



ELSEVIER

Artificial Intelligence in Medicine 703 (2003) 1–26

www.elsevier.com/locate/artmed

**Artificial
Intelligence
in Medicine**

Automated monitoring of medical protocols: a secure and distributed architecture

T. Alsinet^{*}, C. Ansótegui, R. Béjar, C. Fernández, F. Manyà

Department of Computer Science, University de Lleida Jaume II 69, E-25001 Lleida, Spain

Received 15 October 2002; received in revised form 15 October 2002; accepted 18 October 2002

Abstract

The control of the right application of medical protocols is a key issue in hospital environments. For the automated monitoring of medical protocols, we need a domain-independent language for their representation and a fully, or semi, autonomous system that understands the protocols and supervises their application. In this paper we describe a specification language and a multi-agent system architecture for monitoring medical protocols. We model medical services in hospital environments as specialized domain agents and interpret a medical protocol as a negotiation process between agents. A medical service can be involved in multiple medical protocols, and so specialized domain agents are independent of negotiation processes and autonomous system agents perform monitoring tasks. We present the detailed architecture of the system agents and of an important domain agent, the database broker agent, that is responsible of obtaining relevant information about the clinical history of patients. We also describe how we tackle the problems of privacy, integrity and authentication during the process of exchanging information between agents.

© 2003 Published by Elsevier Science B.V.

Keywords: Medical protocols; Specification and representation; Multi-agent system architecture; Distributed monitoring; Secure communication layer

1. Introduction

In recent years physicians have made an important effort to formalize medical protocols (MP) for the diagnosis and treatment of diseases. MPs published by medical associations, like the American Heart Association, the European Resuscitation Council and the Australian Medical Association, can be seen as the consensus opinion of experts based

^{*} Corresponding author. Tel.: +34-973-702-734; fax: +34-973-702-702

E-mail addresses: tracy@eup.udl.es (T. Alsinet), carlos@eup.udl.es (C. Ansótegui), ramon@eup.udl.es (R. Béjar), cesar@eup.udl.es (C. Fernández), felip@eup.udl.es (F. Manyà).

33 on systematic review of the scientific literature. This effort is devoted to improve both
34 decision making in healthcare processes and optimal limited resource management.

35 According to Gordon [14] MPs can be defined as systematically developed statements to
36 assist practitioner and patient decisions about appropriate health care for specific clinical
37 circumstances. To help physicians on the right application of MPs in real hospital
38 environments, the Artificial Intelligence community has considered the challenge of
39 automatic management of MPs. For achieving this objective two different problems have
40 been considered.

41 On the one hand, it is necessary to tackle the automated application and the monitoring
42 problem. That is, to develop a general-purpose system devoted both to coordinate
43 practitioners and patients during a real-time application of a MP and to warn about
44 forbidden actions and decisions. On the other hand, for performing the monitoring task, it is
45 necessary to have a suitable and well-defined representation of MPs. For this purpose we
46 propose a domain-independent language for their representation and so, a user-friendly
47 specification tool for allowing physicians to model MPs and translating the MP models into
48 the representation language.

49 There exist different approaches and tools in the literature for representing and modeling
50 MPs, but not for the automated application and the monitoring problem when considering
51 real hospital environments. In such environments, one has to develop a system able to run
52 in a real-time distributed frame and, because of the nature of the information exchanged in
53 the distributed system, one has also to provide a secure communication layer. Some well-
54 known approaches for modeling MPs are those of DGPM [15], GLIF [26,28]. It should be
55 noticed that those approaches, except the extension of GLIF [28], are not necessarily well-
56 suited to support the automated application and monitoring of the MPs. As to the
57 automated application of the MPs, Boxwala et al. [5] have defined a system based on
58 the GLIF representation model. Their approach is aimed at clinical decision-support and
59 medical education. Other approaches facing this aim are proforma [10,9] and the projects
60 Asgaard [37] and ISAAC [4]. The PRODIGY system [30] is a guideline-based decision-
61 support system designed to assist general practitioners in choosing the appropriate
62 therapeutic action for their patients and has been extended [21] for encoding clinical
63 guidelines for managing patients. Although these approaches help physicians to test and
64 evaluate knowledge of MPs recommendations, they do not consider the integration of a MP
65 agreed by consensus into the medical practice, i.e. the automatic application and
66 monitoring of the MP for a given patient in a hospital environment with existing medical
67 records. In this framework, Terenziani et al. [42] and Alsinet et al. [2,3] have defined
68 modular systems for representing and applying MPs. Both systems [42,2] consist of a user-
69 friendly environment for representing and acquiring MPs involving temporal issues as the
70 treatment of composite, concurrent and repeated actions. However, [42] defines a
71 centralized system architecture (i.e. to be executed in a local way) while [3] defines a
72 multi-agent system architecture for distributed hospital environments.

73 The objective of this paper is to define a system for the assistance and supervision of the
74 real-time application of MPs in distributed hospital environments with computerized-
75 based medical records. A MP specifies possible sequences of composite, concurrent and
76 repeated actions that could be performed when a patient had a particular pathology and,
77 depending on clinical results, establishes which of them are forbidden. In hospital

78 environments, every staff person plays one or more roles. A role specifies a particular
79 service; for instance, infirmary, surgery, radiology or hospital management. So, a MP
80 specifies possible interactions between medical services in front of a particular pathology
81 and forbids some medical decisions depending on the medical history and evolution of the
82 patient. Our system can be used for managing both suggested actions and constrained
83 actions. To be precise, suggestions are modeled as multiple possible actions and constraints
84 as forbidden actions. So, it can be used for managing both medical guidelines and MPs.
85 From now on, we refer to the more general term, i.e. MP.

86 To define a system for monitoring MPs, we make an abstraction of the real world, i.e.
87 hospital environment, in terms of a multi-agent system (MAS) [41,45]. The main idea is to
88 model medical services in hospitals as specialized domain agents and interactions between
89 different services as electronic communication processes. From this point of view, a MP
90 describes a negotiation process between multiple specialized domain agents for treating a
91 particular pathology and specifies behavior rules depending on specific symptoms.

92 To tackle this problem, the first step is to develop a tool and a representation model for
93 hospital environment specification in terms of MASs which, in particular, allow us to
94 model medical services as specialized domain agents and MPs as restricted negotiation
95 processes. In this context, we have implemented a user-friendly environment (called
96 JAFDIS) [2] to graphically represent negotiation processes in MASs based on the notion of
97 dialogical institutions. The second step is to develop a suitable MAS architecture for real-
98 time monitoring MPs in distributed hospital environments. In distributed MASs, agent
99 interactions are performed as electronic communications through a communication layer.
100 Because of the private nature of most of the information exchanged in a hospital
101 environment, the communication layer must provide privacy, integrity and authentication
102 during the process of exchanging information between agents. Therefore, the third step is
103 to define a robust communication interface. Finally, in order to make our system fully
104 compliant with the existing medical records, we develop a database broker agent for
105 automatically recovering the medical history and evolution of patients.

106 The remainder of the paper is organized as follows. [Section 2](#) formalizes the concept of
107 dialogical institution and defines its two basic components, dialogical framework and
108 performative structure. [Section 3](#) shows how a MP can be specified using this formalism.
109 [Section 4](#) describes how a dialogical institution can be modeled using JAFDIS. [Section 5](#)
110 presents the architecture of our MAS for monitoring MPs. [Section 6](#) deals with all
111 necessary security and communication aspects to ensure the privacy of the information
112 exchanged between agents. [Sections 7 and 8](#) describe the functional requirements of two
113 main components of the MAS architecture: the server and supervisor agents of MPs.
114 [Section 9](#) shows the mediator functionalities of the interface between the MAS architecture
115 and the medical services. [Section 10](#) describes the functionalities of the database broker
116 agent. Finally, in [Section 11](#), we summarize our main results and point out our future work.

117 2. Dialogical institutions

118 In this section we provide a new definition of dialogical institution, based on [2], which
119 is *formed by a dialogical framework* and a *performative structure*. This definition allow us

120 to model a traditional institution (say a hospital) and its set of institutional protocols (say
 121 MPs) by detailing what kind of agents exist and what type of messages or illocutions are
 122 interchanged between agents. The dialogical framework describes the general context of
 123 the institution and the performative structure describes the set of protocols. In [29,32,7,35]
 124 the authors propose different approaches for modeling some particular institutions.

125 In what concerns our work, we think of agent-based institutions as a computational
 126 realization of a traditional institution, which intuitively amounts to a set of clearly
 127 established conventions that somehow restrict participating agent's interactions. Therefore,
 128 a dialogical institution can be understood as a set of agents and a set of protocols so that, for
 129 every protocol, there exists a well-defined behavior of the involved agents. An agent is an
 130 entity that plays a specific institution role. Moreover, an agent can play different roles in
 131 different protocols and in different scenes of a same protocol. A protocol is a non-empty set
 132 of well-defined scenes.

133 In a dialogical institution agents are “dialogical entities” that interact with other agents
 134 in a multi-agent context through illocutions. An agent is an entity that expresses illocutions,
 135 reacts to illocutions addressed to it and, furthermore, only the emission and reception of
 136 illocutions constitute its observable behavior.

137

138 **Definition 1 (Dialogical institution).** A dialogical institution is a pair $\langle \mathcal{DF}, \mathcal{PS} \rangle$, where

- \mathcal{DF} is a dialogical framework and
- \mathcal{PS} is a performative structure that defines the set of protocols of the institution in
 141 which the system agents are involved.
 142

143 2.1. Dialogical framework

144 The dialogical framework describes the context of a dialogical institution. On the one
 145 hand, it specifies the basic roles and the agents that take place inside the institution. On the
 146 other hand, it formalizes the communication language of the system agents and the
 147 metalanguage used to model the rules of behavior inside a protocol.
 148

149 **Definition 2 (Dialogical framework).** A dialogical framework is a tuple of the form
 150 $\langle \text{Roles, Agents, } L, \text{CL, TL, ML, } T \rangle$, where

- Roles is a set of different kinds of participants associated with the institution,
- Agents denotes a set of role occurrences,
- L is an object language and describes the syntactic rules for generating the system
 154 identifiers, the agents names, and the illocutions content,
- CL is a communication language and defines the set of illocutions that can be exchanged
 156 between the system agents,
- TL is the transition language that allows one to define the movements of agents between
 158 scenes within a specific protocol,
- ML is the metalanguage used for restricting the agents behavior within a specific
 160 protocol, and
- T is a model of time.

163 At the specification phase, one selects a specific model of time, i.e. a specific time
 164 format, by means of a unique identifier. The monitoring system must provide a suitable
 165 implementation of every model of time. For example, a model of time can be the usual time
 166 format indicating seconds, minutes and hours referring to the Greenwich Meridian Time
 167 together with the set of arithmetic and relational operators. Another possible model of
 168 time¹ is the one based on the notion of days, weeks and months, where the day is the basic
 unit of time.

170 **Definition 3 (Communication language).** The syntax of an illocution of CL is as follows:
 171 illocution_id (agent₁, agent₂, φ , t), where

- illocution_id $\in L$ is the illocution identifier,
- agent₁ \in Agents is the sender of the illocution and agent₂ \in Agents is the recipient of
 174 the illocution,
- $\varphi \in L$ is the content of the illocution, and
- t is the illocution time expressed according to the time model T .

177
 178

179 **Definition 4 (Transition language).** The syntax of a transition of TL is as follows: move
 180 (kind_of_movement, $\{a_1, \dots, a_n\}$, t), where

- move $\in L$ is the label of the transition,
- kind_of_movement $\in L$ is either the label IN $\in L$ or the label OUT $\in L$ and allows one to
 183 model arrivals/departures to/from the scenes within a specific protocol,
- $\{a_1, \dots, a_n\} \subseteq$ Agents describes the set of agents that should move between the scenes
 185 of a specific protocol, and
- t is the transition time expressed according to the time model T .

187
 188

189 At the specification phase, the illocution time and the transition time denote time
 190 variables which are instantiated by the monitoring system when executing a protocol of the
 191 institution. The time variables allow us to specify forbidden actions based on temporal
 constraints and timeout execution errors.

192 2.2. Performative structure

193 The performative structure is a set of protocols. A protocol is a main scene and a set of
 194 interdependent sub-scenes and states that allows us to establish causal and temporal co-
 195 dependencies among them.

196 Intuitively, a protocol of a dialogical institution can be formally specified by means of an
 197 extension of a Petri net. A Petri net [6] has four basic components: places, transitions, arcs,
 198 and tokens. A place represents a state in which the system may be. Tokens indicate the
 199 current states of the system. A transition has zero or more input arcs, coming from its input
 200 places, and zero or more output arcs, going to its output places. A transition is enabled if
 201 there is at least one input token in each of its input places. Any enabled transition is
 202 executed by removing one token from each input place and depositing a token in each

¹ The example of Section 3 uses this model of time.

203 output place. The execution of MPs in our system follows the same basic rules as in the
204 execution of Petri nets.

205 Translating Petri nets into our setting, tokens are associated with agents. Petri nets
206 states are associated with scenes and states of our protocols. We identify three different
207 types of states: atomic, synchronism and parallel states. Petri nets arcs between states are
208 associated with illocutions and agent movements. Specifically, each scene and each state
209 is defined as a set of active agents, where each agent plays an institutional role. Each pair
210 of agents that exchange an illocution are subject to a common atomic interaction
211 represented as an illocution arc between two states. The movement arcs are reserved
212 for representing the arrival/departure of agents to/from a scene. This means that move-
213 ment arcs are only allowed between scenes and synchronism or parallel states. A
214 synchronism state models a synchronized arrival of some agents into the current scene
215 from a different scene. That is, the current scene is suspended until the arrival movements
216 of agents from the other scene are executed. A parallel state models a parallel departure of
217 some agents from the current scene to a different scene. That is, the current scene and the
218 new scene are concurrently executed. Finally, to formalize the interaction between the
219 dialogical institution and the outside we assume the existence of an external scene. The
220 arrival of external agents into a scene is modeled by a synchronism node and the departure
221 is modeled by a parallel node.

222
223 **Definition 5 (Performative structure).** A performative structure of a dialogical institu-
224 tion is a set of protocols. A protocol is a tuple of the form $\langle \text{id_protocol}, S, s_0 \rangle$, where

- $\text{id_protocol} \in L$ is the protocol identifier,
- S is the non-empty set of scenes involved with the protocol,
- $s_0 \in S$ is the initial (main) scene of the protocol. The execution of a protocol starts by
228 executing the initial scene.

229

230

231 **Definition 6 (Scene).** A scene is a tuple of the form

232

233

234

$$\langle \text{id_scene}, \{a_1, \dots, a_n\}, \mathcal{SDG}, \mathcal{BR} \rangle,$$

235

where

- $\text{id_scene} \in L$ is the scene identifier,
- $\{a_1, \dots, a_n\} \subseteq \text{Agents}$ describes the non-empty set of initial agents (i.e. the set of agents
238 that must be active when starting the execution of the scene),
- \mathcal{SDG} is the scene dependence graph,
- \mathcal{BR} is a set of rules that restrict the behavior of agents within the scene. Every behavior
241 rule is defined over the metalanguage ML and has the following general syntax:

241

$$\text{if } C \text{ then } \{ \neg \delta_1, \dots, \neg \delta_n \},$$

244

where C is a Boolean expression built over the language $L \cup \text{CL} \cup \text{TL}$, and δ_i ,
247 $i \in \{1, \dots, n\}$ with $n \geq 1$, is a transition of the scene dependence graph.

247

248

249

All the system restrictions, as the rules of behavior and the initial set of agents, are checked
during the monitoring phase when executing a protocol.

250 **Definition 7 (Scene dependence graph).** A scene dependence graph is a tuple of the form
 251 $\langle W, \delta, w_0, W_T \rangle$, where

- 253 • W is a set of nodes formed by a set of scenes, a set of states and the designated external scene,
- δ is the set of transitions of the dependence graph,
- $w_0 \in W$ is the initial node of the scene dependence graph, and
- 257 • $W_T \subseteq W$ is the non-empty set of final nodes of the scene dependence graph.

258 The execution of a scene starts at the initial node and finishes at a final node.

259
 260 **Definition 8 (State).** A state is a tuple of the form $\langle \text{id_state}, \text{kind_of_state}, \text{timeout} \rangle$, where

- 263 • $\text{id_state} \in L$ is the state identifier,
- 264 • $\text{kind_of_state} \in L$ is either the label $\text{ATOMIC} \in L$, or the label $\text{SYNCHRONISM} \in L$, or the label $\text{PARALLEL} \in L$ (i.e. a state is either an atomic state, or a synchronism state, or a parallel state),
- 266 • timeout is a time constant expressed according to the time model T of the dialogical institution.

267
 268 The initial node of a scene dependence graph is an atomic state. A scene cannot be used
 269 as a final node of a scene dependence graph. When monitoring a protocol, its execution
 270 starts in the initial scene at its initial state. Moreover, the initial set of agents of the scene is
 271 the set of active agents of the initial state. Therefore, the set of active agents for each other
 272 state of the scene is dynamically computed from the initial set and the execution trace.
 273 Finally, the timeout of a state specifies the maximum delay allowed for leaving the state
 274 during the execution phase (i.e. the maximum delay allowed for performing a valid
 275 transition). The time model of the dialogical institution must contain a special time
 276 constant for specifying that there is no timeout constraint.

277 In the following definition S denotes a finite set of scenes, A denotes a finite set of atomic
 278 states, E denotes the designated external scene, Sync denotes a finite set of synchronism
 279 states, and Pll denotes a finite set of parallel states.

280
 281 **Definition 9 (Dependence graph transitions).** A set of transitions δ of a scene dependen-
 282 ce graph is a pair of mappings $\langle \rho_1, \rho_2 \rangle$ between the graph nodes. The mapping ρ_1
 283 specifies the allowed transitions inside the current scene. The mapping ρ_2 specifies the
 284 allowed transitions between the current scene and other scenes of the protocol.

285 The mapping ρ_1 is formalized as follows:

$$286 \rho_1 : (A \cup \text{Pll} \cup \text{Sync}) \times \text{CL} \rightarrow A$$

288 The mapping ρ_2 is formalized as follows:

$$290 \rho_2 : (A \cup \text{Pll} \cup \text{Sync}) \times (S \cup E) \times \text{TL} \rightarrow (\text{Pll} \cup \text{Sync}),$$

292 The mapping ρ_2 must also satisfy the following constraints:

- 293 • If $\rho_2(a, s, \text{transition}) \in \text{Pll}$, where $a \in A \cup \text{Pll} \cup \text{Sync}$, $s \in S \cup E$ and $\text{transition} \in \text{TL}$,
 295 then transition is a departure movement (i.e. an output transition) and

- If $\rho_2(a, s, \text{transition}) \in \text{Sync}$, where $a \in A \cup \text{PII} \cup \text{Sync}$, $s \in S \cup E$ and $\text{transition} \in \text{TL}$, then transition is an arrival movement (i.e. an input transition).

297
298
299
300

The set of constraints imposed by the mappings ρ_1 and ρ_2 are checked during the protocol specification.

301
302
303
304
305
306
307
308
309
310
311

Notice that neither mapping ρ_1 nor mapping ρ_2 restrict the number of transitions for a given state. Therefore, it is possible to reach a state from different source states and from a state one can reach multiple target states. When a state acts as the source of multiple transitions, the state has to be understood as a decision state of the scene. At the specification phase, multiple transitions are constrained by the behavior rules. However, it is during the execution phase of a protocol that one can know which of these multiple transitions are allowed by the behavior rules. Finally, during the execution phase, if more than one transition is allowed, the active agents at the decision state have to select one of them. Moreover, if one maintains a history about agents decisions for protocol executions, this database could be used as knowledge source for refining the protocol specification.

312 3. Specification of a medical protocol for controlling hypertension

313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336

For illustrative purposes, we next describe an example of the use of our modeling language to specify a real medical protocol for detecting and controlling hypertension.² The protocol involves the participation of a patient, a physician and a nurse. After an initial diagnostic phase of the severity of the high pressure condition, the patient must start a controlling phase for checking the high pressure evolution. Since the controlling phase is general enough to exemplify the use of all the elements of our specification language, we concentrate on this phase. Fig. 1 shows an algorithmic description of the controlling phase of the hypertension protocol. Pressure is determined by means of two measures: siastolic blood pressure (SBP) and diastolic blood pressure (DBP). Good pressure levels are assumed when the patient SBP and DBP are less than 140 and 90, respectively. Acceptable pressure levels are assumed when the patient SBP and DBP are in the integer and semi-open interval $[140, 160)$ and $[90, 95)$, respectively. Dangerous pressure levels are assumed when the patient SBP and DBP are greater than 160 and 95, respectively. For good and acceptable pressure levels, patients must follow a periodic basic test, every 3 months, for checking their pressure. For dangerous pressure levels, patients must follow a periodic and exhaustive test, every 2 weeks, for checking their pressure and possible side effects. Usually, the basic test is performed by the nurse and the exhaustive test is performed by the physician. Once the patient has successfully passed three consecutive basic tests, he must follow an annual checkup program.

In our modeling language, the controlling phase of the hypertension protocol is specified by means of an scene, called the controlling scene. Fig. 2 shows the scene dependence graph (SDG) of the controlling scene. Tables 1 and 2 show the illocutions and transitions of the SDG, respectively.

²This protocol has been provided by a clinic of Tarragona, Spain

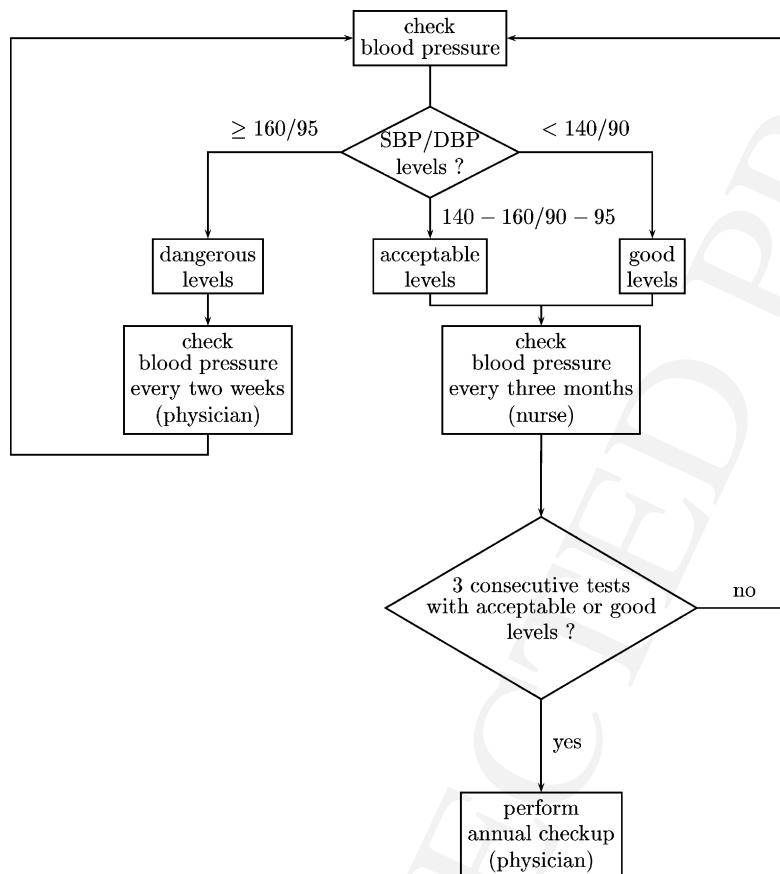


Fig. 1. Algorithm of the controlling phase for the hypertension protocol.

337 The agents involved in the controlling scene are a patient, a physician and a nurse, and
 338 the initial agent of the scene is the nurse. The first action that takes place in the scene is the
 339 synchronization between the nurse and the patient (transition T_1). The nurse is located in
 340 the initial state S_0 and the patient comes from the diagnostic scene (scene D at the SDG). At
 341 the synchronization state $Sync_1$, the nurse checks the blood pressure of the patient.

Table 1
 Illocutions of the SDG of the controlling scene

Identifiers	Illocution
$i_1, i_4, i_9, i_{14}, i_{19}$	Inform (nurse, patient, "dangerous levels")
$i_2, i_5, i_{10}, i_{15}, i_{20}$	Inform (nurse, patient, "acceptable levels")
$i_3, i_6, i_{11}, i_{16}, i_{21}$	Inform (nurse, patient, "good levels")
$i_7, i_8, i_{12}, i_{13}, i_{17}, i_{18}$	Request (nurse, patient, "perform basic test")
i_{22}	Request (nurse, patient, "perform exhaustive test")




Fig. 2. SDG of the controlling scene for the hypertension protocol.

342 Depending on the SBP and DBP levels, the nurse emits either illocution i_1 , or i_2 , or i_3 . The
 343 right illocution is determined by means of behavior rules.

344 On the one hand, when the patient SBP and DBP levels are dangerous, the nurse should
 345 emit illocution i_1 and then illocution i_{22} with a maximum delay of 2 weeks. This fact is
 346 specified by means of a finite timeout attached to state S_1 . Since the exhaustive test must be
 347 performed by the physician, he must come in the scene (transition T_2). After performing
 348 this test, as in state $Sync_1$ and depending on the SBP and DBP levels, the nurse emits either
 349 illocution i_4 , or i_5 , or i_6 , but, in any case, the physician abandons the scene and goes to the
 350 external scene (scene E at the SDG) by means of transition T_3 .

351 On the other hand, when the SBP and DBP levels of the patient are acceptable, the nurse
 352 should emit illocutions i_2 and then i_7 . If levels are good, the nurse should emit illocutions i_3
 353 and then i_8 . In both cases, illocutions must have a maximum delay of 3 months. This fact is
 354 specified by means of a finite timeout attached to states S_2 and S_3 . After performing the
 355 basic test, as in state $Sync_1$ and depending on the SBP and DBP levels, the nurse emits
 356 either illocution i_9 , or i_{10} , or i_{11} . If the patient successfully passes³ two more consecutive
 357 basic tests, the system reaches either state S_{11} , if the SBP and DBP levels for the last basic
 358 test are acceptable, or state S_{12} , if the levels are good. In both cases, the patient must be sent
 359 to the annual checkup scene (scene AR at the SDG) by means of transactions T_4 and T_5 .

360 Finally, the behavior rules attached with state $Sync_1$ are the following ones:

- If $(SBP \geq 160 \text{ and } DBP \geq 95)$ then $\{\neg i_2, \neg i_3\}$.
- If $(SBP \in [140, 160) \text{ and } DBP \in [90, 95))$ then $\{\neg i_1, \neg i_3\}$.
- If $(SBP < 140 \text{ and } DBP < 90)$ then $\{\neg i_1, \neg i_2\}$.

364
 365

The behavior rules attached with states Pll_1 , S_4 , S_7 and S_{10} are defined in an analogous way.

³ A patient successfully passes a basic test whenever the SBP and DBP levels are either acceptable or good.

Table 2
Transitions of the SDG of the controlling scene

Identifiers	Transition
T_1	Move (IN, {patient})
T_2	Move (IN, {physician})
T_3	Move (OUT, {physician})
T_4, T_5	Move (OUT, {patient})

366 4. Modeling a dialogical institution with JAFDIS

367 JAFDIS [2] offers the possibility of working either in graphical mode or in text mode.
 368 That is, the scene dependence graph can be specified either by loading a textual description
 369 in a formal language or by drawing the graph using a graph editor. The graphical mode is
 370 supported by a full-featured graphical user-interface which incorporates MDI technology
 371 (multiple document interface) in order to display the context of all the scenes that configure
 372 a protocol in a common frame.

373 The definition of a new dialogical institution can be performed in two different ways:
 374 static and dynamic. The static way refers to the fact of explicitly specifying the dialogical
 375 framework by means of menu commands before specifying the performative structure. The
 376 dynamic way refers to the fact of implicitly specifying the dialogical framework while
 377 specifying the performative structure.

378 JAFDIS provides a unified frame for specifying all the components of a dialogical
 379 institution. The frame is divided into four areas (see Fig. 3):

- (1) The *menu* and the *toolbar* area.
- (2) The *tabPane* area with two options: S (Scenes) and P (Paint).
 - o The S option shows the hierarchy structure of scenes that make up a protocol.
 - o The P option shows a drawing tool for graphically designing the scenes dependence
 384 graphs.
- (3) The *WorkSpace* area shows each scene dependence graph in a different window.
- (4) The *Debugger* area shows errors generated when checking the consistency of the
 387 scenes dependence graphs.
 388

389 The dialogical framework describes the basic components of the dialogical institution.
 390 JAFDIS provides different dialog boxes for graphically specifying these components.

391 For illustrative purposes, we present the JAFDIS dialog boxes for agents management.
 392 Fig. 4 shows the dialog box for specifying the set of agents. The *Properties* option allows us
 393 to specify their main features (see Fig. 5) identification, description and role. As an agent
 394 can play different roles in the institution and a role can be associated with different agents,
 395 JAFDIS provides a dialog box (see Fig. 6) for specifying the relations between roles and
 396 agents.

397 The performative structure is the set of protocols of the dialogical institution. When
 398 specifying a new protocol the initial scene is created by default. Every scene has a graphical
 399 representation of the scene dependence graph. The basic components of a scene dependen-
 400 ce graph are nodes and transitions. Nodes are atomic states, synchronism states, parallel

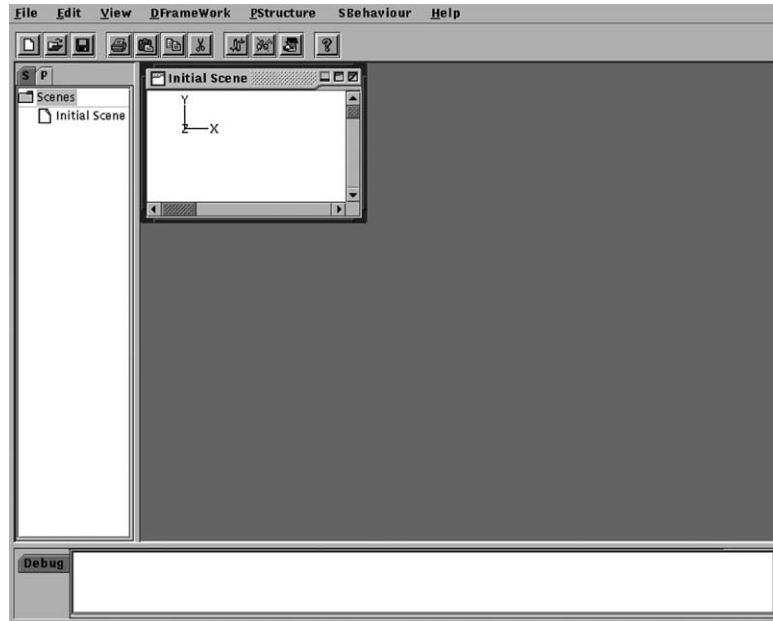


Fig. 3. JAFDIS frame distribution.

401 states and scenes. Transitions are illocutions and movements. Figs. 7 and 8 show the dialog
 402 boxes for specifying nodes and illocutions, respectively. When specifying nodes one can
 403 establish the behavior rules (*R. Behavior* option) that restrict the available transition during
 404 the execution phase and the timeout parameter (*Constraints* option). When specifying
 405 transitions one can establish, during the execution phase, whether the time in which the
 406 transition is performed must be recorded (*mtime* option).

407 Every scene dependence graph must satisfy the following constraints: the scene
 408 dependence graph must be connected, must have at least one initial and one final node,
 409 and every node must be accessible and useful. A node is accessible when there is at least a



Fig. 4. Agents dialog box.



Fig. 5. Agent properties dialog box.




Fig. 6. Agent roles dialog box.

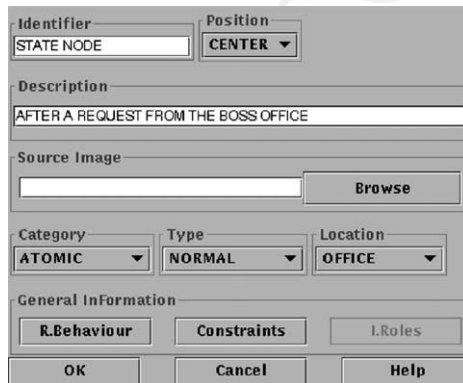


Fig. 7. Nodes properties dialog box.




Fig. 8. Illocution properties dialog box.

410 path from the initial node to this node, and a node is useful when there is at least a path from
411 the node to a final node. A depth-first search algorithm is used to check if the scene
412 dependence graph satisfies these constraints.

413 5. Multi-agent system architecture

414 Nowadays, MASs play a central role in the field of distributed artificial intelligence.
415 MAS applications go from electronic commerce [31,25], aircraft maintenance [38] and
416 business process management [17,19] to e_mail filtering [23]. In the following, an agent
417 has to be understood as a software component that exhibits some aspects of intelligent
418 human behavior [44] and that interacts with its environment. In what concerns our work,
419 we identify three key issues in the design and construction of MASs: negotiation models
420 between autonomous software agents [8,27,45,18,20], agent architectures [33,34,12] and
421 multi-agent architectures [41].

422 In Section 2 we have defined a representation model for specifying hospital institutions in
423 terms of MASs which, in particular, allows us to model medical services as specialized
424 domain agents (SDAs) and MPs as restricted negotiation processes. For identifying the
425 suitable MAS architecture for monitoring MPs we have considered two different approaches:

- 427 (1) To provide SDAs with domain and monitoring knowledge [16]. This approach is based
428 on a layered architecture in which every layer should manage knowledge of a different
429 nature. For instance, the base layer could contain the specific domain knowledge
430 (medical specialty); the middle layer, the specification of MPs related to the specialty;
431 and the upper layer, the monitoring and decision control.
- 432 (2) To distribute the domain and monitoring knowledge [41] among system agents. On
433 the one hand, SDAs could manage the specific domain knowledge and the decision
434 control, i.e. at least a two layer agent architecture. On the other hand, autonomous
435 system agents (ASAs) could control monitoring tasks and manage MP specifications.

436 As a MP can involve multiple medical services, the first approach should keep a same
437 MP specification in multiple SDAs. Moreover, to check forbidden medical decisions could
438 be necessary to know the result of all SDA interactions. So, we can easily see that this
439 approach should maintain a great amount of duplicated information and SDAs should work
440 in a synchronized way, i.e. at every monitoring step all SDAs should know the same
441 information about previous agent interactions. The distributed knowledge approach avoids
442 these problems and allows SDAs to concurrently execute multiple MPs by assigning a
443 different ASA to every monitoring process.

444 In the rest of the paper we make a step towards the formalization of a MAS architecture,
445 based on the distributed knowledge approach, for monitoring MPs by (i) defining a secure
446 communication interface and a certification agent; (ii) providing the system with a server
447 agent and supervisor agents of MPs; and (iii) introducing mediator agents, called
448 interagents, that act as a bridge between SDAs and the rest of the system. SDAs can
449 be of different kinds, i.e. from autonomous software agents to human agents. To support
450 human agents, interagents offer a user-friendly interface accessible from any point of the
451 Internet using a browser.

452 We address the problem of monitoring MPs using a distributed knowledge MAS that
453 allows us to define a distributed environment with secure and robust interactions. Our
454 proposal is based on a computational system composed of several SDAs—one for every
455 medical service in hospital environments—and a set of ASAs that will offer specific
456 services for solving the MPs monitoring problem.

457 Before describing the architecture of the MAS, we identify the fundamental services that
458 the system should provide to agents:⁴

459 *Dynamic location:* The system must be accessible from any point of the Internet (or of
460 the hospital intranet) and allowing agents to interact independently of their physical
461 location (the network address). Thus, we must provide the system with distributed location
462 tables which will be dynamically maintained.

463 *Information confidentiality:* Agent interactions can involve private information about
464 patients. For instance, urine and blood tests, X-rays or biopsy results. Moreover, we have
465 seen that the communication channel of the system is the Internet, which does not
466 guarantee privacy and integrity. Therefore, in order to ensure that only the addressee
467 agent of a message can interpret the content, all messages will be sent encrypted.

468 *Information integrity:* During the communication processes between agents, the system
469 must be able to detect a third party message manipulation. Information manipulation
470 includes such things as insertion, deletion and substitution.

471 *Agents authentication:* In order to ensure the identity of the sender agent during the
472 communication processes between agents, all messages will be digitally signed. Message
473 origin authentication implicitly provides information integrity.

474 *Medical protocols monitoring:* A MP describes a negotiation process between multiple
475 SDAs for treating a particular pathology and forbids some SDA actions depending on the
476 result of previous interactions. So, when executing a MP, the system must supervise all
477 SDA interactions in order to register invalid actions and to assist SDA decisions.

478 *Medical protocols delivery:* As a MP can involve multiple SDAs, the system must
479 manage MP specifications and deliver them with a suitable format.

480 These services are provided by a secure communication interface and some ASAs. SDAs
481 interact with ASAs through interagents allowing SDAs concurrently execute multiple MPs.
482 The secure communication interface ensures that (i) messages are sent encrypted and
483 digitally signed, and (ii) received messages are authentic. The remaining services are
484 provided by the following ASAs:

- 486 • The **certification agent (CA)** acts as a certification authority. The first time an agent
487 interacts with the system it must contact with the CA to obtain its certificate. Agents
need their certificates for signing messages.
- **Supervisor agents (SA)** track interactions between SDAs and verify their validity.
- 490 • The **medical protocols server agent (MPSA)** performs two main system services. On
491 the one hand, it distributes the list of available MPs to SDAs depending on their hospital
492 roles. On the other hand, it distributes MP specifications to SAs.

493 We group agent interactions in two phases. The first one (Fig. 9) includes agents
494 certification, agent addresses distribution and MPs delivery. The second one (Fig. 10)

⁴From now on, by an agent we understand a SDA or an ASA.

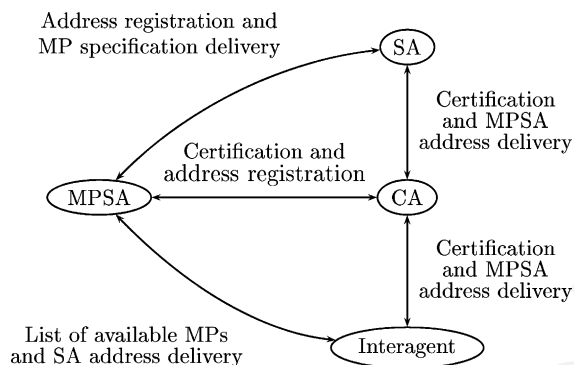


Fig. 9. Certification and information delivery phase.

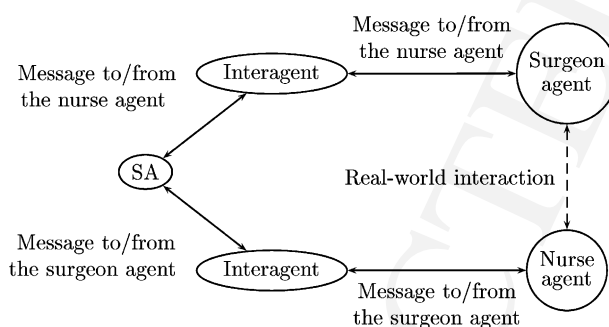


Fig. 10. Monitoring phase.

495 corresponds to the MPs monitoring process. During the first phase, the central component
 496 of the system is the CA. All agents must contact with the CA to obtain their certificates. The
 497 first agent that must contact with the CA is the MPSA. The CA stores the MPSA address. A
 498 SA contacts with the CA for obtaining the MPSA address. At this point the SA provides its
 499 address to the MPSA. This strategy allows us to dynamically distribute the system
 500 monitoring load depending on the number of active SAs. A SDA contacts, through an
 501 interagent, with the CA for obtaining the MPSA address. Next, the interagent contacts with
 502 the MPSA for obtaining the list of available MPs depending on the SDA hospital role. The
 503 SDA decides which MP has to be executed. When all SDAs involved in the MP have
 504 contacted with the MPSA (following the described protocol), the MPSA assigns a SA for
 505 monitoring the MP execution, delivers to interagents the SA address and the monitoring
 506 phase starts. From now on, interagents forward SDAs interactions to the SA.

507 6. Communication interface

508 The private nature of most of the information exchanged in a hospital environment and
 509 the need of authentication led us to use cryptographic tools that we incorporated into a

510 communication interface. Furthermore, these requirements are imposed, in most countries,
511 by governmental laws related to the electronic processing of personal data. An excellent
512 analysis of the types of threats that can impact on a MAS can be found in [43]. Authors also
513 provide architectural solutions to deal with security and trust issues.

514 Our objective, as in most cryptographic systems, has been to provide confidentiality,
515 integrity and authentication to the dialog between agents of our MAS. These features have
516 to be transparent to the user of the communication layer. As it is difficult to find
517 cryptographic classes written in Java, without distribution and license restrictions, we
518 have implemented a complete cryptosystem based on the Secure Socket Layer specifica-
519 tion [11].

520 Our approach is based on a certification authority architecture [22]. This architecture has
521 been used, instead of using a public key server approach, in order to avoid a full storage of
522 the public key of agents and a bottleneck in the system. The potential bottleneck would
523 obey to the fact that an agent, each time it must interact with a new one, must obtain the
524 public key of the new agent through the public key server. The certification authority (in
525 our system, the CA) creates certificates that are used by agents to exchange public keys
526 without contacting a public key server and avoiding congestion scenarios. An agent
527 conveys its public key information to another by transmitting its certificate. Then, the
528 receiver agent verifies that the certificate was created by the CA.

529 The main features of our implementation, outlined in Fig. 11, are the following ones:

- 531 • Communications are symmetrically encrypted by using the triple DES algorithm to
provide confidentiality. The symmetric encryption key is negotiated in every session.
- 533 • The negotiation of the symmetric key is performed by using the public key RSA
534 algorithm. This algorithm allows an authenticated communication between the involved
agents. Therefore, every agent has a public and a private key.
- The public and private keys are generated from the login and password of the agent.
- Two agents transfer one to each other their public key through their certificates. As
537 explained in Section 5, certificates are delivered to agents by the CA.
- For the sake of increasing security, symmetric key negotiation requires a double
539 exchange of nonces.
- A Message Authentication Code (MAC) is used to compute a digest of the plain
541 message. This digest is encrypted with the private key of the sender agent and appended
542 to the plain message in order to provide authentication and integrity.
- The MAC is computed using a Message Digest (MD5) algorithm [36,24].

544 Once explained how the communication among agents achieves certain degree of security,
545 we will explain how certificates are distributed.

547 A certificate for an agent A , say C_A , can be expressed as

$$548 \quad C_A = [ID_A, KU_A, T, E_{KR_{CA}}[D(M)]],$$

550

551 where ID_A is the identification for the agent A , KU_A its public key, T a time stamp and
552 $E_{KR_{CA}}[D(M)]$ is the digest calculated over the three previous components of C_A encrypted
553 with the private key of the CA. This last term provides authentication and integrity when a
554 certificate is requested by an agent, but only if the public key of the CA has been delivered
555 using a secure channel.

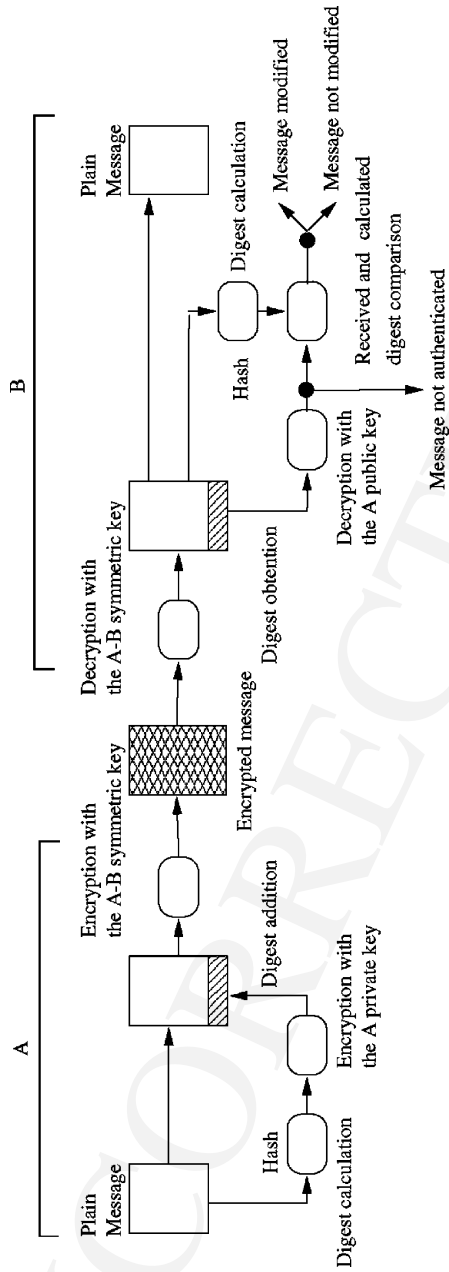


Fig. 11. Block diagram of the secure communication process.

556 The main objective of this implementation has been to establish great similarities with
557 the usual transport level interfaces, i.e. sockets. The use of our interface differs from others
558 in the following points:

- Additional parameters must be specified; for instance, sender keys and certificates.
- Additional error conditions have to be managed; for instance, authentication errors.

561

562 Finally, we describe how an agent should use the secure communication interface. The
563 steps that it must follow are the following ones:

- (1) Open a socket as in conventional transport level interfaces.
- (2) Set up the socket with some parameters: the agent identification, the agent private key and the public key of the CA.
- (3) Execute the handshake protocol. This is the core operation of the secure communication layer. When executing the handshake protocol, certificates of peer agents are exchanged and a session symmetric encryption key is established. Then, two nonces are exchanged following the scheme detailed in [40].
- (4) If the handshake protocol is successfully concluded, the socket is ready to accept messages.

572

573 7. Medical protocols server agent

574 Before specifying the services provided by the medical protocols server agent (MPSA),
575 we summarize the system specification of a MP. A MP is a unique identifier, a non-empty
576 set of scenes and an initial scene.

577 The components of a scene are:

- A unique **identifier**.
- A **scene dependence graph** which specifies a negotiation process between SDAs. States represent points in a negotiation process (SDAs involved in the process and their condition) and transitions between states represent SDA actions. A scene dependence graph contains one initial state and at least one final state. We distinguish two kinds of actions: dialog actions and movement actions. Dialog actions model information exchange between SDAs. Because of the nature of hospital environment, when a SDA is executing a scene of a MP it can be necessary to start a new medical task which will be described by means of a different scene. Therefore, movement actions model arrivals and departures of SDAs to/from scenes.
- A set of **behavior rules** that describes forbidden SDA actions depending on the result of previous actions.
- A non-empty set of **starting agents**. It refers to the set of SDAs that must be present in the system to execute a scene. For instance, to perform a surgical operation at least an anesthesiologist and a surgeon must be present.
- A set of **dynamic joining agents**. It refers to the set of SDAs that could take part in a scene. Their entrance will depend on the monitoring process state. For instance, during a surgical operation, a second surgeon can be requested depending on the evolution of the patient.

580

581

582

583

584

585

586

587

589

591

592

594

595

596

597 The MPSA allows us to define a distributed and scalable MAS in which SDAs can be
598 developed independently of MPs features by:

- 600 • Storing the address of every SA. New SAs can be dynamically created depending on the
601 system monitoring load.
- 602 • Creating new MP instances. A MP is started by a SDA from its set of starting agents (the
603 starting agents of the initial scene). At this point, the MPSA creates a new instance for
604 the MP and blocks it until the rest of starting agents are present. A MP instance contains
605 the MP specification and the monitoring state. When creating a new MP instance the
606 monitoring state is empty.
- 607 • Hanging up MP scene instances. A MP scene can be hanged up by two reasons: a new
608 dynamic joining agent is needed or a SDA involved in the execution of the scene decides
609 to temporally hang it up. In both cases, the MPSA stores the monitoring state provided
610 by the SA and frees it from the monitoring task.
- 611 • Providing SDAs with the list of available MPs depending on SDAs hospital roles. This
612 list will be composed of blocked and hanged up MP scene instances, and all MP that
613 they can start.
- 614 • Providing SAs with MP scene instances. When a blocked or hanged up MP scene
615 instances is ready (all required SDAs are present) the MPSA selects one SA and delivers
616 to it the instance to be monitored. Moreover, as a SA must intercept through interagents
617 all SDA interactions, the MPSA delivers to interagents the SA address and the MP scene
618 instance identifier. This strategy allows us to distribute the system monitoring load
619 between SAs, to move a SDA between MP scenes and to concurrently execute multiple
620 scenes of a same MP or of different MPs.

620 8. Medical protocols supervisor agents

621 The main services offered by medical protocols supervisor agents (SAs) are the
622 supervision and assistance of MP executions. When monitoring MPs, SAs assist SDAs
623 in the sense that they inform SDAs about all the possible and forbidden actions specified.
624 Although SDAs behavior is not restricted by SAs, SAs record SDAs anomalous decisions.
625 In some cases, a MP can specify a timeout for executing critical medical actions; for
626 instance, a biopsy has to be performed before 2 days. Thus, SAs control the elapsed time
627 between SDAs actions.

628 The most important benefit of defining specialized agents for monitoring MPs is that
629 SDAs do not need to be synchronized. To check forbidden and critical SDAs actions, it is
630 necessary to know the monitoring state, which must be the same for every SDA involved in
631 the MP execution. This information is centralized and maintained by SAs in our approach.

632 A more detailed description of SA functionalities is the following one:

- 634 • As we have seen, the number of active SAs is not fixed in the system. Therefore, a SA
635 must be able to supervise the execution of several MP instances. To this end, interagents
636 transparently forward to SAs, using the MP scene instance identifier, interactions
637 between SDAs.

- 638 • For each MP instance, a SA must compute the monitoring state, i.e. for each active scene
639 it must compute the current state by means of the scene dependence graph, the sequence
640 of SDA actions from the initial state to the current state and, if necessary, behavior rules
641 of the current state. To check behavior rules, we have defined a *general data type* that
642 encapsulates data and their operations. The content of a message must belong to a data
643 type that derives from this general data type; for instance, a biopsy result should be
644 provided together with its testing operations. This feature is easily achieved in the object
oriented programming paradigm of Java.
- 646 • After computing a monitoring state, SAs assist SDAs by informing them, through
interagents, about possible and forbidden actions.
- 648 • Finally, SAs must manage arrival and departure actions. On the one hand, when a new
649 dynamic joining agent is required (arrival action), SAs hang up the MP scene instance
650 following the process described in [Section 7](#). On the other hand, a SDA can abandon a
MP scene execution in different ways:
 - 652 (1) A SDA decides to temporally hang up a MP scene execution. The SDA contacts
653 with the MPSA who checks if the MP scene instance is already hanged up. In this
654 case, the SDA is simply added to the list of dynamic joining agents. Otherwise, the
MPSA interacts with the SA for obtaining the MP scene instance.
 - 656 (2) When executing a MP scene, a SDA reaches a final state. The SA informs the
interagent about the exit condition.
 - 658 (3) A SDA can, temporally or definitively, leave the current MP scene to execute a
659 new one. The SA provides the interagent with the identifier of the new scene. The
660 interagent interacts with the MPSA for obtaining the list of blocked and hanged up
661 instances of the MP scene. Finally, the SDA decides to execute one of them or a
new one.

662
663 In the next section we show how these system interaction processes are transparent to
664 SDAs. This allows SDAs to focus on the decision control based on the specific domain
665 knowledge.

666 9. Interagents

667 Two functional requirements of our system are that the monitoring process must be
668 transparent to SDAs and SDAs must be able to concurrently execute multiple MPs.
669 Interagents allow us to achieve such objectives by forwarding SDAs interactions to SAs
670 and assigning a different interagent to each MP scene concurrent execution. In order to
671 execute a MP scene, a SDA interacts with other SDAs through one interagent; this
672 interagent is the same for all interactions in the MP scene.

673 When a SDA is executing a MP scene it can be necessary, temporally or definitively, to
674 interrupt the current scene and then execute a new MP scene. In [Section 7 and 8](#) we have
675 described how the MPSA and SAs interact with interagents to perform entrances to MPs
676 scenes. Now, we describe how interagents perform exits from MPs scenes. As we have
677 seen, the system uniquely identifies every MP scene execution. When a SDA enters to a MP
678 scene execution, the interagent pushes the scene instance identifier and the SA address

679 (provided by the MPSA) onto a stack. When a SDA reaches a final state (the SA informs the
680 interagent about the exit condition), the interagent pops the current MP scene instance
681 information from the stack. After performing a pop operation, if the stack is not empty (i.e.,
682 it contains temporally interrupted executions), the interagent contacts with the MPSA for
683 coming back to the last hanged up instance which is located at the top of the stack.

684 An additional feature of interagents is to provide human SDAs with a user-friendly
685 interface accessible using a web browser. This interface represents the content of SDAs
686 messages in a suitable way and assists human agents in making decisions by showing the
687 possible and forbidden actions in every step when executing a MP scene. The operations
688 encapsulated by the *general data type* include not only operations to check behavior rules,
689 but also operations to graphically represent the involved message information. For
690 instance, biopsy and X-rays results should be shown to the user in a different way.

691 10. Database agent

692 According to [13] the integration of the existing medical records within the guidelines
693 specified by means of a MP is a key issue for automatically executing and monitoring MPs.
694 The effective integration requires consistent information structures and content. However,
695 in real hospital environments we identify two key problems. On the one hand, the
696 information structures and content of the medical records are not uniform among different
697 hospital environments. On the other hand, typically, the MPs specifications involve high
698 level information structures and content while medical records management systems store
699 and retrieve information at a lower level. Therefore, to achieve an effective integration it is
700 necessary to provide a *bridge* between the execution and monitoring system and the
701 medical records management system.

702 In our framework, the bridge is provided as an autonomous SDA called *database broker*
703 *agent*. The functionalities of the database broker agent are the following ones:

- 705 • Provides to the monitoring system a unified frame for exchanging information between
706 the system agents and the database broker agent. To be precise, the database broker
707 agent answers queries by using the already proposed *general data type* that encapsu-
lates data and their operations.
- 709 • Provides a uniform high level language for querying the medical records. For instance,
to obtain the results of biopsies and X-rays.
- 711 • Translates the high level queries into the corresponding set of low level queries for the
712 medical records management system. For instance, the results of biopsies and X-rays
713 could be composed of different information elements as patient identification data,
714 images, indexes results, These elements can be stored with different data models:
715 object oriented, relational, hierarchical, And within a model, the information
716 elements can be stored with different management systems and information structures.
717 Nowadays, we have developed this functionality for management systems with rela-
718 tional data models and with the SQL interrogation language. To interact with different
719 management systems, the SDA should be extended with new modules that implement
low level access to these systems.

720 11. Conclusions and future work

721 In this paper, we have first defined a formal specification language for modeling
722 dialogical institutions which, in particular, can be used to model hospital environments
723 and their MPs. Second, we have defined a graphical framework (called JAFDIS) developed
724 in Java for specifying dialogical institutions according to the specification language. Third,
725 we have defined a MAS architecture for the assistance and supervision of the execution of
726 MPs in hospital environments. The architecture is based on a knowledge distributed
727 approach that allows us to distribute the load among agents and to define a dynamically
728 scalable system. SDAs model specific medical services and manage the involved domain
729 knowledge, and ASAs manage MP specifications and supervise their execution. When
730 executing a MP, the system assists SDAs by reporting to them about possible and forbidden
731 actions and SDAs perform the decision control. For supervising executions, interactions
732 between SDAs are transparently forwarded by interagents to SAs allowing SDAs to
733 concurrently execute multiple MPs. Messages are automatically digitally signed and
734 encrypted by means of the services provided by the secure communication interface.
735 All system components have been implemented in Java and this allows us to define
736 messages as objects that encapsulate data and processing operations. Finally, we have
737 defined a database broker agent for automatically bridging medical records management
738 systems and the monitoring system.

739 For the time being, we have developed a prototype of the MAS architecture which has
740 been implemented in Java. This prototype consists of the database broker agent, the secure
741 communication interface and all required ASAs, i.e. a certification agent, a MP server
742 agent and MP supervisor agents, and interagents.

743 The prototype of the MAS architecture has been developed using a Java based
744 representation of MPs. That is, each MP is represented by a Java class hierarchy where
745 each class represents a scene of the MP. This representation model is automatically
746 generated by the graphical specification tool JAFDIS. The Java based representation of
747 MPs is a general representation model in the sense that one can easily define a translator
748 from other MPs representation models (for instance, GEM [39], GLIF [26,28] or UMLS
749 [1]) to our Java class hierarchy model.

750 Our ongoing research is directed toward two main directions. On the one hand, the
751 development of a user-friendly interface tool between the medical records administrator
752 and the database broker agent. The objective of this tool is to allow the medical records
753 administrator to dynamically maintain the available set of high level queries. On the other
754 hand, the design and development of medical services as autonomous intelligent software
755 agents. For instance, it would be interesting to provide a suitable implementation of an
756 agent for automatically scheduling the resources of the hospital taking into account
757 different quality criteria.

758 Acknowledgements

759 We would like to thank the anonymous referees for their comments and suggestions.
760 This research was partially supported by project SMASH (TIC96-1038-C04-03) and

761 project LOGFAC (TIC2001-1577-C03-03) funded by the Spanish *Ministerio de Ciencia y*
762 *Tecnología*.

References

- 764 [1] Achour SL, Dojat M, Rieux C, Bierling P, Lepage E. UMLS-based knowledge acquisition tool for rule-
765 based clinical decision-support system development. *J Am Med Inform Assoc* 2001;8:351–60.
- 766 [2] Alsinet, T, Béjar, R, Ansótegui, C, Fernández, C, Manyà F. JAFDIS, a Java framework for dialogical
767 institution specification. Technical Report DIEI-98-RT-2, Universitat de Lleida, 1998. Available at [http://](http://fermat.eup.udl.es/~cesar/recerca/jafdis_ev.ps.gz)
768 fermat.eup.udl.es/~cesar/recerca/jafdis_ev.ps.gz.
- 769 [3] Alsinet T, Béjar, R, Fernández, C, Manyà, F. A multi-agent system architecture for monitoring medical
770 protocols. In: Sierra, C, Gini, M, Rosenschein, JS, editors. Proceedings of the Fourth International
771 Conference on Autonomous Agents, Barcelona, Spain. New York, USA: ACM Press; 2000. p. 499–505.
- 772 [4] Boon W, Nardi R, Bertelli S. Computerized guidelines for preventative care in general practice. *Stud*
773 *Health Technol Inform* 1995;16:45–56.
- 774 [5] Boxwala, AA, Greenes, RA, Deibel, SR. Architecture for a multipurpose guideline execution engine. In:
775 Lorenzi, NM, editor. Proceedings of AMIA Annual Symposium, Greenes, PA, USA. Philadelphia, USA:
776 Hanley & Belfus; 1999. p. 701–5.
- 777 [6] Danthine AAS. Protocol representation with finite-state models. *IEEE Trans Commun* 1980;COM-
778 20:632–43.
- 779 [7] Esteva, M, Rodríguez-Aguilar, JA, Arcos, JLL, Sierra, C, García, P. Institutionalising open multi-agent
780 systems. In: Kraus, S, Nakashima, H, Tambe, M, editors. Proceedings of the Fourth International
781 Conference on Multiagent Systems, Boston, MA, USA. Los Alamitos, CA, USA: IEEE Press; 2000.
782 p. 381–2.
- 783 [8] Faratin P, Sierra C, Jennings NR. Negotiation decision functions for autonomous agents. *Int J Robot*
784 *Autonom Agents* 1998;24(3/4):159–82.
- 785 [9] Fox, J, Das, S. Safe, Sound: artificial intelligence in hazardous situations. Menlo Park, CA, USA: AAAI
786 Press; 2000.
- 787 [10] Fox J, Johns N, Rahmazadeh A. Disseminating medical knowledge: the PROforma approach. *Artif Intel*
788 *Med* 1998;14:157–81.
- 789 [11] Freier, AO, Karlton, P, Kocher, PC. The SSL Protocol, Version 3.0, 1996.
- 790 [12] Georgeff, MP, Lansky, AL. Reactive reasoning and planning. In: Forbus, K, Shrobe, H, editors.
791 Proceedings of the Sixth National Conference on Artificial Intelligence, Seattle, WA, USA. Menlo Park,
792 CA, USA: AAAI Press; 1987. p. 677–82.
- 793 [13] Glowinski A. Integrating guidelines and the clinical record: the role of semantically constrained
794 terminologies. *Stud Health Technol Inform* 1995;16:207–18.
- 795 [14] Gordon, C. Practice guidelines and healthcare telematics: towards an alliance. In: Gordon, C, Christensen,
796 JP, editors. Health telematics for clinical guidelines and protocols. Amsterdam, The Netherlands: IOS
797 Press; 1995. p. 3–15.
- 798 [15] Herbert, SI. Informatics for care protocols and guidelines: towards a european knowledge model. In:
799 Gordon, C, Christensen, JP, editors. Health telematics for clinical guidelines and protocols. Amsterdam,
800 The Netherlands: IOS Press; 1995. p. 27–42.
- 801 [16] Huang, J, Jennings, NR, Fox, J. An agent architecture for distributed medical care. In: Wooldridge, M,
802 Jennings, NR, editors. Proceedings of the ECAI'94 Workshop on Gent Theories, Architectures and
803 Languages, Amsterdam, The Netherlands, LNCS 890. Heidelberg, Germany: Springer; 1994. p. 219–232.
- 804 [17] Jennings, NR, Faratin, P, Johnson, MJ, O'Brien, PO, Wiegand, ME. Using intelligent agents to manage
805 business processes. In: Proceedings of the First International Conference on The Practical Application of
806 Intelligent Agents and Multi-Agent Technology, London, UK, 1996, p. 345–60.
- 807 [18] Jennings NR, Faratin P, Lomuscio AR, Parsons S, Wooldridge M, Sierra C. Automated negotiation:
808 prospects, methods and challenges. *Group Decision Negotiation* 2001;10:199–215.
- 809 [19] Jennings NR, Norman TJ, Faratin P. ADEPT an agent-based approach to business process management.
810 *ACM SIGMOD Record* 1998;27(4)(4):32–9.

- 811 [20] N.R. Jennings, S. Parsons, C. Sierra, P. Faratin, Automated negotiation. In: Proceedings of The Fifth
812 International Conference on The Practical Application of Intelligent Agents and Multi-Agent Technology,
813 Manchester, UK, 2000, p. 23–30.
- 814 [21] Johnson, PD, Tu, S, Booth, N, Sugden, B, Purves, IN. Using scenarios in chronic disease management
815 guidelines for primary care. In: Overhage, JM, editor. Proceedings of AMIA Annual Symposium, Los
816 Angeles, CA, USA. Philadelphia, USA: Hanley & Belfus; 2000. p. 389–93.
- 817 [22] Kohnfelder, L. Towards a practical public-key cryptosystem. Bachelor's Thesis, MIT, 1978.
- 818 [23] Maes P. Agents that reduce work and information overload. *Commun ACM* 1996;37(7):31–40.
- 819 [24] Menezes, A, Oorschot, P, Vanstone, S. Handbook of applied cryptography. London, UK: CRC Press;
820 1997.
- 821 [25] Nwana, HS, Rosenschein, J, Sandholm, T, Sierra, C, Maes, P, Guttman, R. Agent-mediated electronic
822 commerce: issues, challenges and some viewpoints. In: Sycara, KP, Wooldridge, M, editors. Proceedings
823 of the Second International Conference on Autonomous Agents, Minneapolis, USA. New York, USA:
824 ACM Press; 1998. p. 189–96.
- 825 [26] Ohno-Machado L, Gennari JH, Murphy SN, Jain NL, Tu SW, Oliver DE, Pattison-Gordon E, Greenes RA,
826 Shortliffe EH, Barnett GO. The guideline interchange format: a model for representing guidelines. *J Am
827 Med Inform Assoc* 1998;5:357–72.
- 828 [27] Parsons SD, Sierra C, Jennings NR. Agents that reason and negotiate by arguing. *J Logic Comput*
1998;8(3):261–92.
- 830 [28] Peleg, M, Boxwala, AA, Ogunyemi, O, Zeng, Q, Tu, SW, Bernstam, E, Ohno-Machado, L, Shortliffe, EH,
831 Greenes, RA. GLIF3: the evolution of a guideline representation format. In: Overhage, JM, editor.
832 Proceedings of AMIA Annual Symposium, Los Angeles, CA, USA. Philadelphia, USA: Hanley & Belfus;
833 2000. p. 645–9.
- 834 [29] Plaza, E, Noriega, P, Sierra C. The conference center as an agent-mediated institution. In: Proceedings of
835 the First International Workshop on Agents in Community Ware, Paris, France, 1998, p. 73–82.
- 836 [30] Purves, IN, Sugden, B, Booth, N, Sowerby, M. The prodigy project—the interactive development of the
837 release one model. In: N.M. Lorenzi, editor. Proceedings of AMIA Annual Symposium, Greenes, PA,
838 USA. Philadelphia, USA: Hanley & Belfus; 1999. p. 359–63.
- 839 [31] Rodríguez, JA, Noriega, P, Sierra, C, Padget, J. A Java-based electronic auction house. In: Proceedings of
840 the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent
841 Technology, 1997, p. 207–224.
- 842 [32] Rodríguez-Aguilar, JA, Martín, FF, García, P, Noriega, P, Sierra, C. Towards a formal specification of
843 complex social structures in multi-agent systems. In: Padget, JA, editor. Collaboration between human and
844 artificial societies: coordination and agent-based distributed computing, LNAI 1624. Heidelberg,
845 Germany: Springer; 1999. p. 284–300.
- 846 [33] Rosenschein S. Formal theories of knowledge in AI and robotics. *New Generation Comput* 1985;
847 3(4):345–57.
- 848 [34] Rosenschein, S, Kaelbling, L. The synthesis of digital machines with provable epistemic properties. In:
849 Halpern, JY, editor. Proceedings of the Conference on Theoretical Aspects of Reasoning About
850 Knowledge. New York, USA: Morgan Kaufmann; 1986. p. 207–224.
- 851 [35] Sabater, J, Sierra, C. Regret: reputation in gregarious societies. In: Andre, E, Sen, S, Frasson, C, Müller,
852 JP, editors. Proceedings of The Fifth International Conference on Autonomous Agents, Montreal, Canada.
853 New York, USA: ACM Press; 2001. p. 194–5.
- 854 [36] Schneier, B. Applied cryptography: protocols, algorithms, and source code in C. 2nd ed. Hoboken, NJ,
855 USA: Wiley; 1995.
- 856 [37] Shahar, Y, Miksch, S, Johnson, P. The asgaard project: a task-specific framework for the application and
857 critiquing of time-oriented clinical guidelines. *Artif Intel Med* 1998;14:29–51.
- 858 [38] Shehory, O, Sycara, K, Sukthankar, G, Mukherjee, V. Agent aided aircraft maintenance. In: Etzioni, O,
859 Müller, JP, editors. Proceedings of the Third International Conference on Autonomous Agents, Seattle,
860 WA, USA. New York, USA: ACM Press; 1999.
- 861 [39] Shiffman RN, Karras BT, Agrawal A, Chen R, Marengo L, Nath S. GEM: a proposal for a more
862 comprehensive guideline document model using xml. *J Am Med Inform Assoc* 2000;7:488–98.
- 863 [40] Stallings, W. Network and Internetwork Security. New Jersey, USA: Prentice-Hall; 1995.

- 864 [41] Sycara K, Decker K, Pannu A, Williamson M, Zeng D. Distributed intelligent agents. *IEEE Expert* 1996;11(6):36–45.
- 866 [42] Terenziani P, Molino G, Torchio M. A modular approach for representing and executing clinical
867 guidelines. *Artif Intel Med* 2001;23:249–76.
- 868 [43] Wong, HC, Sycara, K. Adding security and trust to multi-agent systems. In: *Proceedings of the Agents'99*
869 *Workshop on Deception, Fraud and Trust in Agent Societies*, Seattle, WA, USA, 1999, p. 149–61.
- 870 [44] Wooldridge, M, Jennings, NR, Agent theories, architectures, and languages: a survey. In: *Wooldridge, M,*
871 *Jennings, NR, editors. In: Proceedings of the ECAI'94 Workshop on Agent Theories, Architectures and*
872 *Languages*, Amsterdam, The Netherlands, LNCS 890. Heidelberg, Germany: Springer; 1994. p. 1–32.
- 873 [45] Wooldridge MJ, Jennings NR. Cooperative problem solving. *J Logic Comput* 1999;9(4):563–92.