

Hypotheses Refinement under Topological Communication Constraints *

Gauvain Bourgne, Gael Hette, Nicolas Maudet, and Suzanne Pinson

LAMSADE, Univ. Paris-Dauphine, France

{bourgne,hette,maudet,pinson}@lamsade.dauphine.fr

ABSTRACT

We investigate the properties of a multiagent system where each (distributed) agent locally perceives its environment. Upon perception of an unexpected event, each agent locally computes its favoured hypothesis and tries to propagate it to other agents, by exchanging hypotheses and supporting arguments (observations). However, we further assume that communication opportunities are severely constrained and change dynamically. In this paper, we mostly investigate the convergence of such systems towards global consistency. We first show that (for a wide class of protocols that we shall define), the communication constraints induced by the topology will not prevent the convergence of the system, at the condition that the system dynamics guarantees that no agent will ever be isolated forever, and that agents have unlimited time for computation and arguments exchange. As this assumption cannot be made in most situations though, we then set up an experimental framework aiming at comparing the relative efficiency and effectiveness of different interaction protocols for hypotheses exchange. We study a critical situation involving a number of agents aiming at escaping from a burning building. The results reported here provide some insights regarding the design of optimal protocol for hypotheses refinement in this context.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Theory, Experimentation

Keywords

Negotiation and argumentation, agent communication languages and protocols

*The primary author of this paper is a student.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

Copyright 2007 IFAAMAS.

1. INTRODUCTION

We consider a multiagent system where each (distributed) agent locally perceives its environment, and we assume that some unexpected event occurs in that system. If each agent computes only locally its favoured hypothesis, it is only natural to assume that agents will seek to coordinate and refine their hypotheses by confronting their observations with other agents. If, in addition, the communication opportunities are severely constrained (for instance, agents can only communicate when they are close enough to some other agent), and dynamically changing (for instance, agents may change their locations), it becomes crucial to carefully design protocols that will allow agents to converge to some desired state of global consistency. In this paper we exhibit some sufficient conditions on the system dynamics and on the protocol/strategy structures that allow to guarantee that property, and we experimentally study some contexts where (some of) these assumptions are relaxed.

While problems of diagnosis are among the venerable classics in the AI tradition, their multiagent counterparts have much more recently attracted some attention. Roos and colleagues [8, 9] in particular study a situation where a number of distributed entities try to come up with a satisfying global diagnosis of the whole system. They show in particular that the number of messages required to establish this global diagnosis is bound to be prohibitive, unless the communication is enhanced with some suitable protocol. However, they do not put any restrictions on agents' communication options, and do not assume either that the system is dynamic.

The benefits of enhancing communication with supporting information to make convergence to a desired global state of a system more efficient has often been put forward in the literature. This is for instance one of the main idea underlying the *argumentation-based negotiation* approach [7], where the desired state is a compromise between agents with conflicting preferences. Many of these works however make the assumption that this approach is beneficial to start with, and study the technical facets of the problem (or instead emphasize other advantages of using argumentation). Notable exceptions are the works of [3, 4, 2, 5], which studied in contexts different from ours the efficiency of argumentation.

The rest of the paper is as follows. Section 2 specifies the basic elements of our model, and Section 3 goes on to presenting the different protocols and strategies used by the agents to exchange hypotheses and observations. We put special attention at clearly emphasizing the conditions on the system dynamics and protocols/strategies that will be exploited in the rest of the paper. Section 4 details one of

the main results of the paper, namely the fact that under the aforementioned conditions, the constraints that we put on the topology will *not* prevent the convergence of the system towards global consistency, at the condition that no agent ever gets completely “lost” forever in the system, and that unlimited time is allowed for computation and argument exchange. While the conditions on protocols and strategies are fairly mild, it is also clear that these system requirements look much more problematic, even frankly unrealistic in critical situations where distributed approaches are precisely advocated. To get a clearer picture of the situation induced when time is a critical factor, we have set up an experimental framework that we introduce and discuss in Section 5. The critical situation involves a number of agents aiming at escaping from a burning building. The results reported here show that the effectiveness of argument exchange crucially depends upon the nature of the building, and provide some insights regarding the design of optimal protocol for hypotheses refinement in this context.

2. BASIC NOTIONS

We start by defining the basic elements of our system.

Environment

Let $\mathcal{O}$ be the (potentially infinite) set of possible *observations*. We assume the sensors of our agents to be perfect, hence the observations to be *certain*. Let $\mathcal{H}$ be the set of hypotheses, uncertain and revisable. Let $Cons(h, O)$ be the *consistency relation*, a binary relation between a hypothesis $h \in \mathcal{H}$ and a set of observations $O \subseteq \mathcal{O}$. In most cases, $Cons$ will refer to classical consistency relation, however, we may overload its meaning and add some additional properties to that relation (in which case we will mention it).

The environment may include some dynamics, and change over the course of time. We define below sequences of time points to deal with it:

DEFINITION 1 (SEQUENCE OF TIME POINTS). A sequence of time points $t_1, t_2, \dots, t_n$ from t is an ordered set of time points $t_1, t_2, \dots, t_n$ such that $t_1 \geq t$ and $\forall i \in [1, n-1], t_{i+1} \geq t_i$.

Agent

We take a *system* populated by n agents $a_1, \dots, a_n$. Each agent is defined as a tuple $\langle \mathcal{F}, O_i, h_i \rangle$, where:

- $\mathcal{F}$, the set of facts, common knowledge to all agents.
- $O_i \in 2^{\mathcal{O}}$, the set of observations made by the agent so far. We assume a perfect memory, hence this set grows monotonically.
- $h_i \in \mathcal{H}$, the favourite hypothesis of the agent.

A key notion governing the formation of hypotheses is that of consistency, defined below:

DEFINITION 2 (CONSISTENCY). We say that:

- An agent is consistent ($Cons(a_i)$) iff $Cons(h_i, O_i)$ (that is, its hypothesis is consistent with its observation set).
- An agent a_i consistent with a partner agent a_j iff $Cons(a_i)$ and $Cons(h_i, O_j)$ (that is, this agent is consistent and its hypothesis can explain the observation set of the other agent).

- Two agents a_i and a_j are mutually consistent ($MCons(a_i, a_j)$) iff $Cons(a_i, a_j)$ and $Cons(a_j, a_i)$.
- A system is consistent iff $\forall (i, j) \in [1, n]^2$ it is the case that $MCons(a_i, a_j)$.

To ensure its consistency, each agent is equipped with an abstract reasoning machinery that we shall call the *explanation function* $\mathcal{E}_h$. This (deterministic) function takes a set of observation and returns a single preferred hypothesis ($2^{\mathcal{O}} \rightarrow \mathcal{H}$). We assume $h = \mathcal{E}_h(O)$ to be consistent with O by definition of $\mathcal{E}_h$, so using this function on its observation set to determine its favourite hypothesis is a sure way for the agent to achieve consistency. Note however that an hypothesis does not *need* to be generated by $\mathcal{E}_h$ to be consistent with an observation set. As a concrete example of such a function, and one of the main inspiration of this work, one can cite the Theorist reasoning system [6] —as long as it is coupled with a filter selecting a *single* preferred theory among the ones initially selected by Theorist.

Note also that h_i may only be modified as a consequence of the application $\mathcal{E}_h$. We refer to this as the *autonomy* of the agent: no other agent can directly impose a given hypothesis to an agent. As a consequence, only a new observation (being it a new perception, or an observation communicated by a fellow agent) can result in a modification of its preferred hypothesis h_i (but not necessarily of course).

We finally define a property of the system that we shall use in the rest of the paper:

DEFINITION 3 (BOUNDED PERCEPTIONS). A system involves a bounded perception for agents iff $\exists n_0$ s.t. $\forall t | \cup_{i=1}^N O_i | \leq n_0$. (That is, the number of observations to be made by the agents in the system is not infinite.)

Agent Cycle

Now we need to see how these agents will evolve and interact in their environment. In our context, agents evolve in a dynamic environment, and we classically assume the following system cycle:

1. *Environment dynamics*: the environment evolves according to the defined rules of the system dynamics.
2. *Perception step*: agents get perceptions from the environment. These perceptions are typically partial (e.g. the agent can only see a portion of a map).
3. *Reasoning step*: agents compare perception with predictions, seek explanations for (potential) difference(s), refine their hypothesis, draw new conclusions.
4. *Communication step*: agents can communicate hypotheses and observations with other agents through a defined protocol. Any agent can only be involved in one communication with another agent by step.
5. *Action step*: agents do some practical reasoning using the models obtained from the previous steps and select an action. They can then modify the environment by executing it.

The communication of the agents will be further constrained by topological consideration. At a given time, an agent will only be able to communicate with a number of neighbours. Its connexions with these others agents may

evolve with its situation in the environment. Typically, an agent can only communicate with agents that it can sense, but one could imagine evolving topological constraints on communication based on a network of communications between agents where the links are not always active.

Communication

In our system, agents will be able to communicate with each other. However, due to the aforementioned topological constraints, they will not be able to communicate with any agents at anytime. Who an agent can communicate with will be defined dynamically (for instance, this can be a consequence of the agents being close enough to get in touch). We will abstractly denote by $\mathcal{C}(a_i, a_j, t)$ the communication property, in other words, the fact that agents a_i and a_j can communicate at time t (note that this relation is assumed to be symmetric, but of course not transitive). We are now in a position to define two essential properties of our system.

DEFINITION 4 (TEMPORAL PATH). *There exists a temporal communication path at horizon t_f (noted $\mathcal{L}_{t_f}(a_i, a_j)$) between a_i and a_j iff there exists a sequence of time points $t_1, t_2, \dots, t_n$ from t_f and a sequence of agents $k_1, k_2, \dots, k_n$ s.t. (i) $\mathcal{C}(a_i, a_{k_1}, t_1)$, (ii) $\mathcal{C}(a_{k_n}, a_j, t_{n+1})$, (iii) $\forall i \in [1, n]$, $\mathcal{C}(a_{k_i}, a_{k_{i+1}}, t_i)$*

Intuitively, what this property says is that it is possible to find a “temporal path” in the future that would allow to link agent a_i and a_j via a sequence of intermediary agents. Note that the time points are not necessarily successive, and that the sequence of agents may involve the same agents several times.

DEFINITION 5 (TEMPORAL CONNEXITY). *A system is temporally connex iff $\forall t \forall (i, j) \in [1, n]^2 \mathcal{L}_t(a_i, a_j)$*

In short, a temporally connex system guarantees that any agent will be able to communicate with any other agents, no matter how long it might take to do so, at any time. To put it another way, it is never the case that an agent will be isolated for ever from another agent of the system.

We will next discuss the detail of how communication concretely takes place in our system. Remember that in this paper, we only consider the case of bilateral exchanges (an agent can only speak to a single other agent), and that we also assume that any agent can only engage in a single exchange in a given round.

3. PROTOCOLS AND STRATEGIES

In this section, we discuss the requirements of the interaction protocols that govern the exchange of messages between agents, and provide some example instantiation of such protocols. To clarify the presentation, we distinguish two levels: the *local* level, which is concerned with the regulation of bilateral exchanges; and the *global* level, which essentially regulates the way agents can actually engage into a conversation. At each level, we separate what is specified by the protocol, and what is left to agents’ strategies.

Local Protocol and Strategies

We start by inspecting local protocols and strategies that will regulate the communication between the agents of the system. As we limit ourselves to bilateral communication, these protocols will simply involve two agents. Such protocol will have to meet one basic requirement to be satisfying.

- *consistency (CONS)*—a local protocol has to guarantee the mutual consistency of agents upon termination (which implies termination of course).




Figure 1: A Hypotheses Exchange Protocol [1]

One example such protocol is the protocol described in [1] that is pictured in Fig. 1. To further illustrate how such protocol can be used by agents, we give some details on a possible strategy: upon receiving a hypothesis h_1 ($propose(h_1)$ or $counterpropose(h_1)$) from a_1 , agent a_2 is in state 2 and has the following possible replies: *counterexample* (if the agent knows an example contradicting the hypothesis, or not explained by this hypothesis), *challenge* (if the agent lacks evidence to accept this hypothesis), *counterpropose* (if the agent agrees with the hypothesis but prefers another one), or *accept* (if it is indeed as good as its favourite hypothesis). This strategy guarantees, among other properties, the eventual mutual logical consistency of the involved agents [1].

Global Protocol

The global protocol regulates the way bilateral exchanges will be initiated between agents. At each turn, agents will concurrently send one weighted request to communicate to other agents. This weight is a value measuring the agent’s willingness to converse with the targeted agent (in practice, this can be based on different heuristics, but we shall make some assumptions on agents’ strategies, see below). Sending such a request is a kind of conditional commitment for the agent. An agent sending a weighted request commits to engage in conversation with the target if he does not receive and accept himself another request. Once all request have been received, each agent replies with either an *accept* or a *reject*. By answering with an *accept*, an agent makes a full commitment to engage in conversation with the sender. Therefore, it can only send one accept in a given round, as an agent can only participate in one conversation per time step. When all response have been received, each agent receiving an *accept* can either initiate a conversation using the local protocol or send a *cancel* if it has accepted another request.

At the end of all the bilateral exchanges, the agents engaged in conversation are discarded from the protocol. Then each of the remaining agents resends a request and the process iterates until no more requests are sent.

Global Strategy

We now define four requirements for the strategies used by agents, depending on their role in the protocol: two are concerned with the requestee role (how to decide who the

agent *wishes* to communicate with?), the other two with the responder role (how to decide which communication request to accept or not?).

- *Willingness to solve inconsistencies (SOLVE)*—agents want to communicate with any other agents unless they know they are mutually consistent.
- *Focus on solving inconsistencies (FOCUS)*—agents do not request communication with an agent with whom they know they are mutually consistent.
- *Willingness to communicate (COMM)*—agents cannot refuse a weighted communication request, unless they have just received or send a request with a greater weight.
- *Commitment to communication request (REQU)*—agents cannot accept a weighted communication request if they have themselves sent a communication request with a greater weight. Therefore, they will not cancel their request unless they have received a communicational request with greater weight.

Now the protocol structure, together with the properties COMM+REQU, ensure that a request can only be rejected if its target agent engages in communication with another agent. Suppose indeed that agent a_i wants to communicate with a_j by sending a request with weight w . COMM guarantees that an agent receiving a weighted request will either accept this communication, accept a communication with a greater weight or wait for the answer to a request with a greater weight. This ensures that the request with maximal weight will be accepted and not cancelled (as REQU ensures that an agent sending a request can only cancel it if he accepts another request with greater weight). Therefore at least two agents will engage in conversation per round of the global protocol. As the protocol ensures that a_i can resend its request while a_j is not engaged in a conversation, there will be a turn in which a_j must engage in a conversation, either with a_i or another agent.

These requirements concern request sending and acceptance, but agents also need some strategy of weight attribution. We describe below an altruist strategy, used in our experiments. Being cooperative, an agent may want to know more of the communication wishes of other agents in order to improve the overall allocation of exchanges to agents. A *context request* step is then added to the global protocol. Before sending their chosen weighted request, agents attribute a weight to all agents they are prepared to communicate with, according to some internal factors. In the simplest case, this weight will be 1 for all agent with whom the agent is not sure of being mutually consistent (ensuring SOLVE), other agent being not considered for communication (ensuring FOCUS). The agent then sends a *context request* to all agents with whom communication is considered. This request also provides information about the sender (list of considered communications along with their weight). After reception of all the *context requests*, agents will either reply with a *deny*, iff they are already engaged in a conversation (in which case, the requesting agent will not consider communication with them anymore in this turn), or an *inform* giving the requester information about the requests it has sent and received. When all replies have been received, each agent can calculate the weight of all requests concerning it. It does so

by subtracting from the weight of its request the weight of all requests concerning either it or its target (that is, the final weight of the request from a_i to a_j is $W_{i,j} = w_{i,j} + w_{j,i} - (\sum_{k \in R(i) - \{j\}} w_{i,k} + \sum_{k \in S(i) - \{j\}} w_{k,i} + \sum_{k \in R(j) - \{i\}} w_{j,k} + \sum_{k \in S(j) - \{i\}} w_{k,j})$ where $w_{i,j}$ is the weight of the request of a_i to a_j , $R(i)$ is the set of indice of agents having received a request from a_i and $S(i)$ is the set of indice of agents having send a request to a_i). It then finally sends a weighted *request* to the agents who maximise this weight (or wait for a *request*) as described in the global protocol.

4. (CONDITIONAL) CONVERGENCE TO GLOBAL CONSISTENCY

In this section we will show that the requirements regarding protocols and strategies just discussed will be sufficient to ensure that the system will eventually converge towards global consistency, under some conditions. We first show that, if two agents are not mutually consistent at some time, then there will be necessarily a time in the future such that an agent will learn a new observation, being it because it is new for the system, or by learning it from another agent.

LEMMA 1. *Let $\mathcal{S}$ be a system populated by n agents $a_1, a_2, \dots, a_n$, temporally connex, and involving bounded perceptions for these agents. Let n_1 be the sum of cardinalities of the intersection of pairwise observation sets. ($n_1 = \sum_{(i,j) \in [1,n]^2} |O_i \cap O_j|$) Let n_2 be the cardinality of the union of all agents' observations sets. ($n_2 = |\cup_{i=1}^n O_i|$). If $\neg MCons(a_i, a_j)$ at time t_0 , there is necessarily a time $t' > t_0$ s.t. either n_1 or n_2 will increase.*

PROOF. Suppose that there exist a time t_0 and indices (i, j) s.t. $\neg MCons(a_i, a_j)$. We will use $m_{t_0} = \sum_{(k,l) \in [1,n]^2} \varepsilon_{Comm}(a_k, a_l, t_0)$ where $\varepsilon_{Comm}(a_k, a_l, t_0) = 1$ if a_k and a_l have communicated at least once since t_0 , and 0 otherwise. Temporal connexity guarantees that there exist $t_1, \dots, t_{m+1}$ and $k_1, \dots, k_m$ s.t. $\mathcal{C}(a_i, a_{k_1}, t_1)$, $\mathcal{C}(a_{k_m}, a_j, t_{m+1})$, and $\forall p \in [1, m]$, $\mathcal{C}(a_{k_p}, a_{k_{p+1}}, t_p)$. Clearly, if $MCons(a_i, a_{k_1})$, $MCons(a_{k_m}, a_j)$ and $\forall p, MCons(a_{k_p}, a_{k_{p+1}})$, we have $MCons(a_i, a_j)$ which contradicts our hypothesis ($MCons$ being transitive, $MCons(a_i, a_{k_1}) \wedge MCons(a_{k_1}, a_{k_2})$ implies that $MCons(a_i, a_{k_2})$ and so on till $MCons(a_i, a_{k_m}) \wedge MCons(a_{k_m}, a_j)$ which implies $MCons(a_i, a_j)$).

At least two agents are then necessarily inconsistent ($\neg MCons(a_i, a_{k_1})$, or $\neg MCons(a_{k_m}, a_j)$, or $\exists p_0$ t.q. $\neg MCons(a_{k_{p_0}}, a_{k_{p_0+1}})$). Let a_k and a_l be these two neighbours at a time $t' > t_0$ ¹. The SOLVE property ensures that either a_k or a_l will send a communication request to the other agent at time t' . As shown before, this in turn ensures that at least one of these agents will be involved in a communication. Then there are two possibilities: (*case i*) a_k and a_l communicate at time t' . In this case, we know that $\neg MCons(a_k, a_l)$. This and the CONS property ensures that at least one of the agents must change its

¹Strictly speaking, the transitivity of $MCons$ only ensure that a_k and a_l are inconsistent at a time $t'' \geq t_0$ that can be different from the time t' at which they can communicate. But if they become consistent between t'' and t' (or inconsistent between t' and t''), it means that at least one of them have changed its hypothesis between t'' and t' , that is, after t_0 . We can then apply the reasoning of *case iib*.

hypothesis, which in turn, since agents are autonomous, implies at least one exchange of observation. But then $|O_k \cap O_l|$ is bound to increase: $n_1(t') > n_1(t_0)$.

(case ii) a_k communicates with a_p at time t' . We then have again two possibilities:

(case iia) a_k and a_p did not communicate since t_0 . But then $\varepsilon_{Comm}(a_k, a_p, t_0)$ had value 0 and takes value 1. Hence m_{t_0} increases.

(case iib) a_k and a_p did communicate at some time $t'_0 > t_0$. The CONS property of the protocol ensures that $MCons(a_k, a_p)$ at that time. Now the fact that they communicate and FOCUS implies that at least one of them did change its hypothesis in the meantime. The fact that agents are *autonomous* implies in turn that a new observation (perceived or received from another agent) necessarily provoked this change. The latter case would ensure the existence of a time $t'' > t_0$ and an agent a_q s.t. either $|O_p \cap O_q|$ or $|O_k \cap O_q|$ increases of 1 at that time (implying $n_1(t'') > n_1(t_0)$). The former case means that the agent gets a new perception o at time t'' . If that observation was unknown in the system before, then $n_2(t'') > n_2(t_0)$. If some agent a_q already knew this observation before, then either $O_p \cap O_q$ or $O_k \cap O_q$ increases of 1 at time t'' (which implies that $n_1(t'') > n_1(t_0)$). Hence, $\neg MCons(a_i, a_j)$ at time t_0 guarantees that, either:
 $-\exists t' > t_0$ t.q. $n_1(t') > n_1(t_0)$; or
 $-\exists t' > t_0$ t.q. $n_2(t') > n_2(t_0)$; or
 $-\exists t' > t_0$ t.q. m_{t_0} increases of 1 at time t' .

By iterating the reasoning with t' (but keeping t_0 as the time reference for m_{t_0}), we can eliminate the third case (m_{t_0} is integer and bounded by n^2 , which means that after a maximum of n^2 iterations, we necessarily will be in one of two other cases.) As a result, we have proven that if $\neg MCons(a_i, a_j)$ at time t_0 , there is necessarily a time t' s.t. either n_1 or n_2 will increase. $\square$

THEOREM 1 (GLOBAL CONSISTENCY). *Let $\mathcal{S}$ be a system populated by n agents $a_1, a_2, \dots, a_n$, temporally connex, and involving bounded perceptions for these agents. Let $Cons(a_i, a_j)$ be a transitive consistency property. Then any protocol and strategies satisfying properties CONS, SOLVE, FOCUS, COMM and REQU guarantees that the system will converge towards global consistency.*

PROOF. For the sake of contradiction, let us assume $\exists I, J \in [1, N]$ s.t. $\forall t, \exists t_0 > t$, t.q. $\neg Cons(a_I, a_J, t_0)$. Using the lemma, this implies that $\exists t' > t_0$ s.t. either $n_1(t') > n_1(t_0)$ or $n_2(t') > n_2(t_0)$. But we can apply the same reasoning taking $t = t'$, which would give us $t_1 > t' > t_0$ s.t. $\neg Cons(a_I, a_J, t_1)$, which gives us $t'' > t_1$ s.t. either $n_1(t'') > n_1(t_1)$ or $n_2(t'') > n_2(t_1)$. By successive iterations we can then construct a sequence $t_0, t_1, \dots, t_n$, which can be divided in two sub-sequences $t'_0, t'_1, \dots, t'_n$ and $t''_0, t''_1, \dots, t''_n$ s.t. $n_1(t'_0) < n_1(t'_1) < \dots < n_1(t'_n)$ and $n_2(t''_0) < n_2(t''_1) < \dots < n_2(t''_n)$. One of these sub-sequences has to be infinite. However, $n_1(t'_i)$ and $n_2(t''_i)$ are strictly growing, integer, and bounded, which implies that both are finite. Contradiction. $\square$

What the previous result essentially shows is that, in a system where no agent will be isolated from the rest of the agents for ever, only very mild assumptions on the protocols and strategies used by agents suffice to guarantee convergence towards system consistency in a finite amount of time (although it might take very long). Unfortunately, in many

“critical” situations, it will not be possible to assume this temporal connexity. As distributed approaches as the one advocated in this paper are precisely often presented as a good way to tackle problems of reliability or problems of dependence to a center that are of utmost importance in these critical applications, it is certainly interesting to further explore how such a system would behave when we relax this assumption.

5. EXPERIMENTAL STUDY

This experiment involves agents trying to escape from a burning building. The environment is described as a spatial grid with a set of walls and (thankfully) some exits. Time and space are considered discrete. Time is divided in rounds. Agents are localised by their position on the spatial grid. These agents can move and communicate with other agents. In a round, an agent can move of one cell in any of the four cardinal directions, provided it is not blocked by a wall. In this application, agents communicate with any other agent (but, recall, a single one) given that this agent is in view, and that they have not yet exchanged their current favoured hypothesis. Suddenly, a fire erupts in these premises. From this moment, the fire propagates. Each round, for each cases where there is fire, the fire propagates in the four directions. However, the fire cannot propagate through a wall. If the fire propagates in a case where an agent is positioned, that agent burns and is considered dead. It can of course no longer move nor communicate. If an agent gets to an exit, it is considered saved, and can no longer be burned. Agents know the environment and the rules governing the dynamics of this environment, that is, they know the map as well as the rules of fire propagation previously described. They also locally perceive this environment, but cannot see further than 3 cases away, in any direction. Walls also block the line of view, preventing agents from seeing behind them. Within their sight, they can see other agents and whether or not the cases they see are on fire. All these perceptions are memorised.

We now show how this instantiates the abstract framework presented the paper.

- $\mathcal{O} = \{Fire(x, y, t), NoFire(x, y, t), Agent(a_i, x, y, t)\}$
Observations can then be positive ($o \in P(\mathcal{O})$ iff $\exists h \in \mathcal{H}$ s.t. $h \models o$) or negative ($o \in N(\mathcal{O})$ iff $\exists h \in \mathcal{H}$ s.t. $h \models \neg o$).
- $\mathcal{H} = \{FireOrigin(x_1, y_1, t_1) \wedge \dots \wedge FireOrigin(x_l, y_l, t_l)\}$
Hypotheses are conjunctions of *FireOrigins*.
- $Cons(h, O)$ consistency relation satisfies:
 - *coherence* : $\forall o \in N(O), h \not\models \neg o$.
 - *completeness* : $\forall o \in P(O), h \models o$.
 - *minimality* : For all $h' \in \mathcal{H}$, if h' is coherent and complete for O , then h is preferred to h' according to the preference relation ($h \leq_p h'$).²

²Selects first the minimal number of origins, then the most recent (least preemptive strategy [6]), then uses some arbitrary fixed ranking to discriminate ex-aequo. The resulting relation is a total order, hence minimality implies that there will be a single h s.t. $Cons(O, h)$ for a given O . This in turn means that $MCons(a_i, a_j)$ iff $Cons(a_i), Cons(a_j)$, and $h_i = h_j$. This relation is then transitive and symmetric.

- $\mathcal{E}_h$ takes O as argument and returns $\min_{\leq p}$ of the coherent and complete hypothesis for O

5.1 Experimental Evaluation

We will classically (see e.g. [3, 4]) assess the *effectiveness* and *efficiency* of different interaction protocols.

Effectiveness of a protocol

The proportion of agents surviving the fire over the initial number of agents involved in the experiment will determine the effectiveness of a given protocol. If this value is high, the protocol has been effective to propagate the information and/or for the agents to refine their hypotheses and determine the best way to the exit.

Efficiency of a protocol

Typically, the use of supporting information will involve a communication overhead. We will assume here that the efficiency of a given protocol is characterised by the data flow induced by this protocol. In this paper we will only discuss this aspect wrt. local protocols. The main measure that we shall then use here is the mean total size of messages that are exchanged by agents per exchange (hence taking into account both the number of messages and the actual size of the messages, because it could be that messages happen to be very “big”, containing e.g. a large number of observations, which could counter-balance a low number of messages).

5.2 Experimental Settings

The chosen experimental settings are the following:

- *Environmental topology*— Performances of information propagation are highly constrained by the environment topology. The perception skills of the agents depend on the “openness” of the environment. With a large number of walls the perceptions of agents are limited, and also the number of possible inter-agent communications, whereas an “open” environment will provide optimal possibilities of perception and information propagation. Thus, we propose a topological index (see below) as a common basis to characterize the environments (maps) used during experimentations. The *topological index* (TI) is the ratio of the number of cells that can be perceived by agents summed up from all possible positions, divided by the number of cells that would be perceived from the same positions but *without any walls*. (The closer to 1, the more open the environment). We shall also use two additional, more classical [10], measures: the *characteristic path length*³ (CPL) and the *clustering coefficient*⁴ (CC).
- *Number of agents*— The propagation of information also depends on the initial number of agents involved during an experimentation. For instance, the more agents, the more potential communications there is. This means that there will be more potential for propagation, but also that the bilateral exchange restriction will be more crucial.

³The CPL is the median of the means of the shortest path lengths connecting each node to all other nodes.

⁴characterising the isolation degree of a region of an environment in terms of accessibility (number of roads still usable to reach this region).

Map	T.I. (%)	C.P.L.	C.C.
69-1	69,23	4,5	0,69
69-2	68,88	4,38	0,65
69-3	69,80	4,25	0,67
53-1	53,19	5,6	0,59
53-2	53,53	6,38	0,54
53-3	53,92	6,08	0,61
38-1	38,56	8,19	0,50
38-2	38,56	7,3	0,50
38-3	38,23	8,13	0,50

Table 1: Topological Characteristics of the Maps

- *Initial positions of the agents*— Initial positions of the agents have a significant influence on the overall behavior of an instance of our system: being close from an exit will (in general) ease the escape.

5.3 Experimental environments

We choose to realize experiments on three very different topological indexes (69% for “open” environments, 53% for mixed environments, and 38% for labyrinth-like environments).




Figure 2: Two maps (left: TI=69%, right TI=38%)

We designed three different maps for each index (Fig. 2 shows two of them), containing the same maximum number of agents (36 agents max.) with a maximum density of one agent per cell, the same number of exits and a similar fire origin (e.g. starting time and position). The three different maps of a given index are designed as follows. The first map is a model of an existing building floor. The second map has the same “enclosure”, exits and fire origin as the first one, but the number and location of walls are different (wall locations are designed by an heuristic which randomly creates walls on the spatial grid such that no fully closed rooms are created and that no exit is closed). The third map is characterised by geometrical “enclosure” in which walls location is also designed with the aforementioned heuristic. Table 1 summarizes the different topological measures characterizing these different maps. It is worth pointing out that the values confirm the relevance of TI (maps with a high TI have a low CPL and a high CC. However the CPL and CC allows to further refine the difference between the maps, e.g. between 53-1 and 53-2).

5.4 Experimental Results

For each triple of maps defined as above we conduct the same experiments. In each experiment, the society differs in terms of its initial proportion of involved agents, from 1% to 100%. This initial proportion represents the percentage of involved agents with regards to the possible maximum number of agents. For each map and each initial proportion,

we select randomly 100 different initial agents' locations. For each of those different locations we execute the system one time for each different interaction protocol.

Effectiveness of Communication and Argumentation

The first experiment that we set up aims at testing how effective is hypotheses exchange (HE), and in particular how the topological aspects will affect this effectiveness. In order to do so, we have computed the ratio of improvement offered by that protocol over a situation where agents could simply not communicate (*no comm*). To get further insights as to what extent the hypotheses exchange was really crucial, we also tested a much less elaborated protocol consisting of mere observation exchanges (OE). More precisely, this protocol requires that each agent stores any "unexpected" observation that it perceives, and agents simply exchange their respective lists of observations when they discuss. In this case, the local protocol is different (note in particular that it does not guarantee mutual consistency), but the global protocol remains the same (at the only exception that agents' motivation to communicate is to synchronise their list of observations, not their hypothesis). If this protocol is at best as effective as HE, it has the advantage of being more efficient (this is obvious wrt the number of messages which will be limited to 2, less straightforward as far as the size of messages is concerned, but the rough observation that the exchange of observations can be viewed as a "flat" version of the challenge is helpful to see this). The results of these experiments are reported in Fig. 3.




Figure 3: Comparative effectiveness ratio gain of protocols when the proportion of agents augments

The first observation that needs to be made is that communication improves the effectiveness of the process, and this ratio increases as the number of agents grows in the system. The second lesson that we learn here is that closeness relatively makes communication more effective over non communication. Maps exhibiting a T.I. of 38% are constantly above the two others, and 53% are still slightly but significantly better than 69%. However, these curves also suggest, perhaps surprisingly, that HE outperforms OE in precisely those situations where the ratio gain is less important (the only noticeable difference occurs for rather open maps where T.I. is 69%). This may be explained as follows: when a map is open, agents have many potential explanation

candidates, and argumentation becomes useful to discriminate between those. When a map is labyrinth-like, there are fewer possible explanations to an unexpected event.

Importance of the Global Protocol

The second set of experiments seeks to evaluate the importance of the design of the global protocol. We tested our protocol against a "local broadcast" (LB) protocol. Local broadcast means that *all* the neighbours agents perceived by an agent will be involved in a communication with that agent in a given round —we alleviate the constraint of a single communication by agent. This gives us a rough upper bound upon the possible ratio gain in the system (for a given local protocol). Again, we evaluated the ratio gain induced by that LB over our classical HE, for the three different classes of maps. The results are reported in Fig. 4.




Figure 4: Ratio gain of local broadcast over hypotheses exchange

Note to begin with that the ratio gain is 0 when the proportion of agents is 5%, which is easily explained by the fact that it corresponds to situations involving only two agents.

We first observe that all classes of maps witness a ratio gain increasing when the proportion of agents augments: the gain reaches 10 to 20%, depending on the class of maps considered. If one compares this with the improvement reported in the previous experiment, it appears to be of the same magnitude. This illustrates that the design of the global protocol cannot be ignored, especially when the proportion of agents is high. However, we also note that the effectiveness ratio gain curves have very different shapes in both cases: the gain induced by the accuracy of the local protocol increases very quickly with the proportion of agents, while the curve is really smooth for the global one.

Now let us observe more carefully the results reported here: the curve corresponding to a TI of 53% is above that corresponding to 38%. This is so because the more open a map, the more opportunities to communicate with more than one agent (and hence benefits from broadcast). However, we also observe that curve for 69% is below that for 53%. This is explained as follows: in the case of 69%, the potential gain to be made in terms of surviving agents is much lower, because our protocols already give rather efficient outcomes anyway (quickly reaching 90%, see Fig. 3).

A simple rule of thumb could be that when the number of agents is small, special attention should be put on the local

protocol, whereas when that number is large, one should carefully design the global one (unless the map is so open that the protocol is already almost optimally efficient).

Efficiency of the Protocols

The final experiment reported here is concerned with the analysis of the efficiency of the protocols. We analysis here the mean size of the totality of the messages that are exchanged by agents (mean size of exchanges, for short) using the following protocols: HE, OE, and two variant protocols. The first one is an intermediary *restricted hypotheses exchange* protocol (RHE). RHE is as follows: it does not involve any challenge nor counter-propose, which means that agents cannot switch their role during the protocol (this differs from RE in that respect). In short, RHE allows an agent to exhaust its partner's criticism, and eventually this partner will come to adopt the agent's hypothesis. Note that this means that the autonomy of the agent is not preserved here (as an agent will essentially accept any hypothesis it cannot undermine), with the hope that the gain in efficiency will be significant enough to compensate a loss in effectiveness. The second variant protocol is a *complete observation exchange* protocol (COE). COE uses the same principles as OE, but includes in addition all critical negative examples (*nofire*) in the exchange (thus giving all examples used as arguments by the hypotheses exchanges protocol), hence improving effectiveness. Results for map 69-1 are shown on Fig. 5.




Figure 5: Mean size of exchanges

First we can observe the fact that the ordering of the protocols, from the least efficient to the most efficient, is COE, HE, RHE and then OE. HE being more efficient than COE proves that the argumentation process gains efficiency by selecting when it is needed to provide negative example, which have less impact than positive ones in our specific testbed. However, by communicating hypotheses before eventually giving observation to support it (HE) instead of directly giving the most crucial observations (OE), the argumentation process doubles the size of data exchanges. It is the cost for ensuring consistency at the end of the exchange (a property that OE does not support). Also significant is the fact that the mean size of exchanges is slightly higher when the number of agents is small. This is explained by the fact that in these cases only a very few agents have relevant informations in their possession, and that they will need to communicate

a lot in order to come up with a common view of the situation. When the number of agents increases, this knowledge is distributed over more agents which need shorter discussions to get to mutual consistency. As a consequence, the relative gain in efficiency of using RHE appears to be better when the number of agents is small: when it is high, they will hardly argue anyway. Finally, it is worth noticing that the standard deviation for these experiments is rather high, which means that the conversation do not converge to any “stereotypic” pattern.

6. CONCLUSION

This paper has investigated the properties of a multi-agent system where each (distributed) agent locally perceives its environment, and tries to reach consistency with other agents despite severe communication restrictions. In particular we have exhibited conditions allowing convergence, and experimentally investigated a typical situation where those conditions cannot hold. There are many possible extensions to this work, the first being to further investigate the properties of different global protocols belonging to the class we identified, and their influence on the outcome. There are in particular many heuristics, highly dependent on the context of the study, that could intuitively yield interesting results (in our study, selecting the recipient on the basis of what can be inferred from his observed actions could be such a heuristic). One obvious candidate for longer term issues concern the relaxation of the assumption of perfect sensing.

7. REFERENCES

- [1] G. Bourgne, N. Maudet, and S. Pinson. When agents communicate hypotheses in critical situations. In *Proceedings of DALT-2006*, May 2006.
- [2] P. Harvey, C. F. Chang, and A. Ghose. Support-based distributed search: a new approach for multiagent constraint processing. In *Proceedings of AAMAS06*, 2006.
- [3] H. Jung and M. Tambe. Argumentation as distributed constraint satisfaction: Applications and results. In *Proceedings of AGENTS01*, 2001.
- [4] N. C. Karunatilake and N. R. Jennings. Is it worth arguing? In *Proceedings of ArgMAS 2004*, 2004.
- [5] S. Ontañón and E. Plaza. Arguments and counterexamples in case-based joint deliberation. In *Proceedings of ArgMAS-2006*, May 2006.
- [6] D. Poole. Explanation and prediction: An architecture for default and abductive reasoning. *Computational Intelligence*, 5(2):97–110, 1989.
- [7] I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, 4(18):345–375, 2003.
- [8] N. Roos, A. ten Tije, and C. Witteveen. A protocol for multi-agent diagnosis with spatially distributed knowledge. In *Proceedings of AAMAS03*, 2003.
- [9] N. Roos, A. ten Tije, and C. Witteveen. Reaching diagnostic agreement in multiagent diagnosis. In *Proceedings of AAMAS04*, 2004.
- [10] T. Takahashi, Y. Kaneda, and N. Ito. Preliminary study - using robocuprescue simulations for disasters prevention. In *Proceedings of SRMED2004*, 2004.