In *Proceedings of the Second European Workshop onProbabilistic Graphical Models (PGM-04),* pages 97-104, Leiden, The Netherlands, October 2004

An Efficient Sampling Algorithm for Influence Diagrams

Daniel Garcia-Sanchez & Marek J. Druzdzel
Decision Systems Laboratory
School of Information Sciences
and Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
{dgarcia,marek}@sis.pitt.edu

Abstract

We describe an efficient sampling algorithm for solving influence diagrams that achieves its efficiency by reusing samples for each of the decision strategies. Our algorithm is exhaustive in the sense of computing the expected utility of each of the possible decision strategies. We show how by a parallel evaluation of all strategies we not only save a significant amount of computation but also produce better quality anytime behavior.

1 Introduction

Influence diagrams (IDs) (Howard and Matheson, 1984) are acyclic directed graphs modeling decision problems under uncertainty. An ID encodes three basic elements of a decision: (1) available decision options, (2) factors that are relevant to the decision, including how they interact among each other and how the decisions will impact them, and finally, (3) the decision maker's preferences over the possible outcomes of the decision making process. These three elements are encoded in IDs by means of three types of nodes: decision nodes, typically represented as rectangles, random variables, typically represented as ovals, and value nodes, typically represented as diamonds or hexagons. Most popular type of IDs are those in which both the decision options and the random variables are discrete. A decision node in a discrete ID is essentially a list of labels representing decision options. Each random variable is described by a conditional probability table (CPT) containing the probability distribution over its outcomes conditional on its parents. Each value node encodes a utility function that represents a numerical measure of preference over the outcomes of its immediate predecessors. An ID that contains only random variables is called a Bayesian network (BN) (Pearl, 1988).

An ID is amenable to an algorithmic analysis that can yield for each possible strategy (an assignment of values to all decision nodes conditional on those chance nodes that are observed before the decision nodes, e.g., (Clemen, 1996)) its numerical measure of desirability — expected utility. It is a basic premise of decision theory that a decision maker should maximize his or her gain by choosing the strategy that yields the highest expected utility. Even though this analysis is at its foundations NP-hard (Cooper, 1990), there exist several ingenious algorithms that make it feasible for practical ID models. Before we review these algorithms, it is useful to focus on several general characteristic of ID algorithms that will be helpful in making the review more systematic.

While having an exact result is always preferable, in most cases precision can be traded off for computation, dividing ID algorithms into *exact* and *approximate*. An important subclass of approximate algorithms are *anytime* algorithms, which are algorithms that regardless of the complexity of the problem have an answer available almost immediately. Subsequent computation, which gradually improves the quality of this answer, can be interrupted at any time. The last two distinctions are specific to ID al-

gorithms. An algorithm is *exhaustive* if it computes the expected utility of each of the possible decision strategies or *focused* if its only goal is identification of the optimal strategy. Finally, algorithms can be divided into *direct* and *indirect*, i.e., those that work on IDs directly and those that reduce the IDs solution problem into a series of belief updating problems in its underlying BNs.

There is a sizeable body of literature on the topic of solving IDs. Space limitations prevent us from covering most of it with the exception of those that are most directly relevant to our paper. The oldest among the direct evaluation methods originates from the work of Olmsted (1983), whose algorithm performs a series of node removal and arc reversal operations that successively reduce the ID into a single decision node and a single successor utility node that holds the expected utility of each of the decision options. Shachter (1986) has later formalized this algorithm and proven its correctness. Shenoy (1992) introduced a new approach that transforms the ID into a valuation network that allows for slightly more efficient calculations. Ndilikilikesha (1994) modified Olmsted's algorithm making it computationally as efficient as Shenoy's.

The indirect evaluation methods are inspired by a seminal paper by Cooper (1988), who proposed transforming each decision node in an ID into a chance node with uniform distribution and the (single) value node into a chance node whose probability distribution is a linear transformation of the original utility function into the interval [0,1]. Following this transformation, the algorithm instantiates, one by one, each of the strategies, by observing the states of the nodes that represent the original decision nodes, solves the resulting BN using any belief-updating algorithm, and retrieves the expected utility of the current strategy.

Shachter and Peot (1992) designed an efficient variation on Cooper's algorithm for IDs that have one decision node and one utility node. The algorithm is focused, as it identifies the best decision only. A separate, optional step of the algorithm is required to calculate the expected utility of that decision. Zhang (1998) describes another focused indirect evaluation algorithm that is an improvement on Shachter and Peot's algorithm. Other techniques to solve influence diagrams were offered by Jensen, Jensen and Dittmer (1994) and Madsen and Jensen (1999). Focused algorithms are in general more efficient than exhaustive algorithms, but this efficiency comes at a price. The output of a focused algorithm identifies the best next step of a strategy but it does not put it in a context of other possible steps or other strategies. In particular, the output of a focused algorithm does not show whether the suggested decision is a clear winner or a close runner with other, alternative decisions. In applications that include interaction with human decision makers, exhaustive algorithms offer a considerable amount of insight that helps in choosing the superior decision strategy or assists in refining the model.

Another line of research on algorithms for IDs was started by Horvitz (1989) who identified a class of inference policies for probabilistic reasoning systems that he called incremental refinement policies. He presented algorithms that improve the accuracy of their solutions as a monotonically increasing function of the allocated resources and the available information. Several algorithms have been developed along these lines. Ramoni (1995) proposed an anytime algorithm for solving IDs that incrementally refines the confidence in the expected utility of the decisions. One limitation of this algorithm is that its anytime property holds only for IDs with one decision node. Horsch and Poole (1996) developed another anytime algorithm that constructs a tree structure for each decision node in the ID. Each of the constructed tree structures represents a decision function that is improved incrementally. The authors show how the improvements to this function lead to the optimal decision function. Similarly to Ramoni's algorithm, the algorithm loses its anytime properties if applied to decision problems with more than one decision node. A refinement of this algorithm, focusing on multistage influence diagrams was published later (Horsch and Poole, 1998). Charnes and Shenoy

(2004) propose an approximate sampling algorithm for solving IDs. D'Ambrosio and Burgess (1996) performed experiments comparing a variety of real-time influence diagram algorithms.

This paper shows how sampling algorithms can be utilized very efficiently within an indirect evaluation method. Evaluating all decision strategies on the same set of random samples allows for saving a significant amount of computation and also produces high quality anytime behavior. The algorithm we propose is suitable for very large decision models for which exact algorithms are not feasible. Recent advances in stochastic sampling algorithms for BNs (e.g., (Cheng and Druzdzel, 2000; Yuan and Druzdzel, 2003)) have made it possible to perform inference in very large models with unlikely evidence. Because our algorithm can be based on any simulation algorithm for BNs, it allows for taking advantage of these developments. The algorithm is general in the sense of admitting full IDs with multiple decision and value nodes and containing observations. As a side-product, the algorithm computes also the posterior probability distribution of every chance node conditional on each of the possible strategies. For example, in a system that helps to decide between medication and surgery, the algorithm computes not only the expected utility of both choices, but also the probability of the patient dying in each case. We have found that this information is useful in an interactive modeling environment. The algorithm is, to our knowledge, the only algorithm developed for ID that exhibits anytime behavior when the ID has an arbitrary number of decision and utility nodes. Finally, the algorithm is very efficient and this is achieved by reusing random samples. Practically, only one random number per node is generated for each sample of the network. This random number is used to evaluate the probability distribution of the node conditional on each of the strategies.

All random variables used in this paper are multiple-valued, discrete variables. Lower case letters, such as a, b, or c denote random variables. Bold capital letters, such as $\mathbf{A}$, denote sets of variables. Bold capital $\mathbf{E}$ denotes the set of observed nodes and capital $\mathbf{Q}$ denotes the set of other nodes not contained in $\mathbf{E}$. For any node $n \in \mathbf{C}$, $\mathrm{Pa}(n)$ denotes the set of parents of n, $\mathrm{De}(n)$ denotes the descendants of n, and $\mathrm{O}(n)$ denotes the set of outcomes of n.

2 Some definitions

A Bayesian network $\mathcal{B}$ is defined as a pair $\mathcal{B} = (\mathcal{G}, \Pr)$, where $\mathcal{G}(\mathbf{C}, \mathbf{A})$ is an acyclic directed graph consisting of a set of nodes $\mathbf{C}$ and a set of directed arcs $\mathbf{A}$. Each node $n \in \mathbf{C}$ has associated a conditional probability distribution, specifying the probability of each outcome $o \in \mathrm{O}(n)$ conditional on $\mathrm{Pa}(n)$. Similarly, we denote the joint probability distribution over a set of nodes $\mathbf{J} \subseteq \mathbf{C}$ as $\Pr(\mathbf{X}_J)$.

Given $\mathcal{B} = (\mathbf{C}, \mathbf{A})$ and sets $\mathbf{J}$, $\mathbf{K}$, $\mathbf{L} \subseteq \mathbf{C}$, we say $\mathbf{X}_L$ is irrelevant (Geiger et al., 1990) to $\mathbf{X}_J$ given $\mathbf{X}_K$ if $\Pr(\mathbf{X}_J|\mathbf{X}_K, \mathbf{X}_L) = \Pr(\mathbf{X}_J|\mathbf{X}_K)$. In other words, $\mathbf{X}_L$ is irrelevant to $\mathbf{X}_J$ given $\mathbf{X}_K$ if, once $\mathbf{X}_K$ has been observed, we cannot learn anything about $\mathbf{X}_J$ by also observing $\mathbf{X}_L$. We extend the notion of irrelevance to decision nodes in the following way. Given an ID model $\mathcal{M} = (\mathbf{C}, \mathbf{A}, \mathbf{D}, \mathbf{U})$ and sets $\mathbf{L} \subseteq \mathbf{D}$ and $\mathbf{J}$, $\mathbf{K} \subseteq \mathbf{C}$, we say that $\mathbf{L}$ is irrelevant to $\mathbf{X}_J$ given $\mathbf{X}_K$ if $\Pr(\mathbf{X}_J|\mathbf{X}_K, \mathbf{L}) = \Pr(\mathbf{X}_J|\mathbf{X}_K)$. In other words, L is irrelevant to $\mathbf{X}_J$ given $\mathbf{X}_K$ if, once $\mathbf{X}_K$ has been observed, making a decision in L does not tell us anything new about $\mathbf{X}_J$. Irrelevance for utility nodes $\mathbf{U}$ is defined analogously but with respect to expected utility rather than probability.

An ID $\mathcal{M} = (\mathbf{C}, \mathbf{A}, \mathbf{D}, \mathbf{U})$ is a BN extended with $\mathbf{D}$, a set of decision nodes, and $\mathbf{U}$, a set of utility nodes. The decision nodes are ordered in time, $d_1, \ldots, d_n$, and chance nodes are stratified into sets $\mathbf{E}$, $\mathbf{I}_1$, $\ldots$, $\mathbf{I}_{n+1}$, so that chance nodes $\mathbf{I}_i$ are observed before d_i but after d_{i-1}. The nodes in $\mathbf{E}$ have been observed a-priori while the nodes in $\mathbf{I}_{n+1}$ will never be observed. Let $\mathbf{I} = (\mathbf{I}_1, \ldots, \mathbf{I}_n)$ be nodes that are not evidence but that will be observed before the last decision d_n is made. From here on we assume, without loss of generality, that $\mathbf{E} = \emptyset$. Dealing with evidence in sampling algorithms is quite well understood and is not the focus of this paper. For

convenience, we will sometimes say $n \in \mathcal{M}$ if $n \in \{\mathbf{C} \cup \mathbf{D} \cup \mathbf{U}\}$.

Let us define a partial order Ψ_N over a set of nodes $\mathbf{N} \subseteq \mathcal{M}$ such that, given nodes $n, m \in \mathbf{N}$, if $m \in \mathrm{De}(n)$ then n appears in Ψ_N before m. In other words, in Ψ_N a node appears always before any of its descendants.

We say that $d \in (\mathbf{D}, \mathbf{I})$ is an indexing parent of $n \in \mathcal{M}$ if it precedes n in $\Psi_{(\mathbf{D},\mathbf{C},\mathbf{U})}$ and it is relevant to n. We denote by $\mathrm{IP}(n)$ the set of indexing parents of node n. Less formally, the set of indexing parents of a node n is composed of those decision nodes and those chance nodes that will be observed before making a decision, that have an effect on the distribution of the outcomes of n. A strategy is an instantiation of every node $d \in \mathbf{D}$, conditional on $\mathbf{I}$. We will sometimes use the term strategy when making reference to an instantiation of the set of indexing parents of a node.

For each node $n \in \mathcal{M}$ we can define the Cartesian product of the set of outcomes of its indexing parents:

$$\Phi_n = \times_{p \in \mathrm{IP}(n)} \mathrm{O}(p) \ .$$

In other words, Φ_n contains all the possible combinations of the outcomes of n's indexing parents. Similarly, Φ is the set of indexing parents of the global utility node

$$\Phi = \times_{p \in \bigcup_{u_i \in \mathbf{U}} \mathrm{IP}(u_i)} \mathrm{O}(p) \ .$$

We will now define the cardinality operator, $\|\Phi_n\|$, in the following way

$$\|\Phi_n\| = \begin{cases} |\Phi_n|, & \Phi_n \neq \emptyset \\ 1, & \Phi_n = \emptyset \end{cases} \ .$$

This operator, when applied to any set Φ_n, will return its number of elements or 1 if the set is empty.

Theorem 1 *Given an ID $\mathcal{M} = (\mathbf{C}, \mathbf{A}, \mathbf{D}, \mathbf{U})$ and nodes $c, p \in \{\mathbf{C}, \mathbf{D}, \mathbf{U}\}$, if $p \in \mathrm{Pa}(c)$ then $\mathrm{IP}(p) \subseteq \mathrm{IP}(c)$. In other words, every indexing parent of a node p is also an indexing parent of p's children.*

Proof: By definition, every node $n \in \mathrm{IP}(p)$ is relevant to p. Since $p \in \mathrm{Pa}(c)$, it follows that n must also be relevant to c. □

3 The algorithm

This algorithm solves an ID $\mathcal{M} = (\mathbf{C}, \mathbf{A}, \mathbf{D}, \mathbf{U})$, which means that it calculates the expected utility of each possible strategy and also calculates the distribution of each chance node conditional on each of the strategies. As one of its main features, the algorithm is capable of giving an approximate answer at any time. We describe this algorithm in Figure 1.

Input: An ID $\mathcal{M} = (\mathbf{C}, \mathbf{A}, \mathbf{D}, \mathbf{U})$
Output: For every strategy $\delta_k \in \Phi$ expected utility of δ_k.
(1) For each node $n \in \mathcal{M}$
(1.1) If $n \in \mathbf{U}$, apply Cooper's transformation to n and add n to $\mathbf{C}$
(1.2) Find $\mathrm{IP}(n)$
(1.3) Find $\|\Phi_n\|$
(2) Determine Ψ_M, a partial order over $\mathcal{M}$
(3) For each node $c \in \mathbf{C}$ (in Ψ_M order)
(3.1) Generate a random number
(3.2) For each strategy $\delta_k \in \Phi_c$ (or once if $\Phi_c = \emptyset$)
(3.2.1) Lookup the outcome of every node $p \in \mathrm{Pa}(c)$ given δ_k
(3.2.2) Instantiate the outcome of c given δ_k
(4) If you wish to terminate the algorithm, proceed to Step 5. Otherwise jump to Step 3.
(5) Normalize all $\Pr(c), c \in \mathbf{C}$
(6) For each node $u \in \mathbf{U}$, for each $\delta_k \in \Phi_u$ (or once if $\Phi_u = \emptyset$)
(6.1) Compute $EU(u|\delta_k)$, $\delta_k \in \Phi_u$ by reversing Step 1.1
(7) For each $\delta_k \in \Phi$
(7.1) Compute the global expected utility given δ_k.

Figure 1: The basic algorithm.

Cooper's transformation in Step [1.1] is a two-step operation consisting of replacing all decision nodes and all utility nodes with chance nodes. Transformation of each utility node $u \in \mathbf{U}$ into a chance node, in particular, is performed by transforming the utility function $V(u) = F(p_1, \ldots, p_n)$, $p_1, \ldots, p_n \in \mathrm{Pa}(u)$ into a probability distribution by linearly mapping

the range of F onto the interval [0,1].

In order to achieve [1.2], we can use an algorithm like the one described in Shachter (1998). Basically, for a given node $n \in \mathcal{M}$, it traverses the ID, starting from n, marking each d-connected node it finds along the way. Once the traversing is done, those nodes marked are added to the set $\mathrm{IP}(n)$.

In Step [3], we iterate only through chance nodes, which means that decision nodes are ignored at this stage. Remember that utility nodes have been transformed into chance nodes in Step [1.1]. The Steps [3.1], [3.2.1], [3.2.2], and [3.2.3] basically follow any forward sampling scheme, such as probabilistic logic sampling (Henrion, 1988), likelihood weighting (Fung and Chang, 1989; Shachter and Peot, 1989), importance sampling (Shachter and Peot, 1989), adaptive importance sampling (Cheng and Druzdzel, 2000), or approximate posterior importance sampling (Yuan and Druzdzel, 2003). This produces an unbiased estimator of the distribution of node c. But also note that in [3.2] we are iterating through every $\delta_k \in \Phi_c$ and that in Steps [3.2.1] through [3.2.3] we are assuming 'given δ_k'. The net result of Step [3] is the distribution of each chance node c conditional on each strategy relevant to c:

$$\Pr(c_i|\delta_k), c_i \in \mathbf{C}, \delta_k \in \Phi_c \; . \tag{1}$$

In [3.2.1], we say "Lookup the outcome of every node $p \in \mathrm{Pa}(c)$ given δ_k." What this "outcome" means depends on the type of the parent. If $p \in \mathrm{IP}(c)$, "outcome" means state of p in the current δ_k. On the other hand, if $p \notin \mathrm{IP}(c)$, "outcome" means the outcome of p we computed and annotated in Step [3.2.3]. The fact that we are processing nodes in the order specified by Ψ_M guarantees that p has been dealt with before c, so the outcome of p is available to us. Theorem 1 guarantees that we can always find the state of node p given the current δ_k.

Also, note that we are evaluating in parallel every $\delta_k \in \Phi_c$, reusing the random number generated in [3.1], to compute the distribution of c conditional on every strategy. A complete set of samples, one for every node, describes a state of the world that we effectively subject our decision model to. We use the same state of the world to test each of the possible decision strategies, but we do it in parallel as we continue sampling. This sample reuse achieves two important objectives: (1) it saves a large number of calls to the random number generator, and (2) with every sample generated, it performs the evaluation of all strategies, which is a key to achieving anytime behavior.

In Step [4] we need to decide whether we want to take another sample of the network. The simulation can be interrupted at this point and the algorithm will return its best estimate of the expected utilities of each of the strategies.

Step [6] performs reverse of the transformation applied in [1.1] to utility nodes. Utility nodes, indexed by their indexing parents, contain at this point the expected utility of each possible strategy:

$$EU(u_i|\delta_k), \; u_i \in \mathbf{U}, \; \delta_k \in \Phi_u \; . \tag{2}$$

Step [7] is meant to deal with those cases in which $\mathbf{U}$ has more than one utility node. In this case, a multi-attribute utility function of the form $\mathrm{GEU}(V(u_1), \ldots, V(u_n))$, $u_i \in \mathbf{U}$, should be provided (typically, it is a part of the model). This function is intended to produce a global utility and, for linearly additive utility functions, for example, it is defined as $V(u_1) + \ldots + V(u_n)$, $u_i \in \mathbf{U}$. Our algorithm does not depend in any way on the actual form of this multi-attribute utility function.

4 Empirical illustration

It is quite obvious that the proposed algorithm is more efficient in terms of its use of samples and in terms of reducing the variance in results. We illustrate this (rather than testing it empirically) on Hepar II, a model for diagnosis of liver disorders consisting of 73 nodes (Oniśko et al., 2001). Since Hepar II is a BN, we first performed a modification that transformed it into an ID. Basically, Hepar II consists of nodes representing risk factors, liver disorders, and medical tests and observations. The structure of the graph is causal: arcs go from

risk factors to disorders and from disorders to tests and observations. First, we designated four nodes from among the risk factors to act as decision nodes. These decision nodes represented hypothetical choices a person could have about certain risks factors. For example, alcohol abuse can be considered a risk factor for several liver disorders. Converting this chance node into a decision node represents the choice a person has with respect to drinking. Since each of the transformed decision nodes had two states, we ended up with 16 different strategies to evaluate. The second step was to add utility nodes. We decided to add one utility node to each of the disorder nodes such that the utility of not having a particular disorder was 100 while the utility of having the disorder was 0. For simplicity, we assumed that all disorders as equally severe and created an additive linear multi-attribute utility function with weights $w_i = 1.0$ for each of the individual utility nodes. Since the model contains 8 disorders, the total utility ranges from 800, when the patient suffers from no disorders at all, to 0, when the patient suffers from all eight disorders at the same time.

We compared our algorithm against indirect exhaustive algorithm, with both algorithms based on likelihood weighting as their underlying BN algorithm. The only difference between the two algorithms was sample reuse. We ran 1,000,000 samples of the network, taking measurements of intermediate values of expected utility for each of the 16 possible strategies.

Figure 2 shows the expected utility of each of the 16 strategies as a function of the number of samples. In the simulation, the range of expected utilities goes from 500 up to 800 (each horizontal line on the graph represents an increment of 50). This wide range is due to the variations in the estimates that we obtained during the first one hundred samples. Crossing of trajectories means essentially a change in the ranking of the strategies. The expected utilities converge to their exact values (verified by means of an exact algorithm) as the number of samples increases. Figure 3 shows the same plot for our algorithm. In this simulation, the range of expected utilities goes from 550 up to 775 (each

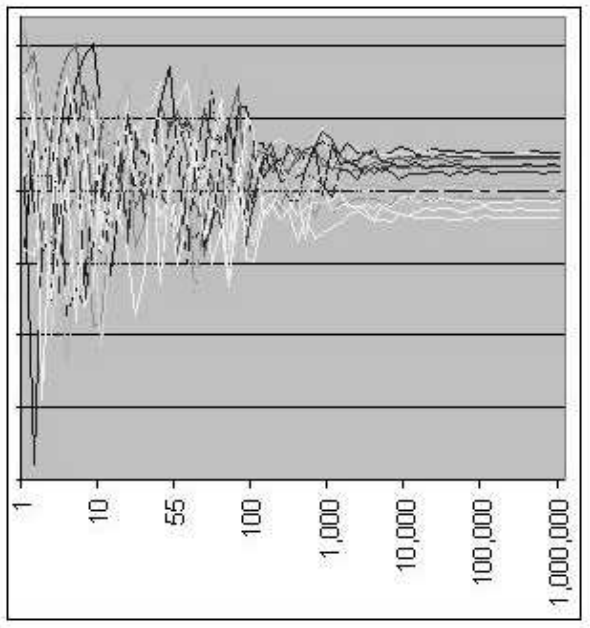

Figure 2: Expected Utility as a function of the number of samples for the crude exhaustive algorithm.

horizontal line also represents an increment of 50). The expected utilities converge to exactly the same values (670 for the worst strategy and 715 for the best one).

Both figures show individual trajectories shifting up and down, but while in the crude exhaustive algorithm each trajectory seems to do so independently of the others, in our algorithm they all shift roughly at the same time and in the same direction. This is a simple consequence of the fact that our simulation tests each strategy against the same sample, i.e., on the same randomly generated state of the world. This greatly reduces variance and our algorithm can rank the strategies after processing a relatively small number of samples. While in Figure 2 we can find trajectory lines crossing even after as many as 90,000 samples taken, in Figure 3 the last crossing of trajectory lines occurs at only 900 samples.

Reuse of samples for each of the 16 strategies reduced the number of calls to the random number generator by 16-fold. There was a 20% to 30% computational time overhead related to

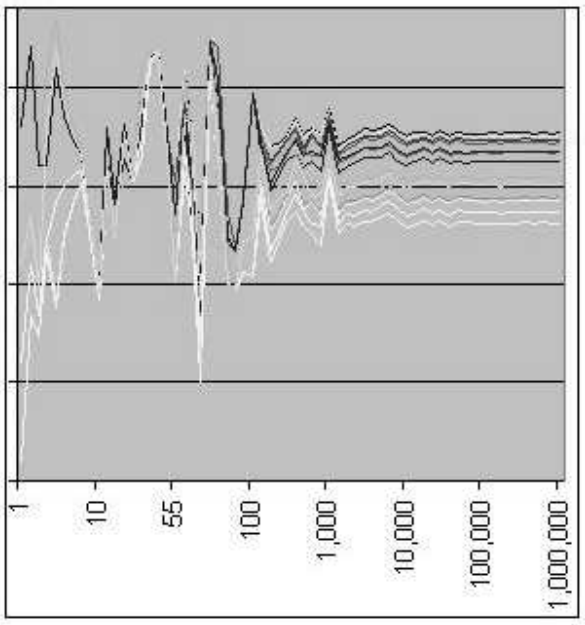

Figure 3: Expected Utility (vertical axis) vs. number of samples (horizontal axis) from our algorithm.

reuse, resulting in roughly a 12-fold reduction of the computation time of our algorithm compared to the crude algorithm.

5 Conclusion

We introduced an approximate anytime sampling algorithm for IDs that computes the expected utilities of all decision strategies. The algorithm is indirect, in the sense of reducing the problem of solving an ID first transforming it into a BN. We have shown how by evaluation of each of the strategies on the same set of samples we not only save much computation but also reduce variance and, hence, produce high quality anytime behavior. The proposed algorithm is, furthermore, amenable to parallelization and a possible further increase in speed.

Our algorithm accommodates any forward sampling scheme. A simple empirical test of the algorithm has shown that it rapidly produces the correct order of strategies and also converges fast to the correct numerical value of the expected utility. We expect that combining this algorithm with stopping rules will lead to efficient algorithms for IDs that gives precision guarantees with respect to both the order and numerical expected utilities of strategies.

Acknowledgments

This research was supported by the National Science Foundation under Faculty Early Career Development (CAREER) Program, grant IRI–9624629, and by the Air Force Office of Scientific Research, grant F49620–03–1–0187. All experimental data have been obtained using SMILE, a Bayesian inference engine developed at the Decision Systems Laboratory and available at `http://www.sis.pitt.edu/~genie`.

References

John M. Charnes and Prakash P. Shenoy. 2004. Multi-stage Monte Carlo method for solving influence diagrams using local computation. *Management Science*, 50(3):405–418.

Jian Cheng and Marek J. Druzdzel. 2000. BN-AIS: An adaptive importance sampling algorithm for evidential reasoning in large Bayesian networks. *Journal of Artificial Intelligence Research*, 13:155–188.

Robert T. Clemen. 1996. *Making Hard Decisions: An Introduction to Decision Analysis*. Duxbury Press, An Imprint of Wadsworth Publishing Company, Belmont, California.

Gregory F. Cooper. 1988. A method for using belief network algorithms to solve decision-network problems. In *Proceedings of the 1988 Workshop on Uncertainty in Artificial Intelligence, UAI–88*, Minneapolis, Minnesota.

Gregory F. Cooper. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2–3):393–405, March.

Bruce D'Ambrosio and Scott Burgess. 1996. Some experiments with real-time decision algorithms. In *Proceedings of the Twelfth Annual Conference on Uncertainty in Artificial Intelligence (UAI–96)*, pages 194–202, San Francisco, CA. Morgan Kaufmann Publishers.

Robert Fung and Kuo-Chu Chang. 1989. Weighing and integrating evidence for stochastic simulation in Bayesian networks. In M. Henrion, R.D. Shachter, L.N. Kanal, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence 5*, pages 209–219. Elsevier Science Publishing Company, Inc., New York, N. Y.

Dan Geiger, Thomas S. Verma, and Judea Pearl. 1990. Identifying independence in Bayesian networks. *Networks*, 20(5):507–534, August.

Max Henrion. 1988. Propagating uncertainty in Bayesian networks by probabilistic logic sampling. In L.N. Kanal, T.S. Levitt, and J.F. Lemmer, editors, *Uncertainty in Artificial Intellgience 2*, pages 149–163. Elsevier Science Publishing Company, Inc., New York, N. Y.

Michael C. Horsch and David Poole. 1996. Flexible policy construction by information refinement. In *Proceedings of the Twelfth Annual Conference on Uncertainty in Artificial Intelligence (UAI–96)*, pages 315–324, San Francisco, CA. Morgan Kaufmann Publishers.

Michael C. Horsch and David Poole. 1998. An anytime algorithm for decision making under uncertainty. In *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–98)*, pages 246–255, San Francisco, CA. Morgan Kaufmann Publishers.

Eric J. Horvitz. 1989. Reasoning under varying and uncertain resource constraints. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence, IJCAI–89*, pages 1121–1127, San MAteo, CA. Morgan Kauffman.

Ronald A. Howard and James E. Matheson. 1984. Influence diagrams. In Ronald A. Howard and James E. Matheson, editors, *The Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA.

Frank Jensen, Finn V. Jensen, and Søren L. Dittmer. 1994. From influence diagrams to junction trees. In *Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–94)*, pages 367–373, San Francisco, CA. Morgan Kaufmann Publishers.

Anders L. Madsen and Finn V. Jensen. 1999. Lazy evaluation of symmetric Bayesian decision problems. In *Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–99)*, pages 382–390, San Francisco, CA. Morgan Kaufmann Publishers.

P. Ndilikilikesha. 1994. Potential influence diagrams. *International Journal of Approximate Reasoning*, 11:251–285.

Scott M. Olmsted. 1983. *On Representing and Solving Decision Problems*. Ph.D. thesis, Stanford University, Department of Engineering-Economic Systems, Stanford, CA, December.

Agnieszka Oniśko, Marek J. Druzdzel, and Hanna Wasyluk. 2001. Learning Bayesian network parameters from small data sets: Application of Noisy-OR gates. *International Journal of Approximate Reasoning*, 27(2):165–182.

Judea Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., San Mateo, CA.

Marco Ramoni. 1995. Anytime influence diagrams. In *Working notes of the IJCAI–95 workshop on "Anytime Algorithms and Deliberation Scheduling"*, pages 55–62, Montreal, Canada.

Ross D. Shachter and Mark A. Peot. 1989. Simulation approaches to general probabilistic inference on belief networks. In M. Henrion, R.D. Shachter, L.N. Kanal, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence 5*, pages 221–231. Elsevier Science Publishing Company, Inc., New York, N. Y.

R. D. Shachter and M. A. Peot. 1992. Decision making using probabilistic inference methods. In *Proceedings of the Eighth Annual Conference on Uncertainty in Artificial Intelligence (UAI–92)*, pages 276–283, San Francisco, CA. Morgan Kaufmann Publishers.

Ross D. Shachter. 1986. Evaluating influence diagrams. *Operations Research*, 34(6):871–882, November–December.

Ross Shachter. 1998. Bayes-Ball: The rational pasttime (for determining irrelevance and requisite information in belief networks and influence diagrams). In *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–98)*, pages 48–487, San Francisco, CA. Morgan Kaufmann Publishers.

Prakash P. Shenoy. 1992. Valuation-based systems for Bayesian decision analysis. *Operations Research*, 40(3):463–484.

Changhe Yuan and Marek Druzdzel. 2003. An importance sampling algorithm based on evidence pre-propagation. In *Proceedings of the 19th Annual Conference on Uncertainty in Artificial Intelligence (UAI-03)*, pages 624–631, San Francisco, CA. Morgan Kaufmann Publishers.

N. L. Zhang. 1998. Probabilistic inference in influence diagrams. In *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–98)*, pages 514–522, San Francisco, CA. Morgan Kaufmann Publishers.