

Multi-Agent Coordination Off-Line: Structure and Complexity

Carmel Domshlak* and Yefim Dinitz

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva, Israel

Abstract

Coordination between processing entities is one of the most widely studied areas in multi-agent planning research. Recently, efforts have been made to understand the formal computational issues of this important area. In this paper, we make a step toward this direction, and analyze a certain class of coordination problems for dependent agents with independent goals acting in the same environment. We assume that a state-transition description of each agent is given, and that preconditioning an agent's transitions by the states of other agents is the only considered kind of inter-agent dependence. Off-line coordination between the agents is considered. We analyze some structural properties of these problems, and investigate the relationship between these properties and the complexity of coordination in this domain. We show that our general problem is provably intractable, but some significant subclasses are in NP and even polynomial.

Introduction

Coordination between processing entities is one of the most widely studied areas in multi-agent planning research. Recently, efforts have been made to understand the formal computational issues in this important area. In this paper, we make an additional step toward this direction. We describe a certain class of coordination problems for agents with independent goals acting in the same environment, that is the strategy of each agent should be planned while taking into account strategies of other agents as planning constraints. In these problems, (i) the set of feasible plans for each agent is defined by a state-transition graph, and (ii) preconditioning an agent's *transitions* by the *states* of other agents is the only considered kind of inter-agent dependence. For this problem class, *off-line* coordination is considered; structural and computational properties are investigated, and both tractable and intractable subclasses are presented. Although the examined class of problems can be seen as restricted, its place in multi-agent systems was already discussed in AI literature, e.g. (Brooks 1986;

Dudek *et al.* 1996; Moses & Tennenholtz 1995). We believe that the formal model we suggest for the presented problem class can serve as a basis for further extensions toward a representation of richer multi-agent worlds yet preserving convenience for computational analysis.

In the area of multi-agent coordination, we identify two main research directions: *multi-agent collaboration* and *synthesizing multi-agent plans*. In multi-agent collaboration, goal achievement is a team activity, i.e., multiple agents collaborate towards the achievement of an overarching goal. During the last decade, various theories for multi-agent collaboration, e.g. (Cohen & Levesque 1991; Grosz & Kraus 1996; Kinny *et al.* 1994), and several generic approaches for corresponding planning, e.g. (Decker & Lesser 1998; Smith 1994), were proposed; for an overview of this area see (Grosz 1995). In addition, a few formal computational results concerning the coalition formation are found for both benevolent (Shehory & Kraus 1996) and autonomous (d'Inverno, Luck, & Wooldridge 1997) agents.

Synthesizing multi-agent plans addresses synchronization between agents that work on independent goals in the same environment. In (Georgeff 1983), centralized merging of preconstructed individual agent plans is presented, and it is based on introducing synchronization actions. For the purpose of efficiency, this work deals only with the problems in which agents are sufficiently decoupled and thus actions of one agent are largely independent of any other agent. In (Ephrati & Rosenschein 1994), a heuristic approach that exploits decomposed search space is introduced. Whereas in (Ephrati & Rosenschein 1994; Georgeff 1983) methods for merging plans at the primitive level are described, in (Clement & Durfee 1999) a strategy for plan merge on different levels of hierarchical plans is presented. In (Alami, Ingrand, & Qutub 1998; Alami *et al.* 1995), a generic approach for distributed, incremental merging of individual plans is suggested: the agents are considered sequentially, and the plan for the currently processed agent is synchronized with the plans chosen for the previously processed agents. Besides, a few more algorithmic approaches to an application-specific plan synchronization are proposed in (Cao, Fukunaga, & Kahng 1997; Pape 1990; Yuta & Premvuti 1992). Note that all these approaches for inter-agent synchronization assume that the personal plans are already constructed, and that the merging

*Currently with the Faculty of Industrial Engineering and Management, Technion, Israel.
Copyright © 2014, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

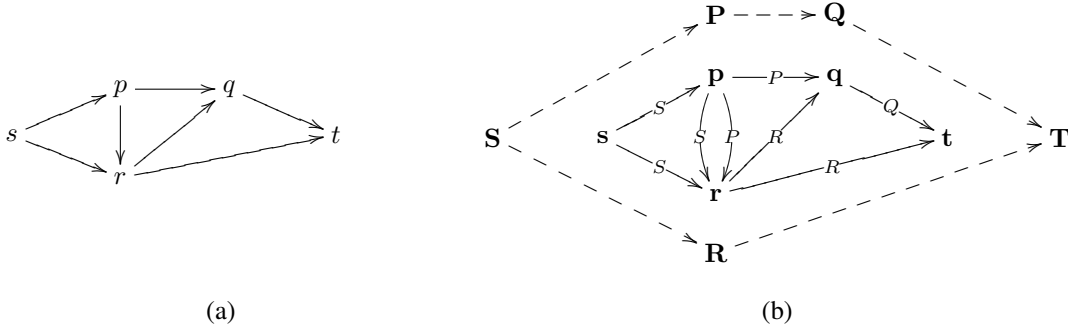


Figure 1: Examples of strategy graphs for: (a) single agent, (b) two-agent group: transitions of X depend on states of Y .

operation is considered only on this fixed set of individual plans (possibly, certain local modifications of the plans may be allowed). This essentially restricts not only the quality, but even the ability to find a coordinated plan.

A wider formal model for multi-agent systems that concerns both determining personal agent plans and their synchronization, is suggested in (Moses & Tennenholtz 1995). With respect to this model, the cooperative goal achievement problem is defined. It can be roughly stated as follows: Given a set of benevolent agents in the same environment, each one with its own, independent goal, does there exist a satisfying system of coordinated personal plans? In (Moses & Tennenholtz 1995), this problem is shown to be PSPACE-complete. Detailed investigation of specific subclasses of the suggested problem can become a natural continuation of this work.

In this paper, we concentrate on a particular class of the cooperative goal achievement problem. In this problem class, which we denote to as STS (state-transition support), an action (transition) of an agent can be constrained by the states of other agents. Following (Brooks 1986; Moses & Tennenholtz 1995), each agent is assumed to be representable as a finite automaton. As in (Clement & Duffee 1999; Ephrati & Rosenschein 1994; Georgeff 1983), the coordination process is assumed to be off-line and centralized. Note that we are not dealing here with many complementary issues of multi-agent systems such as determining agent behavior, collaboration establishing, decision process distributing, etc.

We suggest to represent an STS problem domain using a graph language. The two interrelated graphical structures of our model are as follows:

1. The *strategy (di)graph*: the state-transition graph describing behavior of an agent.
2. The *dependence (di)graph*: description of the dependence between the agents.

The strategy graph compactly captures all alternative personal plans for the corresponding agent: its strategies correspond to the *paths* from the source node to the target node in its strategy graph. Transition of an agent along an edge is conditioned by the *states* of the agents on which it depends, i.e., their locations at certain nodes in their strategy graphs. The reader can find some relation between our

strategy graphs and the interactive automata of (McCarthy & Hayes 1969). The main differences are that we (i) address preconditioning, not interaction between agents, and (ii) consider off-line planning.

The nodes of the dependence graph correspond to the agents. An edge (p, q) appears in the dependence graph if some action of the agent q depends on the state of the agent p . We define subclasses of STS, mainly with respect to structural properties of the dependence graph, and analyze their complexity. The complexity is measured in terms of the total size of the strategy graphs and the size of the dependence graph (which is actually the number of agents). We show that our general problem is provably intractable, but some significant subclasses are in NP and even polynomial. As far as we know, such an analysis of the relation between the structural and the computational properties of a multi-agent coordination problem was not done in previous research. Note that this paper presents only first results of our ongoing research; hence, in particular, we do not discuss relative significance of the presented problem classes, and we do not provide wide final conclusions.

The rest of the paper is organized as follows: Section presents and discusses our model. Section is devoted to the basic case of two agents. Section introduces a classification of STS's subclasses, and presents additional complexity results. For the proofs and some of detailed algorithms we refer to the technical report (Domshlak & Dinitz 2001). Concluding remarks are given in Section , together with some discussion of related issues and future work.

Graphical Model for STS

In this section we define two graphical structures of the STS problem domain. The *strategy graph* G^A captures the alternative personal strategies of an agent A . The agent's states are represented by the nodes of G^A . Each (directed) edge represents a possible transition between the two corresponding states. If a transition is conditioned by a state of some "supporting" agent (or states of a few such agents), then the corresponding edge has a *label* which describes the condition. When a transition can be done under several alternative conditions, this is modeled by multiple edges, each labeled by a single condition. There are two distinguished nodes in G^A , the source node and the target node,

which represent the initial and goal states of A , respectively. The possible agent's strategies are represented by the paths from the source node to the target node. For an illustration see Fig. 1(a), where the possible strategies are: $spqt$, $sprt$, $spqrt$, srt , $srqt$.

The size of this structure is linear in the length of the problem description. Observe that in general, the number of different, potentially applicable, strategies may be exponential in the size of the problem even for an acyclic structure; for a general structure, this number can, formally, be infinite.

Consider a pair of agents X and Y so that X depends on Y in the following sense: each transition of X is preconditioned by a certain state of Y . In graph terms, each edge in G^X is labeled by the name of a node in G^Y . An example is given in Fig. 1(b), where the inner graph models X and the outer one models Y . Note that there are two transitions of X from the state p to the state r which are preconditioned by different states of Y . The following pair of strategies for X and Y is coordinated: The agent X begins from state s and Y begins from state S . Subsequently, (i) X moves to p , (ii) Y moves to P , while X is in p , (iii) X moves to q , (iv) Y moves to Q , (v) X moves to t , and finally, (iv) Y moves to T . We denote this coordinated multi-agent strategy by the sequence $\langle (s,p)(S,P)(p,q)(P,Q)(q,t)(Q,T) \rangle$. The other possible coordinated plans for X and Y are $\langle (s,p)(p,r)(S,R)(r,t)(R,T) \rangle$ and $\langle (s,r)(S,R)(r,t)(R,T) \rangle$. No other coordinated pair of strategies exists. For example, if the agents begin with $\langle (s,p)(S,P)(p,r) \rangle$, then X cannot leave r , since it is impossible for Y to reach R after being in P .

In general, there are several agents A_1, A_2, \dots, A_n with strategy graphs $G^i = G^{A_i}$, $1 \leq i \leq n$. We say that an agent A_i depends on an agent A_j if transitions of A_i have states of A_j as preconditions. Note that a few agents may depend on the same agent, and an agent may depend on a few other agents. Hence, a transition of A_i is preconditioned by states of agents that A_i depends on; in general, a transition can depend on a part of them.

The *dependence graph* \mathcal{G} is a digraph whose nodes are A_i , $1 \leq i \leq n$, and there is an edge from A_j to A_i if agent A_i depends on agent A_j . In what following, we identify nodes of \mathcal{G} with the corresponding agents. Likewise, we identify the nodes and edges of G^i with the states and transitions of A_i , respectively. An edge of G^i has a compound label, each component of which is a state of an immediate predecessor of A_i in \mathcal{G} (i.e., a node in the strategy graph of this predecessor).

Both the strategy graph and the dependence graph can be constructed straightforwardly, given an STS problem. Likewise, the structural properties of the dependence graph, which are investigated in this paper, can be verified in low polynomial time in the size of the graph.

In this paper, *each* transition of an agent is preconditioned by the states of *all* agents that this agent depends on, if the opposite is not declared explicitly. We consider only connected dependence graphs. The reason is that otherwise the set of agents can be divided into disjoint subsets that can be considered independently.

A natural motivation for a pair of agents X and Y as

above is that X fulfills some mission, while Y supports it. In practice, the mission of X may be a movement towards a target, while Y either watches X from certain observation points, or feeds it from certain warehouses, or protects it from some fortified positions. Such a support is usual in military and is wide-spread in various other activities. For example, a worker X assembles some compound product, while a mobile crane Y moves the (half-assembled) product from one working place of X to another. In general, when there are several agents, each agent is supported by a few agents, and in turn, supports a few other agents.

Having read this far the reader may argue that it is possible that we never know an agent's behavior in complete detail, and even if we do, the number of states that represent an agent's behavior may be very large. In this case, exploiting and even constructing strategy graphs seems to be infeasible.

Indeed, it is not realistic to present detailed models of many real-life systems as finite state machines. However, as it is argued in (Moses & Tennenholtz 1995), the representation an agent uses need not present the world in sufficient detail to require more expressive description language. Intuitively, it is possible to distinguish between the physical state of agents and their computational state. It is the computational state that we have difficulty in modeling by finite-state machines. However, the computational state of an agent is not accessible to the other agents. On the other hand, a finite-state description can serve well as a representation of the physical state. A wider discussion of these and other issues of finite-automata-based knowledge representation and reasoning can be found in (Moses & Tennenholtz 1995). Likewise, number of finite-automata-based architectures are discussed in the AI literature, e.g., see (Brooks 1986).

In case that the number of parameters describing the physical state (and thus the size of the state space) is still large, it is often possible to take advantage of the structure of an agent in order to derive a concise description of its state. In particular, an agent may be viewed as being composed of a number of largely independent components, each with its own state. Each such component (such as robot's hand, motor, wheel, etc.) can be represented by an agent, whose physical state is relatively simple, since it describes a particular feature of a complex system. In this case, the strategy graphs of the agents are expected to be of reasonable size.

The Basic Case of Two Agents

In the rest of the paper, we study the complexity of various multi-agent coordination problems based on the model suggested above. In this section we demonstrate our techniques on solving the basic case of agent X depending on agent Y . Given the strategy graphs of X and Y , we build a new graph describing X restricted by Y . As a result, our problem is reduced to the known problem of finding a source-to-target path in this new graph. Such a reduction becomes possible because of the locality of the original restrictions.

We describe a solution of our problem as an interleaved sequence of transitions of agents X and Y , as was shown in Section . Let us analyze the structure of a general coordinated plan Π for these agents. As mentioned above, private

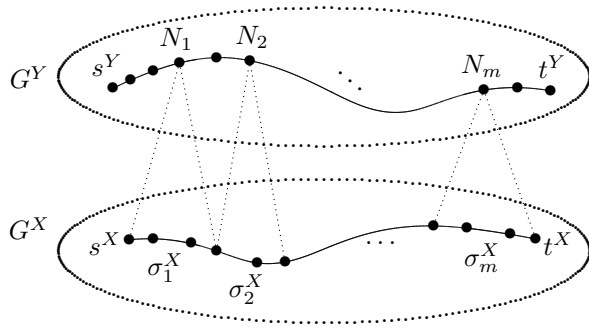


Figure 2: Relation between personal strategies of X and Y

strategies for X and Y are source-to-target paths in G^X and G^Y respectively, and we denote them by σ^X and σ^Y . Let us divide σ^X into intervals $\sigma_1^X, \sigma_2^X, \dots, \sigma_m^X$ of constancy of labels: all edges (transitions) in σ_i^X are labeled (preconditioned) by the same node (state) N_i of Y . The Y 's path σ^Y is also divided by the nodes N_i into some parts. The considered plan is as follows: (1) If $N_1 \neq s^Y$, then agent Y moves from s^Y to state N_1 along the initial part of σ^Y ; (2) Agent X moves along σ_1^X to its end; (3) Y moves from N_1 to N_2 along the next part of σ^Y ; (4) X moves along σ_2^X , etc. Finally, if $N_m \neq t^Y$, then agent Y moves from N_m to t^Y along the final part of σ^Y . For illustration see Figure 2.

Let see what “degrees of freedom” Π has. For simplicity of notation, we denote s^Y by N_0 and t^Y by N_{m+1} . Observe that the replacement of part of σ^Y between N_i to N_{i+1} , $0 \leq i \leq m$, by any other path from N_i to N_{i+1} retains the plan feasible. That is, the only requirement for feasibility of σ^X is the *existence* of such a path for each label change along σ^X . Moreover, replacement of any σ_i^X by another path between its ends such that all its edges are labeled by N_i retains the plan feasible.

Lets now consider building such a coordinated plan Π . We take as a central issue finding an appropriate path σ^X . As mentioned above, the only critical points in its building are the moments of label change along σ^X : if the label changes from N_i to N_{i+1} , then a “connecting” path from N_i to N_{i+1} in G^Y must exist. In order to include into the same framework the cases in which the label of the first edge is not s^Y and the label of the last edge is not t^Y , we extend G^X as follows. We add to G^X a dummy source \tilde{s}^X , with an edge (\tilde{s}^X, s^X) labeled s^Y , and a dummy target \tilde{t}^X , with an edge (t^X, \tilde{t}^X) labeled t^Y , and consider path $\tilde{\sigma}^X$ from \tilde{s}^X to \tilde{t}^X in this extended graph \tilde{G}^X . It is easy to see that existence of connecting paths in G^Y for all critical points in $\tilde{\sigma}^X$ is necessary and sufficient for the feasibility σ^X . Indeed, the twin path σ^Y is the concatenation of all connecting paths, for $0 \leq i \leq m$. Therefore, the only information we need in order to build an appropriate path in \tilde{G}^X is whether two edges may be consequent on a path $\tilde{\sigma}^X$ in the above sense, for all pairs of adjacent edges in \tilde{G}^X .

Information on existence of a path between any pair of nodes N' and N'' of G^Y (reachability of N'' from N') can be obtained by a preprocessing of G^Y . Appropriate

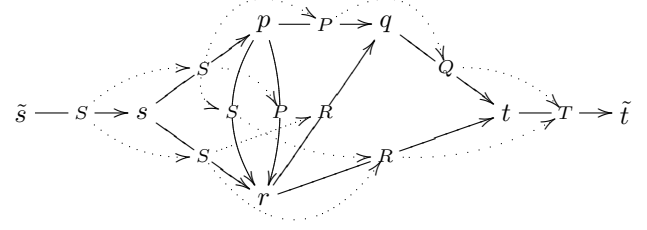


Figure 3: Illustration of \tilde{G}^X with the corresponding $G^{X|Y}$.

algorithms are widely known and are very fast; the paths themselves, if any, are provided by these algorithms as well, e.g. see (Cormen, Leiserson, & Rivest 1990). Hence, we can form the following database for all pairs of consequent edges in \tilde{G}^X : If an edge e' labeled N' enters some node and an edge e'' labeled N'' leaves the same node, then the pair (e', e'') is marked “permitted” if either $N' = N''$ or N'' is reachable from N' in G^Y . For example see Figure 3, where dotted lines show the permitted edge pairs.

In fact, the suggested construction provides a reduction of the problem of finding a path $\tilde{\sigma}^X$ as defined above to a known problem of finding a path in a graph. Let the new graph $G^{X|Y}$ have the edges of \tilde{G}^X as nodes, and let its edges be defined by the permitted pairs of edges in \tilde{G}^X .¹ Now, all we need is a path from the source node (\tilde{s}^X, s^X) to the target node (t^X, \tilde{t}^X) in $G^{X|Y}$. Clearly, such a path defines the corresponding path in \tilde{G}^X immediately.

Figure 4 summarizes the algorithm for finding a coordinated plan for agent X depending on (“supported by”) agent Y . This algorithm is, evidently, polynomial. The reader can “execute” it by himself on the example given in Figure 3. This example is the same as the one in Section , and any of the three feasible coordinated plans listed there can be obtained as the result of such an execution. Thus we have achieved our first result, and in what following, we use the presented approach for analyzing other variants of STS.

Theorem 1 *There exists a polynomial time algorithm that, given a pair of agents X, Y , in which X depends on Y , determines existence of a coordinated plan and finds such a plan, if exists.*

This result can be easily generalized to finding an *optimal* coordinated plan. Assume that the cost function is the sum or the product of weights of states and/or transitions used in the plan. Since in step 4 of the algorithm, an arbitrary path can be chosen, let us choose an optimal path; this suffices for global optimality. Any algorithm for finding a cheapest path in a graph can be used, e.g. see (Cormen, Leiserson, & Rivest 1990).

Now let us require, in addition, *simplicity* of Y 's strategy, i.e., let us forbid Y to visit the same state twice. A motivation for this can be that each state of Y relates to some

¹Such a construction is a variant of the so called “edge graph” known in graph theory; the addition in our case is the exclusion of non-permitted edges from it.

1. Extending strategy graph G^X to \tilde{G}^X .
2. Preprocessing of strategy graph G^Y : finding a path from any node to any other node (in fact, it is sufficient to do this for the node pairs needed in the next phase only).
3. Constructing permitted-edge graph $G^{X|Y}$.
4. Finding a source-to-target path in $G^{X|Y}$, if any, or reporting on non-existence of a coordinated plan, otherwise.
5. Reconstructing from the path found in phase 4 a path $\tilde{\sigma}^X$ in \tilde{G}^X and its abridged variant σ^X in G^X .
6. Building a path σ^Y in G^Y as the concatenation of paths found at phase 2, for all label changes along $\tilde{\sigma}^X$.
7. Outputting properly interleaved strategies σ^X and σ^Y .

Figure 4: Algorithm for the basic STS problem.

consumable resources. It turns out that this seemingly non-essential change in the problem definition, change the problem to be NP-complete. Informally, the reason for this complexity worsening is that the locality of precondition is thus broken. This result is mainly interesting from the theoretical point of view, since it points on a computational sensitivity of our original problem.

Results for Some Other Subclasses of STS

In this section we discuss various additional complexity results for STS. Problem classes are defined mainly by the structure of their dependence graph. Moreover, both results and methods depend crucially on the number and size of possible agent strategies: if both of them are polynomial in the size of problem domain, then we distinguish such a class as *bounded*. Figure 5 summarizes the achieved results, and each box in the figure denotes an STS's subclass.

The notation is as follows: First, each STS subclass is denoted with respect to the form of the dependence graph: \mathbb{C} stands for *directed chain*, \mathbb{F}^\wedge for *directed fork* (graph with exactly one node with a non-zero outdegree), \mathbb{F}^\vee for *directed inverse fork* (graph with exactly one node with a non-zero indegree), \mathbb{P} for *polytree* (digraph whose underlying undirected graph is a tree), and \mathbb{S} for *singly-connected DAG* (directed acyclic graph with at most one directed path between any pair of nodes). Second, the subscript denotes the number of agents: k stands for a constant bound, and n indicates no bounds.

By default, we assume that each transition of an agent is preconditioned by the states of *all* agents that this agent depends on. Alternatively, π in a superscript denotes a possibility of *partial* dependence. σ in a superscript denotes the requirement that all strategies chosen are *simple*. Note that cases \mathbb{C}_2 and \mathbb{C}_2^σ have been discussed in Section .

Coordinating General Agents A natural generalization \mathbb{C}_n of the class \mathbb{C}_2 (as in Section) occurs when there is a dependence chain: agents A_i , $1 \leq i \leq n$, with A_i dependent on A_{i-1} , $i \geq 2$. For simplicity of notation, let us abridge A_i to i in superscripts (e.g., G^i instead of G^{A_i}).

Our approach for \mathbb{C}_n is to iterate the analysis and the algorithm for \mathbb{C}_2 along the dependence chain. A naive scheme is as follows. For processing pair A_3, A_2 , we need a database on *coordinated* reachability in G^2 . Observe that our algo-

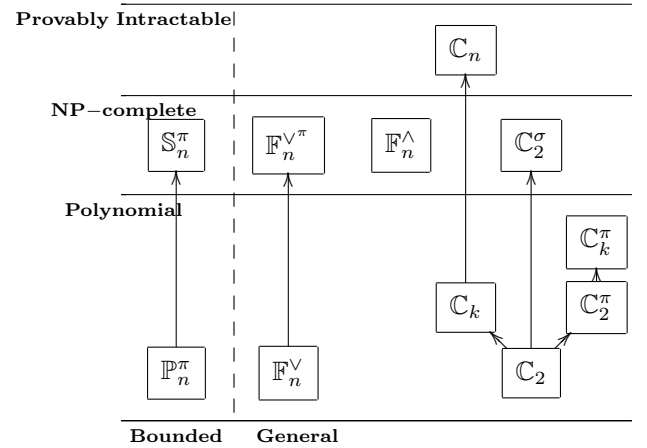


Figure 5: Complexity results for STS

rithm for \mathbb{C}_2 applied to pair A_2, A_1 provides such an information for nodes s^2, t^2 ; clearly, this can be done for any pair of nodes of G^2 . Considering pair A_2, A_1 , let's check reachability and store paths, if any, for *all* pairs of nodes in G^2 , as required. Now, we are as if sure that if we would determine a coordinated plan for pair A_3, A_2 based on the above reachability relation for A_2 , a corresponding strategy for A_1 will be obtained easily. Hence, the situation with A_3, A_2 is similar to that with A_2, A_1 , and thus we can iterate up to A_n . Personal coordinated strategies can be consequently reconstructed, from A_{n-1} to A_1 .

However, this scheme turns to be inconsistent. The point is that, in fact, the coordinated reachability relation in G^2 is on its *edges*, not on its nodes, and it is captured by the edge graph $G^{2|1}$. Therefore, informally, the right approach is to create an edge graph for G^3 basing on $G^{2|1}$, not on G^2 . This idea of recursive edge graph creating can be generalized for dependence chain of arbitrary length. Using this method, we never arrive at a deadend as in the naive scheme, and we can obtain any coordinated plan for this problem. For the detailed algorithm for \mathbb{C}_n , together with an illustrating example of its execution and a counterexample for the naive approach we refer to the technical report (Domshlak & Dinitz 2001).

Main

Repeat *Pruning* while strategy sets are changing.
 If all strategy sets are empty, then report that no coordinated plan exists.
 Otherwise, perform *Construction*.

Pruning

For all edges (A_j, A_i) in \mathcal{G} do:
 For each $\sigma^i \in S^i$ do:
 If *Compatibility-Check* (σ_i, σ_j) returns false for all $\sigma^j \in S^j$ then remove σ^i from S^i .
 For each $\sigma^j \in S^j$ do:
 If *Compatibility-Check* (σ_i, σ_j) returns false for all $\sigma^i \in S^i$ then remove σ^j from S^j .

Compatibility-Check (σ_i, σ_j)

If $l(\sigma^i)[j]$ with neighboring repetitions removed is a subsequence of σ^j ,
 then return true, else return false.

Construction

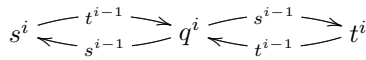
Pick any agent A and any strategy for it.
 Traverse \mathcal{G} undirectly from the node A . For each visited node A' , choose any strategy of A'
 that is compatible with the strategy chosen for the node from which we came to A' .

Figure 6: Algorithm for \mathbb{P}_n^π .

Theorem 2 (I) *There exists an algorithm for \mathbb{C}_n that determines existence of a coordinated plan and finds such a plan, if exists. This algorithm is polynomial in the total size of strategy graphs and exponential in n ; thus \mathbb{C}_k is proved to be polynomial.*

(II) \mathbb{C}_n has instances with exponentially sized minimal solutions.

The proof of the exponential lower bound is by the following example. Each agent A_i , $2 \leq i \leq n$, has the strategy graph as follows:



The strategy graph for A_1 looks the same, except that its edges are unlabeled. This problem instance has a unique minimal coordinated plan of total length $2^{n+1} - 2$ transitions.

By showing that \mathbb{C}_n is provably intractable, theorem 2 emphasizes our motivation for exploiting various structural restrictions on STS in order to find problem classes which are polynomial, e.g., \mathbb{C}_k , or at least belong to NP. In particular, the subclass of \mathbb{C}_n strategy graphs of which are acyclic can be easily shown to belong to NP. However, the question of its exact hardness is still an open question. The same question is open for the decision version of \mathbb{C}_n .

Theorem 3 summarizes the complexity results for some other classes of STS (for notations see Section).

Theorem 3 (I) *There exists a polynomial time algorithm for \mathbb{F}_n^\vee that determines existence of a coordinated plan and finds such a plan, if exists. (II) There exists a polynomial time algorithm for \mathbb{C}_2^π that determines existence of a coordinated plan and finds such a plan, if exists. (III) \mathbb{F}_n^\wedge is NP-complete. (IV) $\mathbb{F}_n^{\vee\pi}$ is NP-complete.*

Proof sketch. (I) In \mathbb{F}_n^\vee , a single agent A_1 depends on a group of independent supporting agents A_2, \dots, A_n . The

algorithm is similar to that for \mathbb{C}_2 . The idea is that several preconditions on the same transition can be checked independently. (II) \mathbb{C}_2^π is an extension of \mathbb{C}_2 in which only *some* of transitions of the supported agent depend on the state of the supporting agent. The algorithm for \mathbb{C}_2^π is similar to that for \mathbb{C}_2 , except that a certain modification of the permitted-edge graph $G^{X|Y}$ is used. (III),(IV) In \mathbb{F}_n^\wedge , a single agent A_1 supports a group of agents A_2, \dots, A_n . The membership in NP can be easily shown for both \mathbb{F}_n^\wedge and $\mathbb{F}_n^{\vee\pi}$. In turn, the proofs of hardness for \mathbb{F}_n^\wedge and $\mathbb{F}_n^{\vee\pi}$ are by polynomial reductions from 3-SAT and PATH WITH FORBIDDEN PAIRS problems, respectively.

Remarks: The algorithm for \mathbb{C}_n can be extended for \mathbb{C}_n^π , similarly to the extension of the \mathbb{C}_2 algorithm for \mathbb{C}_2^π . Results presented for the cases of chain and inverse fork dependence graphs can be generalized to the case of dependence graph being a tree directed to its root.

Coordinating Bounded Agents In this section we consider the class of problems for which the number and size of strategies for each agent is polynomial in the size of the problem domain. In what follows, such agents are referred as *bounded*. Let us denote by S^i the set of allowed strategies σ^i of A_i (i.e., source-to-target paths in G^i). We show that if a dependence graph forms a polytree then STS for bounded agents is polynomial. However, its extension to singly connected directed graphs is already NP-complete.

Consider \mathbb{P}_n^π . Notice that any polytree is an acyclic graph. Denote by $l(\sigma^i)$ the sequence of edge labels along σ^i , and by $l_j(\sigma^i)[j]$ the projection of $l(\sigma^i)$ to the nodes of G^j . Figure 6 presents our algorithm for \mathbb{P}_n^π .

Theorem 4 (I) *There exists a polynomial time algorithm for \mathbb{P}_n^π with bounded agents that determines existence of a coordinated plan and finds such a plan, if exists. (II) \mathbb{S}_n^π for bounded agents is NP-complete.*

Proof sketch. (I) According to the analysis in Section , the condition of *Compatibility-Check* confirms that the checked pair σ^i, σ^j is coordinated. Therefore, convergence of the pruning process ensures that for each agent A : (i) any one of its strategies has at least one supporting strategy of each predecessor of A in \mathcal{G} , and (ii) any one of its strategies supports at least one strategy for any successor of A in \mathcal{G} . Hence, if *Construction* has a personal strategy to begin with, then it is surely successful. Polynomiality holds, since each execution of *Pruning* decreases the total number of personal strategies.

(II) Membership in NP for \mathbb{S}_n^π is straightforward, and the hardness proof is by a polynomial reduction from 3-SAT.

Discussion and Conclusions

In this paper, we concerned coordination for a set of agents that work on independent goals in the same environment. We investigated a particular family of the problems of finding coordinated plans, where actions of an agent depend on the local states of certain other agents. First, we presented a graphical representation model for this family of problems. Then, we classified these problems mainly according to the structural properties of the model. In the main part of the paper, we analyzed the computational properties of various problem classes. We gave polynomial solutions for several classes, while for some other we proved their NP-completeness or even provable intractability. A number of unsolved problems remains for the further research. The suggested approach—to find a sufficiently formal description for some problem classes, and to analyze their computational properties—seems to be prospective in the multi-agent area. Following some additional observations with respect to results presented in this paper.

First, let each agent and its states be considered as a multi-valued variable and its values, respectively, and let the set of agents' transitions be considered as a set of operators over the above variables. In this context, our results concern complexity analysis in the area of classical planning over multi-valued variables (Bäckström & Nebel 1995; Jonsson & Bäckström 1998). Specifically, our results address the planning problems over multi-valued variables and only unary (= single effect) operators. Complexity analysis for planning with unary operators, and *binary* variables can be found in (Domshlak & Brafman 2001).

Second, in our model, each agent is assumed to be assigned to a personal goal, which is independent of the personal goals of all other agents. However, a supporting agent may have no explicit personal goal, while supporting activities of some other agents may be its only destiny. This particular relaxation can be immediately added into the presented model, with no negative impact on the computational properties of the problems. The coordinated group of dependent agents can be viewed as collaborating towards the achievement of some global goal.

Third, in this work we address only groups of fully controlled entities. Therefore, in particular, our results do not concern state-transition settings where agent's strategy depends on uncontrolled environment, like that studied in (Arkin & Balch 1998; Balch *et al.* 1995).

In the future, we plan to continue with analysis of various classes of STS. In addition, we want to examine other forms of dependence between the agents, and other forms of goal(s) definition for a group of agents. This issues will be examined in the context of their impact on the complexity of coordination. We also plan to address related optimization problems.

References

- Alami, R.; Robert, F.; Ingrand, F.; and Suzuki, S. 1995. Multi-robot cooperation through incremental plan-merging. In *International Conference on Robotics and Automation*.
- Alami, R.; Ingrand, F.; and Qutub, S. 1998. A scheme for coordinating multi-robot planning activities and plans execution. In *Proceedings of the 13th European Conference on Artificial Intelligence – ECAI*.
- Arkin, R., and Balch, T. 1998. Cooperative multi-agent robotic systems. In *Artificial Intelligence and Mobile Robots*. MIT Press.
- Bäckström, C., and Nebel, B. 1995. Complexity results for SAS⁺ planning. *Computational Intelligence* 11(4):625–655.
- Balch, T.; Boone, G.; Collins, T.; Forbes, H.; MacKenzie, D.; and Santamaria, J. 1995. Io, Ganymede and Callisto – A multiagent robot trash-collecting team. *AI Magazine* 16(2):39–53.
- Brooks, R. A. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* 2(1):14–23.
- Cao, Y.; Fukunaga, A. S.; and Kahng, A. B. 1997. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots* 4.
- Clement, B. J., and Durfee, E. H. 1999. Top-down search for coordinating the hierarchical plans of multiple agents. In *Proceedings of the Third International Conference on Autonomous Agents*, 252–259.
- Cohen, P., and Levesque, H. 1991. Confirmations and joint actions. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 951–957.
- Cormen, T. H.; Leiserson, C. E.; and Rivest, R. L. 1990. *Introduction to Algorithms*. MIT Press.
- Decker, K., and Lesser, V. 1998. Designing a family of coordination algorithms. In Huhns, M., and M. Singh., eds., *Readings in Agents*. Morgan Kaufmann. 450–457.
- d’Inverno, M.; Luck, M.; and Wooldridge, M. 1997. Cooperation structures. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 600–605.
- Domshlak, C., and Brafman, R. I. 2001. Structure and complexity in planning with unary operators. Technical Report CS-01-05, Dept. of Computer Science, Ben-Gurion Univ.
- Domshlak, C., and Dinitz, Y. 2001. Multi-agent off-line coordination: Structure and complexity. Technical Report CS-01-04, Dept. of Computer Science, Ben-Gurion Univ.

- Dudek, G.; Jenkin, M.; Milios, E.; and Wilkes, D. 1996. A taxonomy for multi-agent robotics. *Autonomous Robots* 3(4):375–397.
- Ephrati, E., and Rosenschein, J. S. 1994. Divide and conquer in multi-agent planning. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 375–380.
- Georgeff, M. 1983. Communication and interaction in multi-agent planning. In *Proceedings of the Third National Conference on Artificial Intelligence*, 125–129.
- Grosz, B., and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence* 86(2):269–357.
- Grosz, B. 1995. Collaborative systems. *AI Magazine* 17(2):67–85.
- Jonsson, P., and Bäckström, C. 1998. State-variable planning under structural restrictions: Algorithms and complexity. *Artificial Intelligence* 100(1–2):125–176.
- Kinny, D. N.; Ljungberg, M.; Rao, A. S.; Sonenberg, E. S.; Tidhar, G.; and Werner, E. 1994. Planned team activity. In *Artificial Social Systems*, volume 830 of *LNCS*. 227–256.
- McCarthy, J., and Hayes, P. 1969. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence* 6.
- Moses, Y., and Tennenholtz, M. 1995. Multi-entity models. *Machine Intelligence* 14:63–88.
- Pape, C. L. 1990. A combination of centralized and distributed methods for multi-agent planning and scheduling. In *Proceedings of the IEEE International Conference on Robotics and Automation – ICRA*, 488–493.
- Shehory, O., and Kraus, S. 1996. Formation of overlapping coalitions for precedence-ordered task execution among autonomous agents. In *Proceedings of the Second International Conference on Multi-Agent Systems*, 330–337. AAAI Press / MIT Press.
- Smith, R. C. 1994. The contract net protocol: High-level communication and control in a distributed problem solving. *IEEE Transactions on Computers* 29(12).
- Yuta, S., and Premvuti, S. 1992. Coordinating autonomous and centralized decision making to achieve cooperative behaviors between multiple mobile robots. In *Proceedings of International Conference on Intelligent Robots and Systems*, 1566–1574.