

# Structured models for multi-agent interactions\*

Daphne Koller                      Brian Milch  
Computer Science Department  
Stanford University  
Stanford, CA 94305-9010  
{*koller,milch*}@cs.stanford.edu

## Abstract

The traditional representations of games using the extensive form or the strategic (normal) form obscure much of the structure that is present in real-world games. In this paper, we propose a new representation language for general multi-player noncooperative games — *multi-agent influence diagrams (MAIDs)*. This representation extends graphical models for probability distributions to a multi-agent decision-making context. The basic elements in the MAID representation are *variables* rather than *strategies* (as in the normal form) or *events* (as in the extensive form). They can thus explicitly encode structure involving the dependence relationships among variables. As a consequence, we can define a notion of *strategic relevance* of one decision variable to another:  $D'$  is *strategically relevant* to  $D$  if, to optimize the decision rule at  $D$ , the decision maker needs to take into consideration the decision rule at  $D'$ . We provide a sound and complete graphical criterion for determining strategic relevance. We then show how strategic relevance can be used to detect structure in games, allowing large games to be broken up into a set of interacting smaller games, which can be solved in sequence. We show that this decomposition can lead to substantial savings in the computational cost of finding Nash equilibria in these games.

## 1 Introduction

Game theory provides a mathematical framework for determining what behavior is rational for agents interacting with each other in a partially observable environment. However, the traditional representations of games are primarily designed to be amenable to abstract mathematical formulation and analysis. As a consequence, the standard game representations, both the normal (matrix) form and the extensive (game tree) form, obscure certain important structure that is often present in real-world scenarios — the decomposition of the situation into chance and decision *variables*, and the dependence relationships between these variables. In this paper, we provide a representation that captures this type of structure. We also show that capturing this structure explicitly has several advantages, both in our ability to analyze the game in novel ways, and in our ability to compute Nash equilibria efficiently.

Our representation is based on the framework of *probabilistic graphical models* [19], used in many disciplines, including statistics and artificial intelligence. Probabilistic graphical models represent the world via a set of *variables*, that take on values in some (discrete or continuous) space. For example, in a simple economic model, we might have a continuous variable for each of several possible goods, indicating its supply at a given time. We might also have a discrete variable representing the amount of rainfall in a region over the last year (e.g., taking the values “drought”, “low”, “normal”, or “high”). Each possible state (or trajectory) of the world is then an assignment of values to these variables. By representing the world in terms of these variables, we can make statements about the relationships between them. For example, we might know that the supply of oranges depends on the rainfall variable,

---

\*This paper is an extended version of a paper titled “Multi-Agent Influence Diagrams for Representing and Solving Games which appears in the Proceedings of the International Joint Conference on Artificial Intelligence, 2001.

whereas the supply of oil does not. The graphical model represents this structure using a directed graph structure, where the nodes represent the variables, and the edges represent the direct dependence of one variable on another. As we discuss in Section 2, these graphs, called *Bayesian networks* or *probabilistic networks* have clear and formal semantics as a representation of a probability distribution over the state space defined by the variables. Furthermore, the graph structure itself makes explicit certain important aspects of the probability distribution, such as the conditional independence properties of the variables in the distribution.

*Influence diagrams* [7] extend Bayesian networks to the decision-theoretic setting, where an agent has to make decisions in accordance with his preferences. In addition to chance variables, influence diagrams contain *decision variables*, which are variables whose value the agent selects as part of his strategy, and *utility variables*, that represent the agent’s preferences.

In this paper, we define *multi-agent influence diagrams (MAIDs)*, which represent decision problems involving multiple agents. We show that MAIDs have clearly defined semantics as noncooperative games, and can be reduced to an equivalent sequence form or normal form game, albeit at the cost of obscuring the variable-level interaction structure that the MAID makes explicit. MAIDs allow us to represent complex games in a natural way, whose size is no larger than that of the extensive form, but which can be exponentially more compact.

We show that MAIDs allow us to define a qualitative notion of dependence between *decision variables*. We define a notion of *strategic relevance*: a decision variable  $D$  strategically relies on another decision variable  $D'$  when, to optimize the decision rule in  $D$ , the decision-making agent needs to take into consideration the decision rule in  $D'$ . This notion provides new insight about the relationships between the agents’ decisions in a strategic interaction. We provide a graph-based criterion, which we call *s-reachability*, for determining strategic relevance based purely on the graph structure. We also provide a polynomial time algorithm, which considers only the graph structure, for computing s-reachability.

The notion of strategic relevance allows us to define a data structure that we call the *relevance graph* — a directed graph that indicates when one decision variable in the MAID relies on another. This graph can be used to provide a natural decomposition of a complex game into interacting fragments, and provide an algorithm that finds equilibria for these subgames in a way that is guaranteed to produce a global equilibrium for the entire game. As a special case, the algorithm generalizes the standard backward induction algorithm for game trees, showing a general subclass of games where backward induction can be applied, even in some games that are not perfect information. We show that our algorithm can be exponentially more efficient than an application of standard game-theoretic solution algorithms, including the more efficient solution algorithms of [20, 10] that work directly on the game tree.

The remainder of this paper is structured as follows. In Section 2, we review some of the key concepts in the framework of probabilistic graphical models that underlie our work. In Section 3, we present the framework of multi-agent influence diagrams, and in Section 4, we relate it to the standard game-theoretic concepts. In Section 5, we define the notion of strategic relevance, and provide a criterion for determining strategic relevance from the graph structure. In Section 6 we show how to exploit strategic relevance to break up a game into smaller games, and compute equilibria more effectively. We conclude in Section 7 with a discussion of some extensions, including additional structure that can be induced from the MAID representation.

## 2 Bayesian Networks

Our work builds on the framework of *Bayesian networks* (also known as probabilistic networks or belief networks) [19] and on its decision-theoretic extension, *Influence diagrams* [7]. In this section, we briefly review the Bayesian network framework, setting up much of the necessary foundations for the remainder of this paper.

A Bayesian network is a graphical representation of a distribution over the joint probability space defined by a set of variables. More precisely, consider a set of variables  $X_1, \dots, X_n$ , where each  $X_i$  takes on values in some finite set  $dom(X_i)$ . Together, the variables define a cross-product space  $\times_{i=1}^n dom(X_i)$ . Our goal is to represent a joint distribution over this joint space. We use  $\mathcal{X}$  to refer to the set of variables  $X_1, \dots, X_n$ , and  $dom(\mathcal{X})$  to refer to their joint domain.

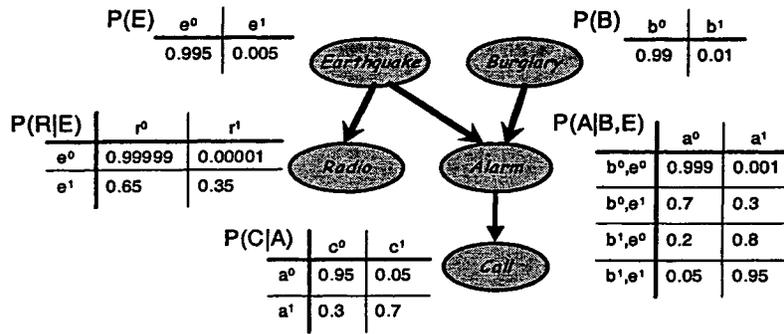


Figure 1: Alarm BN; for these binary variables, we use  $x^1$  as shorthand for  $X = 1$  and  $x^0$  for  $X = 0$ .

The number of states in the joint distribution grows exponentially with the number of variables: If we have  $n$  variables that take on binary values, the total number of states is  $2^n$ . Even for fairly small  $n$ , the representation of this distribution as a list of numbers, one for each state, is impractical.

A Bayesian network (BN) represents the distribution using a graph, whose nodes represent the random variables and whose edges represent (in a very formal sense) direct influence of one variable on another. More precisely, we have the following definition for the Bayesian network representation:

**Definition 1** A Bayesian network  $\mathcal{B}$  over the variables  $X_1, \dots, X_n$  is a pair  $(G, \Pr)$ .  $G$  is a directed acyclic graph with  $n$  nodes labeled  $X_1, \dots, X_n$ . For a node  $X$ , we use  $Pa(X)$  to denote the parents of  $X$  in the graph, i.e., those nodes  $Y$  such that there is a directed edge from  $Y$  to  $X$ .  $\Pr$  is a mapping that associates with each node  $X$  a conditional probability distribution (CPD)  $\Pr(X | Pa(X))$ , which specifies a probability distribution  $\Pr(X | pa)$  over the values of  $X$  for each instantiation  $pa$  of  $Pa(X)$ .

**Example 1** Consider a scenario where we are trying to reason about an alarm installed in a house. The alarm can go off either because of a burglary or because of a minor earthquake. If there is an alarm, then a neighbor might call. If there is an earthquake, the local radio station may report it. There are five variables in this domain, all of which are binary valued:  $A$  (alarm);  $B$  (burglary);  $E$  (earthquake);  $C$  (phone call);  $R$  (radio report). The structure, shown in Fig. 1, represents our intuition about the way the world works. The CPDs associated with the nodes represent the local probability models for the different variables. One such model —  $\Pr(B)$  — would specify the prior probability of a burglary, i.e., it would specify the prior probabilities of the events  $B = 1$  and  $B = 0$ . The probability of the alarm going off is a conditional probability distribution  $\Pr(A | B, E)$ ; it specifies the probability of the alarm going off in any of the relevant circumstances: a burglary and an earthquake ( $B = 1, E = 1$ ), a burglary and no earthquake ( $B = 1, E = 0$ ), etc. One (simplified) choice of CPDs for this domain is shown in Fig. 1(b).

The semantics of a BN is as a joint distribution over  $dom(\mathcal{X})$ :

**Definition 2** A BN  $\mathcal{B} = (G, \Pr)$  over  $X_1, \dots, X_n$  defines a joint distribution over  $X_1, \dots, X_n$  via the chain rule for Bayesian networks:

$$P(X_1, \dots, X_n) = \prod_{i=1}^n \Pr(X_i | Pa(X_i)) \quad (1)$$

A BN can be viewed as a compact representation of a symmetric probability tree. For example, the tree for our Alarm BN would have 32 paths, each with five binary splits, one for each variable. Note that the tree has a lot of duplication over the BN. For example, there are eight branches in the tree all labeled with the same probability  $\Pr(R = 0 | E = 1)$  (one for each assignment to  $B, A, C$ ).

A BN is a full representation of a joint distribution. Hence, it can be used to answer any query that can be answered with a joint distribution. In particular, we can assess the probability distribution over any variable conditioned on any instantiation of values to a subset of others. For example, the probability

of a burglary given that the neighbor called —  $P(B = 1 | C = 1)$  — is 0.1. Now, assume that we turn on the radio and hear a report of an earthquake in our neighborhood. The probability of an earthquake goes up substantially —  $P(E = 1 | C = 1, R = 1) = 0.999$  as compared to  $P(E = 1 | C = 1) = 0.021$ . More interestingly, the probability of a burglary goes *down* —  $P(B = 1 | C = 1, R = 1) = 0.027$ . Note that burglary and earthquake are not mutually exclusive; indeed, they occur independently. However, the earthquake provides an explanation for the neighbor’s phone call; therefore, our basis for believing the alternative explanation of the burglary is no longer as active.

These are examples of conclusions that we can extract from the joint distribution. Of course, explicitly enumerating the joint distribution is a computationally expensive operation, except for the smallest networks. Fortunately, there are algorithms (e.g., [12]) that exploit the sparsity of the network structure to perform this type of reasoning without enumerating the joint distribution explicitly. These algorithms can work effectively even for very large networks, as long as there are not too many direct connections between the variables in the network.

The BN structure induces a set of independence assumptions about the variables  $X_1, \dots, X_n$ . These are assumptions that are a direct consequence of the definition of the joint distribution using the chain rule. They hold for any parameterization  $\Pr$  of the graph  $G$ . These independencies are critical to our analysis, so we review them here.

We first define the basic notion of *conditional independence* of random variables.

**Definition 3** Let  $P$  be a distribution over  $\mathcal{X}$ , and let  $X$ ,  $Y$ , and  $Z$  be three pairwise disjoint subsets of  $\mathcal{X}$ . We say that  $X$  is conditionally independent of  $Y$  given  $Z$  in  $P$ , denoted  $P \models I(X; Y | Z)$ , if, for any  $z \in \text{dom}(Z)$  such that  $P(z) > 0$ , and for any  $x \in \text{dom}(X), y \in \text{dom}(Y)$ , we have that  $P(x | y, z) = P(x | z)$ .

This notion of conditional independence is quite strong, as it implies conditional independence for every assignment of values to the variables involved. Conversely, note that conditional independence is very different from the more standard notion of marginal independence. For example, in the distribution represented by Fig. 1, we have that  $C$  is conditionally independent of  $B$  given  $A$  — once we know whether the alarm sounded or not, knowledge about a possible burglary no longer gives us any information about the chances of a phone call. This independence is very natural; it reflects our intuition that the neighbor decides whether to call based only on hearing the alarm; he has no direct knowledge of the burglary. However,  $B$  and  $C$  are *not* marginally independent: the presence of a burglary makes the phone call more likely.

It turns out that we can use the graph structure to provide a qualitative criterion for determining independence properties of the distribution associated with the graph. At an intuitive level, we can show that probabilistic influence between the variables “flows” along paths in the graph, but can “blocked” somewhere along the path. In our example above, influence can flow from  $B$  to  $C$ , but can be blocked if we condition on  $A$ .

The basic notion in this analysis is that of an *active path*, i.e., one along which influence can “flow”. It helps to first analyze the simplest case, where one node  $X$  can influence another  $Y$  via a third node  $Z$ . As we discussed, this depends on what our evidence is, i.e., what we are conditioning on. Let  $E$  be the set of variables on which we are conditioning; i.e., assume that our evidence is  $E = e$  for some assignment  $e$ , and we are trying to analyze the independencies in the distribution  $P$  conditioned on  $E = e$ . There are four cases to consider.

1. The path has the form  $X \rightarrow Z \rightarrow Y$ . This case is the same as our example  $B \rightarrow A \rightarrow C$ . In this case, we say that the path is *active* if  $Z$  is not observed, i.e.,  $Z \notin E$ . If  $Z \in E$ , we say that the path is *blocked*.
2. The path has the form  $X \leftarrow Z \leftarrow Y$ . As probabilistic dependence is symmetrical, this case is precisely analogous to the first.
3. The path has the form  $X \leftarrow Z \rightarrow Y$ . This case represents a situation where a common cause has two effects, for example, the earthquake causing both the alarm and the news report. It seems fairly intuitive that observing a news report about an earthquake will change our probability in the alarm. However, if we know that there is an earthquake, the news report gives us no additional

information. In light of this intuition, we say that this path is active if  $Z \notin E$  and blocked otherwise.

4. The path is a *v-structure*  $X \rightarrow Z \leftarrow Y$ , like the path from  $B$  through  $A$  to  $E$ . Note that, in the prior distribution  $P$ ,  $B$  and  $E$  are marginally independent. Hence, when  $A$  is not observed, the path is not active, by contrast to the three previous cases. More interestingly, when we observe  $A$ ,  $B$  and  $E$  become correlated. A similar phenomenon occurs if we observe  $C$ . These examples lead us to define a path  $X \rightarrow Z \leftarrow Y$  to be active if  $Z$  or one of  $Z$ 's descendants in  $G$  is in  $E$ , and blocked otherwise.

The definition of active for longer paths is a simple composition of the definition for the paths of length two. Intuitively, influence flows from  $X$  to  $Y$  via a long path if it flows through every intervening node. Thus, we define a general active path as follows:

**Definition 4** Let  $G$  be a BN structure, and  $X_1 - \dots - X_n$  an undirected path in  $G$ . Let  $E$  be a subset of nodes of  $G$ . The path  $X_1 - \dots - X_n$  is active given evidence  $E$  if

- Whenever we have a configuration  $X_{i-1} \rightarrow X_i \leftarrow X_{i+1}$ , then  $X_i$  or one of its descendants in  $E$ ;
- no other node along the path is in  $E$ .

A path which is not active is blocked. We say that  $X$  and  $Y$  are d-separated in  $G$  given  $E$  if every path between them is blocked.

In our Alarm example, we have that  $R$  and  $B$  are d-separated, because the v-structure blocks the only possible path. On the other hand,  $R$  and  $B$  are not d-separated given  $C$ , as observing  $C$  activates the path. But,  $R$  and  $B$  are d-separated given  $E$  and  $C$ , as observing  $E$  blocks the path again. We note that d-separation can be computed in linear time, using, for example, Shachter's Bayes-Ball algorithm [23].

So far, we have defined d-separation based purely on intuitive arguments. However, as shown by Geiger *et al.* [6, 19], these intuitive arguments correspond to provably correct statements about independencies in the distribution. First, we can show that d-separation is *sound*: d-separation in  $G$  implies conditional independence for any Bayesian network  $B = (G, \text{Pr})$ .

**Theorem 1 (Soundness)** Let  $G$  be a Bayesian network structure, and let  $X$  and  $Y$  be nodes in  $G$  and  $E$  a set of nodes such that  $X, Y \notin E$ . If  $X$  and  $Y$  are d-separated in  $G$  given  $E$ , then for any Bayesian network  $B = (G, \text{Pr})$ , we have that  $B \models I(X; Y \mid E)$ .

In other words, the independencies that we derive qualitatively from the graph structure via d-separation hold for every parameterization of the network structure with CPDs.

The d-separation criterion is also *complete*, but in a weaker sense. It is not the case that if  $B \models I(X; Y \mid E)$  then we can necessarily detect that from the network structure. We might choose parameters that create spurious independencies, simply because two different probability expressions happen to be equal to each other. However, as the following theorem shows, if an independence does not follow from the d-separation criterion, then there is at least one counterexample to it.

**Theorem 2 (Completeness)** Let  $G$  be a Bayesian network structure, and let  $X$  and  $Y$  be nodes in  $G$  and  $E$  a set of nodes such that  $X, Y \notin E$ . If  $X$  and  $Y$  are not d-separated in  $G$  given  $E$ , then there exists a Bayesian network  $B = (G, \text{Pr})$  such that  $B \not\models I(X; Y \mid E)$ .

In fact, we can prove an even stronger version: in *almost all* Bayesian networks  $B = (G, \text{Pr})$  (i.e., in all except for a set of measure zero), we have that  $B \not\models I(X; Y \mid E)$ .

Thus, Bayesian networks provide us with a formal framework for representing independence structure in a joint distribution. They allow us to exploit this structure in order to provide a compact representation of complex joint distributions. They also provide us with a qualitative method for determining the presence and absence of independence relations in the joint distribution.

### 3 Multi-Agent Influence Diagrams (MAIDs)

*Influence diagrams* augment the Bayesian network framework with the notions of agents that make decisions strategically, to optimize their utility. Influence diagrams were introduced by Howard [7], and have been investigated almost entirely in a single-agent setting. In this section, we present *multi-agent influence diagrams (MAIDs)*, which extend the influence diagram framework to the multi-agent case. Extending the influence diagram representation from one agent to several agents is fairly straightforward, so there is no point in treating the single-agent and multi-agent cases separately.

We will introduce MAIDs using a simple two-agent scenario:

**Example 2** Alice is considering building a patio behind her house, and the patio would be more valuable to her if she could get a clear view of the ocean. Unfortunately, there is a tree in her neighbor Bob's yard that blocks her view. Being somewhat unscrupulous, Alice considers poisoning Bob's tree, which might cause it to become sick. Bob cannot tell whether Alice has poisoned his tree, but he can tell if the tree is getting sick, and he has the option of calling in a tree doctor (at some cost). The attention of a tree doctor reduces the chance that the tree will die during the coming winter. Meanwhile, Alice must make a decision about building her patio before the weather gets too cold. When she makes this decision, she knows whether a tree doctor has come, but she cannot observe the health of the tree directly. A MAID for this scenario is shown in Fig. 2(a)

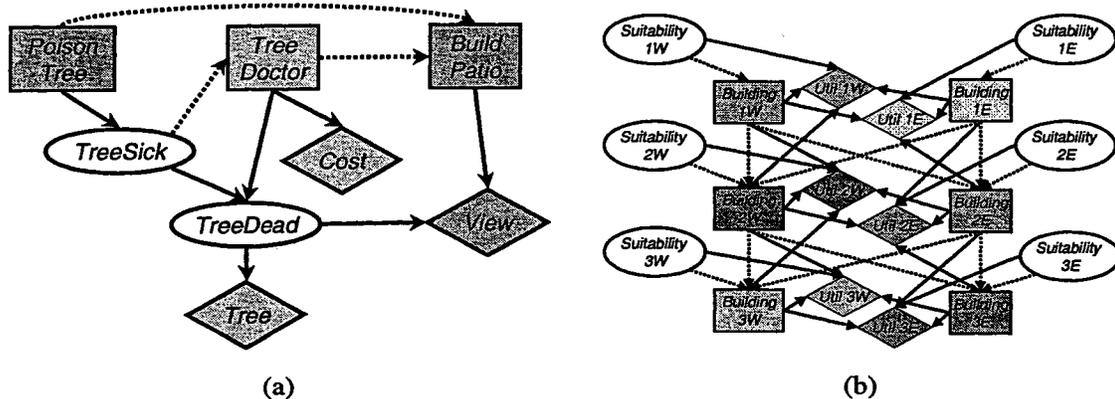


Figure 2: (a) A MAID for the Tree Killer example; Alice's decision and utility variables are in dark gray and Bob's in light gray. (b) A MAID for the Road example with  $n = 3$ .

To define a MAID more formally, we begin with a set  $\mathcal{A}$  of agents. The world in which the agents act is represented by the set  $\mathcal{X}$  of *chance variables*, and a set  $\mathcal{D}_a$  of *decision variables* for each agent  $a \in \mathcal{A}$ . Chance variables correspond to decisions of nature, as in the Bayesian network formalism. They are represented in the diagram as ovals. The decision variables for agent  $a$  are variables whose values  $a$  gets to choose, and are represented as rectangles in the diagram. We use  $\mathcal{D}$  to denote  $\bigcup_{a \in \mathcal{A}} \mathcal{D}_a$ . The agents' utility functions are specified using *utility variables*: For each agent  $a \in \mathcal{A}$ , we have a set  $\mathcal{U}_a$  of utility variables, represented as diamonds in the diagram. The domain of a utility variable is always a finite set of real numbers (a chance or decision variable can have any finite domain). We use  $\mathcal{U}$  to denote  $\bigcup_{a \in \mathcal{A}} \mathcal{U}_a$ . and  $\mathcal{V}$  to denote  $\mathcal{X} \cup \mathcal{D} \cup \mathcal{U}$ .

Like a BN, a MAID defines a directed acyclic graph with its variables as the nodes, where each variable  $X$  is associated with a set of parents  $Pa(X) \subset \mathcal{X} \cup \mathcal{D}$ . Note that utility variables cannot be parents of other variables. For each chance variable  $X \in \mathcal{X}$ , the MAID specifies a CPD  $\Pr(X | Pa(X))$ , as in a BN. For a decision variable  $D \in \mathcal{D}_a$ ,  $Pa(D)$  is the set of variables whose values agent  $a$  knows when he chooses a value for  $D$ . Thus, the choice agent  $a$  makes for  $D$  can be contingent only on these variables. (See Definition 5 below.) For a utility variable  $U$ , the MAID also specifies a CPD  $\Pr(U | \mathbf{pa})$  for each instantiation  $\mathbf{pa}$  of  $Pa(X)$ . However, we require that the value of a utility variable

be a deterministic function of the values of its parents: for each  $\mathbf{pa} \in \text{dom}(Pa(U))$ , there is one value of  $U$  that has probability 1, and all other values of  $U$  have probability 0. We use  $U(\mathbf{pa})$  to denote the value of node  $U$  that has probability 1 when  $Pa(U) = \mathbf{pa}$ . The total utility that an agent  $a$  derives from an instantiation of  $\mathcal{V}$  is the sum of the values of  $\mathcal{U}_a$  in this instantiation. Thus, by breaking an agent's utility function into several variables, we are simply defining an additive decomposition of the agent's utility function [7, 9].

The agents get to select their behavior at each of their decision nodes. An agent's decision at a variable  $D$  can depend on the variables that the agent observes prior to making  $D$  —  $D$ 's parents. The agent's choice of strategy is specified via a set of *decision rules*.

**Definition 5** A decision rule for a decision variable  $D$  is a function that maps each instantiation  $\mathbf{pa}$  of  $Pa(D)$  to a probability distribution over  $\text{dom}(D)$ . An assignment of decision rules to every decision  $D \in \mathcal{D}_a$  for a particular agent  $a \in \mathcal{A}$  is called a strategy.

An assignment  $\sigma$  of decision rules to every decision  $D \in \mathcal{D}$  is called a *strategy profile*. A *partial strategy profile*  $\sigma_{\mathcal{E}}$  is an assignment of decision rules to a subset  $\mathcal{E}$  of  $\mathcal{D}$ . We will also use  $\sigma_{\mathcal{E}}$  to denote the restriction of  $\sigma$  to  $\mathcal{E}$ , and  $\sigma_{-\mathcal{E}}$  to denote the restriction of  $\sigma$  to variables not in  $\mathcal{E}$ .

Note that a decision rule has exactly the same form as a CPD. Thus, if we have a MAID  $\mathcal{M}$ , then a partial strategy profile  $\sigma$  that assigns decision rules to a set  $\mathcal{E}$  of decision variables induces a new MAID  $\mathcal{M}[\sigma]$  where the elements of  $\mathcal{E}$  have become chance variables. That is, each  $D \in \mathcal{E}$  corresponds to a chance variable in  $\mathcal{M}[\sigma]$  with  $\sigma(D)$  as its CPD. When  $\sigma$  assigns a decision rule to every decision variable in  $\mathcal{M}$ , the induced MAID is simply a BN: it has no more decision variables. This BN defines a joint probability distribution  $P_{\mathcal{M}[\sigma]}$  over all the variables in  $\mathcal{M}$ .

**Definition 6** If  $\mathcal{M}$  is a MAID and  $\sigma$  is a strategy profile for  $\mathcal{M}$ , then the joint distribution for  $\mathcal{M}$  induced by  $\sigma$ , denoted  $P_{\mathcal{M}[\sigma]}$ , is the joint distribution over  $\mathcal{V}$  defined by the Bayes net where:

- the set of variables is  $\mathcal{V}$ ;
- for  $X, Y \in \mathcal{V}$ , there is an edge  $X \rightarrow Y$  if and only if  $X \in Pa(Y)$ ;
- for all  $X \in \mathcal{X} \cup \mathcal{U}$ , the CPD for  $X$  is  $\text{Pr}(X)$ ;
- for all  $D \in \mathcal{D}$ , the CPD for  $D$  is  $\sigma(D)$ .

With this probability distribution, we can now write an equation for the utility that agent  $a$  expects to receive in a MAID  $\mathcal{M}$  if the agents play a given strategy profile  $\sigma$ . Suppose  $\mathcal{U}_a = \{U_1, \dots, U_m\}$ . Then:

$$EU_a(\sigma) = \sum_{(u_1, \dots, u_m) \in \text{dom}(\mathcal{U}_a)} P_{\mathcal{M}[\sigma]}(u_1, \dots, u_m) \sum_{i=1}^m u_i \quad (2)$$

where  $\text{dom}(\mathcal{U}_a)$  is the joint domain of  $\mathcal{U}_a$ .

Having defined the notion of an expected utility, we can now define what it means for an agent to optimize his decision at one or more of his decision rules, relative to the a given set of decision rules for the other variables. begin to d

**Definition 7** Let  $\mathcal{E}$  be a subset of  $\mathcal{D}_a$ , and let  $\sigma$  be a strategy profile. We say that  $\sigma_{\mathcal{E}}^*$  is optimal for the strategy profile  $\sigma$  if, in the induced MAID  $\mathcal{M}[\sigma_{-\mathcal{E}}]$ , where the only remaining decisions are those in  $\mathcal{E}$ , the strategy  $\sigma_{\mathcal{E}}^*$  is optimal, i.e., if for all strategies  $\sigma'_{\mathcal{E}}$ :

$$EU_a((\sigma_{-\mathcal{E}}, \sigma_{\mathcal{E}}^*)) \geq EU_a((\sigma_{-\mathcal{E}}, \sigma'_{\mathcal{E}}))$$

Note that, in this definition, it does not matter what decision rules  $\sigma$  assigns to the variables in  $\mathcal{E}$ .

In the game-theoretic framework, we typically consider a strategy profile to represent rational behavior if it is a *Nash equilibrium* [16]. Intuitively, a strategy profile is a Nash equilibrium if no agent has an incentive to deviate from the strategy specified for him by the profile, as long as the other agents do not deviate from their specified strategies.

**Definition 8** A strategy profile  $\sigma$  is a Nash equilibrium for a MAID  $\mathcal{M}$  if for all agents  $a \in \mathcal{A}$ ,  $\sigma_{\mathcal{D}_a}$  is optimal for the strategy profile  $\sigma$ .

The task of finding a Nash equilibrium for a game is arguably the most fundamental task in noncooperative game theory.

## 4 MAIDs and Games

A MAID provides a compact representation of a scenario that can also be represented as a game in strategic or extensive form. In this section, we discuss how to convert a MAID into an extensive-form game. We also show how, once we have found an equilibrium strategy profile for a MAID, we can convert it into a behavior strategy profile for the extensive form game. The word “node” in this section refers solely to a node in the tree, as distinguished from the nodes in the MAID.

There are many ways to convert a MAID into a game tree, all of which are equivalent from a semantic perspective, but which differ dramatically in their computational costs. To understand this point, consider a MAID with  $n$  binary-valued chance variables  $C_1, \dots, C_n$ , with  $C_i$  the sole parent of  $C_{i+1}$ , and a single decision  $D$  whose sole parent is  $C_n$ . One naive approach is to generate a symmetric tree where the root is a chance node splitting on  $C_1$ ; its two children are both chance nodes splitting on  $C_2$ ; etc. Each path down the tree will thus contain  $n$  splits (one for each of  $C_1, \dots, C_n$ ), and then have a final split for  $D$ , belonging to the agent. To preserve the information structure, the tree will have two information sets, one for  $C_n = \text{true}$  and one for  $C_n = \text{false}$ ; thus, all of the  $2^{n-1}$  nodes where  $C_n = \text{true}$  would be in the first information set. The advantage of this tree is that the probabilities associated with the chance-splits are simply extracted from the CPDs in the MAID. The disadvantage is that the size of the tree grows exponentially with the number of chance variables in the MAID, even when this blowup is extraneous.

A alternative approach (based on the construction in [19]) is to generate a tree that has only the minimal set of splits. In our example, we only need a single split for  $C_n$ , and then another split for  $D$ . The probabilities of the  $C_n$  split must now be computed from the MAID using probabilistic inference; i.e., we want to compute  $P(C_n)$  in the joint distribution over  $C_1, \dots, C_n$  defined by the MAID. Similarly, we must now compute the expected utilities at the leaves.

More generally, we need to split on a chance variable before its value is observed by some decision node. Furthermore, we need only split on chance variables that are observed at some point in the process. Thus, the set of variables included in our game tree is  $\mathcal{G} = \mathcal{D} \cup \bigcup_{D \in \mathcal{D}} Pa(D)$ . We present the construction below, referring the reader to [19] for a complete discussion.

We define a total ordering  $\prec$  over  $\mathcal{G}$  that is consistent with the topological order of the MAID: if there is a directed path from  $X$  to  $Y$ , then  $X \prec Y$ . Our tree  $\mathcal{T}$  is a symmetric tree, with each path containing splits over all the variables in  $\mathcal{G}$  in the order defined by  $\prec$ . Each node is labeled with a partial instantiation  $inst(N)$  of  $\mathcal{G}$ , in the obvious way. For each agent  $a$ , the nodes corresponding to variables  $D \in \mathcal{D}_a$  are decision nodes for  $a$ ; the other nodes are all chance nodes. To define the information sets, consider two decision nodes  $M$  and  $M'$  that correspond to a variable  $D$ . We place  $M$  and  $M'$  in the same information set if and only if  $inst(M)$  and  $inst(M')$  assign the same values to  $Pa(D)$ .

Our next task is to determine the split probabilities at the chance nodes. Consider a chance node  $N$  corresponding to a chance variable  $C$ . For each value  $c \in dom(C)$ , let  $N_c$  be the child of  $N$  corresponding to the choice  $C = c$ . We want to compute the probability of going from  $N$  to  $N_c$ . The problem, of course, is that a MAID does not define a full joint probability distribution until decision rules for the agents are selected. It turns out that we can choose an arbitrary fully mixed strategy profile  $\sigma$  for our MAID  $\mathcal{M}$  (one where no decision has probability zero), and do inference in the BN  $\mathcal{M}[\sigma]$  induced by this strategy profile, by computing  $P_{\mathcal{M}[\sigma]}(inst(N_c) \mid inst(N))$ .

The value of this expression does not depend on our choice of  $\sigma$ . To see why this is true, note that if we split on a decision variable  $D$  before  $C$ , then the decision rule  $\sigma_D$  does not affect the computation of  $P_{\mathcal{M}[\sigma]}(inst(N_c) \mid inst(N))$ , because  $inst(N)$  includes values for  $D$  and all its parents. If we split on  $D$  after  $C$ , then  $D$  cannot be an ancestor of  $C$  in the MAID. Also, by the topological ordering of the nodes in the tree, we know that  $inst(N)$  cannot specify evidence on  $D$  or any of its descendants. Therefore,  $\sigma_D$  cannot affect the computation. Hence, the probabilities of the chance nodes are well-defined.

We define the payoffs at the leaves by computing a distribution over the utility nodes, given an instantiation of  $\mathcal{G}$ . For a leaf  $N$ , the payoff for agent  $a$  is:

$$\sum_{U \in \mathcal{U}_a} \sum_{u \in dom(U)} P_{\mathcal{M}[\sigma]}(U = u \mid inst(N)) \cdot u \quad (3)$$

We can also show that the value of (3) does not depend on our choice of  $\sigma$ . The basic idea here is that

$inst(N)$  determines the values of  $D$  and  $Pa(D)$  for each decision variable  $D$ . Hence, the agents' moves and information are all fully determined, and the probabilities with which different actions are chosen in  $\sigma$  are irrelevant. We omit details.

The mapping between MAIDs and trees also induces an obvious mapping between strategy profiles in the different representations. A MAID strategy profile specifies a probability distribution over  $dom(D)$  for each pair  $(D, pa)$ , where  $pa$  is an instantiation of  $Pa(D)$ . The information sets in the game tree correspond one-to-one with these pairs, and a behavior strategy in the game tree is a mapping from information sets to probability distributions. Clearly the two are equivalent.

Based on this construction, we can now state the following equivalence proposition:

**Proposition 1** *Let  $\mathcal{M}$  be a MAID and  $\mathcal{T}$  be its corresponding game tree. Then for any strategy profile  $\sigma$ , the payoff vector for  $\sigma$  in  $\mathcal{M}$  is the same as the payoff vector for  $\sigma$  in  $\mathcal{T}$ .*

The number of nodes in  $\mathcal{T}$  is exponential in the number of decision variables, and in the number of chance variables that are observed during the course of the game. While this blowup is unavoidable in a tree representation, it can be quite significant in certain games. As we now show, a MAID can be exponentially smaller than the extensive game it corresponds to.

**Example 3** *Suppose a road is being built from north to south through undeveloped land, and  $n$  agents have purchased plots of land along the road. As the road reaches each agent's plot, the agent needs to choose what to build on his land. His utility depends on what he builds, on some private information about the suitability of his land for various purposes, and on what is built north, south, and across the road from his land. The agent can observe what has already been built immediately to the north of his land (on both sides of the road), but he cannot observe further north; nor can he observe what will be built across from his land or south of it.*

The MAID representation, shown in Fig. 2(a) for  $n = 3$ , is very compact. There are  $n$  chance nodes, corresponding to the private information about each agent's land, and  $n$  decision variables. Each decision variable has at most three parents: the agent's private information, and the two decisions regarding the two plots to the north of the agent's land. Thus, the size of the MAID is linear in  $n$ . Conversely, any game tree for this situation must split on each of the  $n$  chance nodes and each of the  $n$  decisions, leading to a representation that is exponential in  $n$ . Concretely, suppose the chance and decision variables each have three possible values, corresponding to three types of buildings. Then the game tree corresponding to the Road MAID has  $3^{2n}$  leaves.

A MAID representation is not always more compact. If the game tree is naturally asymmetric, a naive MAID representation can be exponentially larger than the tree.

**Example 4** *Consider, for example, a standard centipede game, a perfect information two-player game. The agents take turns moving, each move consisting of either "right" or "down"; the game ends as soon as any agent decides to move "down". We can view the agents' decisions as variables  $D_1, \dots, D_n$ , each of which takes one of two values. As the game has perfect information, each node  $D_i$  has  $D_1, \dots, D_{i-1}$  as parents. As a consequence, the decision rule for  $D_i$ , if treated naively, requires that we make a decision for each of the exponentially many combinations of values to  $D_1, \dots, D_{i-1}$ . Thus, a naive representation of the MAID strategy grows exponentially in this example, despite the fact that the tree only has  $n$  decisions, and can be represented very compactly.*

It is possible to avoid this problem by using the techniques of Boutilier *et al.* [1]. Rather than representing decision rules as tables, we can represent them as trees, only considering combinations of parents that represent achievable paths in the game. Using ideas along these lines, it is fairly straightforward to provide a transformation from game trees to MAIDs which causes no blowup, i.e., so that the size of the MAID is the same as that of the game tree. The transformation is somewhat technical, and brings no real insight into MAIDs, so we omit it from this paper. However, it has the following important consequence: We can state that the MAID representation of a decision-making situation is no larger than the extensive form representation, and is exponentially smaller in many cases.

## 5 Strategic Relevance

To take advantage of the independence structure in a MAID, we would like to find a global equilibrium through a series of relatively simple local computations. The difficulty is that, in order to determine the optimal decision rule for a single decision variable, we usually need to know the decision rules for some other variables. In Example 2, when Alice is deciding whether to poison the tree, she needs to compare the expected utilities of her two alternatives. However, the probability of the tree dying depends on the probability of Bob calling a tree doctor if he observes that the tree is sick. Thus, we need to know the decision rule for *CallTreeDoctor* to determine the optimal decision rule for *PoisonTree*. In such situations, we will say that *PoisonTree* (strategically) relies on *CallTreeDoctor*, or that *CallTreeDoctor* is relevant to *PoisonTree*. On the other hand, *CallTreeDoctor* does not rely on *PoisonTree*. Bob gets to observe whether the tree is sick, and *TreeDead* is conditionally independent of *PoisonTree* given *TreeSick*, so the decision rule for *PoisonTree* is not relevant to Bob's decision.

We will now formalize this intuitive discussion of strategic relevance. Suppose we have a strategy profile, and we would like to find a decision rule for a single decision variable  $D \in \mathcal{D}_a$  that maximizes  $a$ 's expected utility, assuming the rest of the strategy profile remains fixed.

According to Definition 7, to determine whether  $\delta$  is optimal for  $\sigma$ , we construct the induced MAID where all decision nodes except  $D$  are turned into chance nodes, with their CPDs specified by  $\sigma$ . The decision rule  $\delta$  is optimal for  $\sigma$  if it maximizes  $a$ 's expected utility in this single-decision MAID. The key question that motivates our definition of strategic relevance is the following: What other decision rules are relevant for optimizing the decision rule at  $D$ ?

**Definition 9** Let  $D$  be a decision node in a MAID  $\mathcal{M}$ ,  $\delta$  be a decision rule for  $D$ , and  $\sigma$  be a strategy profile such that  $\delta$  is optimal for  $\sigma$ .  $D$  strategically relies on a decision node  $D'$  in  $\mathcal{M}$  if there is another strategy profile  $\sigma'$  such that  $\sigma'$  differs from  $\sigma$  only at  $D'$ , but no decision rule  $\delta'$  that agrees with  $\delta$  on all parent instantiations  $pa \in \text{dom}(Pa(D))$  where  $P_{\mathcal{M}[\sigma]}(pa) > 0$  is optimal for  $\sigma'$ .

In other words, if a decision rule  $\delta$  for  $D$  is optimal for a strategy profile  $\sigma$ , and  $D$  does not rely on  $D'$ , then  $\delta$  is also optimal for any strategy profile  $\sigma'$  that differs from  $\sigma$  only at  $D'$ . The last clause of this definition is needed to deal with a problem that arises in many other places in game theory — the problem of suboptimal decisions in response to observations that have zero probability (such as observing an irrational move by another agent).

Relevance is a numeric criterion that depends on the specific probabilities and utilities in the MAID. It is not obvious how we would check for strategic relevance without testing all possible pairs of strategy profiles  $\sigma$  and  $\sigma'$ . We would like to find a qualitative criterion that can help us determine strategic relevance purely from the structure of the graph, i.e., a criterion analogous to the d-separation criterion for determining conditional independence in Bayesian networks.

The key insight is based on the following reasoning. Each decision only influences utility nodes that are its descendants, and therefore we need only consider them. As a consequence, to perform our optimization, we care only about the probability distribution over the values of these utility nodes. Hence, only decision nodes that can influence our (subjective) probability distribution over these utility nodes are relevant. Consider, for example, the *CallTreeDoctor* decision variable in our example. To find an optimal decision rule for this variable, we only need to evaluate two probabilistic queries:  $P_{\mathcal{M}[\sigma]}(\text{Tree} \mid \text{TreeSick}, \text{CallTreeDoctor})$  and  $P_{\mathcal{M}[\sigma]}(\text{Cost} \mid \text{TreeSick}, \text{CallTreeDoctor})$ . The second query is obviously trivial since *CallTreeDoctor* is the sole parent of *Cost*; the first can be evaluated without referring to the decision rules for *PoisonTree* or *BuildPatio* (because of the independence relations in the MAID). Thus, if we change  $\sigma$  to another strategy profile  $\sigma'$  that assigns different decision rules to *PoisonTree* or *BuildPatio*,  $\delta$  will still be optimal for  $\sigma'$ .

The problem of determining which nodes' CPDs might affect the evaluation of a probabilistic query is a standard one in the Bayesian network literature, so that we can build on a graphical criterion already defined for Bayesian networks, that of a *requisite probability node*:

**Definition 10** Let  $G$  be a BN structure, and let  $X$  and  $Y$  be sets of variables in the BN. Then a node  $Z$  is a requisite probability node for the query  $P(X \mid Y)$  if there exist two Bayesian networks  $\mathcal{B}_1$  and  $\mathcal{B}_2$  over  $G$ , that are identical except in the CPD they assign to  $Z$ , but  $P_{\mathcal{B}_1}(X \mid Y) \neq P_{\mathcal{B}_2}(X \mid Y)$ .

As we will see, the decision rule at  $D'$  is only relevant to  $D$  if  $D'$  (viewed as a chance node) is a requisite probability node for  $P(\mathcal{U}_D \mid D, Pa(D))$ .

Geiger *et al.* [6] provide a graphical criterion for testing whether a node  $Z$  is a requisite probability node for a query  $P(\mathcal{X} \mid \mathcal{Y})$ . We add to  $Z$  a new “dummy” parent  $\widehat{Z}$  whose values correspond to CPDs for  $Z$ , selected from some set of possible CPDs. Then  $Z$  is a requisite probability node for  $P(\mathcal{X} \mid \mathcal{Y})$  if and only if  $\widehat{Z}$  can influence  $\mathcal{X}$  given  $\mathcal{Y}$ . Based on these considerations, we can define *s-reachability*, a graphical criterion for detecting strategic relevance. Note that, unlike d-separation in Bayesian networks, s-reachability is not necessarily a symmetric relation.

**Definition 11** *A node  $D'$  is s-reachable from a node  $D$  in  $\mathcal{M}$  if there is some utility node  $U \in \mathcal{U}_D$  such that if a new parent  $\widehat{D}'$  were added to  $D'$ , there would be an active path in  $\mathcal{M}$  from  $\widehat{D}'$  to  $U$  given  $Pa(D) \cup \{D\}$ , where a path is active in a MAID if it is active in the same graph, viewed as a BN.*

We can show that s-reachability is sound and complete for strategic relevance (almost) in the same sense that d-separation is sound and complete for independence in Bayesian networks. As for d-separation the soundness result is very strong: without s-reachability, one decision cannot be relevant to another.

**Theorem 3 (Soundness)** *If  $D$  and  $D'$  are two decision nodes in a MAID  $\mathcal{M}$  and  $D'$  is not s-reachable from  $D$  in  $\mathcal{M}$ , then  $D$  does not rely on  $D'$ .*

As for BNs, the result is not as strong in the other direction: s-reachability does not imply relevance in every MAID. We can choose the probabilities and utilities in the MAID in such a way that the influence of one decision rule on another does not manifest itself. However, s-reachability is the most precise graphical criterion we can use: it will not identify a strategic relevance unless that relevance actually exists in some MAID that has the given graph structure. We say that two MAIDs have the same graph structure when the two MAIDs have the same sets of variables and agents, each variable has the same parents in the two MAIDs, and the assignment of decision and utility variables to agents is the same in both MAIDs. The chance and decision variables must have the same domains in both MAIDs, but we allow the actual utility values of the utility variables (their domains) to vary. The CPDs in the two MAIDs may also be different.

**Theorem 4 (Completeness)** *If a node  $D'$  is s-reachable from a node  $D$  in a MAID, then there is some MAID with the same graph structure in which  $D$  relies on  $D'$ .*

Since s-reachability is a binary relation, we can represent it as a directed graph. As we show below, this graph turns out to be extremely useful.

**Definition 12** *The relevance graph for a MAID  $\mathcal{M}$  is a directed graph whose nodes are the decision nodes of  $\mathcal{M}$ , and which contains an edge  $D \rightarrow D'$  if and only if  $D'$  is s-reachable from  $D$ .*

By Theorem 3, if  $D$  relies on  $D'$ , then there is an edge from  $D$  to  $D'$  in the relevance graph.

To construct the graph for a given MAID, we need to determine, for each decision node  $D$ , the set of nodes  $D'$  that are s-reachable from  $D$ . Using an algorithm such as Shachter’s Bayes-Ball [23], we can find this set for any given  $D$  in time linear in the number of nodes in the MAID. By repeating the algorithm for each  $D$ , we can derive the relevance graph in time quadratic in the number of MAID nodes.

Recall our original statement that a decision node  $D$  strategically relies on a decision node  $D'$  if one needs to know the decision rule for  $D'$  in order to evaluate possible decision rules for  $D$ . Although we now have a graph-theoretic characterization of strategic relevance, it will be helpful to develop some intuition by examining some simple MAIDs, and seeing when one decision node relies on another. In the five examples shown in Fig. 3, the decision node  $D$  belongs to agent  $a$ , and  $D'$  belongs to agent  $b$ . Example (a) represents a perfect-information game. Since agent  $b$  can observe the value of  $D$ , he does not need to know the decision rule for  $D$  in order to evaluate his options. Thus,  $D'$  does not rely on  $D$ . On the other hand, agent  $a$  cannot observe  $D'$  when she makes decision  $D$ , and  $D'$  is relevant to  $a$ ’s utility, so  $D$  relies on  $D'$ . Example (b) represents a game where the agents do not have perfect information: agent  $b$  cannot observe  $D$  when making decision  $D'$ . However, the information is “perfect

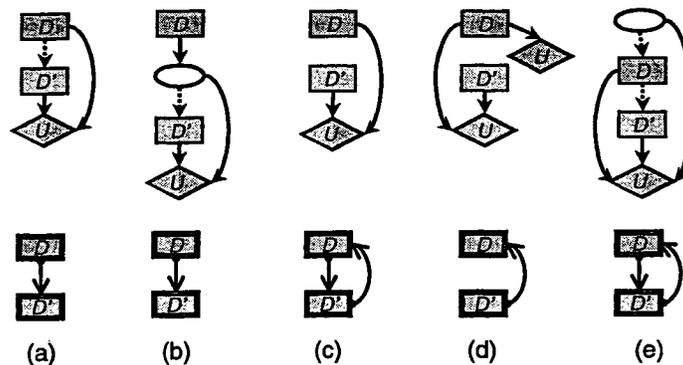


Figure 3: Five simple MAIDs (top), and their relevance graphs (bottom). A two-color diamond represents a pair of utility nodes, one for each agent, with the same parents.

enough”: the utility for  $b$  does not depend on  $D$  directly, but only on the chance node, which  $b$  can observe. Hence  $D'$  does not rely on  $D$ .

Examples (c) and (d) represent scenarios where the agents move simultaneously, and thus neither can observe the other’s move. In (c), each agent’s utility node is influenced by both decisions, so  $D$  relies on  $D'$  and  $D'$  relies on  $D$ . Thus, the relevance graph is cyclic. In (d), however, the relevance graph is acyclic despite the fact that the agents move simultaneously. The difference here is that agent  $a$  no longer cares what agent  $b$  does, because her utility is not influenced by  $b$ ’s decision. In graphical terms, there is no active path from  $D'$  to  $a$ ’s utility node given  $D$ .

One might conclude that a decision node  $D'$  never relies on a decision node  $D$  when  $D$  is observed by  $D'$ , but the situation is more subtle. Consider example (e), which represents a simple card game: agent  $a$  observes a card, and decides whether to bet ( $D$ ); agent  $b$  observes only agent  $a$ ’s bet, and decides whether to bet ( $D'$ ); the utility of both depends on their bets and the value of the card. Even though agent  $b$  observes the actual decision in  $D$ , he needs to know the decision rule for  $D$  in order to know what the value of  $D$  tells him about the chance node. Thus,  $D'$  relies on  $D$ ; indeed, when  $D$  is observed, there is an active path from  $D$  that runs through the chance node to the utility node.

## 6 Computing Equilibria

The computation of a Nash equilibrium for a game is arguably the key computational task in game theory. In this section, we show how the structure of the MAID can be exploited to provide efficient algorithms for finding equilibria in certain games.

The key insight behind our algorithms is the use of the relevance graph to break up the task of finding an equilibrium into a series of subtasks, each over a much smaller game. Since algorithms for finding equilibria in general games have complexity that is superlinear in the number of levels in the game tree, breaking the game into smaller games significantly improves the complexity of finding a global equilibrium.

Our algorithm is a generalization of existing backward induction algorithms for decision trees and perfect information games [25] and for influence diagrams [8]. The basic idea is as follows: in order to optimize the decision rule for  $D$ , we need to know the decision rule for all decisions  $D'$  that are relevant for  $D$ . For example, in the Tree Killer example, in order to optimize *PoisonTree*, we must first decide on the decision rules for *BuildPatio* and *TreeDoctor*. However, we can optimize *TreeDoctor* without knowing the decision rules for either of the other decision variables. Having decided on the decision rule for *TreeDoctor*, we can now optimize *BuildPatio* and then finally *PoisonTree*. In other words, our relevance graph has edges from *PoisonTree* to *BuildPatio* and *TreeDoctor* and from *BuildPatio* to *TreeDoctor*.

We can apply this simple backward induction procedure in any MAID which, like the Tree Killer example, has an acyclic relevance graph. When the relevance graph is acyclic, we can construct a topological ordering of the decision nodes: an ordering  $D_1, \dots, D_n$  such that if  $i < j$ , then  $D_i$  is not  $s$ -reachable from  $D_j$ . We can then iterate backward from  $D_n$  to  $D_1$ , deriving an optimal decision rule for each decision node in turn. Each decision  $D_i$  relies only on the decisions that succeed it in the order, and these will have been computed by the time we have to select the decision rule for  $D_i$ .

The relevance graph is acyclic in all perfect-information games, and in all single-agent decision problems with perfect recall. There are also some games of imperfect information, such as the Tree Killer example, that have acyclic relevance graphs. But in most games we will encounter cycles in the relevance graph. For example, any simple two-player simultaneous move game with two decisions  $D_1$  and  $D_2$ , where both players' payoffs depend on the decisions at both  $D_1$  and  $D_2$ , has a cyclic relevance graph (as in Fig. 3(c)). However, we can often utilize relevance structure even in games where the relevance graph is cyclic.

**Example 5** Consider the relevance graph for the Road example, shown in Fig. 4(a) for  $n = 6$  agents. We can see that we have pairs of interdependent decision variables, corresponding to the two agents whose lots are across the road from each other. Also, the decision for a given plot relies on the decision for the plot directly to the south. However, it does not rely on the decision about the land directly north of it, because this decision is observed. None of the other decisions affect this agent's utility directly, and therefore they are not  $s$ -reachable.

Intuitively, although the last pair of nodes in the relevance graph rely on each other, they rely on nothing else. Hence, we can compute an equilibrium for the pair together, regardless of any other decision rules. Once we have computed an equilibrium for this last pair, the decision variables can be treated as chance nodes, and we can proceed to compute an equilibrium for the next pair.

**Definition 13** A set  $S$  of nodes in a directed graph is a strongly connected component (SCC) if for every pair of nodes  $D \neq D' \in S$ , there exists a directed path from  $D$  to  $D'$ . A maximal SCC is an SCC that is not a strict subset of any other SCC.

The maximal SCCs for the Road example are outlined in Fig. 4(a).

We can find the maximal SCCs of a relevance graph in linear time, by constructing a *component graph* whose nodes are the maximal SCCs of the graph [2]. There is an edge from component  $C$  to component  $C'$  in the component graph if and only if there is an edge in the relevance graph from some element of  $C$  to some element of  $C'$ . The component graph is always acyclic, so we can define an ordering  $C_1, \dots, C_m$  over the SCCs, such that whenever  $i < j$ , no element of  $C_i$  is  $s$ -reachable from any element of  $C_j$ .

We can now provide a divide-and-conquer algorithm for computing Nash equilibria in general MAIDs.

### Algorithm 1

Given a MAID  $\mathcal{M}$

a topological ordering  $C_1, \dots, C_m$  of the SCCs in the relevance graph for  $\mathcal{M}$

1 Let  $\sigma^0$  be an arbitrary fully mixed strategy profile

2 For  $i = 0$  through  $m - 1$ :

3 Let  $\tau$  be a partial strategy profile for  $C_{(m-i)}$  that is a Nash equilibrium in  $\mathcal{M} \left[ \sigma^i_{-C_{(m-i)}} \right]$

4 Let  $\sigma^{i+1} = (\sigma^i_{-C_{(m-i)}}, \tau)$

5 Output  $\sigma^m$  as an equilibrium of  $\mathcal{M}$

The algorithm iterates backwards over the SCC's, finding an equilibrium strategy profile for each SCC in the MAID induced by the previously selected decision rules (with arbitrary decision rules for some decisions that are not relevant for this SCC). In the induced MAID, the only remaining decision nodes are those in the current SCC; all the other decision nodes have been converted to chance nodes. We can find the equilibrium in this MAID by converting it to a game tree, as described in Section 4, and using a standard game-solving algorithm [14] as a subroutine. Note that if the relevance graph is acyclic, each SCC consists of a single decision node. Thus, step 3 involves finding a Nash equilibrium in

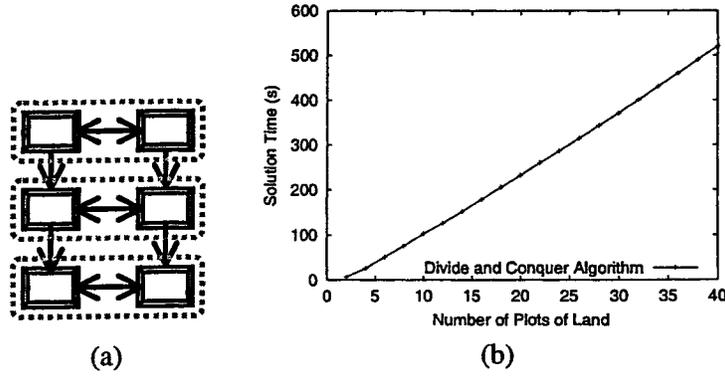


Figure 4: Road example: (a) Relevance graph for  $n = 6$ . (b) Performance results.

a single-player game, which reduces to simply finding a decision rule that maximizes the single agent’s expected utility.

In proving the correctness of Algorithm 1, we encounter a subtle technical difficulty. The definition of strategic relevance (Def. 9) only deals with the optimality of a single decision rule for a strategy profile. But in Algorithm 1, we derive not just single decision rules, but a complete strategy for each agent. To make the leap from the optimality of single decision rules to the optimality of whole strategies in our proof, we must make the standard assumption of *perfect recall* — that agents never forget their previous actions or observations. More formally:

**Definition 14** *An agent  $a$  has perfect recall with respect to a total order  $D_1, \dots, D_n$  over  $\mathcal{D}_a$  if for all  $D_i, D_j \in \mathcal{D}_a$ ,  $i < j$  implies that  $D_i \in Pa(D_j)$  and  $Pa(D_i) \subset Pa(D_j)$ .*

We can now prove the correctness of Algorithm 1.

**Theorem 5** *Let  $\mathcal{M}$  be a MAID where every agent has perfect recall, and let  $C_1, \dots, C_m$  be a topological ordering of the SCCs in the relevance graph for  $\mathcal{M}$ . Then the strategy profile  $\sigma^m$  produced by running Algorithm 1 with  $\mathcal{M}$  and  $C_1, \dots, C_m$  as inputs is a Nash equilibrium for  $\mathcal{M}$ .*

To demonstrate the potential savings resulting from our algorithm, we tried it on the Road example, for different numbers of agents  $n$ . Note that the model we used differs slightly from that shown in Fig. 2(b): In our experiments, each agent had not just one utility node, but a separate utility node for each neighboring plot of land, and an additional node that depends on the suitability of the plot for different purposes. The agent’s decision node is a parent of all these utility nodes. The idea is that an agent gets some base payoff for the building he builds, and then the neighboring plots and the suitability node apply additive bonuses and penalties to his payoff. Thus, instead of having one utility node with  $3^5 = 243$  parent instantiations, we have 4 utility nodes with  $3^2 = 9$  parent instantiations each. This change has no effect on the structure of the relevance graph, which is shown for  $n = 6$  in Fig. 4(a). The SCCs in the relevance graph all have size 2; as we discussed, they correspond to pairs of decisions about plots that are across the road from each other.

Even for small values of  $n$ , it is infeasible to solve the Road example with standard game-solving algorithms. As we discussed, the game tree for the MAID has  $3^{2n}$  leaves, whereas the MAID representation is linear in  $n$ . The normal form adds another exponential factor. Since each agent (except the first two) can observe three ternary variables, he has 27 information sets. Hence, the number of possible pure (deterministic) strategies for each agent is  $3^{27}$ , and the number of pure strategy profiles for all  $n$  players is  $(3^{27})^{(n-2)} \cdot (3^3)^2$ . In the simplest interesting case, where  $n = 4$ , we obtain a game tree with 6561 terminal nodes, and standard solution algorithms, that very often use the normal form, would need to operate on a game matrix with about  $4.7 \times 10^{27}$  entries (one for each pure strategy profile).

Solving the Road game either in its extensive form or in the normal form is infeasible even for  $n = 4$ . By contrast, our divide-and-conquer algorithm ends up generating a sequence of small games, each with

two decision variables. Fig. 4(b) shows the computational cost of the algorithm as  $n$  grows. We converted each of the induced MAIDs constructed during the algorithm into a small game tree, and used the game solver GAMBIT [5] to solve it. As expected, the time required by our algorithm grows approximately linearly with  $n$ . Thus, for example, we can solve a Road MAID with 40 agents (corresponding to a game tree with  $3^{80}$  terminal nodes) in 8 minutes 40 seconds.

## 7 Discussion and Future Work

We have introduced a new formalism, multi-agent influence diagrams (MAIDs), for modeling multi-agent scenarios with imperfect information. MAIDs use a representation where variables are the basic unit, and allow the dependencies between these variables to be represented explicitly, in a graphical form. They therefore reveal important qualitative structure in a game, which can be useful both for understanding the game and as the basis for algorithms that find equilibria efficiently. In particular, we have shown that our divide-and-conquer algorithm for finding equilibria provides exponential savings over existing solution algorithms in some cases, such as the Road example, where the maximal size of an SCC in the relevance graph is much smaller than the total number of decision variables. In the worst case, the relevance graph forms a single large SCC, and our algorithm simply solves the game in its entirety, with no computational benefits.

Although the possibility of extending influence diagrams to multi-agent scenarios was recognized at least fifteen years ago [21], the idea seems to have been dormant for some time. Suryadi and Gmytrasiewicz [24] have used influence diagrams as a framework for learning in multi-agent systems. Milch and Koller [15] use multi-agent influence diagrams as a representational framework for reasoning about agents' beliefs and decisions. However, the focus of both these papers is very different, and they do not consider the structural properties of the influence diagram representation, nor the computational benefits derived from it. Nilsson and Lauritzen [17] have done related work on limited memory influence diagrams, or LIMIDs. Although they mention that LIMIDs could be applied to multi-agent scenarios, they only consider the use of LIMIDs to speed up inference in single-agent settings.

MAIDs are also related to La Mura's [11] game networks, which incorporate both probabilistic and utility independence. La Mura defines a notion of strategic independence, and also uses it to break up the game into separate components. However, his notion of strategic independence is an undirected one, and thus does not allow as fine-grained a decomposition as the directed relevance graph used in this paper, nor the use of a backward induction process for interacting decisions.

This work opens the door to a variety of possible extensions. One obvious direction is to relate MAIDs to existing concepts in game theory, particularly equilibrium refinements. We can show that that the solution found by our algorithm in the case of acyclic relevance graphs is a *perfect Bayesian equilibrium* [4]; it would be interesting to show an analogous result for the more general case.

Another direction relates to additional structure that is revealed by the notion of strategic relevance. In particular, even if the relevance graph is cyclic, it might not be a fully connected subgraph; for example, we might have a situation where  $D_1$  relies on  $D_2$ , which relies on  $D_3$ , which relies on  $D_1$ . Clearly, this type of structure tells us something about the interaction between the decisions in the game. An important open question is to analyze the meaning of these types of structures, and to see whether they can be exploited for computation gain.

Finally, the notion of strategic relevance is not the only type of insight that we can obtain from the MAID representation. We can use a similar type of path-based analysis in the MAID graph to determine which of the variables that an agent can observe before making a decision actually provide relevant information for that decision. In complex scenarios, agents tend to accumulate a great many observations, and the complexity of a decision rule increases exponentially with the number of observed variables. Thus, there has been considerable work on identifying irrelevant parents of decision nodes in single-agent influence diagrams [7, 22, 23]. In this case, we can use d-separation to identify irrelevant parents of a given decision node in time linear in the number of variables. We can use a similar technique in MAIDs, allowing us to eliminate some parents of the decision node. At the level of the associated game tree, this process results in collapsing of several information sets into one (a process called *deflation* in [3, 18]). The structure of the MAID allows us to detect cases when this process can be executed

without any loss to the agents. However, the multi-agent case also raises subtleties that are absent in the single-agent case. In this case, an observed variable that does not directly influence one agent's payoff might nevertheless be relevant, if another agent conditions his behavior on this variable. MAIDs provide a unique framework for studying this question.

**Acknowledgements** This work was supported by Air Force contract F30602-00-2-0598 under DARPA's TASK program and by ONR MURI N00014-00-1-0637 under the program "Decision Making under Uncertainty".

## References

- [1] C. Boutilier, N. Friedman, M. Goldszmidt, and D. Koller. Context-specific independence in Bayesian networks. In *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence (UAI-96)*, pages 115–123, 1996.
- [2] T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT Press, 1990.
- [3] N. Dalkey. Equivalence of information patterns and essentially determinate games. *Annals of Mathematics Studies*, 28:217–243, 1953.
- [4] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [5] GAMBIT software, California Institute of Technology, 2000. <http://www.hss.caltech.edu/gambit/Gambit.html>.
- [6] D. Geiger, T. Verma, and J. Pearl. Identifying independence in Bayesian networks. *Networks*, 20:507–534, 1990.
- [7] R. A. Howard and J. E. Matheson. Influence diagrams. In *Readings on the Principles and Applications of Decision Analysis*, pages 721–762. Strategic Decisions Group, 1984.
- [8] F. Jensen, F.V. Jensen, and S.L. Dittmer. From influence diagrams to junction trees. In *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence (UAI-94)*, pages 367–373, 1994.
- [9] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley & Sons, Inc., 1976.
- [10] D. Koller, N. Megiddo, and B. von Stengel. Fast algorithms for finding randomized strategies in game trees. In *Proceedings of the 26th ACM Symposium on Theory of Computing (STOC-94)*, pages 750–759, 1994.
- [11] P. LaMura. Game networks. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI 2000)*, pages 335–342, 2000.
- [12] S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *J. Royal Stat. Soc. B*, 50(2):157–224, 1988.
- [13] E. Maskin and J. Tirole. Markov perfect equilibrium I: Observable actions. Working Paper 1799, Harvard Institute for Economic Research (HIER), 1997.
- [14] R.D. McKelvey and A. McLennan. Computation of equilibria in finite games. In *Handbook of Computational Economics*, volume 1, pages 87–142. Elsevier Science, Amsterdam, 1996.
- [15] B. Milch and D. Koller. Probabilistic models for agents' beliefs and decisions. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI 2000)*, pages 389–396, 2000.
- [16] J. Nash. Equilibrium points in n-person games. *Proc. National Academy of Sciences of the USA*, 36:48–49, 1950.
- [17] D. Nilsson and S.L. Lauritzen. Evaluating influence diagrams with LIMIDs. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI 2000)*, pages 436–445, 2000.
- [18] A. Okada. Complete inflation and perfect recall in extensive games. *International Journal of Game Theory*, 16(2):85–91, 1987.
- [19] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Francisco, 1988.
- [20] I. V. Romanovskii. Reduction of a game with complete memory to a matrix game. *Soviet Mathematics*, 3:678–681, 1962.
- [21] R. D. Shachter. Evaluating influence diagrams. *Operations Research*, 34:871–882, 1986.
- [22] R. D. Shachter. An ordered examination of influence diagrams. *Networks*, 20:535–563, 1990.
- [23] R. D. Shachter. Bayes-ball: The rational pastime. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 480–487, 1998.
- [24] D. Suryadi and P.J. Gmytrasiewicz. Learning models of other agents using influence diagrams. In *Proceedings of the Seventh International Conference on User Modeling (UM-99)*, pages 223–232, 1999.
- [25] E. Zermelo. Über eine Anwendung der Mengenlehre auf der Theorie des Schachspiels. In *Proceedings of the Fifth International Congress on Mathematics*, 1913.