



Topological parameters for time-space tradeoff

Rina Dechter^{a,*}, Yousri El Fattah^b

^a *Information & Computer Science, University of California, Irvine, CA 92717, USA*

^b *Rockwell Science Center, 1049 Camino Dos Rios, Thousand Oaks, CA 91360, USA*

Received 8 June 1999; received in revised form 13 June 2000

Abstract

In this paper we propose a family of algorithms combining tree-clustering with conditioning that trade space for time. Such algorithms are useful for reasoning in probabilistic and deterministic networks as well as for accomplishing optimization tasks. By analyzing the problem structure, the user can select from a spectrum of algorithms, the one that best meets a given time-space specification. To determine the potential of this approach we analyze the structural properties of problems coming from the circuit diagnosis domain. The analysis demonstrates how the tradeoffs associated with various hybrids can be used for each problem instance. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Time-space; Topological parameters; Bayesian networks; Constraint networks; Automated inference; Optimization tasks; Hybrid algorithms; Empirical evaluation

1. Introduction

Problem solving methods can be viewed as hybrids of two main principles: inference and search. Tree-clustering is an example of an inference algorithm, while the cycle-cutset conditioning is an example of a search method. Tree-clustering algorithms are time and space exponential in the size of their cliques while search algorithms are time exponential but require only linear memory. In this paper we develop a hybrid scheme that uses inference (tree-clustering) and search (cycle-cutset conditioning) as its two extremes and, using a single, structure-based design parameter, permits the user to control the storage-time tradeoff in accordance with its problem domain and the available resources.

* Corresponding author.

E-mail addresses: dechter@ics.uci.edu (R. Dechter), yousri@rsc.rockwell.com (Y. El Fattah).

Indeed, structure-based algorithms for constraint satisfaction and probabilistic reasoning fall into two distinct classes. One class is centered on tree-clustering, the other on cycle-cutset decomposition. Tree-clustering involves transforming the original problem into a tree-like problem [14,22,26] that can then be solved by a specialized efficient tree-solving algorithm [27,28]. The transforming algorithm identifies subproblems that together form a tree, and the solutions to the subproblems serve as the new values of variables in a tree metalevel problem. The metalevel problem is called a *join-tree*. The tree-clustering algorithm is time and space exponential in the tree-width of the problem's graph. A related parameter is the *induced width* which equals the tree-width. We will use both terms interchangeably.

The *cycle-cutset* method, also called *loop-cutset conditioning*, utilizes the problem's structure in a different way. It exploits the fact that variable instantiation changes the effective connectivity of the underlying graph. A cycle-cutset of an undirected graph is a subset of its nodes which, once removed, cuts all of the graph's cycles. A typical cycle-cutset method enumerates the possible assignments to a set of cutset variables and, for each cutset assignment, solves (or reasons about) a tree-like problem in polynomial time. Thus, the overall time complexity is exponential in the size of the cycle-cutset [8,15,29]. Fortunately, enumerating all the cutset's assignments can be accomplished in linear space, yielding an overall linear space algorithm.

The first question is which method, tree-clustering or the cycle-cutset scheme, provides a better worst-case time guarantee. This question was answered by Bertele and Briochi [5] in 1972 and later reaffirmed in [31]. They showed that the minimal cycle-cutset of any graph can be much larger, and is never smaller, than its minimal tree-width. In fact, for an arbitrary graph, $r \leq c + 1$, where c is the size of the minimal cycle-cutset and r is the tree-width [5]. Consequently, for any problem instance the time guarantees accompanying the cycle-cutset scheme are never tighter than those of tree-clustering, and can even be much worse. On the other hand, while tree-clustering requires exponential space (in the induced width) the cycle-cutset requires only linear space.

Since the space complexity of tree-clustering can severely limit its usefulness, we investigate in this paper the extent to which space complexity can be reduced, while reasonable time complexity guarantees are maintained. Is it possible to have the time guarantees of tree-clustering while using linear space? On some problem instances, it is possible, specifically, on those problems whose associated graph has an induced width and a cycle-cutset of comparable sizes (e.g., on a ring, the cutset size is 1 and the tree-width is 2, leading to identical time bounds). We conjecture, however, that any algorithm that has a time bound guarantee exponential in the induced width will, on some problem instances, require exponential space in the induced width.

The space complexity of tree-clustering can be bounded more tightly using the *separator size*, which is defined as the size of the maximum subset of variables shared by adjacent subproblems in the join-tree. Indeed, some of the variants of tree-clustering algorithms send messages over the separators only and therefore can comply with space complexity that is exponential in the separator size only [22,38]. Our investigation employs the separator size to control the time-space tradeoff. The idea is to combine adjacent subproblems joined by a large separator into one bigger cluster or a subproblem so that the remaining separators are of smaller size. Once a join-tree with smaller separators is

generated, its potentially larger clusters can be solved using the cycle-cutset method or any other linear-space scheme.

In this paper we will develop a time-space tradeoff scheme that is applicable to belief network processing, constraint processing, and optimization tasks, yielding a sequence of parameterized algorithms that can trade space for time. With this scheme it will be possible to select from a spectrum of algorithms the one that best meets some time-space requirement. Algorithm tree-clustering and cycle-cutset conditioning are two extremes in this spectrum.

We investigate the potential of our scheme in the domain of combinatorial circuits. This domain is frequently used as an application area in both probabilistic and deterministic reasoning [16,18,40]. We analyze 11 benchmark combinatorial circuits widely used in the fault diagnosis and testing community [6] (see Table 1 ahead). For each circuit, the analysis is summarized in a chart displaying the time-space complexity tradeoffs for diagnosing that circuit. The analysis allows tailoring the hybrid of tree-clustering and cycle-cutset decomposition to the available memory.

Different variants of the tree-clustering transformation algorithms were developed in recent years, both for constraint processing [11] and for probabilistic inference [22–24, 26,29,37,38]. Some of these algorithms compile functions over the clusters and some restrict this information to the separators. Some are directional, query-based while others are symmetrical, compiling answers for a variety of queries. We choose to demonstrate our approach using a directional variant of tree-clustering that is query-based, which we hope simplifies the exposition. The approach is applicable to any variant of tree-clustering and can be carried out off-line, in a compilation mode, as well as online.

Section 2 gives definitions and preliminaries and introduces the time-space tradeoff ideas for belief networks. Sections 3 and 4 briefly extend these ideas to constraint networks and to optimization problems. Section 5 describes the empirical analysis and the results. Section 6 discusses related work and Section 7 gives our conclusions.

The paper assumes familiarity with the basic concepts of tree-clustering and cycle-cutset conditioning and provides only brief necessary background. For more details the reader should consult the references.

2. Probabilistic networks

2.1. Overview

2.1.1. Definitions and notations

Belief networks provide a formalism for reasoning about partial beliefs under conditions of uncertainty. It is defined by a directed acyclic graph over nodes representing random variables of interest (e.g., the temperature of a device, the gender of a patient, a feature of an object, the occurrence of an event). The arcs signify the existence of direct causal influences between the linked variables. The strength of these influences are quantified by conditional probabilities that are attached to each cluster of parents-child nodes in the network. A belief network is a concise description of a complete probability distribution. It uses the concept of a directed graph.

Definition 1 (*Directed graph*). A directed graph $G = \{V, E\}$, where $V = \{X_1, \dots, X_n\}$ is a set of elements and $E = \{(X_i, X_j) \mid X_i, X_j \in V\}$ is a set of edges. If an arc $(X_i, X_j) \in E$, we say that X_i points to X_j . For each variable X_i , $pa(X_i)$ is the set of variables pointing to X_i in G , while $ch(X_i)$ is the set of variables that X_i points to. The family of X_i includes X_i and its parent variables. A directed graph is acyclic if it has no directed cycles. In an undirected graph the direction of the arcs is ignored: (X_i, X_j) and (X_j, X_i) are identical. An undirected graph is chordal if every cycle of length at least 4 has a chord. A clique is a subgraph that is completely connected, and a maximal clique of graph is a clique that is not contained in any other clique of the graph.

Definition 2 (*Belief networks*). Let $X = \{X_1, \dots, X_n\}$ be a set of random variables over multi-valued domains, D_1, \dots, D_n . A belief network is a pair (G, P) where $G = (X, E)$ is a directed acyclic graph over the nodes X and $P = \{P_i\}$ are the conditional probability matrices over the families of G , $P_i = P(X_i \mid pa(X_i))$. An assignment $(X_1 = x_1, \dots, X_n = x_n)$ can be abbreviated as $x = (x_1, \dots, x_n)$. The belief network represents a probability distribution over X having the product form

$$P(x_1, \dots, x_n) = \prod_{i=1}^n P(x_i \mid x_{pa(X_i)}),$$

where $x_{pa(X_i)}$ denotes the projection of a tuple x over $pa(X_i)$. An evidence set e is an instantiated subset of variables. A moral graph of a belief network is an undirected graph generated by connecting the tail variables of any two head-to-head pointing arcs in G and ignoring the arrows. A belief network is a polytree if its underlying undirected (unmoralized) graph has no cycles (namely, it is a tree).

Definition 3 (*Induced width, induced graph*). An ordered graph is a pair (G, d) where G is an undirected graph and $d = X_1, \dots, X_n$ is an ordering of the nodes. The width of a node in an ordered graph is the number of its earlier neighbors. The width $w(d)$ of an ordering d , is the maximum width over all nodes. The induced width of an ordered graph, $w^*(d)$, is the width of the induced ordered graph obtained by processing the nodes recursively, from last to first; when node X is processed, all its earlier neighbors are connected. This process is also called “triangulation”. The induced (triangulated) graph is clearly chordal. The induced width of a graph, w^* , is the minimal induced width over all its orderings (for more information see [9,13]).

Example 1. Fig. 1 shows a belief network’s acyclic graph and its associated moral graph. The width of the graph in Fig. 1(b) along the ordering $d = A, B, C, D, G, E, F, H$ is 3. Since the moral graph in Fig. 1(b) is chordal no arc is added when generating the induced ordered graph. Therefore, the induced width w^* of the graph is also 3, because no smaller induced width can be attained.

The most common task over belief networks is to determine posterior beliefs of some variables. Other important tasks are *mpe*: finding the most probable explanation given a set of observations, and *map*: finding the maximum a posteriori hypotheses given evidence. Both tasks are relevant to abduction and diagnosis [29]. It is well known that such tasks

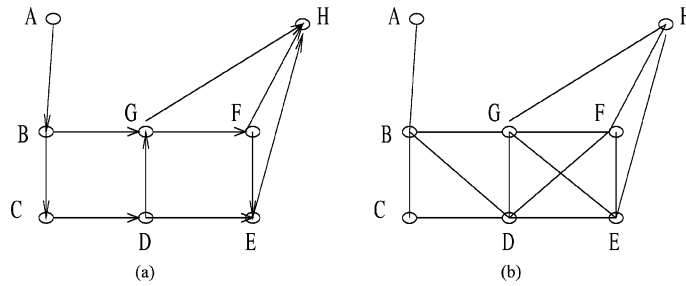


Fig. 1. (a) A belief network and (b) its moral graph.

can be answered effectively for singly-connected polytrees by a belief propagation algorithm [29] that can be extended to multiply-connected networks by either tree-clustering or loop-cutset conditioning.

2.1.2. Tree-clustering

The most widely used method for processing belief networks is join-tree clustering. The algorithm transforms the original network into a tree of subproblems called a *join-tree*. Tree-clustering methods have two parts. In the first part the structure of the newly generated tree problem is decided, and in the second part the conditional probabilities between the subproblems (viewed as high-dimensional variables) is determined. The structure of the join-tree is determined primarily by graph information, embedding the graph in a tree of cliques as follows. First the moral graph is embedded in a chordal graph by adding some edges. This is accomplished by picking a variable ordering $d = X_1, \dots, X_n$, then, moving from X_n to X_1 , recursively connecting all the earlier neighbors of X_i in the moral graph yielding the induced ordered graph. Its *induced width* $w^*(d)$, as defined earlier, is the maximal number of earlier neighbors each node has.

Clearly, each node and its earlier neighbors in the induced graph are a clique. The maximal cliques, indexed by their latest variable in the ordering, can be connected into a *clique-tree* and can serve as the subproblems (or clusters) in the final join-tree. The clique-tree is created by connecting every clique C_i to an earlier clique C_j , called its parent, with whom it shares a maximal number of variables. Clearly, the induced width $w^*(d)$ equals the size of the maximal clique minus 1.

Once the join-tree structure is determined, each conditional probability table (CPT) is placed in a clique containing all its arguments. The marginal probability distributions for each clique can then be computed by multiplying all the CPTs and normalizing, and subsequently the conditional probabilities between every clique and its parent clique can be derived. Tree-clustering, therefore, is time and space exponential in the size of the maximal clique, namely, exponential in the moral graph's induced width (plus 1). In Fig. 1(b), the maximal cliques of the chordal graph are $\{(A, B), (B, C, D), (B, D, G), (G, D, E, F), (H, G, F, E)\}$, resulting in the join-tree structure given in Fig. 3(a).

Clearly, the established connection between induced width and complexity motivates finding an ordering with a smallest induced width, a task known to be hard [1,2]. However, useful greedy heuristics as well as approximation algorithms were developed in the last

Algorithm directional join-tree clustering (DTC)

Input: A belief network (G, P) , where G is a DAG and $P = \{P_1, \dots, P_n\}$, evidence e .

Output: the belief of X_1 given evidence e .

1. Generate a join-tree clustering of G , identified by its cliques C_1, \dots, C_t . Place each P_i and each observation in one clique that contains its arguments. (One's favorite method for structuring into tree-clustering can be used.)
2. Impose directionality on the join-tree, namely create a rooted directed tree whose root is a clique containing the queried variable. Let $d = C_1, \dots, C_t$ be a breadth-first ordering of the rooted clique-tree, let $S_{p(i)}$ and $C_{p(i)}$ be the parent separator and the parent clique of C_i , respectively.
3. From $i \leftarrow t$ downto 1 do
4. (Processing clique C_i):

Let $\lambda_1, \lambda_2, \dots, \lambda_j$ be the functions in clique C_i , and where C_i denotes also its set of variables.

 - For any observation $X_j = x_j$ in clique C_i substitute x_j in each function over x_j in the clique.
 - Let $U_i = C_i - S_{p(i)}$ and let u_i be an assignment to U_i . Compute function

$$\lambda = \sum_{u_i} \prod_{i=1}^j \lambda_i.$$
 Put λ in parent clique $C_{p(i)}$.
5. **Return** (processing root clique, C_1) Let $\lambda_1, \dots, \lambda_j$ be functions in C_1

$$Bel(x_1) = \alpha \sum_{u_1} \prod_i \lambda_i$$

$$\alpha \text{ is a normalizing constant and } u_1 \text{ is an assignment to } U_1 = C_1 - \{X_1\}.$$

Fig. 2. Algorithm directional join-tree clustering.

decade, taking into account not only graph information, but also the variance in the variables' domain sizes [4,5,14,24,26,29,39].

A tighter bound on the space complexity of tree-clustering may be obtained using the *separator size*. The separator size of a join-tree is the maximal size of the intersections between any two cliques, and the *separator size of a graph* is the minimal separator size over all the graph's join-trees [7].

Algorithm *directional join-tree clustering* (DTC), presented in Fig. 2, is a query-based variant of join-tree clustering. It records functions on separators only. The algorithm can be viewed as rephrasing one phase (the collect phase) in both the Shafer–Shenoy architecture [36,38], as well as the one-phase propagation in the Hugin architecture. The Hugin architecture has its roots in the method proposed by Lauritzen and Spiegelhalter [26] for computing marginals of probability distributions. It was proposed by Jensen et al. [22] and is incorporated in the software product Hugin.

Step 1 of algorithm DTC computes the join-tree by first determining the cliques and connecting them in a tree-structure where each clique has a parent clique and a parent separator. Once the structuring part of the join-tree is determined, each CPT is placed in one clique that contains its arguments. For example, given the join-tree structure in Fig. 3(a), the cliques contain the CPTs as follows. Clique AB contains $P(B|A)$; BCD contains $P(C|B)$ and $P(D|C)$; BDG contains $P(G|B, D)$; $GDEF$ contains $P(E|D, F)$ and $P(F|G)$; and finally, clique $GEFH$ contains $P(H|G, F, E)$. Subsequently, algorithm DTC processes the cliques recursively from leaves to the root. Processing a clique involves computing the product of all the probabilistic functions that reside in that clique and then summing over all the variables that do not appear in its parent separator. The computed

function (over the parent separator) is added to the parent clique. Computation terminates at the root clique.¹

As shown ahead, the algorithm tightens the bound on space complexity using its separator size. This modification to its space management can be applied to any variant of tree-clustering which is not necessarily query-based. In summary,

Theorem 1 (Time-space complexity of join-tree clustering). *Given a belief network having n variables, whose moral graph can be embedded in a clique-tree having induced width r and separator size s , the time complexity for determining beliefs and the mpe by a join-tree clustering algorithm (e.g., by DTC) is $O(n \cdot \exp(r))$, while its space complexity is $O(n \cdot \exp(s))$.*

Proof. It is well known that the time and space complexity of join-tree clustering is bounded by the induced width (clique sizes) of the graph, r . The only thing that needs to be shown, therefore, is that the tighter bound on space complexity is valid. Consider the processing of a clique (step 4 in Algorithm DTC). Let C be the variables in the clique, let S be the variables in the separator with its parent clique and let $U = C - S$. In step 4 we compute $\lambda = \sum_U \prod_{i=1}^j \lambda_i$. Namely, λ is a function defined over S , because all the variables in U are eliminated by summation. Namely, for each assignment s to S , $\lambda(s)$ can be computed in linear space as follows: we initialize $\lambda(s) \leftarrow 0$, and then for every assignment u to U we compute the running sum: $\lambda(s) \leftarrow \lambda(s) + \prod_i \lambda_i(s, u)$. \square

Clearly $s \leq r$. Note that since in the example of Fig. 3(a) the separator size is 3 and the induced width is also 3, we do not gain much space-wise by the modified algorithm. There are, however, many cases where the separator size is much smaller than the induced width.

Algorithm directional join-tree clustering (DTC) can be adapted for the task of finding the most probable explanation (mpe) by replacing the summation operation by maximization. Although it is presented as an online algorithm, when evidence variables are known in advance, the algorithm can exploit this information and simplify the structure of the compiled join-tree off-line. This can be done by treating the evidence variables as assigned variables.

2.1.3. Cycle-cutset conditioning

Belief networks may be processed also by cycle-cutset conditioning [29]. A subset of nodes of an undirected graph is called a *cycle-cutset* if removing all the edges incident to nodes in the cutset makes the graph cycle-free. A subset of nodes of an acyclic directed graph is called a *loop-cutset* if removing all the outgoing edges of nodes in the cutset results in a polytree [29,30]. A minimal cycle-cutset (respectively, minimal loop-cutset) is such that if one node is removed from the set, the set is no longer a cycle-cutset (respectively, a loop-cutset).

Algorithm cycle-cutset-conditioning (also called cycle-cutset decomposition or loop-cutset conditioning) is based on the observation that assigning a value to a variable

¹ We disregard algorithmic details that do not affect asymptotic worst-case analysis here.

changes the effective connectivity of the network. Graphically, this amounts to removing all outgoing arcs from the assigned variables. Consequently, an assignment to a subset of variables that constitute a loop-cutset means that belief updating, conditioned on this assignment, can be carried out in the resulting polytree [29]. Multiply-connected belief networks can therefore be processed by enumerating all possible instantiations of a loop-cutset and solving each conditioned network using the polytree algorithm. Subsequently, the conditioned beliefs are combined using a weighted sum where the weights are the probabilities of the joint assignments to the loop-cutset variables conditioned on the evidence. Pearl [29] showed that weight computation is not more costly than enumerating all the conditioned beliefs.

This scheme was later simplified by Peot and Shachter [30]. They showed that if the polytree algorithm is modified to compute the probability of each variable-value proposition *conjoined* with the evidence, rather than *conditioned* on the evidence, the weighted sum can be replaced by a simple sum. In other words:

$$P(x|e) = \alpha P(x, e) = \alpha \sum_c P(x, e, c).$$

If $\{X\} \cup C \cup E$ is a loop-cutset (note that C and E denote subsets of variables) then $P(x, e, c)$ can be computed efficiently using a propagation-like algorithm on polytrees. Consequently the complexity of the cycle-cutset scheme is exponential in the size of C where $C \cup \{X\} \cup E$ is a loop-cutset. Some additional improvements are presented in [15]. In summary,

Theorem 2 [29,30]. *Given a belief network (G, P) having n variables, and family sizes bounded by $|F|$, and a loop-cutset bounded by c , belief updating and mpe can be computed in time $O(\exp(|F|) \cdot \exp(c))$ and in linear space.²*

2.2. Trading space for time

Assume now that we have a problem whose join-tree has induced width r and separator size s but space restrictions do not allow the necessary $O(\exp(s))$ memory required by tree-clustering. One way to overcome this problem is to collapse cliques joined by large separators into one big cluster. The resulting join-tree has larger subproblems but smaller separators. This yields a sequence of tree-decomposition algorithms parameterized by the sizes of their separators.

Definition 4 (*Primary and secondary join-trees*). Let T be a clique-tree embedding of the moral graph of G . Let s_0, s_1, \dots, s_n be the sizes of the separators in T listed in strictly descending order. With each separator size s_i , we associate a tree decomposition T_i generated by combining adjacent clusters whose separator sizes are strictly greater than s_i . $T = T_0$ is called the primary join-tree, while T_i , when $i > 0$, is a secondary join-tree. We denote by r_i the largest cluster size in T_i .

² Another bound often used is $O(n \cdot \exp(c+2))$ where the “2” in the exponent comes from the fact that belief updating on binary trees is linear in the size of the CPTs between pairs of variables which are at least $O(k^2)$ where k bounds the domains of the variables.

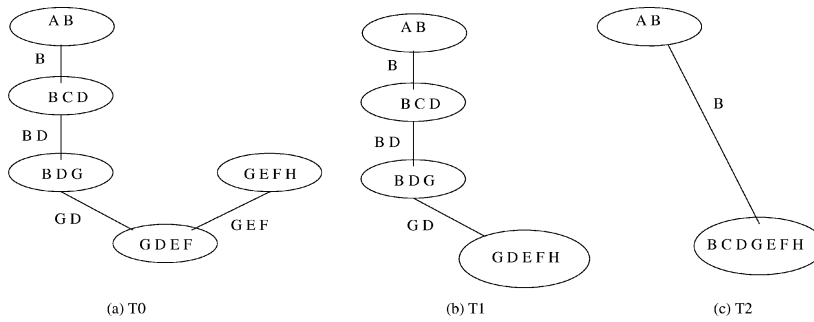


Fig. 3. A tree-decomposition with separators equal to (a) 3, (b) 2, and (c) 1.

Note that as s_i decreases, r_i increases. Clearly, from Theorem 1 it follows that

Theorem 3. Given a join-tree T over n variables, having separator sizes s_0, s_1, \dots, s_t and corresponding secondary join-trees having maximal clusters, r_0, r_1, \dots, r_t , belief updating and mpe can be computed using any one of the following time and space bounds $O(n \cdot \exp(r_i))$ time, and $O(n \cdot \exp(s_i))$ space (i ranges over all of the secondary join-trees), respectively.

Proof. For each i , a secondary tree T_i is a structure underlying a possible execution of directional join-tree clustering. From Theorem 1 it follows that the time complexity is bounded exponentially by the corresponding clique size (e.g., r_i) and space complexity is bounded exponentially by the corresponding separator size, s_i . \square

Example 2. If in our example we allow only separators of size 2, we get the join-tree T_1 in Fig. 3(b). This structure suggests that we can update beliefs and compute mpe in time which is exponential in the largest cluster, 5, while using space exponential in 2. If space considerations allow only singleton separators, we can use the secondary tree T_2 in Fig. 3(c). We conclude that the problem can be solved, either in $O(k^4)$ time (k being the maximum domain size) and $O(k^3)$ space using the primary tree T_0 , or in $O(k^5)$ time and $O(k^2)$ space using T_1 , or in $O(k^7)$ time and $O(k)$ space using T_2 .

We know that finding the smallest induced width of a graph (or finding a join-tree having smallest cliques) is NP-complete [2,35]. Nevertheless, many greedy ordering algorithms provide useful upper bounds. We denote by w_s^* the smallest induced width among all the tree embeddings of G whose separators are of size s or less. Finding w_s^* may be hard as well, however. We can conclude that: Given a belief network BN , for any $s \leq n$, if $O(\exp(s))$ space can be used, then belief updating and mpe can potentially be computed in time $O(\exp(w_s^* + 1))$.

2.2.1. Using the cycle-cutset scheme within cliques

Finally, instead of executing a brute-force algorithm to compute the marginal distributions over the separators in each clique (see step 4 in Algorithm DTC), we can use the

loop-cutset scheme. Given a clique C_p with a separator parent set S_p , step 4 computes a function defined over the separator, by

$$\lambda_p = \sum_{u_p} \prod_{i=1}^j \lambda_i,$$

where $U_p = C_p - S_p$. This seems to suggest that we have to enumerate explicitly all tuples over C_p . However, we observe that when computing λ_p for a particular value assignment of the separator x_s , those assignments can be viewed as cycle-breaking values in the graph. So, when the separator constitutes a loop-cutset, then for each tuple x_p the sum can be computed in linear time, either by propagation over the resulting polytree or by an equivalent variable elimination procedure [10].

If the instantiated separator set does not cut all loops in the cluster we can add additional nodes from the clique until we get a full loop-cutset. If the resulting loop-cutset (containing the separator variables) has size c_s , the clique's processing is time exponential in c_s only and not in the full size of the clique.

In summary, given a join-tree decomposition, in every clique we can choose a loop-cutset of clique C_i that is a minimal subset of variables, which together with its parent separator set constitute a loop-cutset of the subnetwork defined over C_i . Note that the cutset is defined relative to the original subgraph that is determined by the clique. We conclude:

Theorem 4. *Let n be the number of variables in a belief network. Given a constant $s \leq n$, let T_s be a clique-tree whose separator size has size s or less, and let c_s^* be the maximum size of a minimal cycle-cutset in any subgraph defined by the cliques in T_s . Then belief assessment and mpe can be computed in space $O(n \cdot \exp(s))$ and in time $O(n \cdot \exp(c_s^*))$, where $c_s^* \geq s$, while c_s^* is smaller than the clique size.*

Proof. Since computation in each clique is done by the cycle-cutset conditioning, time is exponentially bounded by c_s^* , the maximal cycle-cutset over all the cliques of T_s . The space complexity remains exponential in the maximum separator size s . Since for every clique, the loop-cutset we select contains its parent separator, we clearly have $c_s^* > s$. \square

Algorithm STC is presented in Fig. 4. We conclude with an example that demonstrates the time-space tradeoff when using cycle-cutset in each clique.

Example 3. Considering the join-trees in Fig. 3, if we apply the cycle-cutset scheme inside each subnetwork defined by each clique, we get no improvement in the bound for T_0 because the largest loop-cutset size in each cluster is 3 since it always exceeds the largest separator. (Remember also that once a loop-cutset is instantiated, processing the simplified network by propagation or by any efficient method is $O(k^2)$.) However, when using the secondary tree T_1 , we can reduce the time bound from $O(k^5)$ to $O(k^4)$ (while still using only $O(\exp(2))$ space) because the cutset size of the largest subgraph restricted to $\{G, D, E, F, H\}$ is 2; in this case the separator $\{G, D\}$ is already a loop-cutset and, therefore, when applying algorithm cutset-conditioning to this subnetwork the overall time complexity is now $O(k^4)$. When applying conditioning to the clusters in T_2 , we get

Algorithm space-based join-tree clustering (STC(s))

Input: A belief network (G, P) , where G is a DAG and $P = \{P_1, \dots, P_n\}$, a space parameter s , evidence e .

Output: The belief of X_1 given evidence e .

1. Generate a join-tree clustering of G , and call it T_0 .
2. Generate the secondary join-tree by combining any two adjacent cliques whose separator is strictly larger than s . Let C_1, \dots, C_t be the cliques in the resulting secondary join-tree. Place each P_i and each observation in one clique that contains its arguments.
3. Impose directionality on the secondary join-tree. Let $d = C_1, \dots, C_t$ be a breadth-first ordering of the rooted clique-tree, let $S_{p(i)}$ and $C_{p(i)}$ be the parent separator and the parent clique of C_i , respectively.
4. For $i \leftarrow t$ downto 1 do (Processing clique C_i with cycle-cutset):
 - Find a subset of variables $I_i \subseteq C_i$ s.t. $I_i \cup S_{p(i)}$ is a loop-cutset of the subgraph of G restricted to nodes C_i .
 - Let $\lambda_1, \lambda_2, \dots, \lambda_j$ be the functions in clique C_i ,
 - For any observation $X_j = x_j$ in C_i assign x_j to each function.
 - For every assignment \bar{x} of $S_{p(i)}$ do,
 - $\lambda_i(\bar{x}) \leftarrow 0$.
 - For every assignment \bar{y} of I_i do ($U_i = C_i - I_i - S_{p(i)}$)
 - Using the cycle-cutset scheme compute:

$$\lambda(\bar{x}, \bar{y}) \leftarrow \sum_{\{u_i | S_{p(i)} = \bar{x}, I_i = \bar{y}\}} \prod_{m=1}^j \lambda_m.$$
 - $\lambda_i(\bar{x}) \leftarrow \lambda_i(\bar{x}) + \lambda(\bar{x}, \bar{y})$.
 - Put λ_i in parent clique $C_{p(i)}$.
5. **Return** (processing root clique, C_1),
 - $Bel(x_1) = \alpha \sum_{u_1} \prod_m \lambda_m$,
 - α is a normalizing constant.

Fig. 4. Algorithm space-based join-tree clustering.

a time bound of $O(k^5)$ with just $O(k)$ space because the loop-cutset of the subnetwork over $\{B, C, D, G, E, F, H\}$ has three nodes only, $\{B, G, E\}$. In summary, the dominating tradeoffs (when considering only the exponents) are between an algorithm based on T_1 that requires $O(k^4)$ time and quadratic space and an algorithm based on T_2 that requires $O(k^5)$ time and $O(k)$ space.

3. Constraint networks

Constraint networks have proven successful in modeling mundane cognitive tasks such as vision, language comprehension, default reasoning, and abduction, as well as in applications such as scheduling, design, diagnosis, and temporal and spatial reasoning. In general, constraint satisfaction tasks are computationally intractable.

Definition 5 (*Constraint network*). A *constraint network* consists of a finite set of variables $X = \{X_1, \dots, X_n\}$, each associated with a domain of discrete values, D_1, \dots, D_n , and a set of constraints, $\{R_1, \dots, R_t\}$. A constraint C_i has two parts: (1) the subset of variables

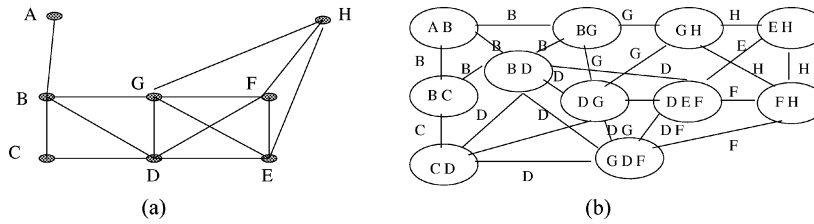


Fig. 5. Primal (a) and dual (b) constraint graphs.

$S_i = \{X_{i_1}, \dots, X_{i_{j(i)}}\}$, on which the constraint is defined, called its *scope*, and (2) a *relation*, rel_i , defined over S_i : $rel_i \subseteq D_{i_1} \times \dots \times D_{i_{j(i)}}$, whose tuples are all the compatible value assignments. The *scheme* of a constraint network is the set of scopes on which constraints are defined. An assignment of a unique domain value to each member of some subset of variables is called an *instantiation*. A consistent instantiation of *all* the variables that does not violate any constraint is called a *solution*. Typical queries associated with constraint networks are to determine whether a solution exists and to find one or all solutions. A *primal constraint graph* represents variables by nodes and associates an arc with any two nodes residing in the same constraint. A *dual constraint graph* represents each constraint scope by a node and associates a labeled arc with any two nodes whose scopes share variables. The arcs are labeled by the shared variables.

Fig. 5(a) and (b) present the primal and the dual graphs of a constraint problem having the scheme $\{AB, BC, CD, BD, BG, DG, GDF, GH, DEF, EH, FH\}$. Tree-clustering for constraint networks is similar to join-tree clustering for probabilistic networks. In fact, the structuring part is identical. Once the join-tree structure is determined, each constraint is placed in a clique (or a cluster) that contains its scope and then each clustered subproblem can be solved independently.³ The *time and space complexity* of tree-clustering is governed by the time and space required to generate the constraints (relations) of each clique in the join-tree. This process is exponential in the clique's size, and therefore in the problem's induced width w^* [9,14]. Since the graph in Fig. 5(a) is identical to the graph in Fig. 1(b), it possesses the same clique-tree embeddings.

Refining the clustering method for constraint networks can be done just as we did for probabilistic networks, thus tree-clustering in constraint networks obeys similar time and space complexities. The directional version of join-tree clustering for finding a solution to a set of constraints is given in Fig. 6. We can show:

Theorem 5 (Time-space complexity of tree-clustering, [14]). *Given a constraint problem over n variables whose constraint graph can be embedded in a clique-tree having induced width r and separator size s , the time complexity of tree-clustering for deciding consistency and for finding one solution is $O(n \cdot \exp(r))$ and its space complexity is $O(n \cdot \exp(s))$. The time complexity for generating all solutions is $O(n \cdot \exp(r) + |\text{solutions}|)$, also requiring $O(n \cdot \exp(s))$ memory.*

³ By "solving" we mean listing the relation of all solutions.

Algorithm directional tree-clustering for CSPs

Input: A set of constraints R_1, \dots, R_l over $X = \{X_1, \dots, X_n\}$, having scopes S_1, \dots, S_l , respectively, and its primal constraint graph G .

Output: A solution to the constraint problem.

1. Generate a join-tree clustering of G , identified by its cliques C_1, \dots, C_l . Place each R_i in one clique that contains its scope.
2. Impose directionality on the join-tree, namely create a rooted directed tree whose root is any clique. Let $d = C_1, \dots, C_l$ be a breadth-first ordering of the rooted clique-tree, let $S_{p(i)}$ and $C_{p(i)}$ be the parent separator and the parent clique of C_i , respectively.
3. From $i \leftarrow l$ downto 1 do
4. (Processing clique C_i):
 Let R_1, R_2, \dots, R_j be the constraints in clique C_i . Let U_i be the set of variables in clique C_i .
 - Solve the subproblem in C_i and call the set of solutions ρ_i . Project this set of solutions on the parent separator. Let $\rho_{S_{p(i)}}$ be the projected relation. Formally, (using the database projection operator \prod and join operator \bowtie) compute $\rho_{S_{p(i)}} \leftarrow \prod_{S_{p(i)}} \bowtie_{k=1}^j R_k$
 Put $\rho_{S_{p(i)}}$ in parent clique $C_{p(i)}$.
5. **Return** a solution generated in a backtrack-free manner going from the root clique towards the leaves.

Fig. 6. Algorithm directional tree-clustering for constraints.

When the space required by clustering is beyond the available resources, tree-clustering can be coerced to yield smaller separators and larger subproblems, as we have seen earlier for processing belief networks. This leads to a conclusion similar to Theorem 3.

Theorem 6. *Given a constraint network over n variables whose constraint graph can be embedded in a primary clique-tree having separator sizes s_0, s_1, \dots, s_k , whose corresponding maximal clique sizes in the secondary join-trees are r_0, r_1, \dots, r_k , then deciding consistency and finding a solution can be accomplished using any of the time and space complexity bounds $O(n \cdot \exp(r_i))$ and $O(n \cdot \exp(s_i))$, respectively.*

Proof. Analogous to Theorem 3. \square

Finally, similar to belief networks, any linear-space method can replace backtracking for solving each of the subproblems defined by the cliques. One possibility is to use the cycle-cutset scheme. The cycle-cutset method for constraint networks (like in belief networks) enumerates the possible solutions to a set of cycle-cutset variables and, for each consistent cutset assignment, solves the restricted tree-like problem in polynomial time. Thus, the overall time complexity is bounded by $O(n \cdot k^{c+2})$, where c is the cutset size, k is the domain size, and n is the number of variables [8]. Therefore,

Theorem 7. *Let G be a constraint graph over n variables and let T be a corresponding join-tree with separator size s or less. Let c_s be the largest minimal cycle-cutset⁴ in any*

⁴ As before, the cycle-cutset contains the separator set.

subproblem in T . Then the problem can be solved in space $O(n \cdot \exp(s))$ and in time $O(n \cdot \exp(c_s + 2))$, where $c_s \geq s$.

Example 4. Applying the cycle-cutset method to each subproblem in T_0, T_1, T_2 (see Fig. 3) yields the same time-space tradeoffs as for the belief network case.

A special case of Theorem 7, observed before in [13,17], occurs when the graph is decomposed into non-separable components (i.e., when the separator size equals 1).

Corollary 1. *If G has a decomposition to non-separable components such that the size of the maximal cutsets in each component is bounded by c , then the problem can be solved in $O(n \cdot \exp(c))$ time, using linear space.*

4. Optimization tasks

Clustering and conditioning are applicable also to optimization tasks defined over probabilistic and deterministic networks. An optimization task is defined relative to a real-valued criterion or cost function associated with every instantiation. In the context of constraint networks, the task is to find a consistent instantiation having optimal cost. Applications include diagnosis and scheduling problems. In the context of probabilistic networks, the criterion function denotes a utility or a value function, and the task is to find an assignment to a subset of decision variables that maximize the expected criterion function. Applications include planning and decision making under uncertainty. If the criterion function is decomposable, its structure can be incorporated into the corresponding graph to subsequently be exploited by either tree-clustering or cycle-cutset conditioning.

Definition 6 (*Decomposable criterion function* [3,25]). A criterion function over a set X of n variables $\{X_1, \dots, X_n\}$ having value domains D_1, \dots, D_n is additively decomposable relative to a scheme Q_1, \dots, Q_t where $Q_i \subseteq X$ iff

$$f(x) = \sum_{i \in T} f_i(x_{Q_i}),$$

where $T = \{1, \dots, t\}$ is a set of indices denoting the subsets of variables $\{Q_i\}$ and x is an instantiation of all the variables. The functions f_i are the components of the criterion function and are specified, in general, by stored tables.

Definition 7 (*Constraint optimization, augmented graph*). Given a constraint network over a set of n variables $X = X_1, \dots, X_n$ and a set of constraints C_1, \dots, C_t having scopes S_1, \dots, S_t , and given a criterion function f decomposable into $\{f_1, \dots, f_l\}$ over Q_1, \dots, Q_l , the constraint optimization problem is to find a consistent assignment $x = (x_1, \dots, x_n)$ such that the criterion function $f = \sum_i f_i$, is maximized. The *augmented constraint graph* contains a node for each variable and an arc connecting any two variables that appear either in the same scope of a constraint or in the same functional component of the criterion function.

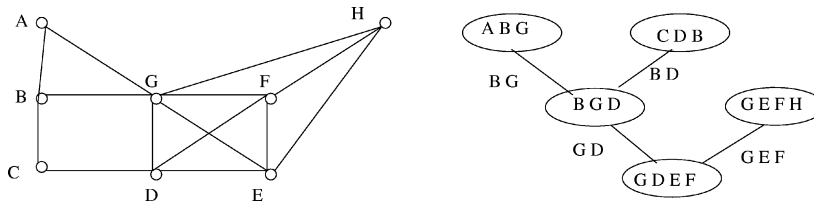


Fig. 7. An augmented moral graph for the utility function $f(a, b, c, d, e, f, g, h) = a \cdot g + c^2 + 5d \cdot e \cdot f$.

Since constraint optimization can be performed in linear time when the augmented constraint graph is a tree, both join-tree clustering and cutset conditioning can be extended to non-tree structures [32] in the usual manner. We can conclude:

Theorem 8 (Time-space of constraint optimization [32]). *Given a constraint optimization problem over n variables whose augmented constraint graph has a cycle-cutset of size c , and whose augmented graph can be embedded in a clique-tree having tree-width r and separator size s , the time complexity of finding an optimal consistent solution using tree-clustering is $O(n \cdot \exp(r))$ and the space complexity $O(n \cdot \exp(s))$. The time complexity for finding a consistent optimal solution using the cycle-cutset conditioning is $O(n \cdot \exp(c))$ while its space complexity is linear.*

In a similar manner, the structure of the criterion function can augment the moral graph when computing the maximum expected utility (MEU) of some decisions in a general influence diagram [20]. For more details see [12].

Once we have established the graph that guides tree-clustering and conditioning, the same principle of trading space for time becomes applicable and will yield a collection of parameterized algorithms governed by the primary and secondary clique-trees and cycle-cutsets of the augmented graphs as we have seen before. For completeness sake we restate the full theorem:

Theorem 9. *Given a constraint network over n variables and given an additively decomposable criterion function f , if the augmented constraint graph relative to the criterion function can be embedded in a clique-tree having separator sizes s_0, s_1, \dots, s_k , and corresponding maximal clique sizes r_0, r_1, \dots, r_k and corresponding maximal cycle-cutset sizes c_0, c_1, \dots, c_k , then finding an optimal solution can be accomplished using any one of the following bounds on the time and space: If a brute-force approach is used for processing each subproblem the bounds are $O(n \cdot \exp(r_i))$ time and $O(n \cdot \exp(s_i))$ space. If cycle-cutset conditioning is used for each cluster, the bounds are $O(n \cdot \exp(c_i))$ time and $O(n \cdot \exp(s_i))$ space, where $c_i \geq s_i$.*

Example 5. Consider the following criterion function defined over the constraint network in Fig. 5(a)

$$u(a, b, c, d, e, f, g, h) = a \cdot g + c^2 + 5d \cdot e \cdot f.$$

Table 1
ISCAS'85 benchmark circuit characteristics

Circuit name	Circuit function	Total gates	Input lines	Output lines
C17		6	5	2
C432	Priority decoder	160 (18 EXOR)	36	7
C499	ECAT	202 (104 EXOR)	41	32
C880	ALU and control	383	60	26
C1355	ECAT	546	41	32
C1908	ECAT	880	33	25
C2670	ALU and control	1193	233	140
C3540	ALU and control	1669	50	22
C5315	ALU and selector	2307	178	123
C6288	16-bit multiplier	2406	32	32
C7552	ALU and control	3512	207	108

Here the augmented graph will have one additional arc connecting nodes A and G (see Fig. 7(a)), resulting in a primary clique-tree embedding in Fig. 7(b) that differs from the tree in Fig. 3(a). As a result one has to consider the clique ABG instead of the original clique AB . Thus, applying join-tree clustering to the primary tree yields time complexity $O(\exp(4))$ and space complexity $O(k^3)$. If only binary functions can be recorded, we will need to combine clique $(GDEF)$ with $(GEFH)$ yielding a clique of size 5. Using cycle-cutset conditioning, this results in time complexity of $O(k^4)$ as well, while using only $O(k^2)$ space. If this space requirement is too heavy, we need to solve the whole problem as one cluster using cycle-cutset conditioning which, in this case, requires $O(k^5)$ time and linear space.

5. Empirical framework and results

The motivation for the experiments is twofold. One, to analyze the structural parameters of tree-clustering and the cycle-cutset on real-life instances. Two, to gain further understanding of how time-space tradeoff can be exploited to alleviate space bottlenecks. We analyzed empirically benchmark combinatorial circuits, widely used in the fault diagnosis and testing community [6] (see Table 1). The experiments allow us to assess in advance the complexity of diagnosis and abduction tasks on those circuits, and to determine the appropriate combination of tree-clustering and cycle-cutset methods to perform those tasks for each instance. None of the circuits are trees and they all have considerable node fanout as shown in the schematic diagram of circuit C432 in Fig. 8.

A directed acyclic graph (DAG) is computed for each circuit. The graph includes a node for each variable in the circuit. For every gate in the circuit, the graph has an edge directed from each gate's input to the gate's output. The nodes with no parents (children) in the

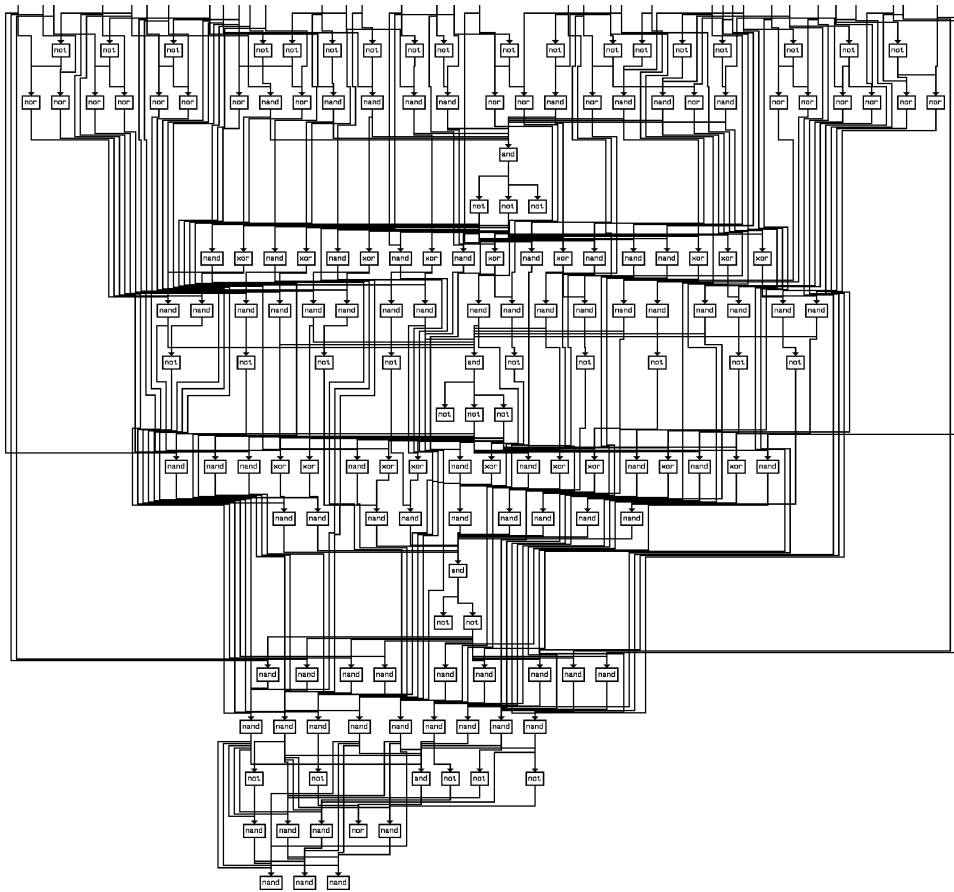


Fig. 8. Schematic of circuit C432: 36 inputs, 7 outputs, and 160 components.

Table 2

Number of nodes and edges for the primal graphs of the circuits

Circuit	C17	C432	C499	C880	C1355	C1908	C2670	C3540	C5315	C6288	C7552
#nodes	11	196	243	443	587	913	1426	1719	2485	2448	3719
#edges	18	660	692	1140	1660	2507	3226	4787	7320	7184	9572

DAG are the primary inputs (outputs) of the circuit. The primal graph for each circuit is then computed as the moral graph for the corresponding DAG. Table 2 gives the number of nodes and edges of the primal graph for each circuit.

Tree-clustering is performed on the primal graphs as usual: by selecting an ordering for the nodes, then triangulating the graph and identifying its maximum cliques. There are many possible heuristics for ordering the nodes with the aim of obtaining a join-tree with small cliques.

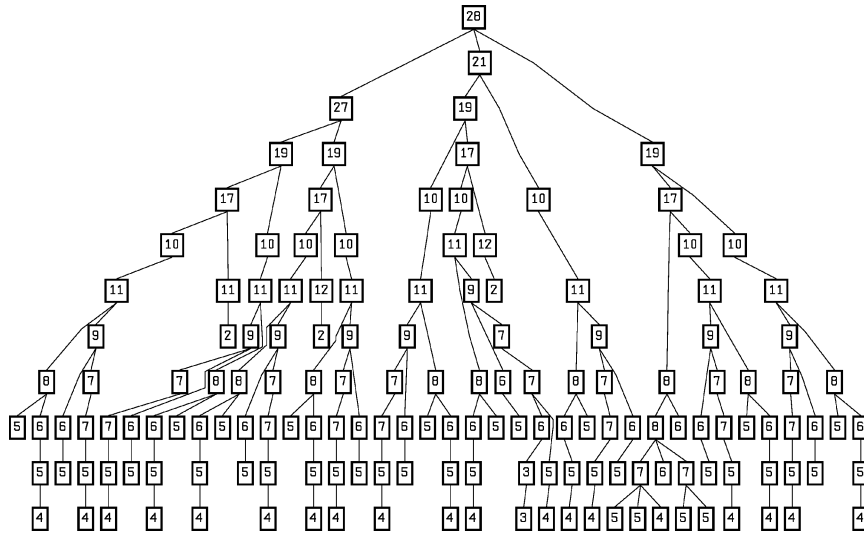


Fig. 9. Primary join-tree (157 cliques) for circuit C432 (196 variables); the maximum separator size is 23.

Four ordering heuristics were considered:

- (1) maximum-cardinality ordering,
- (2) minimum-width ordering,
- (3) minimum-degree ordering, and
- (4) causal ordering.

The max-cardinality ordering is generated from first to last by picking the first node arbitrarily and then repeatedly selecting the unordered node that is adjacent to the maximum number of already ordered nodes. The min-width ordering is computed from last to first by repeatedly selecting the node having the least number of neighbors in the graph, removing the node and its incident edges from the graph, and continuing until the graph is empty. The min-degree ordering [5] is exactly like min-width except that we connect neighbors of selected nodes, and the causal ordering is just a topological sort of the directed graph for the circuit.

The structural portion of tree-clustering was implemented using each of the four orderings on each of the benchmark circuits of Table 1 and we observed that the min-degree ordering was by far the best, yielding the smallest cliques sizes and separators. Our evaluation of the performance of the orderings is consistent with the results in [23]. Therefore, we report here the results only relative to the *min-degree ordering*. For results on the other orderings see [12].

5.1. Results: Primary join-trees

For each primary join-tree generated, three parameters are computed:

- (1) the size of cliques,
- (2) the size of cycle-cutsets in each of the subgraphs defined by the cliques, and
- (3) the size of the separator sets.

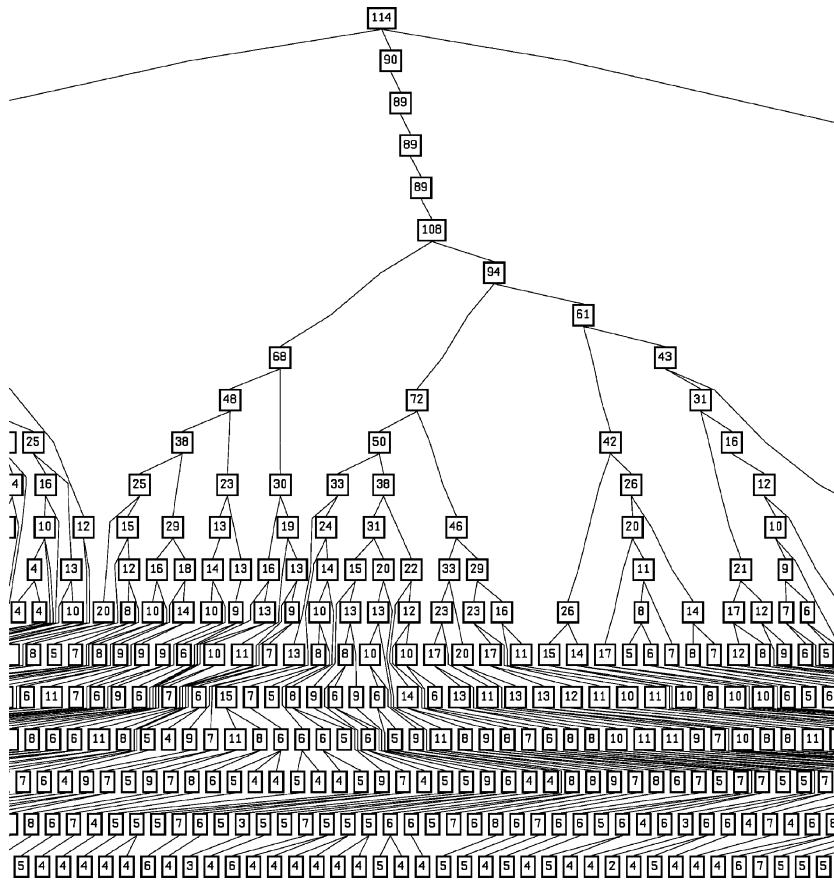


Fig. 10. Part of the primary join-tree (1419 cliques) for circuit C3540 (1719 variables) showing the descendants of the root node down to the leaves; the maximum separator size is 89.

The nodes of the join-tree are labeled by the cliques (or clusters) sizes. In this section we present the results on two circuits C432 and C3540, having 196 and 1719 variables, respectively. Results on other circuits are summarized in [12].

Figs. 9 and 10 present information on the primary join-trees. Fig. 9 shows that the clique sizes range from 2 to 28. The root node has 28 nodes and the descendant nodes have strictly smaller sizes. The depth of the tree is 11 and all nodes whose distance from the root is greater than 6 have sizes strictly less than 10. The leaves have sizes ranging from 2 to 6. The corresponding numbers for the primary join-tree of the larger circuit C3540 are shown in Fig. 10.

Figs. 11 and 12 provide additional details, showing the frequencies of cliques sizes, separator sizes, and cutset sizes for both circuits. These figures (and all the corresponding figures in [12]) show that the structural parameters are skewed with the vast majority of the parameters having values below the midpoint (the point dividing the range of values from smallest to largest).

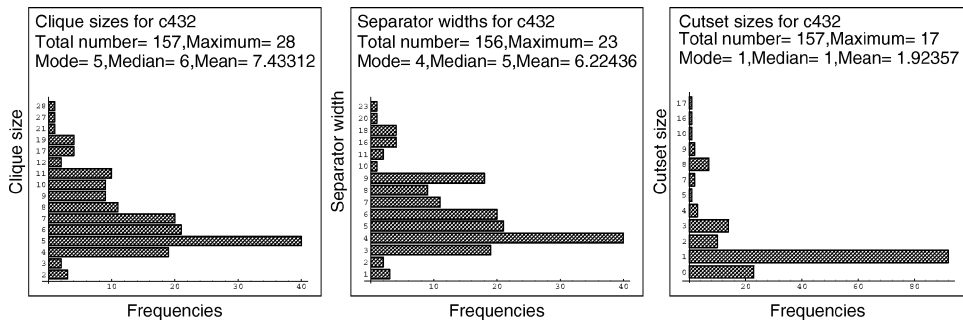


Fig. 11. Histograms of the clique sizes, the separator sizes, and the cutset sizes of the primary join-tree for circuit C432 (196 variables).

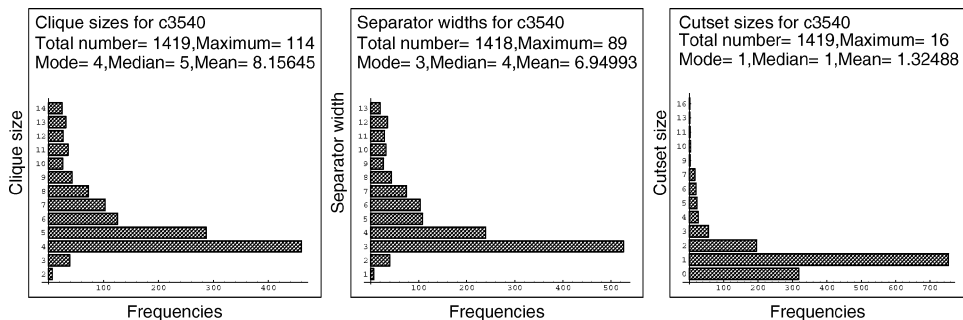


Fig. 12. Histograms of the clique sizes (0.9th quantile range), the separator sizes (0.9th quantile range), and the cutset sizes of the primary join-tree for circuit C3540 (1719 variables).

We see in Fig. 11 that the number of cliques is 157 and that the clique sizes range from 2 to 28. The mode is 5, the median is 6 and the mean is 7.433. Forty cliques out of the total 157 have size 5, and only 23 out of 157 have size greater than 9. The separator sizes are in the range from 1 to 23. The mode is 4, the median is 5, and the mean is 6.224. Out of the total 156 separator sizes, 40 have size 4 and only 13 have sizes greater than 10. The cutset sizes are in the range from 0 to 17. The mode is 1, the median is 1, and the mean is 1.923. Out of 157 cliques, 23 have cutset size 0. This means that the projection of the primal graph on each of those 23 cliques is already acyclic.

The corresponding figures for C3540 can be read from Fig. 12, showing the 0.9th quantile distribution of the separator sizes. Like the cliques, 90% of the separator sizes are small (between 1 and 13) and the remaining 10% span a broad range of values (from 14 to 89). For the cutset sizes, we note that 318 cutsets out of the total 1419 have size 0, namely the projection on each of those 318 cliques is already acyclic. We also note that 753 out of 1419 cliques have singleton cutsets. Only 47 out of 1419 cutsets have sizes greater than 5.

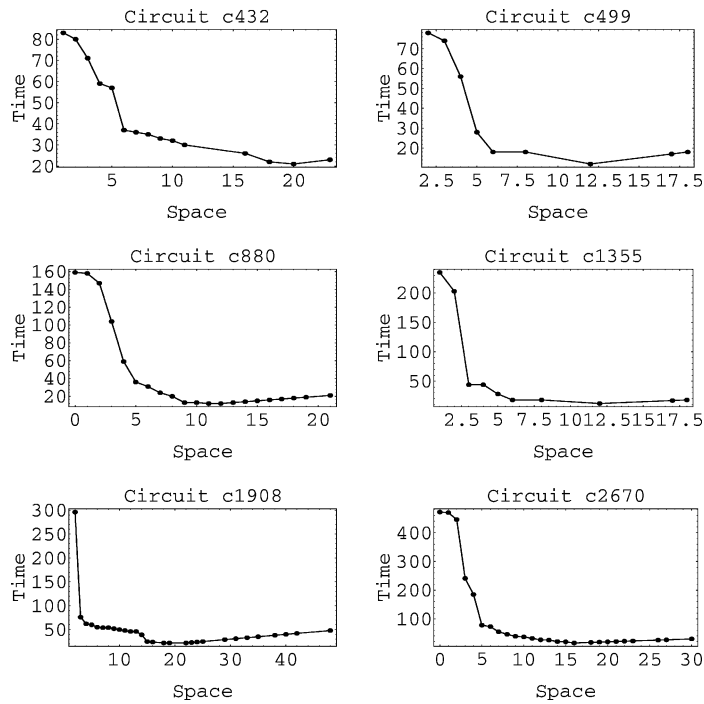


Fig. 13. Time-space tradeoff for C432 (196 variables), C499 (243 variables), C880 (443 variables), C1355 (587 variables), C1908 (913 variables), and C2670 (1426 variables). Time is measured by the maximum of the separator size and the cutset size, and space by the maximum separator size.

5.2. Results: Hybrid clustering + conditioning

As we see, some cliques and separators require memory space exponential in 23 for circuit C432 and exponential in 89 for circuit C3540. This is clearly not feasible. We will next evaluate the potential of the tradeoff scheme proposed in this paper.

Let s_0, c_0 be the maximum cutset and separator sizes of the primary join-tree T_0 obtained by tree-clustering. Let s_0, s_1, \dots, s_n be the sizes of the separators in T_0 listed from largest to smallest. As explained earlier, with each separator size, s_i , we associate a tree decomposition T_i generated by combining adjacent clusters whose separators' sizes are strictly larger than s_i . We denote by c_i the largest cutset size in any cluster of T_i .

We estimate the time-space bounds for each circuit based on the graph parameters observed using our tree decomposition scheme. Fig. 13 gives a chart presenting bounds for time versus space for each circuit. Each point in the chart corresponds to a specific secondary join-tree decomposition T_i and has the space complexity measured by the separator size, s_i , and the time complexity by the cutset size.

Each chart in Fig. 13 can be used to select the algorithm from the spectrum of conditioning + clustering algorithms that best meets a given time-space specification. They show the gradual effect of lowering space on the time required by a corresponding

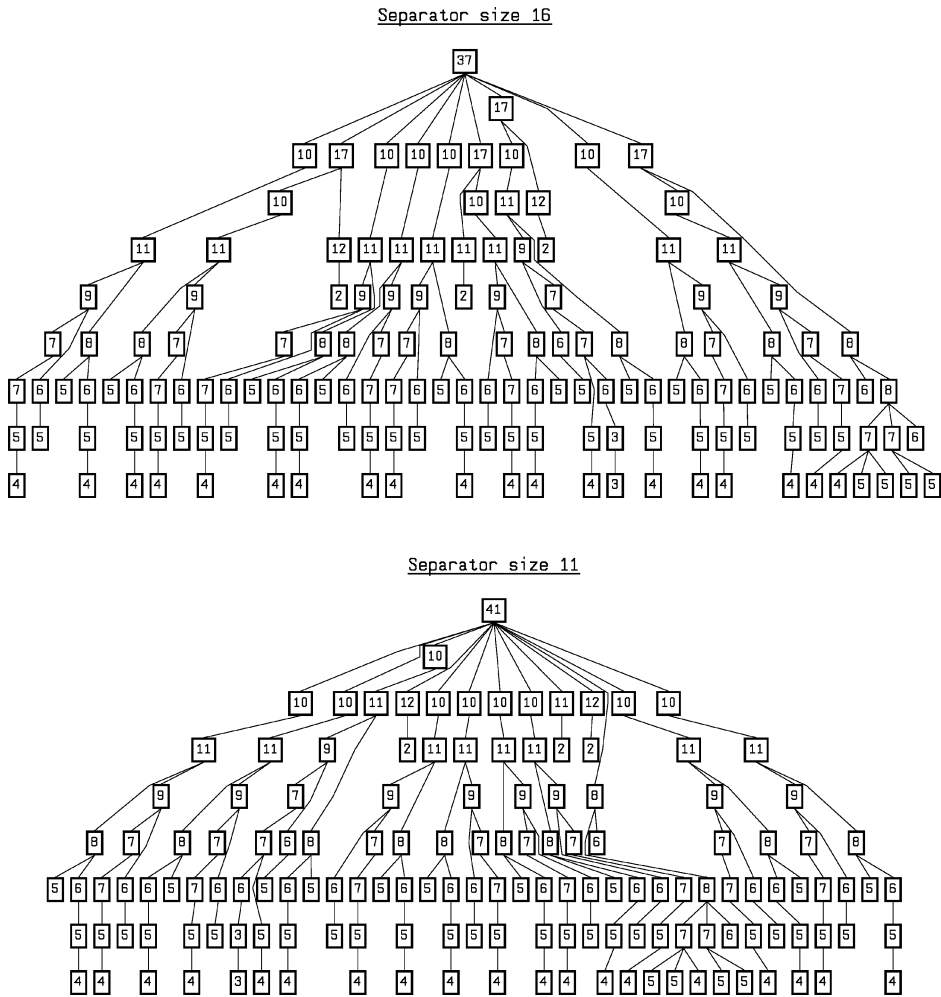


Fig. 14. Secondary trees for C432 with separator sizes 16 and 11.

clustering+conditioning algorithm. For example, circuit C432 (Fig. 13) shows the separator size (space) which is initially 23 (for the primary join-tree) gradually reduced down to 1 in a series of secondary trees. The figure demonstrates that reducing the separator size (to meet the space restrictions) increases the worst-case time complexity of the hybrid algorithm. The time increases because of the large clusters contained in the secondary join-tree and the corresponding increase in the size of cutsets.

Note that the charts in Fig. 13 all display a “knee” phenomenon in the time-space tradeoff where time increases only slightly for a wide range of space reduction beyond which further reduction in space causes significant rise in the time bound.

Figs. 14 and 15 display the structure of secondary join-trees for C432. The primary join-tree for the circuit is shown in Fig. 9. The secondary trees are indexed by the separator

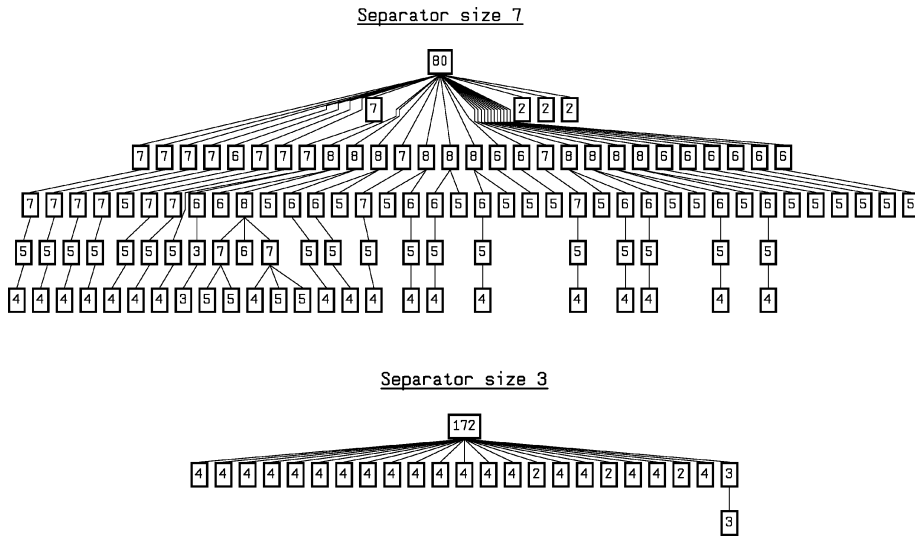


Fig. 15. Secondary trees for C432 with separator sizes 7 and 3.

sizes of the primary tree which range from 1 to 23 (Fig. 11). As the separator decreases the maximum clique size increases, and both the size and the depth of the tree decrease. Like the primary join-tree, each secondary join-tree also has a skewed distribution of the clique sizes. Note that the clique size for the root node is significantly larger than for all other nodes, and is increasing as the separator decreases.

6. Related work

The cycle-cutset scheme for probabilistic inference was introduced by Pearl [29] and for constraint networks by Dechter [8]. It was further improved and extended for probabilistic reasoning by Peot and Shachter [30], Darwiche [7], and Diez [15].

In subsequent years the cycle-cutset scheme was recognized as a special case of *conditioning*, namely, value assignments to a subset of variables creates subproblems that can be solved by any means. While the cycle-cutset scheme requires that the conditioning set will be large enough so that the resulting subproblem is singly-connected, any size of conditioning set can be used, yielding simplified problems that can be solved by tree-clustering or by any other method. This idea of extending the combination of conditioning and tree-clustering beyond the cycle-cutset scheme appears in [21] for constraint networks and in [19,31] for probabilistic networks. In [21] various heuristic are presented, aiming at creating a hybrid algorithm having improved time performance. In [31], the issue of reducing the space of tree-clustering by combination with conditioning is also briefly addressed. The latter paper includes an alternative proof to the (worst-case) time superiority of tree-clustering over the cycle-cutset method. In [19] the idea is applied to the Pathfinder system, when the conditioning set is restricted to the set of diseases.

Finally, in [33,34] a scheme for combining conditioning and variable elimination for propositional theories is outlined and analyzed. It is shown that although the worst-case time guarantee of a hybrid cannot be superior to tree-clustering (nor to a variable elimination scheme), for some problem classes a hybrid algorithm can have a better time performance than both pure clustering and pure search. The work in this paper provides an alternative hybrid scheme between conditioning and variable elimination.

7. Summary and conclusions

Problem solving methods can be viewed as hybrids of two main principles: inference and search. Tree-clustering is an example of an inference algorithm, while the cycle-cutset is an example of a search method. Tree-clustering algorithms are time and space exponential in the size of their cliques, while search algorithms are time exponential in the number of variables but require only linear memory. In this paper we developed a structure-based hybrid scheme that uses tree-clustering and cycle-cutset conditioning as its two extremes and, using a single design parameter, permits the user to control and tailor the storage-time tradeoff in accordance with the problem domain and the available resources

Specifically, we have shown that constraint processing and belief network processing obey a structure-based time-space tradeoff that allows tailoring a hybrid of tree-clustering and cycle-cutset conditioning to certain time and space requirements. The same tradeoff is obeyed by optimization problems when augmenting the graph by arcs reflecting the structure of the criterion function. Our analysis presents a spectrum of algorithms that allows a rich time-space performance balance applicable across a variety of tasks.

The structural parameters of interest are:

- (1) the size of cliques in a join-tree, namely, the induced width or tree-width,
- (2) the size of cycle-cutsets in each of the subgraphs defined by the cliques, and
- (3) the size of the separator sets.

To demonstrate the applicability of our scheme to real-life domains, we studied the structural parameters of 11 benchmark circuits widely used in the fault diagnosis and testing community [6]. We observed that the join-trees of the circuits all shared the unexpected property that few cliques are distinctly large and the majority of clique sizes are relatively small. Also, the distributions of all the structural parameters are skewed. This observation has an important practical implication. Although the primary join-tree obtained by tree-clustering may require too much space, a major portion of the tree can be solved without any space problem.

Our analysis should be qualified, however. All the results present worst-case guarantees of the corresponding algorithm. It is still not clear that the bounds are tight nor that they correlate with average-case performance. This analysis should be extended in the future to include implementation and testing of the involved algorithms.

References

- [1] S. Arnborg, A. Proskourowski, Linear time algorithms for NP-hard problems restricted to partial k -trees, *Discrete Appl. Math.* 23 (1989) 11–24.

- [2] S.A. Arnborg, Efficient algorithms for combinatorial problems on graphs with bounded decomposability—A survey, *BIT* 25 (1985) 2–23.
- [3] F. Bacchus, P. van Run, Dynamic variable ordering in CSPs, in: *Principles and Practice of Constraints Programming (CP'95)*, Cassis, France, Lecture Notes in Comput. Sci., Vol. 976, Springer, Berlin, 1995, pp. 258–277.
- [4] A. Becker, D. Geiger, A sufficiently fast algorithm for finding close to optimal junction trees, in: *Proc. Uncertainty in AI (UAI-96)*, Portland, OR, 1996, pp. 81–89.
- [5] U. Bertele, F. Brioschi, *Nonserial Dynamic Programming*, Academic Press, New York, 1972.
- [6] F. Brglez, H. Fujiwara, A neutral netlist of 10 combinatorial benchmark circuits and a target translator in Fortran, in: *Proc. IEEE International Symposium on Circuits and Systems*, 1996.
- [7] A. Darwiche, Conditioning algorithms for exact and approximate inference in causal networks, in: *Proc. Uncertainty in Artificial Intelligence (UAI-95)*, Montreal, Quebec, 1995, pp. 99–107.
- [8] R. Dechter, Enhancement schemes for constraint processing: Backjumping, learning and cutset decomposition, *Artificial Intelligence* 41 (1990) 273–312.
- [9] R. Dechter, Constraints networks, in: *Encyclopedia of Artificial Intelligence*, 2nd edn., 1992, pp. 276–285.
- [10] R. Dechter, Bucket elimination: A unifying framework for probabilistic inference algorithms, in: *Proc. Uncertainty in Artificial Intelligence (UAI-96)*, Portland, OR, 1996, pp. 211–219.
- [11] R. Dechter, Bucket elimination: A unifying framework for reasoning, *Artificial Intelligence* 113 (1999) 41–85.
- [12] R. Dechter, Y. El Fattah, Topological parameters for time-space tradeoff, Technical Report, Information and Computer Science, University of California, Irvine, CA, 1999.
- [13] R. Dechter, J. Pearl, Network-based heuristics for constraint satisfaction problems, *Artificial Intelligence* 34 (1987) 1–38.
- [14] R. Dechter, J. Pearl, Tree clustering for constraint networks, *Artificial Intelligence* 38 (1989) 353–366.
- [15] F.J. Diez, Local conditioning in Bayesian networks, *Artificial Intelligence* 87 (1996) 1–20.
- [16] Y. El Fattah, R. Dechter, Diagnosing tree-decomposable circuits, in: *Proc. IJCAI-95*, Montreal, Quebec, 1995, pp. 1742–1748.
- [17] E.C. Freuder, A sufficient condition for backtrack-free search, *J. ACM* 29 (1) (1982) 24–32.
- [18] H. Geffner, J. Pearl, An improved constraint propagation algorithm for diagnosis, in: *Proc. IJCAI-87*, Milan, Italy, 1987, pp. 1105–1111.
- [19] G.F. Cooper, H.J. Suermondt, D.E. Heckerman, A combination of cutset conditioning with clique-tree propagation in the path-finder system, in: *Proc. Uncertainty in Artificial Intelligence (UAI-91)*, Los Angeles, CA, 1991, pp. 245–253.
- [20] R.A. Howard, J.E. Matheson, Influence diagrams, in: R.A. Howard, J.E. Matheson (Eds.), *The Principles and Applications of Decision Analysis*, Vol. 2, Strategic Decision Group, Menlo Park, CA, 1984.
- [21] P. Jegou, Cyclic clustering: A compromise between tree-clustering and the cycle-cutset method for improving search efficiency, in: *Proc. European Conference on AI (ECAI-90)*, Stockholm, Sweden, 1990, pp. 369–371.
- [22] F.V. Jensen, S.L. Lauritzen, K.G. Olesen, Bayesian updating in causal probabilistic networks by local computation, *Computational Statistics Quarterly* 4 (1990) 269–282.
- [23] U. Kjærulff, Triangulation of graphs—Algorithms giving small total state space, Technical Report 90-09, Department of Mathematics and Computer Science, Aalborg University, Denmark, 1990.
- [24] U. Kjærulff, Nested junction trees, in: *Proc. Uncertainty in Artificial Intelligence (UAI-97)*, Providence, RI, 1997, pp. 294–301.
- [25] D.H. Krantz, R.D. Luce, P. Suppes, A. Tversky, *Foundations of Measurements*, Academic Press, New York, 1976.
- [26] S.L. Lauritzen, D.J. Spiegelhalter, Local computation with probabilities on graphical structures and their application to expert systems, *J. Roy. Statist. Soc. B* 50 (2) (1988) 157–224.
- [27] A.K. Mackworth, E.C. Freuder, The complexity of some polynomial network consistency algorithms for constraint satisfaction problems, *Artificial Intelligence* 25 (1985) 65–74.
- [28] J. Pearl, Fusion propagation and structuring in belief networks, *Artificial Intelligence* 19 (3) (1986) 241–248.
- [29] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann, San Mateo, CA, 1988.
- [30] M.A. Peot, R.D. Shachter, Fusion and propagation with multiple observations in belief networks, *Artificial Intelligence* 48 (1991) 299–318.

- [31] S.K. Anderson, R.D. Shachter, P. Solovitz, Global conditioning for probabilistic inference in belief networks, in: Proc. Uncertainty in Artificial Intelligence (UAI-94), Seattle, WA, 1994, pp. 514–522.
- [32] A. Dechter, R. Dechter, J. Pearl, Optimization in constraint networks, in: Influence Diagrams, Belief Nets and Decision Analysis, Wiley, New York, 1990, pp. 411–425.
- [33] I. Rish, R. Dechter, To guess or to think? Hybrid algorithms for SAT, in: Proc. Principles of Constraint Programming (CP-96), 1996, pp. 555–556.
- [34] I. Rish, R. Dechter, Resolution vs. SAT: Two approaches to SAT, *J. Approx. Reason. (Special Issue on SAT)* 24 (2000) 225–275.
- [35] D.G. Corneil, S.A. Arnborg, A. Proskourowski, Complexity of finding embeddings in a k -tree, *SIAM J. Discrete Math.* 8 (1987) 277–284.
- [36] T. Schmidt, P.P. Shenoy, Some improvements to the Shenoy-Shafer and Hugin architectures for computing marginals, *Artificial Intelligence* 102 (1998) 323–333.
- [37] G. Shafer, *Probabilistic Expert Systems*, SIAM, Philadelphia, PA, 1996.
- [38] G.R. Shafer, P.P. Shenoy, Probability propagation, *Ann. Math. Artificial Intelligence* 2 (1990) 327–352.
- [39] K. Shoiket, D. Geiger, A practical algorithm for finding optimal triangulations, in: Proc. AAAI-97, Providence, RI, 1997, pp. 185–190.
- [40] S. Srinivas, A probabilistic approach to hierarchical model-based diagnosis, in: Working Notes of the Fifth International Workshop on Principles of Diagnosis, New Paltz, NY, 1994, pp. 305–311.