INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team MAIA

MAchine Intelligente Autonome

Lorraine

THEME 3A

Activity Report

2003

Table of contents

1. Team

Key words: Autonomous Agents and Multi-Agent Systems, Distributed Artificial Intelligence, Decision Theory, Reasoning under Uncertainty, Learning, Planning, Diagnosis.

Head of the Project

François Charpillet [Director of Research, INRIA]

Administrative Assistant

Martine Kuhlmann [CNRS, part-time]

Research Scientist

Alain Dutech [Chargé de Recherche, INRIA]

Faculty Members

Christine Bourjot [Maître de Conférences, U. Nancy 2]
Anne Boyer [Maître de Conférences, U. Nancy 2]
Vincent Chevrier [Maître de Conférences, U. Nancy 1]
Jean-Paul Haton [Professeur IUF, U. Nancy 1, member of l'Institut Universitaire de France (Chair of Modélisation informatique des processus perceptifs et cognitifs), also member of the project team Parole]
Alexis Scheuer [Maître de Conférences, U. Nancy 1]

Associate member

Laurent Piraud

Post-doctorates

Olivier Buffet [ATER, U. Nancy 1]
Iadine Chadès [ATER, U. Nancy 2, till october]
Romaric Charton [ATER, U. Nancy 2]
Frank Gechter [ATER, l'IUFM de Lorraine]
Laurent Jeanpierre [ATER, U. Nancy 1, till october]
Bruno Scherrer [ATER, U. Nancy 1, till october]

Ph.D. Students

Vincent Thomas [MENRT scholarship]
Simon Le Gloannec [INRIA/CNRS scholarship, co-directed by A. Mouaddib, Professor, University of Caen]
Daniel Szer [INRIA scholarship]
Raghav Aras [INRIA scholarship]
Cherif Smaili [Egide scholarship, DRT student U. Nancy 2]

Project Technical Staff

Régis Lhoste [Engineer INRIA]
Eric Lucchese [Engineer U. Nancy 1]
Loïc Pelissier [Engineer INRIA]
Cedric Rose [Engineer associated with the development of Dialhemo]

2. Overall Objectives

The long term objective of the MAIA project is to understand how an artificial intelligence software system can be designed based on the basic idea that surprisingly intelligent behavior can arise from fairly simple autonomous and distributed entities called "agents"[1], having very simple control mechanisms but interacting with a complex environment. The team proposes to pursue this idea in two ways: first, to engineer the behaviors given specifications of the environment, agent architecture and objective; and second, to engineer the behaviours by imitating natural systems.

[1] An agent as defined by Russell and Norvig in their book entitled *Artificial Intelligence: A modern approach* is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators.

- The first approach makes use of the decision-theoric approach (using Markov decision processes or Bayesian networks) to planning and learning. In particular, the team focus on learning and planning in partially observable environments, and on learning and planning to coordinate collaborative multi-agent systems.
- The second part of the project focuses on the design of agent based systems according to the principles observed in nature. Natural systems display robustness, adaptability, flexibility, self-organisation and self regulation capabilities which are appealing for designing complex systems.

The MAIA team addresses systems made of one, several or large number of agents constructed by one or both of the two above-mentioned approaches. The most difficult case but the most appealing one is when we face with large number of agents. In this case we are interested in the study of emergent collective intelligence of groups of simple agents, i.e. swarm intelligence. As argued by Eric Bonabeau, Marco Dorigo and Guy Theraulaz in their book entitled *Swarm Intelligence: From Natural to Artificial Systems*, “at a time when the world become so complex that no human being can understand it, ... when software systems become so intractable that they can no longer be controlled, swarm intelligence offers an alternative way of designing intelligent systems, in which autonomy, emergence and distributed functioning replace control, programming and centralisation.”

The basic research results obtained by MAIA group in sequential or distributed decision making, the strong knowhow in using mathematical tools such as Hidden Markov Model, Markov Decision process, bayesian networks, combined with a strong experience in centralized or distributed learning, and a good knowledge in recent advance in “agent technology” open to the MAIA group many opportunities for addressing a wide range of application domains. Among them the team concentrates its effort on developing medical applications, web services or robotic application. The mobile robots are a perfect test bed for a variety of learning and planning approach, and will be crucial for studying emergent collective behavior. The team has more than ten robots, which will make a fascinating test domain. Work on robots as a test domain contributes directly to the field of intelligent robots; but robots are also useful to test a variety of high level principles that may be applied in other projects. A good example of such transfer is the work done for designing the Diatelic system, a telemedecine system which monitors the hydration of kidney disease people.

3. Scientific Foundations

3.1. The decision theoric approach

3.1.1. Our approach

An agent as defined by Russell and Norvig in their book entitled *Artificial Intelligence: A modern approach* is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators. This view makes Markov decision processes (**MDPs**) a good candidate for formulating agents. It is probably why MDPs have received considerable attention in recent years by the artificial intelligence (AI) community. They have been adopted as a general framework for planning under uncertainty and reinforcement learning.

Formally, a Markov Decision Process is a four tuple $\langle S, A, P, R \rangle$, where :

- S is the state space,
- A is the action space,
- P is the state-transition probability function that models the dynamic of the system. P(s, a, s’) is the probability of transitioning from s to s’ on taking action a in s,
- R is the reward function. R(s, a, s’) stands for the reward obtained from taking action a in state s, and transitioning to state s’. Here s, $s' \in S$, $a \in A$.

This framework makes it possible to model the interaction between a rational agent which perceives any state of S, and which acts using any action A upon the environment E in which the agent is integrated. The environment can be considered as a Markov decision process which evolves over time when an agent acts on it. When the system is in a given state s, an action a being chosen by A, the probability for the system to end in state s' is given by P(s, a, s'). After each transition, the environment generates a reward r(s, a, s'). The problem is then to find a mapping $\pi(s, a)$ between states and actions. Such a mapping is called a policy.

In the general case, the information that an agent can acquire on the environment is generally restricted to *observations* obtained for example using sensors. These observations only give partial information on the state in which the agent stands. Therefore, the agent can confound two different states (the aliasing problem). Thus, The decision process has hidden state, and the issue of finding an optimal policy is no more a Markov problem N-MDP. **POMDPs** are an example of N-MDPs. More formally, a POMDP is a tuple $\langle S, A, P, R, \Omega, O\rangle$ where

- S is a finite set of states.
- A is a finite action set.
- P is a transition probability. P(s, a, s') is a rational representing the probability of transitioning from s to s' on taking action a in s. Here s, $s' \in S$, $a \in A$.
- R is a reward function. R(s, a, s') is a rational representing the reward obtained from taking action a in state s, and transitioning to state s'. Here s, $s' \in S$, $a \in A$.
- Ω is a finite set of observations.
- O is a table of observation probabilities. O(s, a, s', o) is a rational representing the probability of transitioning from s to s' on taking action a in s while observing o. Here s, $s' \in S$, $a \in A$, $o \in \Omega$.

Considering an agent several issues can be addressed, and are addressed by our group :

1. Which behavior should an agent adopt in order to optimally achieve its objectives ? This issue raises up the problem of finding the optimal behavior (policy) i.e the policy π which maximises a given performance criteria that the agent can get over time when it executes this policy. There are three popular performance criteria to evaluate a policy: expected reward to target, discounted cumulative reward, the average expected reward per stage.
2. How can an agent learn what to do in each state of the environment from reward by a trial error based process. This learning approach is termed reinforcement learning. The agent in this approach learns to behave (find the best policy) by interacting with its environment so as to maximize the expected reward signal that it can expect to get in the future. It is, as stated by Sutton and Barto, an approach for understanding and automating goal-directed learning and decision-making that is quite different from supervised learning. Contrary to supervised learning, in interactive problems learning it's in most cases impossible to get examples of good behaviors for all situations in which an agent has to act. Trade off between exploration and exploitation is one of the major issues to address in this approach. An other important issue is our work on learning in partially observable environments. In this case the agent must remember something about its past actions and observations. Figuring out what to remember about past and how to behave given that information is very difficult.
3. What is the best representation of the environment and other agents so that an agent can achieve the above objectives.
4. how to design the intern architecture of an agent in term of perceptions, internal states, communication capabilities, decision capabilities, aso.
5. how to address non markovian processes, non stationary processes, aso.

Hidden Markov Models (HMM) are a particular case of POMDP that the MAIA team is studying too. A HMM is a POMDP in which there is no action and no reward. Given a proper description of the *state* of a system, it is possible to model it as a Markov chain. The dynamics of the systems is modeled as *transition probabilities* between states. The information that an external observer of the system can acquire on it can be modeled using *observations* which only give partial information on the state of the system. The problem is then to find the most likely state given a sequence of observations.

While substantial progress has been made in planning and control of single agents, a similar formal treatment of multi-agent systems has been lacking. Some preliminary work has been reported, but it generally avoids the central issue in multi-agent systems: agents typically have different information and different knowledge about the overall system and they cannot share all this information all the time. To address the problem of coordination and control of collaborative multi-agent systems, we are conducting both analytical and experimental research aimed at understanding the computational complexity of the problem and developing effective algorithms for solving it. The main objectives of the project are:

- Develop a formal foundation for analysis, algorithm development, and evaluation of different approaches to the control of collaborative multi-agent systems that explicitly captures the notion of communication cost.
- Identify the complexity of the planning and control problem under various constraints on information observability and communication costs.
- Gain a better understanding of what makes decentralized planning and control a hard problem and how to simplify it without compromising the efficiency of the model.
- Develop new general-purpose algorithms for solving different classes of the decentralized planning and control problem.
- Demonstrate the applicability of the new techniques to realistic applications and develop evaluation metrics suitable for decentralized planning and control.

In formalizing coordination, we take an approach based on distributed optimization, in part because we feel that this is the richest such framework: it handles coordination problems in which there are multiple and concurrent goals of varying worth, hard and soft deadlines for goal achievement, alternative ways of achieving goals that offer a trade off between the quality of the solution and resources. Equally important is the fact that this decision-theoretic approach allows us to model explicitly the effects of environmental uncertainty, incomplete and uncertain information and action outcome uncertainty. Coping with these uncertainties is one of the key challenges in designing sophisticated coordination protocols. Finally, a decision-theoretic framework is the most natural one for quantifying the performance of coordination protocols from a statistical perspective.

3.1.2. Related work and position with respect to national and international laboratories

As far as probabilistic planning is concerned, since the mid-1990s, models based on Markov decision processes (MDPs) have been increasingly used by the AI research community, and more and more researchers in this domain are now referring MDPs. Our research group has participated (for instance, notably in co-organizing workshops of AAAI, IJCAI and ECAI in association with ARC INRIA LIRE and with Philippe Chassaing of the Omega project) and continues to participate in and contribute to the development of this domain through collaborations with Shlomo Zilberstein (on two NSF-INRIA projects, INRIA associated team) and with NASA (on a project entitled *Self-directed cooperative planetary rovers*) in association with Shlomo Zilberstein and Victor Lesser of the University of Massachussetts, Eric Hansen of the Mississippi State University, Richard Washington of the NASA Ames Research Center and Abdel-illah Mouaddib of CRIL, Lens.

We have been using the strengths of the basic theoretical properties of the two major approaches for learning and planning that we follow, to design exact algorithms that are able to deal with practical problems of high complexity. Instances of these algorithms include the "JLO" algorithm for Bayesian networks, the

"Q-learning", "TD(λ)" and "Witness" algorithms for problems based in the Markov decision process (MDP) formalism etc. While it is true that the majority of this work has been done in the United States, the French research community is catching up quickly by developing further this domain on it's own. MAIA has been involved directly in making substantial contributions to this development, notably through our active participation in the (informally formed) group of French researchers working on MDPs. Thus, today there are quite a few research teams in France working on MDPs. To name a few, Toulouse-based teams such as IRIT, CERT, INRA, LAAS etc, the Greyc at Caen, and certain Paris-based researchers such as Rémi Munos, Olivier Sigaud etc.

Most of the current work is focused on finding approximate algorithms. Besides applying these algorithms to a multi-agent system (MAS) framework, we have also been focusing on reducing the complexity of implementing these algorithms by making use of the meta-knowledge available in the system being modeled. Thus in implementing the algorithms, we seek temporal, spatial and structural dynamics or functions of the given problem. This is very time-effective in finding approximate solutions of the problem. Moreover, we are seeking ways to combine rigorously these two forms of learning, and to then use them for applications involving planning or learning for agents located in an environment.

3.2. Natural inspired approach

3.2.1. Objectives

One of the issues of the MAIA project is to study and conceive collective intelligence systems based on numerous and simple agents.

Such systems are characterized by a decentralized control (no agent has a knowledge of the whole system) and agents have limited (possibly no) representation of themselves, of the others and of the environment. This implies that agents behaviors are based upon stimulus-response rules, decision-making is based on limited information about the environment and on limited internal states and they do not refer to explicit deliberation.

This work deals with *Swarm Intelligence*; thus the collective observed complexity is out of the individual simplicity and is the consequence of successive actions and interactions of agents through the environment. The complex phenomenon produced must be interpreted as a solution to a problem. Such systems involve two different levels of description: one for invidual behavior (with no reference to the global phenomena) and one to express collective phenomena.

Designing such systems to solve a problem consists in providing a model for the the environment, the agent behaviors and the dynamics of the whole such as the society is able to fulfil its requirements with a reasonable efficiency by providing some collective properties. One of the difficulty is to establish a link between simple individual behaviors and complex collective properties.

The design problem can be summarized as the two following questions:

1. Considering a global expected property or behavior, how to build individual behaviors and system dynamics in order to obtain it?
2. Considering a set of individual behaviors and a system dynamics, how to predict (or guarantee) the global property?

Currently no theory, nor methodology can adress this issue without making very restrictive hypotheses. We will contribute to this purpose.

3.2.2. Our approach : Principles learnt from self-organised natural systems (biology, physics)

Societies of living entities in their dynamic environment are good candidates for formulating robust multi-agents systems. Of the two principal approaches to the study of multi-agent systems (MAS), we have been inclined to adopt for our research, the approach of "collective" systems which emphasize the notions of interactions and organization.

Natural systems like insects, mammals societies or in a broadest sense physical systems provide inspiration to make a link between individual behaviors and collective response. Moreover, such an approach lead to solution which are:

1. dynamic: natural systems organize themselves according to their current context and are able to continuously adapt to its variations
2. decentralized: in such system there is no central entity which controls the overall activities
3. simple: model of individuals behavior is simple.

Our approach consist in to benefit of principles learnt from self-organised natural systems and to transpose them in order to design computer models.

We organize our research in three complementary steps:

1. understanding self-organise by studying examples of such systems
2. transposing principles found in example systems to solve problems
3. providing a framework to help in formalize, generalize and abstract such principles.

- The first step is to model existing self-organized phenomena and thus have a better understanding of the underlying mechanisms. For instance, social phenomena in biology provide a lot of examples in which a collection of simple, situated entities (as ants) can collectively exhibit complex properties which can be interpreted as a collective response to an environmental problem. We have worked with biologists and provided several models of self organized activities in case of spiders and rats.
- Since individual models and system dynamics are established, the next step consists in transposing them in order to solve a given problem. The transposition corresponds to encode the problem such as to be an input for the swarm mechanism; to adapt the swarm mechanism to the specificities of the problem, if necessary to improve it for efficiency purpose; and then to interpret the collective result of the swarm mechanism as a solution of the problem.
- The final step aims at providing a framework to (try to) face the following issues:

    i. Is it possible to describe such mechanisms in order to easily adapt and reuse them to several different instances of problem (generic description)?

    ii. If such a generic description of a system is available, is it possible to assess the behaviour of the system i n order to derive properties that will be conserved in its instantiations (a priori assessment)?

3.2.3. Related work in the national / international research community

This is evident in the numerous collaborations that we have undertaken with (mainly) francophone researchers of this field as well as in the kinds of research groups we associate and work with:

- The research group MARCIA (from 1993 under the aegis of GDR-PRC IA , which V. Chevrier co-directed from 1995 to 1997), which interests itself in studying auto-organization, among other things
- the research group 'Interaction' of the AFCET-AFIA (from 1994 to 1996)

- and the research group 'Colline' (under the aegis of GDR I3 and the AFIA) since 1997.

Similarly, the works of R. Foisel and A. Dury also concern models of interaction and organization.

The approach that we have adopted for the design of multi-agent systems is based on the notion of auto-organization, and it notably also includes the study of their emerging properties. Even if the research community working in this specific sub-domain is even smaller, it is a growing one, and it is developping interestingly, especially through the work being done at Iremia (at the University of Réunion), at IRIT, at the University of Neuchâtel (Switzerland) and in certain other laboratories of the EU (D. Van Parunak, R. Brooks for example).

Some of these researchers have taken inspiration from biological models to envisage the emerging properties. Principally, this current work is inspired by ant-colony models (such as at LIP6, at Lirmm or at the IRIDIA, Brussels). We consider the use of models of spider colonies or societies as an original contribution from us toward this study, it having never been used before. It must be mentioned that this field has been influenced to a considerable extent by the work of Jean-Louis Deneubourg of CENOLI (Brussels) which concerns phenomena involving auto-organization in such colonies and the mechanisms of interaction by pheronomes in ant-colonies.

4. Application Domains

In order to achieve its basic research program, the MAIA team has developed and is developing a strong known-how in sequential or distributed decision making, in using mathematical tools such as Markov decision processes, hidden Markov models or Bayesian Networks, in centralized or distributed learning that is appropriate for the development of real applications such as

- monitoring the hydration state of patients suffering from kidney disease
- elderly fall prevention
- infection skin prevention around the catheter used for peritoneal dialyses
- e-maintenance
- collaborative filtering
- ambient intelligence
- learning mediation strategies for heterogeneous agents cooperation

5. Software

- A planning algorithm has been developed to compute robotic paths along which curvature is continuous, and its derivative with respect to the arc length is bounded. Compared to classical optimal paths made of line segments and circular arcs, it also contains Clothoids (or Cornu's Spiral) to connect the previous two, with a curvature varying linearly (with respect to the arc length) from zero to the maximum. As a consequence, tracking error (distance between actual and predicted positions) when following such a path is more than fifteen times smaller than when following a classical optimal path. While a first prototype of this planner has been develop in 1997, it used an analytical, computational and graphical library set called Leda. As this library was not open, suffered compatibility problems with previous versions, and at last is not free any more, the planner had to be completely redeveloped. A new (open-source) version will be shortly available, and is already presented at http://www.loria.fr/~scheuer/SCPLS.
- MAIAToolbox - It is a toolbox integrating most of the algorithms developed so far by the Maia team members. A graphical tool makes it possible to build a new application by connecting small boxes (representing an elementary algorithm) to each other.

- a license of the Diatelic software and the corresponding patent has been sold to Diatelic Sa a start-up created for the deployement of the dialic system. The Diatelic project address the hydration monitoring of patients suffering from kidney disease. In addition to optimising the transmission of data via intra/internet, the DIATELIC system has a diagnosis system designed to detect anomalies, which can appear progressively and imperceptibly. This system uses a system based on decision theory, which enables the interpretation using basic data (weight, blood pressure, differential blood pressure, orthostatic blood pressure, peritoneal ultra filtration, etc.). Based on such data the system triggers alarms based on thresholds, which are pre-established or calculated using the averages of the previous 15 days. The nephrologists can personalize the system by changing the patient profile to a certain extent. This can be done easily by telling the system when the diagnosis it provides is not relevant to the nephrologists' diagnosis. A learning procedure adapts the patient profile accordingly. the Diatelic project who's the intelligent agent has been designed by our research team leads to the creation of Diatelic SA. An experimentation concerning about 150 patients in east of France is ongoing.

6. New Results

6.1. Resource-Bounded Reasoning in Intelligent Systems

Over the past seven years, a very fruitful research collaboration has been established between the MAIA group and the RBR group at the University of Massachusetts, directed by Prof. Shlomo Zilberstein. The collaboration was conceived at a meeting that took place in 1995 at the International Joint Conference on Artificial Intelligence in Montreal. During this meeting, we identified a high degree of overlap between our interests, research projects, and solution techniques. These common interests relate to the development of planning and monitoring techniques for autonomous systems that can operate in real-time and can cope with uncertainty and limited computational resources. At the time, the U.S. team investigated a solution technique based on "anytime algorithms" and the French team investigated the "progressive processing" model. Since then, we have worked together on both of these models and exploited the synergy to improve their applicability and effectiveness. This year, some work related to this collaboration has been published.

6.1.1. Optimal Sequencing of Contract Algorithms

Participants: S. Zilberstein, F. Charpillet, P. Chassaing.

In [11], we address the problem of building an interruptible real-time system using non-interruptible components. Some artificial intelligence techniques offer a trade off between computation time and quality of results, but their run-time must be determined when they are activated. These techniques, called contract algorithms, introduce a complex scheduling problem when there is uncertainty about the amount of time available for problem-solving. We show how to optimally sequence contract algorithms to create the best possible interruptible system with or without stochastic information about the deadline. These results extend the foundation of real-time problem-solving and provide useful guidance for embedding contract algorithms in applications.

6.1.2. Application of Progressive Processing to Control Planetary Rovers

Participants: A. Mouaddib, F. Charpillet, S. Le Gloannec, S. Zilberstein.

Planetary rovers are small unmanned vehicles equipped with cameras and a variety of scientific sensors. They have proved to be a cost-effective mechanism in space exploration. Recent rover missions by NASA, such as Sojourner's Mars exploration, have suffered from total reliance on ground-based commanding and employed on-board autonomy only to safely follow up linked commands. The inherent uncertainty that characterizes exploration of new environments and the limited communication bandwidth and time delays increase the risk of execution failures and rover downtime. We have addressed this problem using the progressive processing model. Our approach is to equip the rovers with pre-compiled control polices for making fast decisions on such issues as: how to best perform a given task given a set of alternatives; when the quality of the result is satisfactory; how to react to failure; how many times to retry to perform a certain operation; and how to best allocate limited resources to the entire set of activities over a certain window of operation.

The purpose of our work is to enable a Mars rover to take decisions under resource uncertainty. This autonomous agent has to collect some useful information for scientists in limited time. Unfortunately, It's not only limited by time, but also in the amount of energy available, and the disk space where he saves the data collected. Our approach is based on Markov Decision Processes (MDP), which allows to plan and to take decisions under uncertainty. We represent the rover's states by two features: the amount of resources remaining and the quality of the data already collected. The main problem of this approach is the number of generated states. We seek techniques to reduce the complexity of what are known as large MDPs. We are currently working on state aggregation and state space decomposition [22]. So far, we have worked on dynamic progressive processing, in which we assumed that the task structure is known in advance. The agent must be able to choose at any given time from among a set of tasks, which one will be the best for it. We adapted the Dynamic Progressive Processing so that it can handle multiple resources.

6.2. Designing single agent using Markov decision processes and reinforcement learning

6.2.1. Automatic Generation of an Agent's Basic Behaviors

Participants: Olivier Buffet, Alain Dutech, François Charpillet.

We are trying to develop a methodology for the automated design of agents (in the framework of Markov decision processes) in the case where the global task can be decomposed into simpler –possibly concurrent sub-tasks. Our main idea [13] is to automatically combine basic behaviors using RL methods. This led us to propose two complementary mechanisms. The first mechanism builds a global policy using a weighted combination of basic policies (which are reusable), the weights being learned by the agent (using Simulated Annealing in our case). An agent designed this way is highly scalable as, without further refinement of the global behavior, it can automatically combine several instances of the same basic behavior to take into account concurrent occurrences of the same subtask. The second mechanism aims at creating new basic behaviors for combination. It is based on an incremental learning method that builds on the approximate solution obtained through the combination of older behaviors.

6.2.2. Reinforcement Learning in Partially Observable Markov Decision Processes

Participants: Alain Dutech, Manuel Samuelides.

In this work a new algorithm that extends the Reinforcement Learning framework to Partially Observed Markov Decision Processes (POMDP) is presented. The main idea of our method [10]is to build a state extension, called exhaustive observable, which allow us to define a next processes that is Markovian. We prove that solving this new process, to which classical RL methods can be applied, brings an optimal solution to the original POMDP. We apply the algorithm derived from that proof to several examples to test its performance and robustness.

6.2.3. Dynamic decomposition of Markov Decision Processes

Participants: Daniel Szer, François Charpillet.

The core question of this research is the dynamic definition and resolution of sub-goals in MDPs, and this for two reasons. On the one hand we argue, that most complex decision problems can be solved better with a hierarchical approach than with a flat one. Our goal was it to discover such a hierarchy autonomously while learning. On the other hand, solutions to sub-problems which are often called macros, can later be reused in other tasks. Discovery of sub-goals is essentially guided by the discovery of repeated environmental patterns.

We chose two different problems to test our approach [24]. The first one is a robot soccer task, where one robot has to push a ball into a goal area. The second is the so called taxi problem: A taxi has to pick up a person at a certain location and carry her to another place. Useful macros such as go-to-position, pick-up-person-at, push-ball-to, deliver-person-at could be found and reasonably combined in a hierarchical manner.

6.2.4. Representation issues

Participant: Bruno Scherrer.

The aim of this work is to provide a sound framework for addressing a difficult problem: the automatic construction of an autonomous agent's modular architecture. We briefly present two apparently uncorrelated frameworks: Autonomous planning through Markov Decision Processes and Kernel Clustering [7][32][23]. Our fundamental idea is that the former addresses autonomy whereas the latter allows to tackle self-organizing issues. Relying on both frameworks, we show that modular self organization can be formalized as a clustering problem in the space of MDPs. We derive a modular self-organizing algorithm in which an autonomous agent learns to efficiently spread planning problems over m initially blank modules with $m < n$.

6.2.5. Designing mobile robot behaviour using MDP

Participant: Alexis Scheuer.

In order to move autonomously, a mobile robot needs to plan a safe and feasible motion: collision with obstacles must be avoided, and motion constraints (due to the robot's mechanics) have to be respected. Complexity of this problem increases with the density of obstacles, and with the precision of the model of the robot's mechanics. The improvement of computation frequency in today's computer allows to define more precise robot's models: for example, while car-like robots' motions in the 90's corresponded to constant orientation of the steering wheels (optimal motions follow line segments and circular arcs), it is now possible to plan motions corresponding to constant reorientation's velocity of these wheels. This introduces a set of curves (named Clothoids or Cornu's Spiral) connecting line segments and circular arcs, the resulting path being followed more precisely: tracking error (distance between actual and predicted positions) is divided by more than 15.

We are now interested in taking into account the various imprecisions(due to captors and motors) in the search of the best path. While explicit analytic modeling of these imprecisions is very complex and leads to expensive computations, they can easily be represented in probabilist methods, such as Markov Decision Processes(MDP). However, the complexity of the new robot's model increases drastically the computation time of such a method, due to the number of states considered (it is equivalent to e^d, where d is the dimension of the state space, which grows from 3 to 5). To avoid this complexity, a hybrid method was developed during a Master of Research work in 2002: this method is based on a MDP dealing with imprecisions and collision avoidance through a simplified robot's model, whose results are improved analytically with a better robot's model. The main difficulty of this work was the definition of the simplified robot's model: it has to be simple enough to allow a fast computation, while restricting the robot's motion in a way appropriate to the analytical improvement, without limiting too strongly the motion. The resulting algorithm generates in a few seconds a motion which is qualitatively similar to the motion generated in nearly an hour by a MDP using a precise model of the robot.

We are now interested in the tracking of the computed motion. A second Master of Research work proposed in 2003 [30] an innovative analytical approach, using a hierarchy of A* search (with two different levels of precision in the robot's model) to find in real-time (less than 10 ms) the best control value according to the computed motion, the actual obstacles and more precise motion constraints of the robot. Once again, we would like to take also into account the various imprecisions in the tracking process. Analytic approach being still too complex, we consider to search for the best control value using a (possibly hybrid) MDP. To achieve real-time computation, this MDP will not only use the motion computed by the hybrid MDP, but all the state/action/value graph computed during planning, and improve or correct this graph according to a more precise robot's model and a more recent obstacles'map.

This research is currently proposed as a subject for a Master of Research work in 2004.

6.2.6. A medical filtering agent using HMM

Participants: François Charpillet, Laurent Jeanpierre.

Diatelic is in one of the first interactive and intelligent public telemedicine systems applied to the treatment and monitoring of dialysis patients. This system, resulting from the collaboration between LORIA and ALTIR

(Lorrain Association for Renal Failure Treatments), was baptized DIATELIC (Interactive and Cooperative Monitoring of Dialysis Patients). In addition to optimizing the transmission of data via intra/internet, the DIATELIC system has a diagnosis system designed to detect anomalies, which can appear progressively and imperceptibly. This system uses a system based on decision theory (Hidden Markov Model), which enables the interpretation using basic data (weight, blood pressure, differential blood pressure, orthostatic blood pressure, peritoneal ultra filtration, etc.). Based on such data the system triggers alarms based on thresholds, which are pre-established or calculated using the averages of the previous 15 days. The nephrologists can personalize the system by changing the patient profile to a certain extent. This can be done easily by telling the system when the diagnosis it provides is not relevant to the nephrologist' diagnosis. A learning procedure adapts the patient profile accordingly. The Diatelic system is now developed and maintained by a start up Diatelic SA, which is currently deploying the system for all people in Lorraine suffering from kidney disease and treated using peritoneal dialysis.

Several publications [20][21][31] have been made this year on the project. In a first article, we explain how the decision model's semantics can be constrained to facilitate its interaction with physicians and nurses. In particular, we show that the use of fuzzy filters allows for expressing the observation model in an understandable way, from a medical point of view. Simultaneously, this expression can reduce the model's complexity by modeling precisely the target problem while using few parameters.

In a second article, we show how the use of personalized Markov models can enhance the quality of a medical diagnosis system. This kind of models is particularly adapted to the long-term monitoring of chronic pathologies. More precisely, we explain a particular modeling method that allows for enforcing strong medical semantics within each model parameter. This semantics allow physicians for interacting naturally with the system and to customize each patient profile. This customization allows for an even better monitoring of the patients' health level while taking into account patient-specific pathologies, while maintaining a very low diagnosis complexity.

6.3. A Formal Study of Coordination and Control of Collaborative Multi-Agent Systems

This project develops a decision-theoretic framework for planning and control of multi-agent systems by formalizing the problem as decentralized Markov process. It applies to a wide range of application domains in which decision-making must be performed by multiple collaborating agents such as information gathering, distributed sensing, coordination of multiple robots, as well as the operation of complex human organizations.

6.3.1. Coordination through Communication in Multi-agent Reinforcement Learning

Participants: Daniel Szer, François Charpillet.

Our research concentrates on multi-agent coordination techniques in the case of markovian environments. If we take the example of a team of robots that have to accomplish a given task, then we expect the robots to coordinate their behavior in such a way that will optimize the solution of the problem. While such coordination can be easily accomplished with a centralized controller, decentralized coordination is a much more challenging issue. In game theory, where the notion of a "team" does not exist since each player seeks to maximize only its own rewards, the notion of an equilibrium solution has been introduced to characterize a set of strategies such that no player can do better by unilaterally changing its strategy. The same notion of an equilibrium solution has also been recently introduced to the cooperative multi-agent case. Unfortunately these solutions cannot be guaranteed to be optimal in any sense for the team as a whole.

In our work we argue that one form of coordination can be inter-agent communication. More specifically, we seek for learning algorithms that make use of communication in order to provide optimal solutions for the whole team instead of just equilibrium solutions. Surprisingly, this has not been done much yet. The algorithms we develop should furthermore respect some important constraints: Communication usually entails high costs. The optimal solution to the problem should therefore be found using the minimal communication overhead possible. For reasons of computability, we also suppose that the agents are “independent learners”, which means that they consider the effects of their own actions only.

We have proposed an algorithm for deterministic environments that converges to the teams optimal solution using a communication strategy called mutual notification. The algorithm is inspired by the single-agent Q-learning algorithm. Most of the time each agent learns its behavior locally. At some times however, when an agent detects a possible improvement of its current behavior, it notifies all other agents about this detection in order to decide together whether or not it would also constitute a global improvement for the team. If this is the case, the agents is allowed to update its behavior, otherwise it is not. The algorithm has proved to give encouraging results, and we are currently focusing on solutions for the more general case of stochastic environments.

6.3.2. Decentralized control in the pursuit domain

Participants: Iadine Chadès, Bruno Scherrer, Raghav Aras, François Charpillet.

We consider centralized off line planning and decentralized execution for an MAS composed of several agents. The agents attempt to capture a moving prey that they pursue in a discrete world, i.e., a world formed of cells. Our objective was to endow these agents with optimal individual “behavior” to enable them to work together in trapping the prey. For solving this problem using the Markov decision process we have introduced a general iterative heuristic [2][15][14]: at each step one chooses a sub-group of agents and update their policies to optimize the task given the rest of agents have fixed plans. We have analyzed this process in a general purpose and have shown how it can be applied to Markov Decision Processes, Partially Observable Markov Decision Processes and Decentralized Partially Observable Markov Decision Processes.

In an other work [28], we devised a hierarchy of tasks which constitute the main task (namely, capturing the prey). With such a hierarchy, we set up each sub-task as an MDP, and using dynamic programming found an optimal individual policy for each sub-task. The hierarchical division of the task was both physical (division of the discrete world into zones) and conceptual (defining stages of the task or situations in which it could be executed). The resultant collection of MDPs was very weakly coupled. The idea was to make each agents use a policy according to its current situation. Since the agents were homogeneous, they all used the same set of policies. Offline, these individual optimal policies were combined by a central designer after resolving conflicting conditions. Online, each agent, using the set of policies, accomplished a current sub-task choosing an appropriate policy, taking its decision independently of others. Our approach enabled us to address the above problem in a large domain with a variable number of agents.

6.3.3. Multi-Agent Reinforcement Learning

Participants: Alain Dutech, Olivier Buffet, François Charpillet.

In [18], a new reinforcement learning (RL) methodology for the design of reactive multi-agent systems is presented. Although dealing with realistic situated agents with local perception does not belong to the framework where convergence of RL algorithms is guaranteed, in our method each agent learns individually its local behavior. The progressive aspect of learning, which challenges the agents with sub-tasks of growing complexity, allows us to go beyond the classical limitation of RL in this context. Our general framework is validated in a simulated environment where the agents have to coordinate themselves to reach a global goal.

6.3.4. Application to Learning of Mediation Strategies for Heterogeneous Agents Cooperation

Participants: Romaric Charton, Anne Boyer, F. Charpillet.

Our works aim to control adaptive services that are general public, multimedia and interactive. We focused on information retrieval services where users have needs that can be satisfied by an information source. Various interaction problems arise when users are novice or unable to formulate their requests. We defined a coordinator agent to help users formalizing their requests and to provide them with relevant results from the information source. The coordinator has to find the best strategy: a sequence of interactions that optimizes the user’s satisfaction and the information retrieval costs. It has also to manage the uncertainty and the incomplete knowledge related to the other agents and to the media reliability. We modeled the interaction sequence as a Markov Decision Process (MDP) where the coordination strategy is computed as a stochastic policy [16][17].

To define the MDP states, we represent the partial request related to the user goal and the results returned by the information source. Requests and objects are placed in a multidimensional space with an attribute based referential. To reduce the complexity, we abstracted the state space and only kept the binding states of the request attributes and a quantity indicator for the results. The coordinator can ask question to the user to obtain, propose or confirm the value of an attribute. It can also request results from the source and ask the user to select one of them. Since the transition and reward functions of the MDP are unknown, the coordination strategy is obtained with a Q-Learning reinforcement learning method. We use a set of rewards that model the estimated user satisfaction, the coordination length and the costs to get the results. Our approach was experimented on a classical problem of flight booking service. As the learning is unrealistic without a training period, we used a naive user simulator. The results showed that the coordinator is able learn a coordination strategy and to adapt itself to a simulated user. The optimum strategy is reached given the answer generator but becomes difficult to obtain for more complex tasks. Besides, the strategy is rigid compared to human beings diversity. Our future works will be oriented to find a more compact state space that would allow hierarchical decompositions. We also await more realistic strategies from training on real interaction traces that will be collected in the future.

6.4. Natural-inspired approach

6.4.1. A new pattern of specialisation inspired from rats group

Participants: Christine Bourjot, Vincent Chevrier, Vincent Thomas.

Confronting groups of rats to a situation in which they have an increasing difficulty to reach food leads to the emergence of a social structure. Two main profiles appear : supplier composed of rats which access directly to food, and non-carrier rat which access food by stealing it from carrier rats. The regularity of this differentiation and the stability of the phenomenon let biologists assume that an underlying mechanism is responsible for the differentiation.

Hamelin [25], the simulation model we developed, reproduces by this phenomenon and is based on a reactive multi-agent system: each rat is described as a reactive agent and the constraints on the group is coded in the environment. The different behaviours are computed according to a coupling of two existing models: the response thresholds model for modelling direct access to resource and dominance relationships model for modelling stealing attempts. The execution of this simulation leads to a differentiation among initially identical agents. The differentiation results are analogue to the biological observations: the specialisation is characterized by two profiles (supplier and non-carrier) each one corresponding to a specific way to access the resource. We have proposed with this simulation a first approach to organize a community of agents with difficulties to access a resource. This approach is based on an auto-organized process and the agents are very simple but manage nevertheless to coordinate themselves during runtime. The system can be considered as a decentralized collective learning system. The next step is to extract and to generalize the learning principle.

6.4.2. Transposing the principles of collective web weaving in spiders

Participants: Christine Bourjot, Vincent Chevrier, Vincent Thomas.

Thanks to cooperation with biologists we have proposed an original simulation model to reproduce the collective behaviour of weaving in social spiders species. Based on this model of collective weaving, we transposed it to obtain an approach for region detection in gray level images. Furthermore, we showed that this model extends the existing repertoire of swarm mechanisms [8]. The main difference with existing mechanisms is the possibility in the spider model to integrate non-local information in local decision. However, the transposition suffers of some withdraws that avoid the comparison with existing methods in image processing domain. This year, we extend the model by introducing a population regulation mechanism inspired from living ecosystems. A group in charge of the detection of a region is able to self-regulate its population and detects faster the region.

6.4.3. A multi-agent based filtering algorithm for target tracking

Participants: François Charpillet, Vincent Chevrier, Frank Gechter.

This year, we proposed a new method to deal with the localization/tracking issue based on reactive multi-agent systems [19][4]. We designed a model of collective problem solving, applied it and undertook several experiments to assess the work both in term of problem solving process and of application results. The proposed system is able to focus on a variable number of targets moving in a structured and partially known environment and to accomplish its task by using a variable number of sensors, static or mobile in the real world.

The originality of our proposal is to make the relationship between local interactions involving the agents and the expected collective properties through the design of a complete, physically coherent, environment which, both, formalize the problem to resolved and impose the agent's behaviours. Those local behaviours, obtained by combining attraction and repulsion forces, lead, on a collective point of view, to two opposite organizational structures: the agglomeration of the agents on potential targets and a homogeneous distribution of them on target-less areas. We implemented this model using the Madkit toolkit and undertook several experiments with the soccer robots available in the team. Experiments demonstrated the relevance of such an approach and pointed out interesting properties of the problem solving process:

- A variable number of targets can be tracked,
- A variable number of sensors can be used,
- The data fusion process can be achieved reactively through the environment.

In addition, we compared this system with a standard Kalman filtering method. Even if results are not so accurate, our approach requires less knowledge concerning the dynamic of the targets and even concerning the observation process.

8. Other Grants and Activities

8.1. Regional initiatives

8.1.1. TAM ACTION

This action is part of the PRST (Programme de Recherche Scientifique et Technique) "Software Intelligence". The aim of TAM (TéléAssistance Mobile/ remote and mobile assistance) is to develop systems composed of entities that are able of certain autonomy (of perception, decision and or action) and that can achieve collectively or not a given task in a dynamic, noisy and uncertain environment (as it can be in the industrial world). Several research labs of the Lorraine are involved in TAM: CRAN, LORIA and Supelec-Metz.

8.1.2. Telemedecine Action

This action is part of the PRST (Programme de Recherche Scientifique et Technique) "Software Intelligence" in the project TOAI (teleoperation and intelligent assistance). The aim of this action is to develop a basic research program for telemedecine. Telemedicine, like many other emerging technologies, does not yet have any unique definition. Telemedicine is a term indicating, in a general way, the use of an electronic communication network for the transmission of data and information concerning diagnosis or medical treatments. As such, a simple transmission of medical data by fax (lab reports, medical mail, etc.) or by telephone enters into this definition. However, most commonly, the term telemedicine indicates a more systematic deployment (in the medical field) of communication networks (both hardware and software components), including eventually the Internet domain, for the transmission of data by several means (e.g., images, sounds, files, numeric data). From this point of view, the field of telemedicine remains quite vast and encompasses specific medical activities. We therefore suggest classifying various applications of telemedicine not in function of the type of data and their mode of transmission (as it has often been done in the past), but rather in function of the potential users.

- Health-care practitioners' information exchange. In this case, the quality and integrity of the information transmitted is the major issue. For a precise and rapid long distance analysis it is essential to be able to provide the proper software and hardware infrastructure enabling a friendly interaction of medical experts with information made available by the machinery support. Telemedicine services between professionals is developing rapidly and can be applied to a large variety of domains such as the videoconference, tele-training, hospital-home networks, tele-experts, tele-diagnosis, etc.
- Information exchange between patients and health-care practitioners. The heterogeneity of the population concerned and the large number of potential users require easily accessible (also in terms of costs) and user-friendly systems, which can be adapted to patients who might not have experience and familiarity with the deployment of software and hardware facilities. Although, this class of applications and services has been traditionally neglected, in the last few years, the situation has been evolving quite rapidly. Patient empowerment aims indeed at building an infrastructure for better involvement of patients in their own health care.

The focus of the our action work will indeed consider this latter category of "telemedicine" services. In particular, we will work on defining and providing an effective approach for proactive and adaptive treatment and monitoring of chronic patients. The existing systems mentioned in the literature, essentially enable the passive transmission of medical data from the patients to the physicians allow different medical parameters to be monitored without requiring any particular action from the patient. In order to move toward a more active, proactive and interactive role of patients these systems need to embed new facilities enabling personalization of the offered services in a flexible and adaptive way. The idea is to (1) facilitate the access to existing information, (2) improve the capabilities (functionality of available systems) to process such information, and (3) increase the chances of a better coordination between patients and practitioners.

8.2. National initiatives

8.2.1. RNTS PROJECT PARACHUTE

Participants: Loic Pélissier, J. Semoune, François Charpillet, Anne Boyer.

Parachute is a RNTS project dedicated to the falls prevention by elderly people. The partners involved in this project are either research departments (Université Technologique de Troyes or the MAIA team from Loria), medical institution (CHU Nancy-Brabois, Institut Régional de Réadaptation, Institut de Myologie de Paris, Laboratoire de Physiologie neuromusculaire) or industrial partners (ACDM Concept -Nancy-, Application Electronique de Champagne) with the endorsement of the Région Champagne Ardennes, Téfal, Association Française contre les Myopathies and the CCAS of Nancy

Falls by elderly people are a major problem of health services because of their frequency and their medical and social consequences. In France, about two millions of people over sixty five year's old fall every year. Falls are responsible for an important morbidity (50 000 fractures of the thighbone upper extremity each year in France), a mortality evaluated to 13 000 death per year (it is more than the death due to road accidents). As falls are a major risk, it is a key challenge to design systems able to detect falls as fast as possible in order to alert the relevant assistance service. Our task in MAIA is to create a non invasive monitoring of the elderly subject's locomotive behavior in his/her environment in order to isolate parameters which values will serve as an alert signal. Our approach of an intelligent monitoring for fall prevention is based on an ecological analysis of the walk. It will thus be a question of in-place walk's telemonitoring. MAIA is then working on the development of a behavioral model based on the observations made via camera. The model relies on a notion of signature (set of features characterizing the walk behavior).

A first feasibility study was led in LORIA. About twenty people were filmed walking in a corridor by a single standard CCD camera with a fixed position. We used a standard algorithm for motion detection. From these data, we determined some indicators like the trajectory on the ground, the mean speed, the immediate speed, the acceleration, the height and the orientation of the person. This set of parameters, used as a first estimation of a signature, gave us encouraging results as for the precision and the reproducibility of the calculated parameters.

The project began on October 2003, with a two years duration. Reference: Loic Pelissier. DRT Université Nancy 2.

8.2.2. RNTS DEPIC project

Participants: J.P. Thomesse, François Charpillet.

The partners involved in the project are: LPM Insa Lyon, LORIA, ALTIR, Nancy, CALYDIAL, Lyon, CHU Nancy, CHU Lyon.

This project aims at designing a new sensor for the prevention of skin infection around the catheter used for peritoneal dialysis. The sensor is under design by the micro system team (André Dittmar and George Delhomme) of the *Laboratoire de physique de la matière de l'INSA de Lyon*. It is made up of 19 temperature sensors put on a circular that is placed on the skin of the patient for a measure. The idea is that a particular signature could be characteristic of infection. Our team is involved in the interpretation of the signal delivered by the sensor.

8.2.3. INRIA Development Action DialHémo: Telemedicine System for Patients undergoing CAPD

Participants: Cédric Rose, Cherif Smaili, J. Semoune, François Charpillet, Anne Boyer.

The goal of this project is to develop a telemonitoring and diagnosis aid system for the patients undergoing continuous "hémodialyse" at home, or at the hospital. The input data to the system are returned by the dialysis machine (duration of the session, volume of blood) or by other sensors (weight, blood pressure). The results of data analysis will be available to nephrologists. Organizations which combine their efforts in the DialHemo project are : the research project MAIA at LORIA, the DIATELIC society, technology company stemming from LORIA, partners of the medical and health world in Lorraine (CHU Nancy, ALTIR), the GAMBRO society, world leader on the market of dialysis machines.

Developments follow three research themes :

- Definition of a model for the dynamics of hydration of a patient taking into account orders of dialysis given by the doctor and the evolution of physiological parameters of the patient. A model will be built by machine learning from a data base reference enriched along the project : experimentation is necessary to identify useful parameters and critical values (medical follow-up of the patients), and to validate the human interface ; it will be refined and adapted to the medical practice of the regular doctor. The alarm system mainly relies on the diagnosis aid system.
- Connexion of the diagnosis aid system to dialysis machines.
- Publication of the results, attempt of standardization and exploitation outside Lorraine. The evolution of the software will allow to optimize protocols of follow-up for each patient/nephrologist.

The project began in september 2003 and is planned for 2 years.

8.3. European initiatives

8.3.1. KVM

Participants: Anne Boyer, Vincent Chevrier.

The Eureka project KVM (Knowledge Valorisation Matrix) gathers two companies Kappa Ice (Paris) and HTTP Software (London) and the INRIA through three research projects of the LORIA (MAIA, ECOO and Orpailleur). Its aim is to enhance the value of the knowledge of an enterprise thanks to a tool based on a specific knowledge cartography. To achieve this, the KVM software includes a tool for representing the enterprise knowledge together with a representation methodology according to three referentials; and a tool for strategic diagnosis and decision. MAIA concentrates this year on providing tools to automate reasoning and help for strategic decision and diagnosis. Work steps first focused on the formalization of the audit process and the identification of needs, then on the determination of the reasoning model(s) and the subparts of the audit process that can be automated. This lead to the development of a prototype that focused on the diagnosis from the k-elts based on an expert system to represent knowledge. This project ended in March 2003.

8.3.2. The Ozone Project (IST-2000-30026)

Participant: Alain Dutech.

The partners of the project are Philips Electronics Nederland B.V, T-Systems Nova, Berkom, IMEC, Philips Research France (PRF), EPICTOID, Technical University of Eindhoven, INRIA, Thomson Multimedia, LORIA (équipes Parole, Langue et Dialogue, MAIA)

This project is entering its final phase with the specification and the implementation of many demonstrators. The MAIA team is more specifically involved in two demonstrators evolving around an electric automated Cybercar. In these demonstrators, users must interact with the Ozone system by the way of different modality and we are developing an intelligent Modality Adviser that select the best modality according to the user, the context, the quality of the media, the application, etc. Our work is based on stochastic methods with adaptive capabilities (thought learning) and, for comparison, rule based methods. The final demonstrator is expected in April 2004.

8.3.3. IST project ELIN(IST 2000-30188)

Participants: Régis Lhoste, Anne Boyer.

The IST 2000-30188 project Elin (Electronic Newspaper)) gathers two media companies - Diari Segre from Spain and Corren from Sweden - and several organizations such as Alamo (Spain), FhG (Germany), FZI (Germany), LIU (Sweden), UPC (Spain) and INRIA through the MAIA research team. ELIN objective is to improve current electronic newspapers by enhancing consumer's experience by introducing interactive features, advanced personalization and ubiquity of use and by improve the productivity in the publishing chain by introducing better tools and practices on content authoring and management. MAIA is in charge of the definition and the development of a smart agency based on a Jade Agent Platform to perform user adaptation through collaborative filtering.

8.3.4. ITEA PROTEUS project

Participants: François Charpillet, Alain Dutech [WP2 leader], Loic Pélissier.

Partners: AKN, Arc Informatique, BiKit, Cegelec, Til technologies, Ifak, IML Fraunhofer Institut, Lab, Lifc, Loria, Pertinence, Schneider Electric, LIP6, Vartec.

The ambition of PROTEUS is to engineer a change in the landscape of today's maintenance support tools. The project will provide a fully integrated platform that is able to support any broad e-maintenance strategy. Maintenance is considered an integral part of global EAO policies (Enterprise Asset Optimisation) currently being implemented by a growing number of industrial organizations. Predictive maintenance requires the harmonious integration of:

1. Continuous remote monitoring of equipment throughout its lifetime.
2. Maintenance and Repair Operation (MRO) process management, grouping logistic actions to improve the efficiency of remote access to technical documentation and knowledge repositories, on-line use of modeling packages, decision help tools and human experts.
3. Comprehensive data presentation and synthesis involving direct delivery of operational information, including supervisory and decision level, asset management panel and maintenance contract management. Within PROTEUS, we are developing a European generic software architecture for web-based e-maintenance centers, targeting the transportation, energy and other industries. We aim to improve efficiency by bringing expertise via Internet directly to the user site. We will promote a de-facto form of standardization through extensive use of new data-structuring technologies (XML

- Extensible Mark-up Language), application integration techniques and Internet related technologies. This will reduce maintenance process costs (time to diagnosis and duration of intervention), and also prevent failures through early monitoring of field equipment (condition-based predictive maintenance).

8.4. International initiatives

8.4.1. Self-Directed Cooperative Planetary Rovers

Sponsored by NASA Aerospace Technology Enterprise

Shlomo Zilberstein, Principal Investigator

Co-Investigators: Eric Hansen, Victor Lesser, and Rich Washington

Collaborators: Francois Charpillet, and Abdel-Illah Mouaddib

Research Assistants: Raphen Becker, Daniel Bernstein, Max Horstmann, and Zhengzhu Feng

This project focuses on the question of how to best utilize the rover's resources in the face of the above difficulties. Our approach is to equip the rovers with pre-compiled control polices for making fast decisions on such issues as: how to best perform a given task given a set of alternatives; when the quality of the result is satisfactory; how to react to failure; how many times to retry to perform a certain operation; and how to best allocate limited resources to the entire set of activities over a certain window of operation. To achieve these goals we are developing and evaluating several fundamental technologies, focused on the basic need to carefully manage the limited computational resources, power, and communication capabilities of the rover.

8.4.2. A Formal Study of Coordination and Control of Collaborative Multi-Agent Systems

Sponsored by National Science Foundation

Shlomo Zilberstein, Principal Investigator

Co-Investigators: Victor Lesser

Collaborators: Francois Charpillet, and Abdel-Illah Mouaddib

Research Assistants: Raphen Becker, Daniel Bernstein, and Zhengzhu Feng Postdoctoral Fellow: Claudia Goldman

This project develops a decision-theoretic framework for planning and control of multi-agent systems by formalizing the problem as decentralized Markov process. Three approaches to communication are studied based on (1) a cost/benefit analysis of the amount of communication, (2) search in policy space, and (3) transformations of the more tractable centralized policies into decentralized policies.

8.4.3. Developing a decision-theoretic framework for planning and control of collaborative multi-agent systems : INRIA associated team

This year the collaboration between the MAIA group at INRIA, directed by Dr. François Charpillet, and the RBR group at the University of Massachusetts, directed by Prof. Shlomo Zilberstein has been funded by INRIA as a associated team. This association of the two research teams has been focused on the development of a decision-theoretic framework for planning and control of collaborative multi-agent systems by formalizing the problem as decentralized control of a Markov process. The overall goal is to develop sophisticated coordination strategies that stand on a formal footing. This enables us to better understand the strengths and limitations of existing heuristic approaches to coordination and, more importantly, to develop new approaches based on these more formal underpinnings. There is a wide range of application domains in which decision-making must be performed by a number of distributed agents that are trying to achieve a common goal. This includes information-gathering agents, distributed sensing, coordination of multiple distributed robots, decentralized control of a power grid, autonomous space exploration systems, as well as the operation of complex human organizations. These domains require the development of a strategy for each decision maker assuming that decision makers will have limited ability to communicate when they execute their strategies, and therefore will have different knowledge about the global situation.

8.5. Visiting scientists

- Brahim Chaib head of DAMAS group from Laval university at Quebec visited our group for two months (November, December).
- Shlomo Zilberstein, Dan Bernstein and Sherief Abdallah from the University of Massachusetts at Amherst came each for a week

9. Dissemination

9.1. Leadership within scientific community

- François Charpillet is a member of the committee of the RTP 11: Information et intelligence: *raisonner et décider* of the STIC department at CNRS.
- Jean-Paul Haton is the Chairman of ASTI, Association française des sciences et technologies de l'information.
- Maia is still a leading force in the PDMIA group (Processus Decisionnels de Markov et Intelligence Artificielle) and took a great part in the annual meeting of the group. This year, the group held its annual meeting in Caen. This meeting took the form of a very informal workshop with scientific presentations and many open discussions. The most interesting talks and presentation are in the process of being compiled into a special issue of the french electronic journal JEDAI. The next meeting will most probably take place next year in Paris.

9.2. Conference and workshop committees, invited conferences

9.2.1. Workshop organization and invited conferences

- We organized at Nancy the 6th European Workshop on Reinforcement Learning. This workshop is an important venue for the European research community interested in this subject. The workshop was held in Nancy on September 4-5, 2003. More than forty participants - all active members of the reinforcement learning research community - attended the meeting. Two members of the associated team from the University of Massachusetts attended as well. During the 2-day workshop, participants presented and discussed more than 20 presentations of high-quality scientific work. Most subjects of current research in Reinforcement Learning where covered. Interaction between participants was a key factor contributing to the success of the workshop, which concluded with a broad open discussion. Participants agreed that another session of the workshop should be organized within two years.
- invited conference *Colloque Seniors et Nouvelles technologies, Autonomie et place dans la cité*, Nancy, François Charpillet *Parachute, un système de prévention des chutes*.

9.2.2. Program committees, editorial board

- François Charpillet has been member of the program committee of the *Journées Francophones sur les Systèmes Multi-Agents* and of the *Secondes Journées Francophones Modèles Formels de l'Interaction*,
- Alain Dutech and François Charpillet have been members of the program committee of the 6th European Workshop on Reinforcement Learning.
- François Charpillet has been a reviewer of the International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August, 2003.
- Vincent Chevrier has been reviewer for the RFIA 2004 conference
- Jean-Paul Haton is member of the editorial board of the following journals: Journal of Intelligent Manufacturing, Engineering Intelligent Systems, Revue I3.

9.3. PHD and HDR reviewing

- François Charpillet has been Ph.D. reviewer of Stéphane Cardon of the University of Artois. The thesis was entitled *Résolution de processus décisionnels de Markov totalement observables de grandes tailles*.
- François Charpillet has been Ph.D. reviewer of Hosam Hanna of the University of Caen. The thesis was entitled *La planification des tâches et des ressources sous incertitude dans un système multi-agent*.
- Vincent Chevrier has been Ph.D. reviewer of Fenintsoa Andriamasinoro from Université de La Réunion. The thesis was entitled *Proposition d'un modèle d'agents hybrides basé sur la motivation naturelle*.
- Vincent Chevrier has been reviewer for an HDR from Université de La Réunion. The work was entitled *Simulation multi-agents : architecture, ingénierie et expérimentations*.
- Vincent Chevrier has been member of the Ph.D. commitee of Stamatios Nicolis from Cenoli - Université Libre de Belgique. The thesis was entitled *Dynamique du recrutement alimentaire et de l'agrégation chez les insectes sociaux*.
- Jean-Paul Haton has been reviewer for the PHD or HDR of H. Harb (Ecole Centrale de Lyon), N. Arous (ENIT-Tunis), D. BenAyed (ENIT-Tunis), Y. Mami (ENST-Paris), F. Rouai (ENSI-Tunis), F. Bimbot (Université de Rennes), F. Leber (Université Henri-Poincaré).

9.4. Award

Laurent Jeanpierre received the "Best Regional PHD" award for his work on the Diatelic Project.

9.5. Responsibilities

- Anne Boyer is *chargée de mission* for the "global information system" of Université Nancy 2.
- Anne Boyer is responsible of the INRIA action called Dialhemo.
- Anne Boyer is member of the *commission de spécialistes*, University Nancy 2.
- Christine Bourjot is member of the scientific council of the *Reseau Grand Est des Sciences Cognitives*.
- François Charpillet and Vincent Chevrier are members of the steering commitee of the research theme "TeleOperation and Intelligent Assistant" in the PRST "Software Intelligence".
- Vincent Chevrier is member of the laboratory council, of the steering committee of the project *rencontre avec les STIC* of INRIA, of the editorial board of the collection "Exposés scientifiques multimedia" of INRIA, and of the editorial board of the "Lettre du LORIA".
- Vincent Chevrier is the moderator of the mailing list of the French spoken community on multi-agent systems.
- Jean-Paul Haton is a member of the "Commission de spécialistes" of the university of Metz.
- A. Scheuer represents MAIA team at LORIA's CoMIn (Commission des Moyens Informatiques), commission between computer users and services.
  He is also a member of LORIA's CHS (*Commission Hygiène et Sécurité*), commission concerned by the prevention of work accidents. Finally, he is currently member of the *commission de spécialistes* of computer science at Université Henri Poincaré (Nancy I).

10. Bibliography

Doctoral dissertations and "Habilitation" theses

[1] O. BUFFET. *Une double approche modulaire de l'apprentissage par renforcement pour des agents intelligents adaptatifs.* Thèse d'université, Université Henri Poincaré - Nancy 1, Septembre, 2003, http://www.loria.fr/publications/2003/A03-T-341/A03-T-341.ps.

[2] I. CHADÈS. *Planification distribuée dans les systèmes multi-agents à l'aide de processus décisionnels de Markov.* Thèse d'université, Henri Poicaré Nancy 1, Jan, 2003.

[3] R. CHARTON. *Des agents intelligents dans un environnement de communication multimédia : vers la conception de services adaptatifs.* Thèse d'université, Université Henri-Poincaré, Dec, 2003.

[4] F. GECHTER. *Modèle multi-agent réactif pour la fusion de percepts : application à la localisation et au suivi en robotique mobile.* Thèse d'université, UHP NANCY I, Oct, 2003.

[5] R. LHOSTE. *Archimaj, une plateforme multi-agents pour la recherche d'information.* Diplome de recherche technologique, Université Nancy 2, système d'information des organisations, Nov, 2003.

[6] L. PELISSIER. *Les nouvelles technologies au service des personnes agées en perte d'autonomie.* Diplome de recherche technologique, Université Nancy 2, système d'information des organisations, Dec, 2003.

[7] B. SCHERRER. *Apprentissage de représentation et auto-organisation modulaire pour un agent autonome.* Thèse d'université, Université Henri Poincaré - Nancy I, Jan, 2003, http://www.loria.fr/publications/2003/A03-T-023/A03-T-023.ps.

Articles in referred journals and book chapters

[8] C. BOURJOT, V. CHEVRIER, V. THOMAS. *A new swarm mechanism based on social spiders colonies : from web weaving to region detection.* in « Web Intelligence and Agent Systems : An International Journal - WIAS », number 1, volume 1, Mar, 2003, pages 47-64.

[9] A. DUTECH, F. GARCIA. in « Proceedings of the Sixth European Workshop on Reinforcement Learning », Sep, 2003.

[10] A. DUTECH, M. SAMUELIDES. *Apprentissage par renforcement pour les processus décisionnels de Markov partiellement observés..* in « Revue d'Intelligence Artificielle », number 4, volume 17, Sep, 2003, pages 559-589.

[11] S. ZILBERSTEIN, F. CHARPILLET, P. CHASSAING. *Optimal Sequencing of Contract Algorithms.* in « Annals of Mathematics and Artificial Intelligence », number 1, volume 39, Sep, 2003, pages 1-18.

Publications in Conferences and Workshops

[12] O. BUFFET, A. DUTECH. *A Self-Made Agent Based on Action-Selection.* in « Sixth European Workshop on Reinforcement Learning (EWRL-6), Nancy, France », pages 47-48, Sep, 2003.

[13] O. BUFFET, A. DUTECH, F. CHARPILLET. *Automatic Generation of an Agent's Basic Behaviors.* in « Second International Joint Conference on Autonomous Agents and Multi-Agent Systems - AAMAS'03, Melbourne, Victoria, Australie », ACM press, W. ROSENSCHEIN, YOKOO, editors, pages 875-882, Jul, 2003.

[14] I. CHADÈS, F. CHARPILLET. *Modèle de conception de SMA coopératifs par planification réactive.* in « Secondes Journées Francophones Modèles Formels de l'Interaction - MFI'2003, Lille, France », May, 2003.

[15] I. CHADÈS, B. SCHERRER, F. CHARPILLET. *Planning Cooperative Homogeneous Multiagent System Using Markov Decision Processes.* in « 5th International Conference on Enterprise Information Systems - ICEIS 2003, Angers, France », Kluwer, Apr, 2003.

[16] R. CHARTON, A. BOYER, F. CHARPILLET. *Faire coopérer des agents hétérogènes par apprentissage de médiation.* in « Secondes Journées Francophones Modèles Formels de l'Interaction - MFI'03, Lille, France », Cépaduès-Editions, A. HERTZIG, B. C. DRAA, P. MATHIEU, editors, pages 61-70, May, 2003, http://www.loria.fr/publications/2003/A03-R-064/A03-R-064.ps.

[17] R. CHARTON, A. BOYER, F. CHARPILLET. *Learning of Mediation Strategies for Heterogeneous Agents Cooperation.* in « 15th IEEE International Conference on Tools with Artificial Intelligence - ICTAI'2003, Sacramento, Californie, USA », IEEE Computer Society, B. WERNER, editor, pages 330-337, Nov, 2003.

[18] A. DUTECH, O. BUFFET, F. CHARPILLET. *Apprentissage par renforcement pour la conception de Systèmes Multi-Agents Réactifs.* in « Journées Francophones sur les Systèmes Multi-Agents 2003 - JFSMA'03, Hammamet, Tunise », Hermès, pages 219-231, Nov, 2003.

[19] F. GECHTER, V. CHEVRIER, F. CHARPILLET. *Une architecture réactive pour la localisation en robotique mobile.* in « Journées Francophones des Systèmes Multi-Agents - JFSMA'2003, Hammamet, Tunisie », HERMÈS, editor, Nov, 2003.

[20] L. JEANPIERRE, F. CHARPILLET. *Remote monitoring of chronical pathologies using personalized Markov models.* in « First International Conference on Computational Intelligence in Medicine and Healthcare 2003 - CIMED 2003, Sheffield, Angleterre », Barrie W Jervis, Jul, 2003.

[21] L. JEANPIERRE. *Diagnostic automatique par des modèles stochastiques flous. Application à la surveillance de maladies chroniques.* in « 23ème Séminaire de la Société Francophone de Biologie Théorique - SFBT'2003, St-Flour, France », Jun, 2003, Communication orale.

[22] S. LE GLOANNEC. *Planification et raisonnement progressif.* in « Troisièmes Journées Nationales sur Processus Décisionnel de Markov et Intelligence Artificielle - PDMIA'03, Caen, France », Jun, 2003.

[23] B. SCHERRER. *Modular self-organization for a long-living autonomous agent.* in « Eighteenth International Joint Conference on Artificial Intelligence - IJCAI'03, Acapulco, Mexico », Aug, 2003.

[24] D. SZER. *Définition autonome de sous-problèmes dans l'apprentissage par renforcement.* in « Troisièmes Journées Nationales sur Processus Décisionnel de Markov et Intelligence Artificielle - PDMIA'03, Caen, France », Jun, 2003.

[25] V. THOMAS, C. BOURJOT, V. CHEVRIER. *Du collectif pour la satisfaction individuelle : un modèle d'auto-*

organisation. in « Journées Francophones sur les Systèmes Multi-Agent - JFSMA'2003, Hammamet, Tunisie », pages 261-265, Nov, 2003.

[26] V. THOMAS. *Apprentissage collectif par renforcement d'inspiration biologique ? Le système hamelin.* in « Troisièmes Journées Nationales sur Processus Décisionnel de Markov et Intelligence Artificielle - PDM&IA 2003, Caen, France », Jun, 2003.

[27] J.-P. THOMESSE, F. CHARPILLET. *Diatélic et ParaChute, deux systèmes de télémédecine.* in « Seniors et Nouvelles Technologies, Nancy, France », Nov, 2003.

Internal Reports

[28] R. ARAS. *Decentralized control in the pursuit domain.* Stage de DEA, LORIA INRIA-Lorraine, Jul, 2003.

[29] O. BUFFET, A. DUTECH, F. CHARPILLET. *Etude de différentes combinaisons de comportements adaptatives au sein d'un agent.*. Rapport de recherche, LORIA, Nov, 2003.

[30] J.-P. DRON. *Exécution de mouvement en robotique mobile.* Stage de DEA, Sep, 2003.

[31] L. JEANPIERRE. *Automated Medical Diagnosis With Stochastic Models : Monitoring Chronic Diseases.* Rapport de recherche, Oct, 2003.

[32] B. SCHERRER. *Modular self-organization for a long-living autonomous agent.* Rapport technique, Apr, 2003, http://www.loria.fr/publications/2003/A03-R-053/A03-R-053.ps.

[33] V. THOMAS. *Les relations de dominance : de la biologie aux systèmes markoviens.* Comité de thèse, MAIA, Dec, 2003.

Miscellaneous

[34] L. JEANPIERRE. *Cours de réseau - DEUST AMMILOR 2002.* Jul, 2003, 10 présentations PowerPoint 2002..