

Scaling Up AND/OR Abstraction Sampling

Kalev Kask, Bobak Pezeshki, Filjor Broka, Alexander Ihler and Rina Dechter

University of California, Irvine, Irvine, CA 92697, USA

{kkask, pezeshkb, fbroka, ihler, dechter}@ics.uci.edu

Abstract

Abstraction Sampling (AS) is a recently introduced enhancement of Importance Sampling that exploits stratification by using a notion of abstractions: groupings of similar nodes into abstract states. It was previously shown that AS performs particularly well when sampling over an AND/OR search space; however, existing schemes were limited to “proper” abstractions in order to ensure unbiasedness, severely hindering scalability. In this paper, we introduce AOAS, a new Abstraction Sampling scheme on AND/OR search spaces that allow more flexible use of abstractions by circumventing the properness requirement. We analyze the properties of this new algorithm and, in an extensive empirical evaluation on five benchmarks, over 480 problems, and comparing against other state of the art algorithms, illustrate AOAS’s properties and show that it provides a far more powerful and competitive Abstraction Sampling framework.

1 Introduction

An important task in probabilistic graphical model inference is estimating the partition function (Z). Monte-Carlo methods have traditionally been used for generating Z estimates, with Importance Sampling (IS) being a primary framework [Rubin et al., 2007; Liu et al., 2015; Gogate and Dechter, 2011]. In IS, samples represent a single full configuration of the variables, with each given an *importance weight*. A newly introduced framework, Abstraction Sampling (AS) [Broka et al., 2018], inspired by [Knuth, 1975; Chen, 1992], expands on IS by allowing each sample to represent *multiple* configurations rather than just one. These samples of multiple configurations are called *probes*. In Abstraction Sampling, probes are generated by grouping nodes together into *abstract states* based on an *abstraction function*, and then picking a representative node from each group to build a sampled subtree. As such, each probe is a subtree of the full search tree. Each probe is reweighted based on its chosen representative nodes and contributes to a Monte Carlo estimate of the partition function of the full tree.

In particular, Abstraction Sampling smoothly interpolates between standard IS (where all of a variable’s states are

abstracted into a single abstract state) and standard search (where only *provably equivalent* states are abstracted into a single abstract state in the search tree). As such, Abstraction Sampling can be considered a generalization of Stratified Sampling [Rubin et al., 2007; Rizzo, 2007] combined with compact search [Dechter and Mateescu, 2007].

[Broka et al., 2018] defined AS over OR and AND/OR search spaces and provided an analysis of complexity and conditions for unbiasedness. They introduced abstraction functions for which AS significantly improved performance over the baseline of Importance Sampling *and was shown to be highly competitive against two state-of-the-art schemes: Weighted Mini-Bucket Importance Sampling* [Liu et al., 2015] and *IJGP-SampleSearch* [Gogate and Dechter, 2011].

However, a serious limitation of the existing AND/OR AS algorithm is that abstractions need to adhere to a property called *properness* to ensure unbiased estimation. This restriction (to be described) inhibits the algorithm’s ability to group nodes into abstract states and thus to limit the size of its probes, *severely hindering scalability*.

Contributions. 1. We introduce AOAS, a new Abstraction Sampling scheme over AND/OR search spaces. This new scheme lifts the properness restriction altogether, significantly enhancing the scalability of Abstraction Sampling for AND/OR spaces. 2. We provide new analysis and extensive empirical evaluation of all Abstraction Sampling schemes, and highlight how the new AOAS algorithm empowers formation of abstractions. 3. We compare against a new highly competitive scheme, Dynamic Importance Sampling (DIS), and show that AOAS is largely superior. We also strengthen the empirical results on instances lacking exact answers by placing the estimates within bounds produced by DIS.

2 Background and Definitions

A **graphical model**, such as a Bayesian or a Markov network [Pearl, 1988; Darwiche, 2009; Dechter, 2013], can be defined by a 3-tuple $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \Phi)$, where $\mathbf{X} = \{X_i : i \in V\}$ is a set of variables indexed by a set V and $\mathbf{D} = \{D_i : i \in V\}$ is the set of finite domains of values for each X_i . Each function $\psi_\alpha \in \Phi$ is defined over a subset of the variables called its scope, X_α , where $\alpha \subseteq V$ are the indices of variables in its scope and D_α denotes the Cartesian product of their domains, so that $\psi_\alpha : D_\alpha \rightarrow R^{\geq 0}$. The **primal graph** $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ of a graphical model associates each variable with a node

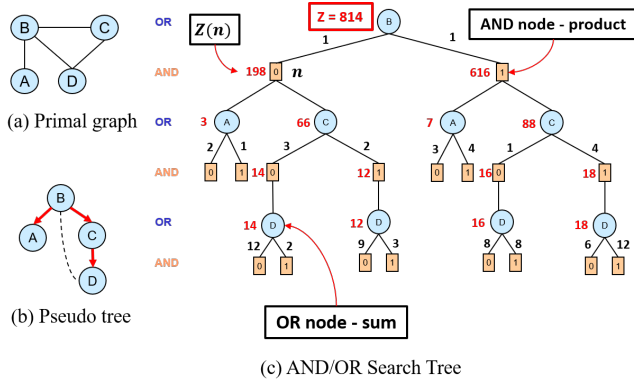


Figure 1: Example of an AND/OR search space. (a) A primal graph, (b) A possible pseudo tree for primal graph a), (c) An AND/OR search tree guided by pseudo tree b). To calculate the value of an OR node, we sum the values of its children multiplied by their arc weights, and for an AND node we multiply the values of its children. Partition function of the model is the resulting value of the root node.

($\mathbf{V} = \mathbf{X}$), while arcs $e \in \mathbf{E}$ connect nodes whose variables appear in the scope of the same local function. Graphical models can be used to represent a global function, often a probability distribution, defined by $Pr(X) \propto \prod_{\alpha} \psi_{\alpha}(X_{\alpha})$. An important task is to compute the normalizing constant, also known as the partition function, $Z = \sum_X \prod_{\alpha} \psi_{\alpha}(X_{\alpha})$.

Search Spaces of Graphical Models

A graphical model can be transformed into a weighted state space graph. In an OR search space, which is constructed layer-by-layer relative to a variable ordering, paths from the root to the leaves represent full configurations - or assignments to all variables - where each successive level corresponds to an assignment of the next variable in the ordering. A graphical model can also be transformed into a more compact AND/OR search space by capturing its conditional independencies, thus facilitating more effective algorithms [Dechter and Mateescu, 2007].

An AND/OR search space is defined relative to a *pseudo tree* of a primal graph. A **pseudo tree** $\mathcal{T} = (\mathbf{V}, \mathbf{E}')$ of a primal graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ is a directed rooted tree that spans \mathcal{G} such that every arc of \mathcal{G} not in \mathbf{E}' is a back-arc in \mathcal{T} connecting a node to one of its ancestors (Figure 1(a),(b)). A variable is a **branching variable** if it has multiple children in \mathcal{T} .

Given a pseudo tree \mathcal{T} of a primal graph \mathcal{G} , the *AND/OR search tree* $T_{\mathcal{T}}$ guided by \mathcal{T} has alternating levels of OR nodes corresponding to variables, and AND nodes corresponding to an assignment from its domain with edge costs extracted from the original functions [Dechter and Mateescu, 2007]. Let n be an AND node in $T_{\mathcal{T}}$, also denoted n_X if X is the last variable of its partial configuration. Each arc into an AND node n has a cost $c(n)$ defined to be the product of all factors ψ_{α} in \mathcal{M} that are instantiated at n but not before.

A **solution tree** is a subtree of $T_{\mathcal{T}}$ satisfying: (1) it contains the root of $T_{\mathcal{T}}$; (2) if an OR node is in the solution tree, exactly one of its AND child nodes is in the solution tree; (3) if an AND node is in the tree then all of its OR children are in the solution tree. Finally, its leaves are leaves of $T_{\mathcal{T}}$.

The cost of a solution tree is the product of its arc-costs and is equal to the cost of the corresponding full configuration of the graphical model [Dechter and Mateescu, 2007].

Value of A Node

The partition function, Z , of a graphical model \mathcal{M} can be obtained by computing the value of the AND/OR search tree T , $V(T)$, by summing the costs of all its solution trees. This can be done recursively from leaves to root [Dechter and Mateescu, 2007] by associating a *value function* $V(n)$ with every node n in the AND/OR search tree. The **value of a node**, $V(n)$, is defined as the sum cost of all partial solution-trees rooted at n . In particular, $V(n_X)$ denotes the value of an AND node n_X of variable X . $V(Y_{n_X})$ denotes the value of an OR node Y that is the child of an AND node n_X , Y being the child of X in the pseudo tree. The value function can be computed by the following recursive expression:

$$V(n_X) = \prod_{Y \in ch_{\mathcal{T}}(X)} V(Y_{n_X}) \quad (1)$$

where

$$V(Y_{n_X}) = \sum_{n_Y \in ch_Y(n_X)} c(n_Y) \cdot V(n_Y) \quad (2)$$

and where $ch()$ denotes child variables either in the pseudo-tree or the search tree itself (depending on the context). Here, $ch_Y(n_X)$ are the child AND nodes of Y descended from AND node n_X . Note that $V(T) = V(\text{root})$, where "root" is the root node of $T_{\mathcal{T}}$. It follows that $V(\text{root}) = Z_{\mathcal{M}}$, the partition function of the underlying model (see Figure 1c).

But what is the meaning of $V(n)$ for an arbitrary AND node n w.r.t the partition function of \mathcal{M} ? Let $path(n)$ denote the configuration of variables on the path from the root to n in $T_{\mathcal{T}}$, let $Z_{|path(n)}$ be Z conditioned on the assignment to $path(n)$, and let $Out(path(n))$ be the set of OR nodes (corresponding to variables) emanating from $path(n)$, which are OUTside $path(n)$. Then Z , when restricting to the partial configuration of $path(n)$, is obtained by multiplying $g(n)$, the cost of $path(n)$, by $V(n)$ and by all the values from OR nodes branching out of $path(n)$ (the latter referred to as the value of the "ancestor branching").

It is easy to see that

Lemma 1.

$$Z_{|path(n)} = g(n) \cdot V(n) \cdot \prod_{Y \in Out(n' \in path(n))} V(Y_{n'}) \quad (3)$$

where $V(Y_{n'})$ is as defined in EQ. 2.

Defining the product factor in Eq. 3, which we call *ancestral branching value* as

$$R(n) = \prod_{Y \in Out(n' \in path(n))} V(Y_{n'}) \quad (4)$$

we get that $Z_{|path(n)} = g(n) \cdot V(n) \cdot R(n)$, an expression that will be used in the algorithm for estimating the partition function at nodes of the search tree.




Figure 2: An AND/OR search tree and possible probes generated by abstraction sampling. Nodes are abstracted based on having the same domain value (denoted by having the same color). (a) A full AND/OR tree, representing all 16 solutions, (b) Boxed in green is the ancestor branching subtree for the path $\rightarrow(B=0)\rightarrow(C=1)$, (c) A probe generated by proper abstractions, representing 8 solutions, (d) An invalid probe generated by improper abstractions with nodes abstracted across different sub-branches under the branching variable B (contains 4 *partial* configurations, and no full solution), (e) A valid probe generated by AOAS with improper abstractions, representing 4 solutions.

Example 1. An example of expression (3) is given in Figure 2(b). Consider the path to node $n = (B=0, C=1)$ marked in red. There is only one branching variable ancestor (B) and, in this case, the only $Out(path(n))$ node is the A node emanating from under ($B = 0$). Thus, $Z_{|path(B=0, C=1)} = g(B = 0, C = 1) \cdot V(B = 0, C = 1) \cdot V(B = 0, A)$.

Stratified Importance Sampling.

Abstraction sampling builds on Importance Sampling and Stratified Sampling. *Importance Sampling* (IS) is a Monte Carlo scheme for approximating likelihood queries over graphical models. *Stratified Sampling* is a variance reduction technique for sampling a search space by first dividing it into disjoint strata. This can be used with importance sampling to further reduce variance. In *Stratified Importance Sampling*, we first divide the sample space into k strata of equal area under the distribution q , then choose re-weighted representatives from each strata. In order to maximize reduction in variance, the variance between strata should be maximized (see [Rizzo, 2007]).

3 AND/OR Abstraction Sampling

Abstraction Sampling (AS) algorithms [Broka *et al.*, 2018] emulate Stratified Importance Sampling with the modification that stratification and selection of representatives is done variable-by-variable along a search tree. Guided by an abstraction function a that groups nodes of a variable together, Abstraction Sampling generates a compact representative subtree of the full search tree $T_{\mathcal{T}}$ called a **probe**. An abstraction function, a , is formally defined as $a : T_{\mathcal{T}} \rightarrow I^+$, where I^+ are integers, and it partitions the nodes in $T_{\mathcal{T}}$, layer by layer. From each partition (i.e. abstraction), a representative node is stochastically chosen based on a proposal p and is reweighted accordingly. AS can use any Importance Sampling proposal. For OR search trees we use p which is proportional to $w(n) \cdot g(n) \cdot h(n)$, where $h(n)$ is a heuristic that provides an upper bound on $V(n)$, $g(n)$ is the cost of $path(n)$, and $w(n)$ is the importance weight. The ORAS algorithm from [Broka *et al.*, 2018], which operates on an OR search tree, is unbiased for any Importance Sampling proposal. However, it was shown that using a heuristic-based proposal function can significantly impact convergence based on its accuracy. When $h(n) = h^* = V(n)$ the proposal is ex-

act and only one sample is needed, as is the case for general Importance sampling.

pAOAS

Abstraction sampling algorithms for AND/OR search spaces operate by expanding and abstracting an AND/OR probe and can traverse the pseudo-tree \mathcal{T} in a depth-first or breadth-first manner. However, performing these abstractions naively may leave partial (invalid) configurations within the resulting probe (ex. Figure 2(d)), potentially leading to biased estimates. To avoid this issue, [Broka *et al.*, 2018] forced abstractions to be *proper* with the scope of abstractions restricted to nodes descending from the same AND node of the most recent branching variable ancestor. A resulting such probe can be seen in Figure 2(c). More formally, consider a variable W whose most recent branching ancestor in \mathcal{T} is U . When performing proper abstractions for W , only the AND nodes of W that are in the common subtree rooted under the same AND node of U can be grouped together [Broka *et al.*, 2018]. This ensure expansion to only valid probes. Unfortunately, properness limits compactness of the sampled AND/OR trees as it constrains the set of AND nodes that can be abstracted together. In particular, instead of bounding the size of the probe by $O(nm)$ AND nodes (as with ORAS), the number of AND nodes in a probe constructed via proper abstractions is bounded only by $O(n \cdot m^{b+1})$, where n is the number of variables, b bounds the number of branching variables along any path and m bounds the maximum number of abstract states per variable. This means that proper AND/OR abstraction schemes are not scalable whenever $b \gg 0$, unless $m = 1$, in effect limiting us to basic Importance Sampling. Various superficial schemes were used in [Broka *et al.*, 2018] to bound probe sizes in this case.

AOAS

Our solution to the above is, AOAS (Algorithm 1), a new AND/OR Abstraction Sampling algorithm. We will show the algorithm to be unbiased and effective far beyond the earlier versions. Given a pseudo tree \mathcal{T} of a graphical model $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \Phi)$ that is rooted at a dummy variable D , a heuristic function $h(n)$ providing upper bound on $V(n)$, and an abstraction function a , AOAS traverses the pseudo tree \mathcal{T} variable-by-variable in a depth-first manner using the traversal to guide the generation of a partial search tree. The algo-

Algorithm 1: AOAS.

Input: Graphical model $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \Phi)$, a pseudo tree \mathcal{T} for \mathcal{M} rooted at a dummy singleton variable D , an abstraction function a , heuristic function h . For any node n , $g(n)$ = its path cost, $w(n)$ = its importance weight, and $\hat{V}(n)$ = its estimated value (initialized to $h(n)$).

Output: $\hat{Z}_{\mathcal{M}}$, an estimate of the partition function of \mathcal{M}

```
1 Function AOAS ( $\mathcal{T}, h, a$ )
2 begin
3    $PROBE \leftarrow n_D, g(n_D), w(n_D), r(n_D), \hat{V}(n_D) \leftarrow 1$ 
4    $STACK \leftarrow push(\text{empty stack}, D)$ 
5   while  $STACK$  is not empty do
6      $X \leftarrow top(STACK)$ 
7     if  $X$  has unvisited children in  $\mathcal{T}$  then
8        $Y \leftarrow$  the next unvisited child of  $X$ 
9       foreach  $n_X \in PROBE$  do
10         $PROBE \leftarrow PROBE$  expanded from
11          $n_X$  to  $Y$ 
12         $F'_Y \leftarrow$  newly added AND nodes of
13          $Y \in PROBE$ 
14        foreach  $n_Y \in F'_Y$  do
15           $w(n_Y) \leftarrow w(n_X)$ 
16           $g(n_Y) \leftarrow g(n_X) \cdot c(n_Y)$ 
17           $r(n_Y) \leftarrow$ 
18             $r(n_X) \cdot \prod_{\{S \neq Y \in ch_{\mathcal{T}}(X)\}} \hat{V}(S_{n_X})$ 
19        end
20      end
21       $A \leftarrow \{A_i \mid A_i = \{n_Y \in PROBE \mid a(n) = i\}\}$ 
22      foreach  $A_i \in A$  do
23        foreach  $n \in A_i$  do
24           $p(n) \leftarrow \frac{w(n) \cdot g(n) \cdot h(n) \cdot r(n)}{\sum_{m \in A_i} w(m) \cdot g(m) \cdot h(m) \cdot r(m)}$ 
25        end
26         $n_{Y_i} \propto_p A_i$ ; // randomly select
27         $w(n_{Y_i}) \leftarrow w(n_{Y_i}) / p(n_{Y_i})$ 
28         $\hat{V}(n_{Y_i}) \leftarrow 1$ 
29         $PROBE \leftarrow PROBE \setminus A_i \cup \{n_{Y_i}\}$ 
30      end
31       $push(STACK, Y)$ 
32    else
33       $pop(STACK), W \leftarrow top(STACK)$ 
34       $PROBE \leftarrow PROBE$  s.t. all  $n_W$  without
35       descendants are pruned
36      foreach  $n_W$  in  $PROBE$  do
37         $\hat{V}(n_W) \leftarrow \hat{V}(n_W) \cdot$ 
38           $\sum_{n_X \leftarrow child(n_W)} \hat{V}(n_X) \cdot c(n_X) \cdot \frac{w(n_X)}{w(n_W)}$ 
39      end
40      if  $X = D$  then  $\hat{Z}_{\mathcal{M}} = \hat{V}(D)$ ;
41    end
42  end
43  return  $\hat{Z}_{\mathcal{M}}$ 
44 end
```

algorithm uses a data structure referred to as *PROBE* to store the generated subtree of the full search tree $T_{\mathcal{T}}$. As it progresses forward in the traversal, it expands and abstracts nodes within *PROBE*. Expansion is done by adding nodes of the newly visited variable across the tree in a breadth-first manner. During backtracking, the algorithm prunes nodes and updates values.

More specifically, in the forward step (lines 8-28), *PROBE* is extended from all current leaf AND nodes of a variable X to the next variable Y (child of X in \mathcal{T}) in a breadth-first manner. The newly added AND nodes $n = n_Y$ are new leaves in *PROBE*. Each newly generated node n inherits the weight of its parent AND node n_X and has its path cost $g(n)$ computed. $r(n)$ is also computed (line 15), where it estimates $R(n)$ in the expression for $Z_{|path(n)}$ in EQ. (3, 4). Note that $r(n)$ is simply inherited from $r(n_X)$ if the parent variable X is not a branching variable. Otherwise, the $r(n_X)$ from its parent is multiplied by the estimated values contributed by its sibling branches, yielding an $r(n)$ which bounds the ancestor branching value for n .

Once nodes are expanded, we perform abstractions (lines 18-27). Nodes are partitioned into abstract states according to an abstraction function a , and for each abstract state, proposals are generated from which a representative node is stochastically selected. The proposal function at node n is proportional to the quantity $w(n) \cdot g(n) \cdot h(n) \cdot r(n)$. The portion $g(n) \cdot h(n) \cdot r(n)$ estimates $Z_{|path(n)}$ (EQ. 3, 4), where $V(n)$ is approximated by $h(n)$ and $R(n)$ is estimated by $r(n)$. We multiply by $w(n)$ to account for all the mass n represents.

When backtracking to a variable W from variable X (lines 30-35), we prune any n_W that has no descendant n_X , thus preventing the creation of invalid probes (line 31). For the surviving nodes, we back up the values from their children from which we have backtracked (line 33). Once the algorithm backtracks to the dummy node, it will have backed up the estimated value of the full model, which is returned.

AOAS Trace

Figure 3 shows a step-by-step trace of AOAS. We follow a DFS traversal of the pseudo-tree \mathcal{T} from Figure 1(b) and sample from the corresponding AND/OR search tree $T_{\mathcal{T}}$ (Figure 2(a)). The abstraction function we use groups nodes that have the same domain value, depicted by the color of the nodes. When performing abstractions, we box nodes that are being abstracted in gray. A red "X" marks a pruned subtree.

Starting with variable B (Figure 3(a)), each node belongs to a different abstraction and is therefore kept. Next, we expand to A and abstract across its nodes (Figure 3(b)). Not restricted to *proper* abstractions, we partition across *all* nodes of A , regardless of whether they fall under $B=0$ or $B=1$. We see two nodes in each abstract state (denoted by the red and blue coloring). Next we calculate their respective proposals (line 21). Note that the proposal of each node n relies on $r(n)$ (line 15), which captures the values of the nodes in its $Out(path(n))$, in this case nodes of C . Since the nodes of C have not been expanded yet, we use their heuristic values as an approximation of their values. We then stochastically choose a representative from each abstract state (line 23). Suppose that both red and blue representatives are stochas-

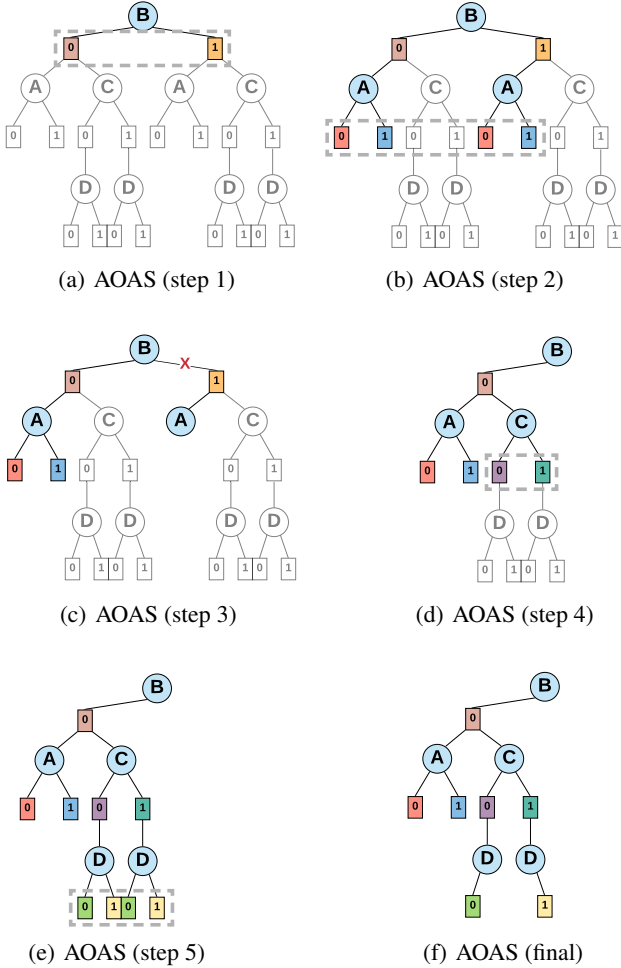


Figure 3: A sample trace of AOAS based on the AND/OR tree in Figure 2(a). Nodes with the same domain values are abstracted (also indicated by node color). A red “X” indicates pruning. Transparent nodes indicate portions of the tree not pruned and yet to be explored.

tically chosen from under $B=0$ (Figure 3(c)). Since A has no descendant, we backtrack to B , updating its node values (line 33) and performing a pruning step (line 31). In pruning, we remove AND nodes of B that do not extend to AND nodes of A , and thus prune $B=1$ (denoted by the red “X” in Figure 3(c)), in order to ensure formation of proper AND/OR probes. Finally, we expand and abstract C and D (Figures 3(d)-3(f)). The $r(n)$ for D ’s nodes is inherited from the $r(n_C)$ of its respective n_C parent. We backtrack from D to the root updating values (no further pruning was necessary). The result is a valid probe (Figure 3(f)) containing four solutions: $(B=0, A=0, C=0, D=0)$, $(B=0, A=0, C=1, D=1)$, $(B=0, A=1, C=0, D=0)$, and $(B=0, A=1, C=1, D=1)$. We estimate the partition function by computing $\hat{V}(B)$.

Properties of AOAS

Algorithm AOAS does away with the properness requirement of pAOAS while ensuring unbiasedness. The proof is based on the fact that there is always a corresponding equivalent

OR sampling algorithm, (e.g., using a DFS ordering of the pseudo-tree), such that there is a one-to-one correspondence between AND nodes in AND/OR sampling and equivalent OR sampling, and ensuring also that the proposal for the matching nodes is the same, implying that the two algorithms produce the same estimate. Since ORAS is known to be unbiased for any proposal, our claim follows. A key to this is that our choice of proposal, which includes $r(n)$, captures the estimated \hat{Z} of the full configuration in the same way that would be done in the equivalent OR case. In other words we can show the following:

THEOREM 1 (simulation of AOAS with ORAS). *Given a graphical model $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \Phi)$, pseudo-tree \mathcal{T}_{AO} of \mathcal{M} , an abstraction a^{AO} , and a proposal function p , then there exists a variable ordering d defining a chain-like pseudo-tree \mathcal{T}_{OR} of \mathcal{M} , and an abstraction function a^{OR} defined on \mathcal{T}_{OR} such that for every execution trace of AOAS on \mathcal{T}_{AO} using a^{AO} and proposal $p \cdot r$ and producing a corresponding $\hat{Z}_{\mathcal{M}}$ estimate, there is an execution trace of ORAS on \mathcal{T}_{OR} using a^{OR} and proposal p producing the same $\hat{Z}_{\mathcal{M}}$ estimate.*

THEOREM 2 (unbiasedness). *Given a graphical model $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \Phi)$, algorithm AOAS provides an unbiased estimate for the partition function of \mathcal{M} .*

Complexity. It is easy to see that, like for ORAS, probe sizes for AOAS will not be larger than $O(m \cdot n)$, where n is the number of variables and m bounds the number of abstract states per variable. Note that theorem 1 does not imply that AOAS on \mathcal{T} is the same as ORAS on the DFS ordering of \mathcal{T} with the same a . In fact, the former can lead to more informed probes that express more solutions as we show next.

Contrasting Scalability. Figures 4(b), 4(c), and 4(d) show probes generated by ORAS, pAOAS , and AOAS, respectively, sampling over a search tree guided by the pseudo tree in Figure 4(a). ORAS ends up expanding 26 nodes, 8 of which are retained in the final probe constituting only two solutions: $(X, Y, Z, T, R, L=0, M=0)$ and $(X, Y, Z, T, R, L=0, M=1)$. pAOAS inevitably expands and keeps the entire AND/OR search tree of 42 nodes! However, since it samples a compact AND/OR search space, this nevertheless corresponds to a relatively many captured solutions (128 in this case). AOAS only expands 20 nodes (fewer than both of the other schemes), and results in a probe size of 11 (comparable to ORAS). Further more, its final probe corresponds to 16 solutions, far more than that by ORAS. Thus, we see that AOAS is much more equipped than pAOAS to limit its probe size whilst able to take advantage of an AND/OR search space.

Cutset Sampling. An improvement to any IS sampling scheme can be made by sampling over a subset of the variables using *cutset decomposition* called *cutset sampling* [Dechter, 2013; Bidyuk and Dechter, 2007]. Using the weighted mini-bucket heuristic it is possible to identify when we reach a node for which the heuristic is exact (i.e. $h(n) = h^*(n) = V(n)$). When this occurs, there is no longer a need to sample below the node as the value of the subtree is known. Utilizing this method, the number of variables sampled is bounded by what is known as the *i-cutset*, which in turn bounds the size of the probes. Consequently, this leads

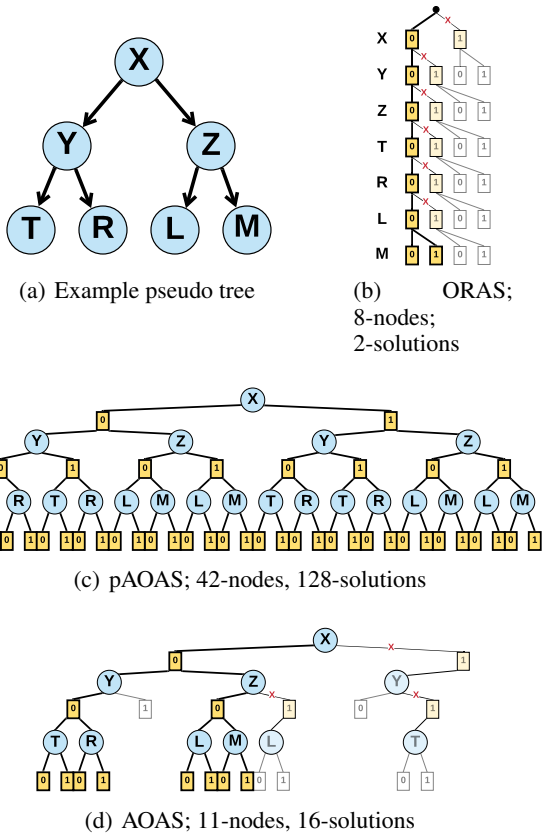


Figure 4: Contrasting Scalability. AS probes sampled from a search space corresponding to the pseudo tree in (a). Nodes are abstracted based on having the same domain value. White nodes were discarded through the abstraction process. Light-yellow nodes are selected representatives that are later pruned. Pruning is depicted by a red “X”. Dark-yellow nodes constitute the final probe. (b) a probe by ORAS; (c) a probe by pAOAS; (d) a probe by AOAS

to a greater number of samples and improves accuracy, as we reference in our empirical analysis.

4 Empirical Evaluation

Algorithms. Experiments were run on three classes of AS algorithms: AOAS, pAOAS, and ORAS (using a depth-first ordering of the pseudo tree). The AS algorithms were also tested against high performing state-of-the-art Dynamic Importance Sampling (DIS) [Lou *et al.*, 2019] using an “*equal-time*” policy. All algorithms were implemented in C++. All experiments were run for 1 hr on a 2.66 GHz processor with 8 GB of memory (24 GB for the Linkage-Type4 benchmark).

Abstraction Functions. We design abstractions to reduce variance within abstract states by using the notion of a variable’s context. Formally, the **context** of a variable X identifies a subset $C(X)$ of its ancestor in a pseudo-tree \mathcal{T} whose assignment uniquely determines the AND/OR subtree below it [Dechter and Mateescu, 2007]. Intuitively, the context of a variable is its ancestors that directly affect its value. Thus, abstractions based on a subset of the context aim to group nodes

Problem	Size	Total	$\in Bnds$	$AOAS \geq$	$AOAS >$
DBN	small	66	62	57	47
	large	48	40	38	35
Grids	small	8	5	5	2
	large	19	7	7	6
Linkage ¹	large	82	82	82	82
Pedigree	small	24	24	24	19
Promedas	small	65	58	49	29
	large	173	165	141	113

Table 1: How often AOAS estimates fall within DIS probabilistic bounds ($\in Bnds$), were comparable/better than DIS’s ($AOAS \geq$), and were strictly better than DIS’s ($AOAS >$)

based on having similar values. AOAS and ORAS utilize relaxed context-based (**RelCB**) and randomized context-based (**RandCB**) abstractions as in [Broka *et al.*, 2018]. RelCB is parametrized by a level j , selecting the closest $j-1$ variables from a variable’s context (ie. its *relaxed context*) plus itself. It abstracts nodes of the same domain value that also share the same assignment to the relaxed context. This yields k^j abstract states at each level, assuming domain size of k . RelCB-0 vacuously groups *all* nodes of a variable into a single abstract state and is referred to as a *Knuth Abstraction*. The randomized scheme, RandCB, is parameterized by a level d determining the number of abstractions per level, leading to a varying number of nodes in each abstract state. pAOAS utilizes a variant of RelCB that maintains *properness*. It is parameterized by levels j and k resulting in a j -level abstraction for nodes having no more than k branching variables in the path to the root and 0-level abstractions below it.

Heuristics. To inform the sampling proposal, Weighted Mini-Bucket Elimination (WMBE) [Dechter and Rish, 2003; Liu and Ihler, 2011] is used as a heuristic. The i -bound parameter controls the strength of WMBE, where higher i -bounds generally lead to stronger heuristics and, thus, better proposals at the expense of higher computation and memory. We standardize our experiments by using i -bound 10.

Benchmarks. We perform high-throughput experiments on over 480 problems from five well known benchmarks: DBN, Grids, Linkage-Type4, Pedigree, and Promedas. For brevity, we show detailed aggregated statistics on only large problem instances, thus excluding Pedigree, whose problems were all small and results relatively uniform across all algorithms. Average statistics for the benchmarks can be found in Figure 5.

Performance Measure. To evaluate the performance of the various algorithms, we calculate error as: $error = \log_{10} \hat{Z} - \log_{10} Z^*$, where Z is the partition function, $\log_{10} \hat{Z}$ is the \log_{10} of the experimentally obtained Z estimate, and $\log_{10} Z^*$ is the reference $\log_{10} Z$ value. When the exact Z value is unknown, an empirical estimate based on an average over 100×1 hr of abstraction sampling is used as the reference. We verified that 98% of these estimates fell within the 95% probabilistic bounds determined by DIS.

i-Bound = 10										DIS log(eri) = -39.463 (3600 sec)													
Bmk	Sz	Graph Scheme	Context Scheme	Abs	n*	log(eri)	error distr.			#probes	#nodes/probe	Bmk	Sz	Graph Scheme	Context Scheme	Abs	n*	log(eri)	error distr.			#probes	#nodes/probe
											0.5 2 10												
DBN	(LARGE, n:1000, d:2, w:78, h:79/215)	AOAS	RelCB	0	48	-3.251	15	29	42	5.40E+05	434	(LARGE, n:3432, d:2, w:27, h:220/3431)	AOAS	RelCB	0	19	-167.080	0	0	4	1.92E+06	3326	
				4	48	-4.045	5	20	41	5.52E+04	6352				4	19	-67.780	0	4	7	2.34E+05	43449	
				8	48	-5.180	5	12	40	4.08E+03	91614				8	19	-49.468	2	4	9	2.07E+04	407963	
		RandCB	16	48	-2.358	8	31	45	5.20E+04	6249	16		19	-79.864	1	3	7	2.82E+05	32965				
			256	48	-2.606	10	30	45	1.38E+04	4131	256		19	-77.505	0	3	5	1.67E+05	66376				
			0	48	-3.123	14	31	42	6.35E+05	434	0		19	-161.895	0	1	4	4.10E+06	3326				
	pAOAS	RelCB (k=5)	2	48	-3.591	10	23	42	2.00E+05	1690	2	19	-129.787	0	3	4	6.83E+03	2849697					
			4	48	-3.712	7	22	43	5.80E+04	6352	4	3	-2.972	0	1	3	4.00E+01	1.44E+08					
			16	48	-2.328	12	31	44	5.41E+04	6276	16	3	-4.025	0	1	3	1.50E+02	41914557					
	ORAS-DFS	RelCB	0	48	-3.632	11	27	42	4.30E+05	434	0	19	-176.694	0	0	4	6.65E+05	6844					
			4	48	-4.377	4	19	42	3.42E+04	6586	4	19	-96.276	0	3	6	6.32E+04	78212					
			8	48	-5.284	2	11	40	2.66E+03	91881	8	19	-60.563	0	4	7	5.10E+03	886885					
RandCB	16	48	-2.704	7	28	44	3.71E+04	4621	16	19	-96.497	0	0	4	1.02E+04	75095							
	256	48	-3.312	9	22	45	7.38E+03	30289	256	19	-95.435	0	1	5	1.47E+03	786178							

(a) DBN

(b) Grids

i-Bound = 10										DIS Unable to Solve (3600 sec)													
Bmk	Sz	Graph Scheme	Context Scheme	Abs	n*	log(eri)	error distr.			#probes	#nodes/probe	Bmk	Sz	Graph Scheme	Context Scheme	Abs	n*	log(eri)	error distr.			#probes	#nodes/probe
											0.5 2 10												
Linkage-Type4	(LARGE, n:10822, d:1.8, w:24, h:581/5745)	AOAS	RelCB	0	7	-32.487	0	0	0	5.52E+05	829	(LARGE, n:1209, d:2, w:27, h:114/1193)	AOAS	RelCB	0	130	-4.781	10	46	111	2.10E+07	171	
				4	26	-14.174	1	1	10	1.92E+05	16472				4	145	-4.106	24	62	129	2.43E+06	2127	
				8	41	-11.090	3	10	26	1.20E+04	463468				8	139	-4.520	19	53	121	1.81E+05	23935	
		RandCB	16	26	-18.582	0	0	9	2.50E+05	10532	16		154	-4.142	26	71	140	1.89E+06	2495				
			256	27	-18.221	0	1	9	1.74E+05	18523	256		173	-3.531	29	78	161	4.83E+05	12845				
			0	8	-27.742	0	0	0	1.14E+07	759	0		131	-4.664	18	46	112	3.31E+07	171				
	pAOAS	RelCB (k=5)	2	13	-27.201	0	0	1	6.96E+04	280176	2	158	-5.158	14	59	132	1.03E+06	11266					
			4	11	-25.427	0	0	0	8.90E+02	34177733	4	145	-5.178	10	46	126	4.61E+04	925471					
			16	12	-24.204	0	0	1	4.24E+03	8820722	16	162	-3.306	35	89	150	5.58E+04	534634					
	ORAS-DFS	RelCB	256	14	-30.930	0	0	0	2.58E+03	24532116	256	171	-2.050	34	107	168	2.36E+03	22671070					
			0	7	-29.820	0	0	0	1.42E+06	2520	0	128	-5.479	9	38	105	9.69E+06	361					
			4	15	-20.802	0	0	2	5.14E+04	26863	4	134	-5.261	10	36	111	1.78E+06	2566					
RandCB	8	16	-18.971	0	0	5	2.44E+03	355404	8	134	-5.073	14	37	118	1.96E+05	27440							
	16	13	-22.843	0	0	1	1.05E+04	23723	16	144	-4.600	18	50	128	1.66E+05	4897							
	256	12	-27.151	0	0	0	1.18E+02	288516	256	173	-2.821	15	70	171	2.66E+03	104091							

(c) Linkage-Type4

(d) Promedas

Figure 5: Aggregated statistics. Displayed are the number of problems solved (n^*), average $\log_{10}Z$ error ($\log(eri)$), count of problems solved within an error threshold ($error\ distr.$), average number of probes ($\#probes$), and average size of a probe ($\#nodes/probe$). Color bars visually show the magnitude of the values, and darker colors show greater values. Red n^* cells indicate an algorithm’s inability to solve relatively many problems. Lines in bold indicate the best performing algorithms. Each benchmark also displays the average number of nodes (n), domain size (d), tree width (w), and AND/OR and OR search tree height (h) of its problems.

4.1 Results

For extended results, please view the full paper¹.

Aggregation Tables. Data from experiments on the same benchmark using a specific AS scheme and abstraction level are combined in the aggregation tables (Figure 5). Importance Sampling corresponds to RelCB-0. At the top of each table (in red) we also report the average $\log_{10}Z$ error for DIS. To help the reader identify AS algorithm configurations that performed particularly well, rows are bolded where (1) a relatively large number of problems are solved and (2) the error is within a factor of 0.2 of the minimum error within the set of algorithms solving the greatest number of problems.

Representative Plots. Figure 6 provides a representative plot from the Grids benchmark. The legend includes probe statistics and $\log_{10}Z$ error values for each algorithm. A

DIS plot is also overlaid onto each subplot for comparison. Although plots vary between benchmarks and problem instances, here we attempt to show a plot that captures the main trend of our data. In the majority of plots, we noticed AOAS converged towards the reference Z value faster, and had overall better performance than both ORAS-DFS and pAOAS. This was also true compared to DIS in many cases. We can also see the anytime nature of the algorithms as they continue to converge towards the reference Z value over time.

AOAS vs DIS Table. Table 1 compares AOAS RandCB-256 and DIS. For each benchmark (partitioned into small and large instances), we tally how often AOAS’s Z estimates: (i) fall within DIS probabilistic bounds, (ii) were comparable or better than DIS’s², or (iii) were strictly better than DIS’s. Note that in order to compare Abstraction Sampling to DIS

¹<https://www.ics.uci.edu/~dechter/publications.html>

²comparable means falling within ± 0.1 or ± 0.5 of the DIS $\log_{10}Z$ value, for small and large problems respectively

grid80x80.f10.wrap

Graph Type: MARKOV, N: 6400, cliques: 19200, K(min): 2, K(max): 2, K(avg): 2.0, Scope Size (max): 2, Fxn Size (max): 4




Figure 6: Plots of AS and DIS on a representative large Grid problem. The dashed line marks the reference $\log_{10}Z$ value. The legend marks the algorithm and abstraction level, number of probes ($\#p$), size of a probe ($\#n/p$), and $\log_{10}Z$ error (*est. error*).

fairly, we selected only a single set of Abstraction Sampling parameters (AOAS RandCB-256).

4.2 Analysis

AOAS Performance. The aggregation tables (Figure 5) show that for each benchmark (except Promedas) AOAS algorithms are among the top performers. Furthermore, AOAS stands out as the *only* best performing algorithm when solving both large Grids and Linkage-Type4 problems. Note that, for Linkage-Type4 (Figure 5(c)), AOAS with RelCB-8 is able to produce estimates for 41 problems, where the best of any other scheme, ORAS-DFS, is only able to produce estimates for 16 problems.

AOAS vs. pAOAS. A noticeable empirical difference between AOAS and pAOAS is their respective abilities to control their probe sizes with more refined abstractions. For example, within the Grids aggregation table (Figure 5(b)) both AOAS and pAOAS RelCB-0 have average probe size of 3326, however, upon increasing the abstraction level of each individually, probe sizes for pAOAS explode. *This is a crucial difference as being able to temper the growth of probe sizes is instrumental to the scalability of the algorithms with increasing abstraction levels.*

AOAS vs. ORAS-DFS. Despite the relationship stated in Theorem 1, AOAS regularly outperforms ORAS-DFS across the same abstraction level (see Figure 5). We believe that this is due to the smaller probe size of AOAS as compared with ORAS-DFS (for the same abstraction schema), which we believe is impacted significantly by the cutset effect [Mateescu and Dechter, 2005]. Namely, the number of i -cutset variables in the AND/OR space is smaller than that of the corresponding the OR space and, as noted earlier, since the heuristic is exact below the cutset, only the cutset variables are sampled by the abstraction sampling schemes.

Comparing with Non Abstraction Sampling Schemes. [Broka *et al.*, 2018] already showed that earlier versions of AS are highly competitive against Weighted Mini-Bucket Importance Sampling [Liu *et al.*, 2015] and IJGP-SampleSearch

[Gogate and Dechter, 2011]. By superseding the performance of the earlier AS algorithms, we also further solidify superiority over these non-AS state-of-the-art schemes. In addition, we also test against a new competitive sampling scheme, *Dynamic Importance Sampling* (DIS) [Lou *et al.*, 2017; Lou *et al.*, 2019], which produces probabilistic bounds as well as estimates in an anytime manner. Through Table 1 we see that not only do AOAS estimates fall within the bounds produced by DIS in most cases, but also that the estimates are comparable or better than that of DIS. This is particularly true of hard problems. We also note that unlike AS algorithms (AOAS in particular), DIS was unable to generate estimates for Linkage-Type4 problems. Furthermore, the aggregated errors for AS are often smaller than that of DIS's (Table 5).

5 Summary and Conclusion

The paper presents a new algorithm for abstraction sampling over AND/OR search spaces that can accommodate any abstraction function. In particular it freed AND/OR abstraction sampling from its earlier handicapping restriction of properness and opens up the horizon for far more scalable and effective performance. We provided analysis and extensive empirical evaluation over 5 benchmarks with large and hard problems showing clearly that the new AOAS algorithm is overall superior to some of the most competitive state-of-the-art IS schemes and to previous AS schemes. We illustrated that, like previous AND/OR AS schemes, AOAS maintains the ability to exploit the decomposition expressed in AND/OR search spaces, yet it also has far better control of probe-size thus making it uniquely scalable. This gives AOAS the power to more smoothly interpolate between (stochastic) sampling and (systematic) search [Broka *et al.*, 2018]. With our findings, we can now turn our attention to advancing the scheme along its central components including the development of good abstraction functions beyond context-based abstractions, understanding the desired strengths of a heuristic function and how it should be balanced with the strength of abstractions, and finally, how we should fit an abstraction level to a problem instance.

References

- [Bidyuk and Dechter, 2007] Bozhena Bidyuk and Rina Dechter. Cutset sampling for bayesian networks. *J. Artif. Intell. Res. (JAIR)*, 28:1–48, 2007.
- [Broka *et al.*, 2018] Filjor Broka, Rina Dechter, Alexander T. Ihler, and Kalev Kask. Abstraction sampling in graphical models. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, pages 632–641, 2018.
- [Chen, 1992] P.-C. Chen. Heuristic sampling: A method for predicting the performance of tree searching programs. *SIAM Journal on Computing*, 21:295–315, 1992.
- [Darwiche, 2009] A. Darwiche. *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press, 2009.
- [Dechter and Mateescu, 2007] Rina Dechter and Robert Mateescu. AND/OR search spaces for graphical models. *Artificial Intelligence*, 171(2-3):73–106, 2007.
- [Dechter and Rish, 2003] Rina Dechter and Irina Rish. Mini-buckets: A general scheme for bounded inference. *J. ACM*, 50(2):107–153, 2003.
- [Dechter, 2013] Rina Dechter. *Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2013.
- [Gogate and Dechter, 2011] Vibhav Gogate and Rina Dechter. Samplesearch: Importance sampling in presence of determinism. *Artif. Intell.*, 175(2):694–729, 2011.
- [Knuth, 1975] D.E. Knuth. Estimating the efficiency of backtracking algorithms. *Math. Comput.*, 29:1121–136, 1975.
- [Liu and Ihler, 2011] Qiang Liu and Alexander T. Ihler. Bounding the partition function using holder’s inequality. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 849–856, 2011.
- [Liu *et al.*, 2015] Qiang Liu, John W Fisher III, and Alexander T Ihler. Probabilistic variational bounds for graphical models. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1432–1440, Montreal, Canada, 2015. Curran Associates, Inc.
- [Lou *et al.*, 2017] Qi Lou, Rina Dechter, and Alexander T. Ihler. Dynamic importance sampling for anytime bounds of the partition function. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 3199–3207, 2017.
- [Lou *et al.*, 2019] Qi Lou, Rina Dechter, and Alexander Ihler. Interleave variational optimization with monte carlo sampling: A tale of two approximate inference paradigms. 2019.
- [Mateescu and Dechter, 2005] Robert Mateescu and Rina Dechter. And/or cutset conditioning. In *International Joint Conference on Artificial Intelligence (Ijcai-2005)*, 2005.
- [Pearl, 1988] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.
- [Rizzo, 2007] Maria L. Rizzo. *Statistical computing with R*. Chapman & Hall/CRC, 2007.
- [Rubinstein and Kroese, 2007] Reuven Y. Rubinstein and Dirk P. Kroese. *Simulation and the Monte Carlo Method (Wiley Series in Probability and Statistics)*. 2 edition, 2007.

6 Supplemental Information

In this section, we provide supplemental material to help visualize, elaborate on, and emphasize the importance of the contributions of our work.

6.1 Additional Algorithm Traces

Here we provide a step-by-step trace of two additional algorithms: `pAOAS`, the previous AND/OR AS scheme which uses proper abstractions, and `naive-npAS`, a naive AS algorithm that uses non-proper abstractions and generates invalid probes. For each trace, we follow a DFS traversal of the pseudo-tree \mathcal{T} from Figure 1(part b), sampling from the corresponding full AND-OR search tree $T_{\mathcal{T}}$ (Figure 2(a)). For our example, we use an abstraction function that groups together nodes that have the same domain value. For each variable, the abstract state they belong to can be visualized based on the color of the AND nodes in their depicted trees. When performing abstractions, we box nodes that are in the same abstraction scope, and thus can be abstracted together should they be part of the same abstract state. We also depict pruning of nodes by showing a red "X" through the incoming arc of the subtree being pruned.

pAOAS trace. In Figure 7 we show the trace of `pAOAS`, which adheres to proper abstractions to guarantee valid probes. We start with the expansion and abstraction of the nodes of the branching variable B (Figure 7(a)). Since each node belongs to a different abstraction group, each is kept as a representative. Next we expand the probe to variable A and abstract its nodes (Figure 7(b)). Among A 's AND nodes, we see two red nodes and two blue nodes. If we were applying non-proper abstractions, the two nodes of each color could be abstracted together, with only one being kept as a representative (the other discarded). However, since we are adhering to properness, only nodes occurring in the same subtree rooted under the *same* AND node of branching variable B can be abstracted together. (This scope is depicted by the dashed box in the figures). Thus, due to the scope of abstractions enforced by properness, the two red nodes cannot be abstracted together, and neither can the two blue nodes. Thus, again, each node is kept! The same applies to variable C when expanded (Figure 7(c)); none of the purple/teal nodes can be removed. When expanding and processing variable D (Figure 7(d)), the pairs rooted under the same AND node of B can be abstracted together. We see the result in Figure 7(e), where under $B = 0$ both representatives of D kept are the left-most nodes, and under $B = 1$, the green node kept is from under $C = 0$ and the teal node kept is from under $C = 1$. Since the subtree under $B = 0, C = 1$ does not extend to an assignment to D (and thus would yield a partial configuration), the subtree under $C = 1$ is pruned. The resulting probe (Figure 7(f)) is a valid probe encoding 8 solutions: $(B = 0, A = 0, C = 0, D = 0)$, $(B = 0, A = 0, C = 0, D = 1)$, $(B = 0, A = 1, C = 0, D = 0)$, $(B = 0, A = 1, C = 0, D = 1)$, $(B = 1, A = 0, C = 0, D = 0)$, $(B = 1, A = 0, C = 1, D = 1)$, $(B = 1, A = 1, C = 0, D = 0)$, and $(B = 1, A = 1, C = 1, D = 1)$. Through this simple example, We can already see how properness severely hindered our ability to limit probe size, a major drawback of the former

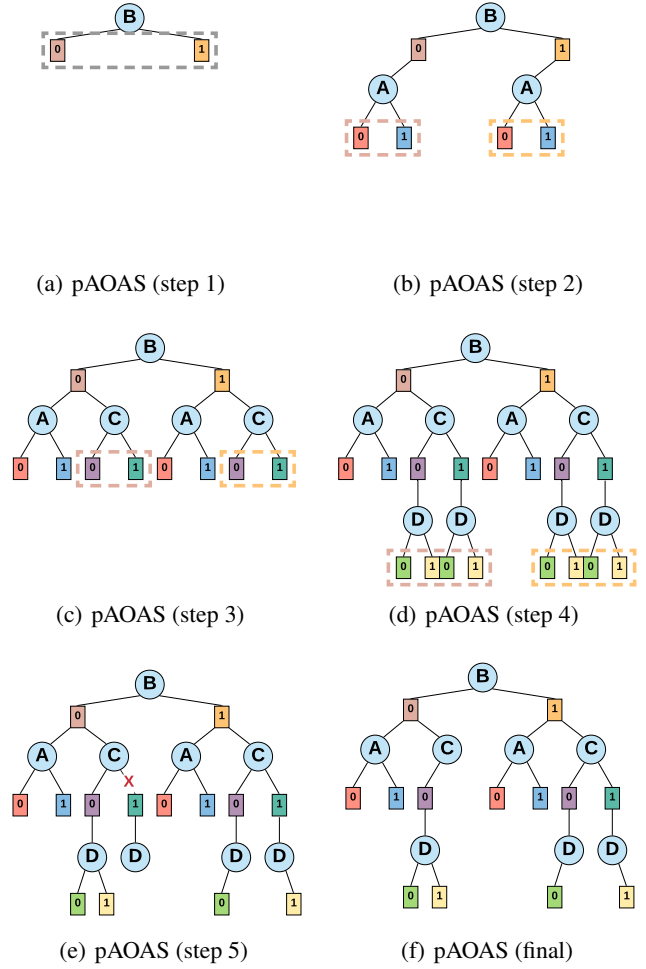




Figure 7: Example pAOAS processing.



`pAOAS` algorithm.

naive-npAS trace. In Figure 8 we show a trace of a naive AS algorithm that uses non-proper abstractions, `naive-npAS`. As with `pAOAS`, we first expand and abstract the nodes of variable B (Figure 8(a)). Next, we expand the probe to A and abstract across its AND nodes (Figure 8(b)). Note that since we are not adhering to properness, the scope of our abstraction is *all* the AND nodes of A , regardless of which subtree under the branching variable B they fall in. Assuming both red and blue node representatives are from under $B = 0$, we next expand and abstract C similarly (Figure 8(c)). This time, by chance, the chosen representatives of C 's AND nodes are all under $B = 1$ (Figure 8(d)), and we expand the probe to D . The representative green node of D is chosen under $B = 1, C = 0$, and that of the yellow node under $B = 1, C = 1$ (Figure 8(e)). *At this point, we can see we have an invalid probe.* Namely, the probe contains no full configurations, rather only four partial configurations: $(B = 0, A = 0)$, $(B = 0, A = 1)$, $(B = 1, C = 0, D = 0)$, $(B = 1, C = 1, D = 1)$. Thus, if we were to prune to remove incomplete configurations (Figures 8(e) and 8(f)), the




(a) Naive non-proper AS (step 1)

(b) Naive non-proper AS (step 2)



(c) Naive non-proper AS (step 3)

(d) Naive non-proper AS (step 4)



(e) Naive non-proper AS (step 5)



(f) Naive non-proper AS (final)

Figure 8: Example invalid naive non-proper AS processing.

entire probe would be pruned, and all work done resulting in no usable probe.

6.2 Additional Results

On the following pages, we include more comprehensive results from our experimentation.

AOAS vs. DIS Comparison . We provide two additional comparisons of AOAS algorithms with DIS. As mentioned in the main paper, we notice that different parameterizations work better with some problem types as compared to others. We leave exploring this phenomenon to future work.

Full Aggregation Tables . We provide our full aggregation tables for each of the benchmark problem types, which includes both small and large problems, and more granular data reporting. As with the aggregation tables in the paper, we have highlighted best performing algorithms.

Representative Plots . We provide five more representative plots, one for each benchmark type. Although plots vary

between problem types, we attempt to show more interesting plots and those that also capture the main trend of our data. Note that as DIS was unable to solve any of the Linkage-Type4 problems, it is thus is not present in the representative plot.

Table 2: How often AOAS *RelCB-8* estimates: fall within DIS probabilistic bounds (*in Bounds*), were comparable/better than DIS's ($AOAS \geq$), and were strictly better than DIS's ($AOAS >$)

Benchmark	Size	Total	in Bounds	$AOAS \geq$	$AOAS >$
DBN	Small	66	54	51	46
	Large	48	36	34	31
Grids	Small	8	6	5	1
	Large	19	12	14	13
Linkage-Type4	Large	82	82	82	82
Pedigree	Small	24	24	24	17
Promedas	Small	65	61	45	30
	Large	173	118	90	69

Table 3: How often AOAS *RandCB-16* estimates: fall within DIS probabilistic bounds (*in Bounds*), were comparable/better than DIS's ($AOAS \geq$), and were strictly better than DIS's ($AOAS >$)

Benchmark	Size	Total	in Bounds	$AOAS \geq$	$AOAS >$
DBN	Small	66	64	53	47
	Large	48	41	39	37
Grids	Small	8	5	4	1
	Large	19	10	7	6
Linkage-Type4	Large	82	82	82	82
Pedigree	Small	24	24	24	20
Promedas	Small	65	63	59	33
	Large	173	149	128	106

		i-Bound = 10				60 sec								3600 sec								
Problem Type	Sz.	Graph Scheme	Context Scheme	Abstr. Level	Agg. Exp.	log(err)	error distribution					#probes	#nodes/probe	Agg. Exp.	log(err)	error distribution					#probes	#nodes/probe
							0.5	1	2	5	10					0.5	1	2	5	10		
DBN	(small, 780, 2, 28, 30/66)	AOAS	RelCB	0	66	-0.951	48	53	56	61	66	1.43E+05	136	66	-0.692	52	56	58	62	66	8.90E+06	136
				4	66	-1.751	40	43	47	59	65	1.20E+04	1898	66	-1.430	43	46	50	60	66	7.54E+05	1898
				8	66	-2.158	31	38	47	58	64	8.98E+02	22254	66	-1.530	43	45	49	60	65	5.69E+04	22254
			RandCB	16	66	-1.026	47	51	54	61	66	1.23E+04	1769	66	-0.473	54	60	62	63	66	7.72E+05	1769
				256	66	-1.236	41	49	52	60	66	6.57E+03	3942	66	-0.669	55	57	59	61	66	4.10E+05	3942
				0	66	-0.901	48	51	58	61	66	1.72E+05	136	66	-0.413	53	57	60	61	66	1.08E+07	136
		pAOAS	RelCB (k=5)	2	66	-0.975	51	54	55	61	65	5.22E+04	510	66	-0.465	54	59	61	63	66	3.30E+06	510
				4	66	-1.768	39	43	46	59	65	1.48E+04	1898	66	-1.371	44	46	51	60	66	9.32E+05	1898
				16	66	-0.881	44	53	57	61	66	1.13E+04	1746	66	-0.531	54	58	60	63	66	7.08E+05	1746
			RandCB	256	66	-1.294	43	53	55	60	65	6.18E+03	3994	66	-0.557	53	58	61	62	66	3.84E+05	3994
				0	66	-1.188	46	52	55	60	65	1.30E+05	136	66	-0.646	54	57	59	62	66	8.09E+06	136
				4	66	-1.639	40	44	47	60	65	1.06E+04	1898	66	-1.471	43	47	49	60	66	6.62E+05	1898
		ORAS -BFS	RelCB	8	66	-2.115	31	38	46	59	64	8.30E+02	22254	66	-1.523	44	46	49	60	65	5.25E+04	22254
				16	66	-1.561	42	49	51	59	63	7.96E+03	1539	66	-0.839	52	55	58	60	66	4.95E+05	1539
				256	66	-1.475	41	51	53	59	65	3.54E+03	3651	66	-0.748	50	56	57	61	66	2.20E+05	3651
			RandCB	0	66	-1.066	50	52	55	61	65	1.30E+05	136	66	-0.695	51	57	58	61	66	8.35E+06	136
				4	66	-1.294	39	44	47	59	65	1.07E+04	1898	66	-1.348	43	47	51	60	66	6.81E+05	1898
				8	66	-2.330	29	37	47	58	63	8.51E+02	22254	66	-1.584	42	45	49	60	65	5.48E+04	22254
		ORAS -DFS	RelCB	16	66	-1.106	45	52	54	60	66	8.02E+03	1543	66	-0.576	52	56	61	62	66	5.13E+05	1543
				256	66	-1.198	36	48	54	60	65	3.62E+03	3640	66	-0.587	53	55	60	62	66	2.29E+05	3640
				0	66	-1.726	39	44	47	59	65	1.07E+04	1898	66	-1.348	43	47	51	60	66	6.81E+05	1898
			RandCB	4	48	-5.150	5	8	18	29	40	8.37E+03	434	48	-3.251	15	19	29	39	42	5.40E+05	434
				4	48	-6.767	3	4	10	24	34	8.18E+02	6352	48	-4.045	5	11	20	39	41	5.52E+04	6352
				8	48	-9.147	1	1	7	17	27	6.19E+01	91614	48	-5.180	5	6	12	27	40	4.08E+03	91614
	AOAS	RelCB	16	48	-3.896	5	8	20	34	44	7.84E+02	6249	48	-2.358	8	17	31	42	45	5.20E+04	6249	
			256	48	-5.142	4	5	14	30	40	2.12E+02	41146	48	-2.606	10	16	30	41	45	1.38E+04	41146	
			0	48	-5.233	4	8	20	32	39	9.48E+03	434	48	-3.123	14	23	31	38	42	6.35E+05	434	
		RandCB	2	48	-5.756	4	6	15	27	39	3.03E+03	1690	48	-3.591	10	14	23	39	42	2.00E+05	1690	
			4	48	-5.859	3	5	11	28	39	8.74E+02	6352	48	-3.712	7	15	22	37	43	5.80E+04	6352	
			16	48	-3.946	4	6	19	34	44	8.09E+02	6275	48	-2.328	12	20	31	42	44	5.41E+04	6275	
	pAOAS	RelCB (k=5)	256	48	-4.557	4	6	18	31	43	2.17E+02	40788	48	-2.591	11	14	29	41	46	1.41E+04	40788	
			0	48	-5.858	3	5	14	30	38	5.88E+03	434	48	-3.547	9	16	28	37	42	3.95E+05	434	
			4	48	-6.910	4	5	12	24	34	4.77E+02	6586	48	-4.014	5	13	23	35	42	3.14E+04	6586	
		RandCB	8	48	-9.254	2	2	6	15	29	3.61E+01	91881	48	-5.464	3	5	15	25	39	2.39E+03	91881	
			16	48	-5.088	3	6	17	29	39	5.16E+02	4643	48	-3.115	8	14	23	41	44	3.36E+04	4643	
			256	48	-6.206	2	6	13	29	37	1.05E+02	30339	48	-3.786	6	12	22	36	43	6.87E+03	30339	
	ORAS -BFS	RelCB	0	48	-5.209	4	9	16	30	40	6.28E+03	434	48	-3.632	11	17	27	38	42	4.30E+05	434	
			4	48	-7.225	2	3	7	25	34	4.92E+02	6586	48	-4.377	4	13	19	33	42	3.42E+04	6586	
			8	48	-8.627	2	3	9	21	30	3.85E+01	91881	48	-5.284	2	5	11	29	40	2.66E+03	91881	
		RandCB	16	48	-5.293	5	7	14	30	41	5.39E+02	4617	48	-2.704	7	15	28	41	44	3.71E+04	4617	
			256	48	-7.305	2	5	11	25	32	1.08E+02	30293	48	-3.312	9	13	22	37	45	7.38E+03	30293	
			0	48	-5.209	4	9	16	30	40	6.28E+03	434	48	-3.632	11	17	27	38	42	4.30E+05	434	
	ORAS -DFS	RelCB	4	48	-7.225	2	3	7	25	34	4.92E+02	6586	48	-4.377	4	13	19	33	42	3.42E+04	6586	
			8	48	-8.627	2	3	9	21	30	3.85E+01	91881	48	-5.284	2	5	11	29	40	2.66E+03	91881	
			16	48	-5.293	5	7	14	30	41	5.39E+02	4617	48	-2.704	7	15	28	41	44	3.71E+04	4617	
		RandCB	256	48	-7.305	2	5	11	25	32	1.08E+02	30293	48	-3.312	9	13	22	37	45	7.38E+03	30293	

Figure 9: DBN Aggregation Table

		i-Bound = 10				60 sec										3600 sec									
Problem Type	Sz.	Graph Scheme	Context Scheme	Abstr. Level	Agg. Exp.	log(err)	error distribution					#probes	#nodes/probe	Agg. Exp.	log(err)	error distribution					#probes	#nodes/probe			
							0.5	1	2	5	10					0.5	1	2	5	10					
Grids	(small, 250, 2, 17, 49/249)	AOAS	RelCB	0	8	-5.824	5	5	5	5	6	5.96E+05	203	8	-4.507	5	5	5	6	6	3.58E+07	203			
				4	8	-1.118	6	6	7	7	8	5.03E+04	2797	8	-0.452	6	6	7	8	8	2.99E+06	2797			
			8	8	-0.632	6	6	7	8	8	3.02E+03	33706	8	-0.579	6	6	7	8	8	1.79E+05	33717				
			16	8	-1.889	5	5	6	7	7	6.46E+04	2223	8	-0.812	5	6	7	8	8	3.86E+06	2223				
		RandCB	256	8	-1.848	5	5	6	6	8	4.78E+04	3296	8	-1.194	6	6	6	7	8	2.86E+06	3296				
			0	8	-5.166	5	5	5	6	6	7.68E+05	203	8	-3.943	6	6	6	6	7	4.70E+07	203				
		pAOAS	RelCB (k=5)	2	8	-4.243	5	5	5	6	7	1.20E+04	80031	8	-3.325	5	5	5	6	7	7.16E+05	80031			
				4	8	-3.483	4	4	4	6	7	7.76E+02	44853263	8	-3.026	5	5	5	6	7	4.59E+04	44853263			
			16	8	-4.298	5	6	6	6	6	2.37E+03	5846955	8	-3.430	5	5	5	6	7	1.41E+05	5842977				
			256	8	-4.336	5	5	5	5	6	1.58E+03	12138610	8	-3.202	5	5	5	6	7	9.33E+04	12175504				
		ORAS -BFS	RelCB	0	8	-5.283	5	5	5	6	6	8.39E+05	163	8	-4.440	5	5	5	6	7	5.02E+07	163			
				4	8	-5.075	5	5	5	5	6	7.86E+04	2495	8	-4.172	5	5	5	6	6	4.70E+06	2495			
			8	8	-4.092	4	5	5	6	7	5.55E+03	37743	8	-3.488	5	5	5	6	7	3.30E+05	37743				
			16	8	-3.675	4	5	5	6	7	3.32E+04	2201	8	-2.467	5	5	6	7	7	1.97E+06	2201				
		RandCB	256	8	-4.458	4	4	4	5	6	1.74E+04	11893	8	-2.582	5	5	5	6	7	1.04E+06	11892				
			0	8	-5.983	5	5	5	5	6	5.49E+05	451	8	-4.561	5	5	5	6	6	3.30E+07	451				
		ORAS -DFS	RelCB	4	8	-2.108	4	5	5	7	7	5.69E+04	5010	8	-1.835	5	5	6	7	8	3.43E+06	5012			
				8	8	-1.647	5	5	6	7	7	4.51E+03	54340	8	-0.633	6	6	6	8	8	2.69E+05	54369			
			16	8	-2.633	5	5	5	7	7	2.22E+04	4621	8	-1.674	5	6	6	7	8	1.33E+06	4626				
			256	8	-2.402	5	5	5	7	8	1.41E+04	17266	8	-1.138	6	6	6	7	8	8.50E+05	17266				
	(LARGE, 3432, 2, 27, 220/3431)	AOAS	RelCB	0	19	-192.560	0	0	0	2	4	3.19E+04	3326	19	-167.080	0	0	0	4	4	1.92E+06	3326			
				4	19	-81.156	0	0	2	4	5	3.81E+03	43450	19	-67.780	0	1	4	6	7	2.34E+05	43449			
			8	19	-63.547	0	1	1	4	6	3.36E+02	408047	19	-49.468	2	3	4	5	9	2.07E+04	407963				
			16	19	-95.589	0	0	1	3	4	4.64E+03	32962	19	-79.864	1	1	3	3	7	2.82E+05	32965				
		RandCB	256	19	-95.610	0	0	1	3	4	2.74E+03	66377	19	-77.505	0	2	3	4	5	1.67E+05	66376				
			0	19	-180.151	0	0	0	2	4	6.60E+04	3326	19	-161.895	0	0	1	3	4	4.10E+06	3326				
		pAOAS	RelCB (k=5)	2	19	-145.545	0	0	0	3	4	1.12E+02	2849697	19	-129.787	0	1	3	3	4	6.83E+03	2849697			
				4	0	-	-	-	-	-	-	-	-	3	-2.972	0	1	1	2	3	4.00E+01	1.44E+08			
			16	3	-8.292	0	0	0	0	2	2.33E+00	41465226	3	-4.025	0	0	1	2	3	1.50E+02	41914557				
			256	0	-	-	-	-	-	-	-	-	0	-	-	-	-	-	-	-	-	-			
		ORAS -BFS	RelCB	0	19	-201.692	0	0	0	0	4	2.77E+04	2827	19	-172.224	0	0	0	4	4	1.67E+06	2827			
				4	15	-141.622	0	0	1	1	3	2.44E+03	36662	19	-136.514	0	0	0	3	4	1.16E+05	45112			
			8	11	-56.152	0	0	0	0	4	2.11E+02	263539	19	-140.661	0	0	2	4	4	7.34E+03	719763				
			16	11	-49.413	0	0	0	2	4	1.12E+03	15073	19	-143.674	0	0	0	2	4	3.90E+04	40976				
		RandCB	256	11	-50.398	0	0	0	0	3	8.07E+01	192189	11	-41.096	0	0	0	2	4	4.81E+03	192244				
			0	19	-200.206	0	0	0	1	3	1.07E+04	6844	19	-176.694	0	0	0	4	4	6.65E+05	6844				
		ORAS -DFS	RelCB	4	19	-118.794	0	0	0	2	4	1.02E+03	78207	19	-96.276	0	1	3	4	6	6.32E+04	78212			
				8	11	-29.495	0	0	1	4	4	1.44E+02	335559	19	-60.563	0	2	4	5	7	5.10E+03	886885			
			16	11	-39.309	0	0	0	3	4	2.89E+02	27582	19	-96.497	0	0	0	4	4	1.02E+04	75095				
			256	11	-46.133	0	0	0	0	2	4.21E+01	260459	19	-95.435	0	0	1	3	5	1.47E+03	786178				

Figure 10: Grids Aggregation Table

		i-Bound = 10				60 sec										3600 sec									
Problem Type	Sz.	Graph Scheme	Context Scheme	Abstr. Level	Agg. Exp.	log(err)	error distribution					#probes	#nodes/probe	Agg. Exp.	log(err)	error distribution					#probes	#nodes/probe			
							0.5	1	2	5	10					0.5	1	2	5	10					
Linkage-Type4	(LARGE, 10822, 1.8, 24, 581/5745)	AOAS	RelCB	0	4	-26.087	0	0	0	0	0	3.15E+03	1076	7	-32.487	0	0	0	0	0	5.52E+05	829			
				4	18	-18.032	0	0	0	2	6	1.72E+03	19089	26	-14.174	1	1	1	7	10	1.92E+05	16472			
			8	26	-8.107	1	2	3	10	19	1.15E+02	525633	41	-11.090	3	5	10	17	26	1.20E+04	463468				
			16	20	-22.868	0	0	0	1	4	2.96E+03	11472	26	-18.582	0	0	0	4	9	2.50E+05	10532				
		RandCB	256	21	-25.033	0	0	0	0	2	2.04E+03	19661	27	-18.221	0	1	1	4	9	1.74E+05	18523				
			0	5	-31.016	0	0	0	0	0	6.63E+04	1009	8	-27.742	0	0	0	0	0	1.14E+07	759				
		pAOAS	RelCB (k=5)	2	7	-29.935	0	0	0	0	0	3.61E+02	449799	13	-27.201	0	0	0	0	1	6.96E+04	280176			
				4	6	-30.513	0	0	0	0	0	1.37E+01	20495116	11	-25.427	0	0	0	0	0	8.90E+02	34177733			
			16	9	-29.958	0	0	0	0	0	3.70E+01	11831679	12	-24.204	0	0	0	0	1	4.24E+03	8820722				
			256	6	-31.092	0	0	0	0	0	3.25E+01	12692016	14	-30.930	0	0	0	0	0	2.58E+03	24532116				
		ORAS -BFS	RelCB	0	4	-28.691	0	0	0	0	0	3.96E+03	1132	7	-30.465	0	0	0	0	0	2.78E+06	870			
				4	4	-25.843	0	0	0	0	0	1.38E+03	18919	10	-23.361	0	0	0	0	0	6.33E+04	18615			
			8	2	-26.171	0	0	0	0	0	1.80E+01	574416	7	-20.932	0	0	0	0	1	1.82E+03	475657				
			16	2	-15.949	0	0	0	0	0	7.40E+01	21214	8	-23.358	0	0	0	0	0	2.10E+04	16660				
		RandCB	256	3	-20.784	0	0	0	0	0	5.00E+00	273550	9	-24.565	0	0	0	0	0	3.96E+02	266340				
			0	3	-28.178	0	0	0	0	0	1.51E+03	3292	7	-29.820	0	0	0	0	0	1.42E+06	2520				
		ORAS -DFS	RelCB	4	7	-25.801	0	0	0	0	0	2.81E+02	33456	15	-20.802	0	0	0	0	2	5.14E+04	26863			
				8	3	-24.080	0	0	0	0	0	2.60E+01	436144	16	-18.971	0	0	0	0	5	2.44E+03	355404			
			16	2	-18.805	0	0	0	0	0	9.75E+01	26231	13	-22.843	0	0	0	0	1	1.05E+04	23723				
			256	1	-31.484	0	0	0	0	0	6.00E+00	78862	12	-27.151	0	0	0	0	0	1.18E+02	288516				

Figure 11: Linkage-Type4 Aggregation Table

Problem Type		i-Bound = 10				60 sec							3600 sec									
Sz.	Graph Scheme	Context Scheme	Abstr. Level	Agg. Exp.	log(err)	error distribution					#probes	#nodes/ probe	Agg. Exp.	log(err)	error distribution					#probes	#nodes/ probe	
						0.5	1	2	5	10					0.5	1	2	5	10			
Pedigree	(small, 840, 2, 19, 88/681)	AOAS	RelCB	0	24	-0.015	24	24	24	24	24	2.32E+05	279	24	0.004	24	24	24	24	24	1.41E+07	279
				4	24	0.003	24	24	24	24	24	2.18E+04	6130	24	-0.001	24	24	24	24	24	1.31E+06	6134
				8	24	-0.034	24	24	24	24	24	1.64E+03	184372	24	0.001	24	24	24	24	24	9.73E+04	184496
			RandCB	16	24	-0.005	24	24	24	24	24	3.22E+04	3389	24	-0.002	24	24	24	24	24	1.96E+06	3389
				256	24	-0.006	24	24	24	24	24	2.31E+04	5485	24	-0.003	24	24	24	24	24	1.39E+06	5481
				0	24	-0.018	24	24	24	24	24	4.70E+05	279	24	-0.001	24	24	24	24	24	2.86E+07	279
		pAOAS	RelCB (k=5)	2	24	-0.016	24	24	24	24	24	7.45E+03	289407	24	0.004	24	24	24	24	24	4.47E+05	301596
				4	20	-0.048	19	20	20	20	20	4.29E+02	11540373	20	-0.021	20	20	20	20	20	2.57E+04	11592761
				16	24	-0.015	23	24	24	24	24	1.04E+03	4860239	24	0.002	24	24	24	24	24	6.20E+04	4410215
			RandCB	256	24	-0.031	23	24	24	24	24	7.02E+02	10309437	24	-0.001	24	24	24	24	24	4.7E+04	1.02E+07
				0	24	-0.003	24	24	24	24	24	2.36E+05	254	24	0.001	24	24	24	24	24	1.45E+07	254
				4	24	-0.042	23	24	24	24	24	1.82E+04	5538	24	-0.013	24	24	24	24	24	1.11E+06	5543
	ORAS -BFS	RelCB	8	24	-0.082	24	24	24	24	24	1.17E+03	245933	24	-0.018	24	24	24	24	24	7.01E+04	245920	
			16	24	-0.026	22	24	24	24	24	7.83E+03	3674	24	-0.013	24	24	24	24	24	4.76E+05	3680	
			256	24	-0.010	22	24	24	24	24	1.76E+03	24155	24	-0.004	24	24	24	24	24	1.05E+05	24117	
		RandCB	0	24	-0.037	24	24	24	24	24	5.00E+04	1075	24	-0.009	24	24	24	24	24	3.09E+06	1075	
			4	24	-0.034	24	24	24	24	24	5.80E+03	11564	24	-0.004	24	24	24	24	24	3.54E+05	11584	
			8	24	-0.014	23	23	24	24	24	5.72E+02	266530	24	-0.003	24	24	24	24	24	3.44E+04	266512	
	ORAS -DFS	RelCB	16	24	-0.086	21	24	24	24	24	1.52E+03	8952	24	-0.003	24	24	24	24	24	24	9.25E+04	8967
			256	24	-0.208	20	23	24	24	24	3.99E+02	40587	24	0.099	23	23	24	24	24	2.40E+04	40839	
			0	24	-0.037	24	24	24	24	24	5.00E+04	1075	24	-0.009	24	24	24	24	24	3.09E+06	1075	

Figure 12: Pedigree Aggregation Table

Problem Type		i-Bound = 10				60 sec							3600 sec											
Sz.	Graph Scheme	Context Scheme	Abstr. Level	Agg. Exp.	log(err)	error distribution					#probes	#nodes/ probe	Agg. Exp.	log(err)	error distribution					#probes	#nodes/ probe			
						0.5	1	2	5	10					0.5	1	2	5	10					
Promedas	(small, 619, 2, 15, 62/611)	AOAS	RelCB	0	64	-0.675	40	50	57	62	64	2.22E+06	85	65	-0.231	59	62	64	64	65	1.28E+08	86		
				4	64	-0.463	39	52	59	64	64	1.77E+06	828	65	-0.327	57	61	63	64	65	1.02E+08	841		
				8	64	-0.616	40	48	55	64	64	1.67E+06	8537	65	-0.262	54	59	64	65	65	9.59E+07	8757		
				RandCB	16	65	-0.399	49	55	62	64	65	1.64E+06	1116	65	-0.111	60	64	64	65	65	9.71E+07	1116	
					256	65	-0.452	47	55	63	65	65	1.62E+06	2357	65	-0.230	58	62	64	65	65	9.51E+07	2357	
					0	64	-0.461	44	51	59	64	64	3.82E+06	85	65	-0.291	58	63	63	64	65	2.20E+08	86	
			pAOAS	RelCB (k=5)	2	65	-0.719	39	45	56	64	65	3.11E+06	4139	65	-0.249	57	61	64	64	65	1.78E+08	4135	
					4	65	-0.862	38	46	54	64	64	3.03E+06	187691	65	-0.361	51	60	64	64	65	1.73E+08	184671	
					16	65	-0.418	49	55	62	64	65	2.98E+06	96663	65	-0.183	59	63	64	65	65	1.71E+08	95997	
				RandCB	256	65	-0.444	46	55	62	65	65	2.96E+06	580315	65	-0.184	59	62	64	65	65	1.71E+08	575424	
					0	64	-0.597	44	50	57	62	64	3.84E+06	89	65	-0.284	57	60	64	64	65	2.20E+08	91	
					4	64	-0.970	36	42	52	63	64	3.42E+06	832	64	-0.283	52	57	64	64	64	1.90E+08	832	
		ORAS -BFS	RelCB	8	64	-1.256	32	33	47	61	64	3.12E+06	9696	64	-0.324	51	58	62	64	64	1.81E+08	9695		
				16	65	-0.814	40	49	55	64	64	3.29E+06	1304	65	-0.410	52	58	63	64	65	1.82E+08	1304		
				256	65	-0.724	42	46	56	65	65	3.15E+06	7325	65	-0.321	51	59	64	65	65	1.80E+08	7327		
			RandCB	0	64	-0.919	37	40	54	63	64	3.25E+06	295	65	-0.422	56	59	62	64	64	1.86E+08	295		
				4	64	-0.768	40	46	55	64	64	3.30E+06	1384	65	-0.393	56	59	63	64	65	1.85E+08	1401		
				8	64	-0.783	38	43	54	63	64	2.98E+06	10952	65	-0.450	49	57	60	64	65	1.71E+08	11162		
		ORAS -DFS	RelCB	16	65	-0.955	40	47	53	64	64	3.13E+06	3422	65	-0.427	49	57	62	64	65	1.80E+08	3423		
				256	65	-0.908	36	42	53	64	65	3.12E+06	23242	65	-0.389	46	55	61	65	65	1.79E+08	23204		
				0	64	-0.919	37	40	54	63	64	3.25E+06	295	65	-0.422	56	59	62	64	64	1.86E+08	295		
			(LARGE, 1209, 2, 27, 114/1193)	AOAS	RelCB	4	123	-5.601	10	14	31	68	105	4.42E+04	1938	145	-4.106	24	36	62	100	129	2.43E+06	2127
						8	108	-5.922	3	7	18	54	93	3.35E+03	21540	139	-4.520	19	32	53	98	121	1.81E+05	23935
						16	132	-5.713	10	18	36	83	111	3.39E+04	2343	154	-4.142	26	45	71	113	140	1.89E+06	2495
	256	173				-6.624	15	20	40	106	143	8.08E+03	12848	173	-3.531	29	45	78	125	161	4.83E+05	12845		
	0	108				-6.100	7	12	22	62	86	5.84E+05	164	131	-4.664	18	30	46	85	112	3.31E+07	171		
	4	103				-5.752	6	7	22	77	111	1.85E+04	10960	158	-5.158	14	31	59	107	132	1.03E+06	11266		
	pAOAS	RelCB (k=5)	4		103	-5.736	1	5	15	52	90	9.30E+02	937740	145	-5.178	10	19	46	95	126	4.61E+04	925471		
			16		132	-4.321	11	20	44	88	121	9.73E+02	561360	162	-3.306	35	58	89	127	150	5.58E+04	534634		
			256		165	-4.342	9	18	53	119	154	4.10E+01	17905502	171	-2.050	34	65	107	155	168	2.36E+03	22671070		
		RandCB	0		100	-6.506	1	3	16	50	78	3.20E+05	201	126	-4.580	15	26	41	85	111	1.80E+07	213		
			4		76	-6.240	2	5	12	35	61	4.23E+04	2069	113	-5.029	7	17	30	69	98	2.31E+06	2254		
			8		59	-8.510	0	2	2	13	40	4.03E+03	23630	101	-5.935	3	5	17	52	85	2.14E+05	25923		
	ORAS -BFS	RelCB	16	93	-6.691	5	9	13	38	73	9.22E+03	2875	126	-5.059	13	20	40	82	108	4.92E+05	3124			
			256	173	-4.571	10	15	28	111	164	3.87E+02	42472	173	-2.710	17	36	71	152	172	2.31E+04	42472			
			0	95	-6.938	4	5	13	39	73	1.62E+05	365	128	-5.479	9	18	38	80	105	9.69E+06	361			
		RandCB	4	107	-7.436	2	4	11	48	78	3.16E+04	2379	134	-5.261	10	19	36	83	111	1.78E+06	2566			
			8	98	-7.098	2	4	12	44	77	3.54E+03	24865	134	-5.073	14	20	37	81	118	1.96E+05	27440			
			16	119	-6.627	5	8	24	63	100	2.81E+03	4751	144	-4.600	18	32	50	98	128	1.66E+05	4897			
	256	173	-5.050	5	7	28	101	163	4.30E+01	103824	173	-2.821	15	34	70	151								

rbm_21

Graph Type: MARKOV, N: 42, cliques: 483, K(min): 2, K(max): 2, K(avg): 2.0, Scope Size (max): 2, Fxn Size (max): 4




Figure 14: Plots of AS and DIS on a small DBN problem. The dashed line marks the reference $\log_{10} Z$ value. The legend marks the algorithm and abstraction level, number of probes ($\#p$), size of a probe ($\#n/p$), and $\log_{10} Z$ error (*est. error*).

grid40x40.f15

Graph Type: MARKOV, N: 1600, cliques: 4720, K(min): 2, K(max): 2, K(avg): 2.0, Scope Size (max): 2, Fxn Size (max): 4

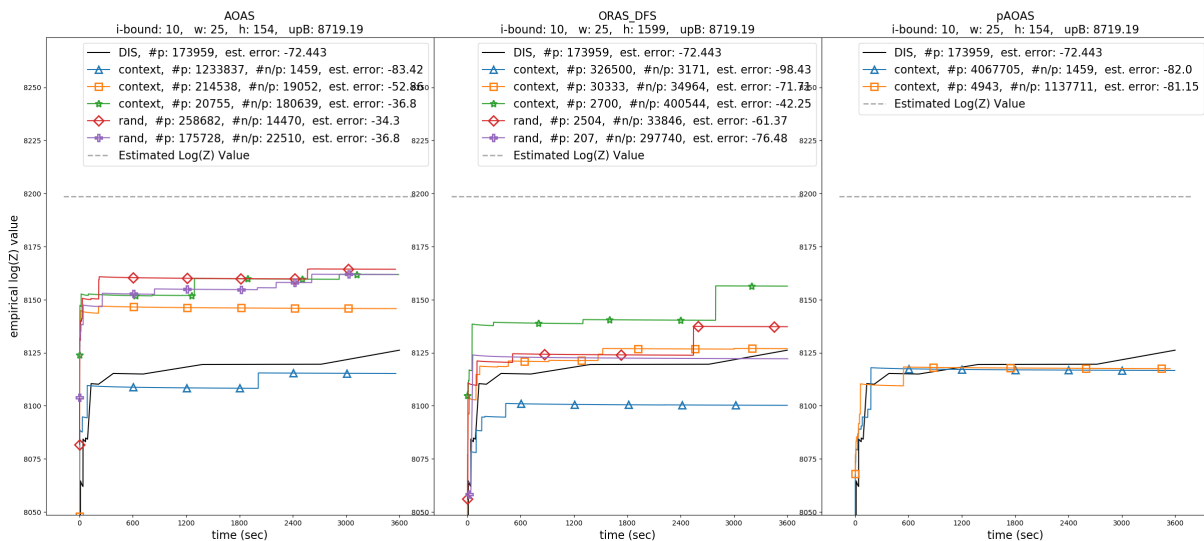


Figure 15: Plots of AS and DIS on a large Grids problem. The dashed line marks the reference $\log_{10} Z$ value. The legend marks the algorithm and abstraction level, number of probes ($\#p$), size of a probe ($\#n/p$), and $\log_{10} Z$ error (*est. error*).

pedfile_Families_Type4_110_15

Graph Type: MARKOV, N: 6639, cliques: 6639, K(min): 1, K(max): 5, K(avg): 1.8, Scope Size (max): 4, Fxn Size (max): 250




Figure 16: Plots of AS and DIS on a large Linkage-Type4 problem. The dashed line marks the reference $\log_{10} Z$ value. The legend marks the algorithm and abstraction level, number of probes ($\#p$), size of a probe ($\#n/p$), and $\log_{10} Z$ error (*est. error*).

pedigree13

Graph Type: MARKOV, N: 1077, cliques: 1077, K(min): 1, K(max): 3, K(avg): 1.9, Scope Size (max): 4, Fxn Size (max): 54

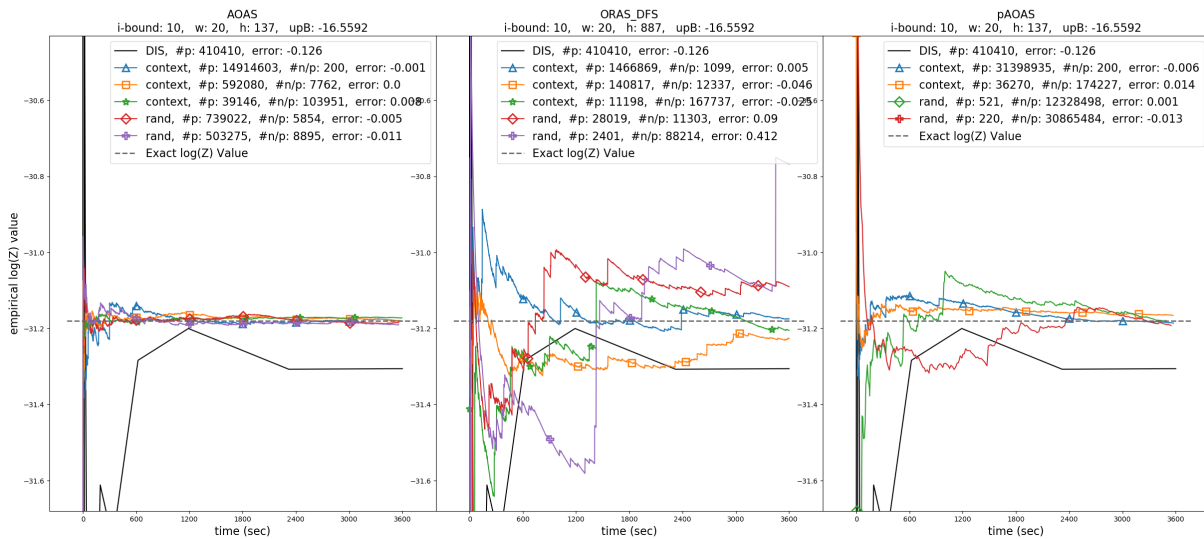


Figure 17: Plots of AS and DIS on a small Pedigree problem. The dashed line marks the reference $\log_{10} Z$ value. The legend marks the algorithm and abstraction level, number of probes ($\#p$), size of a probe ($\#n/p$), and $\log_{10} Z$ error (*est. error*).

or_chain_46.fg

Graph Type: MARKOV, N: 1214, cliques: 1214, K(min): 2, K(max): 2, K(avg): 2.0, Scope Size (max): 8, Fxn Size (max): 0




Figure 18: Plots of AS and DIS on a large Promedas problem. The dashed line marks the reference $\log_{10}Z$ value. The legend marks the algorithm and abstraction level, number of probes ($\#p$), size of a probe ($\#n/p$), and $\log_{10}Z$ error (*est. error*).