

An FPRAS for Model Counting for Non-Deterministic Read-Once Branching Programs

Kuldeep S. Meel   

University of Toronto, Canada

Alexis de Colnet   

TU Wien, Austria

Abstract

Non-deterministic read-once branching programs, also known as non-deterministic free binary decision diagrams (nFBDD), are a fundamental data structure in computer science for representing Boolean functions. In this paper, we focus on #nFBDD, the problem of model counting for non-deterministic read-once branching programs. The #nFBDD problem is #P-hard, and it is known that there exists a quasi-polynomial randomized approximation scheme for #nFBDD. In this paper, we provide the first FPRAS for #nFBDD. Our result relies on the introduction of new analysis techniques that focus on bounding the dependence of samples.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases Approximate model counting, FPRAS, Knowledge compilation, nFBDD

Digital Object Identifier 10.4230/LIPIcs.ICDT.2025.30

Related Version Full Version: <https://arxiv.org/abs/2406.16515>

Funding Meel acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), funding reference number RGPIN-2024-05956; de Colnet is supported by the Austrian Science Fund (FWF), ESPRIT project FWF ESP 235. This work was done in part while de Colnet was visiting the University of Toronto.¹

Acknowledgements This research was initiated at Dagstuhl Seminar 24171 on “Automated Synthesis: Functional, Reactive and Beyond” (<https://www.dagstuhl.de/24171>). We gratefully acknowledge the Schloss Dagstuhl – Leibniz Center for Informatics for providing an excellent environment and support for scientific collaboration.

1 Introduction

Read-once branching programs or binary decision diagrams are fundamental data structures in computer science used to represent Boolean functions. Their variants have been discovered multiple times across various sub-fields of computer science, and consequently, they are referred to by many acronyms [15, 9, 23]. In this paper, we focus on non-deterministic read-once branching programs, also known as non-deterministic free binary decision diagrams (nFBDD).

We study the following computational problem:

#nFBDD. Given a non-deterministic read-once branching program B over a Boolean set of variables X , compute the number of models of B , i.e., the number of assignments over X that B maps to 1.

¹ The authors decided to forgo the old convention of alphabetical ordering of authors in favor of a randomized ordering, denoted by . The publicly verifiable record of the randomization is available at <https://www.aeaweb.org/journals/policies/random-author-order/search>



© Kuldeep S. Meel  and Alexis de Colnet 
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Database Theory (ICDT 2025).

Editors: Sudeepa Roy and Ahmet Kara; Article No. 30; pp. 30:1–30:21

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

From a database perspective, #nFBDD is an important problem owing to the recent connections between query evaluation and knowledge compilation [13, 14, 21, 20, 2, 1]. The field of knowledge compilation has its origins in the artificial intelligence community, where functions represented in input languages are compiled into target languages that can support queries tractably (often viewed as polynomial time) [10]. The typical queries of interest are satisfiability, entailment, enumeration, and counting.

The target languages in the context of databases have been variants of binary decision diagrams, also referred to as branching programs, and circuits in decomposable negation normal form (DNNF) [2, 1]. A binary decision diagram is a representation of a Boolean function as a directed acyclic graph where the nodes correspond to variables and the sinks correspond to values, i.e., 0 or 1. One of the most well-studied forms is the ordered binary decision diagram (OBDD), where the nodes correspond to variables and, along every path from root to leaf, the variables appear in the same order [9]. A generalization of OBDD is nOBDD, where internal nodes can also represent disjunction ($\vee$) gates.

nFBDD are a natural generalization of nOBDDs, as they do not impose restrictions on the ordering of Boolean variables. Since nFBDD do not impose such restrictions, they are known to be exponentially more succinct than nOBDD; that is, there exist functions for which the smallest nOBDD is exponentially larger than the smallest nFBDD [4]. From this viewpoint, nFBDDs occupy a space between nOBDD and DNNF circuits, as they are exponentially more succinct than nOBDDs, while DNNFs are quasi-polynomially more succinct than nFBDD [8, 4].

In the context of databases, the connection between knowledge compilation and query evaluation has been fruitful, leading to the discovery of both tractable algorithms and lower bounds. Of particular note is the application of the knowledge compilation paradigm in query evaluation on probabilistic databases [14], Shapley value computation [11], the enumeration of query answers, probabilistic graph homomorphism [5], counting answers to queries [7]. The knowledge compilation-based approach involves first representing the database task as a query over a Boolean function and then demonstrating that the corresponding Boolean function has a tractable representation in a given target language, which also supports the corresponding query in polynomial time [3]. For example, in the context of query evaluation over probabilistic databases, one can show that the problem of probabilistic query evaluation can be represented as a weighted model counting problem over nOBDD when the underlying query is a path query [5]. Since there is a fully polynomial-time randomized approximation scheme (FPRAS) for the problem of model counting over nOBDD [6], it follows that probabilistic query evaluation for regular path queries over tuple-independent databases admits an FPRAS [5]. In the context of aggregation tasks, the underlying query is often model counting and its variants [19].

The aforementioned strategy makes it desirable to have target languages that are as expressive as possible while still supporting queries such as counting in polynomial time. In this context, the recent flurry of results has been enabled by the breakthrough of Arenas, Croquevielle, Jayaram, and Riveros, who showed that the problem of #nOBDD admits an FPRAS [6]. As mentioned earlier, nOBDD imposes a severe restriction on variable ordering, i.e., along every path from root to leaf, the variable ordering remains the same. nFBDD generalizes nOBDD by alleviating this restriction, thereby enabling succinct representations for several functions that require exponentially large nOBDD. Since nFBDD generalize nOBDD, the #P-hardness of #nFBDD, the problem of model counting over nFBDD, immediately follows. Accordingly, in light of the recent discovery of FPRAS for #nOBDD, an important open question is whether there exists an FPRAS for #nFBDD. The best known

prior result provides a quasi-polynomial time algorithm owing to reduction of nFBDD to DNNF (and accordingly $(+, \times)$ -programs) [12]. As noted in Section 2.1, the techniques developed in the context of design of FPRAS for #nOBDD do not extend to handle the case for #nFBDD and therefore, design of FPRAS for #nFBDD would require development of new techniques.

The primary contribution of our work is to answer the aforementioned question affirmatively, which is formalized in the following theorem.

► **Theorem 1.** *Let B be an nFBDD over n variables, $\varepsilon > 0$ and $\delta > 0$. Algorithm `approxMCnFBDD`(B, ε, δ) runs in time $O(n^5 \varepsilon^{-4} \log(\delta^{-1}) |B|^6 \log |B|)$ and returns `est` with the guarantee that $\Pr[\text{est} \in (1 \pm \varepsilon) |B^{-1}(1)|] \geq 1 - \delta$.*

Organization of the paper. We start with background on nFBDD in Section 2. The different components of the FPRAS are described in Section 4 and the analysis is split in three parts: in Section 5 we introduce the key concept of derivation paths, in Section 6 we describe the particular framework for the analysis, and in Section 7 we go through the proof of the FPRAS guarantees. For space reason, the proofs of several intermediate results are deferred to the appendix.

2 Background

Given a positive integer n and m an integer less than n , $[n]$ denotes the set $\{1, 2, \dots, n\}$ and $[m, n]$ the set $\{m, m+1, \dots, n\}$. For a, b and ε three real numbers with $\varepsilon > 0$, we use $a \in (1 \pm \varepsilon)b$ to denote $(1 - \varepsilon)b \leq a \leq (1 + \varepsilon)b$, similarly, $a \in \frac{b}{1 \pm \varepsilon}$ stands for $\frac{b}{1 + \varepsilon} \leq a \leq \frac{b}{1 - \varepsilon}$. We sometimes use the special value ∞ and in particular that $\frac{1}{\infty}$ equals 0.

Boolean variables take value 0 (*false*) or 1 (*true*). An assignment α to a set X of Boolean variables is mapping from X to $\{0, 1\}$. We sometimes see α as a set $\{x \mapsto \alpha(x) \mid x \in X\}$. We denote by $\alpha_\emptyset$ the empty assignment, which corresponds to the empty set. The set of assignments to X is denoted $\{0, 1\}^X$. A Boolean function f over X is a mapping $\{0, 1\}^X$ to $\{0, 1\}$. The *models* of f are the assignments mapped to 1 by f . When not explicit, the set of variables assigned by α is denoted by $\text{var}(\alpha)$. When $\text{var}(\alpha') \cap \text{var}(\alpha) = \emptyset$, we denote by $\alpha \cup \alpha'$ the assignment to $\text{var}(\alpha') \cup \text{var}(\alpha)$ consistent with both α and α' . For S and S' two sets of assignments, we write $S \otimes S' = \{\alpha \cup \alpha' \mid \alpha \in S, \alpha' \in S'\}$.

nBDD. A *binary decision diagram* (BDD) is a directed acyclic graph (DAG) with a single source node q_{source} , two *sinks* labeled 0 and 1, and where each internal node is labeled by a Boolean variable x and has two outgoing edges: the 0-edge going to the 0-child q_0 and the 1-edge going to the 1-child q_1 (potentially $q_0 = q_1$). Internal nodes are called *decision nodes* and are written $\text{ite}(x, q_1, q_0)$ (*if $x = 1$ then go to q_1 else go to q_0*). A path in the DAG contains a variable x if it contains a node $\text{ite}(x, q_1, q_0)$. Every variables assignment α corresponds to the unique path that starts from the source and, on a decision node $\text{ite}(x, q_1, q_0)$ follows the $\alpha(x)$ -edge. *Non-deterministic BDD* (nBDD) also admit *guess nodes*: unlabeled nodes with arbitrarily many children. When a path for an assignment reaches a guess node it can pursue to any child, so several paths can correspond to the same assignment in an nBDD. For q a node in an nBDD B , $\text{var}(q)$ denotes the set of variables labeling decision nodes reachable from q (including q) in the usual sense of graph reachability. We note $\text{var}(B) = \text{var}(q_{\text{source}})$. B computes a Boolean function over $\text{var}(B)$ whose models are the assignments for which at least one path reaches the 1-sink. Every node q of B is itself the source of an nBDD and therefore represents a function over $\text{var}(q)$ whose set of models we note $\text{mod}(q)$. So $B^{-1}(1) = \text{mod}(q_{\text{source}})$. The function computed by an nBDD is also that computed by the

circuit obtained replacing every decision node $ite(x, q_1, q_0)$ by $(\neg x \wedge q_0) \vee (x \wedge q_1)$ and every guess node with children $q_1, \dots, q_k$ by $q_1 \vee \dots \vee q_k$. Thus, we call $\vee$ -nodes the guess nodes in this paper. The size of an nBDD B , denoted by $|B|$, is its number of edges.

nFBDD. An nBDD is *free* (nFBDD) when every path contains every variable at most once. There are also called in the literature *read-once non-deterministic branching programs* (1-NBP). Note that in an nFBDD, variables may appear in different order depending on the path. When the order of occurrence of the variables is consistent among all paths of the nFBDD we say that we have an *ordered* nBDD (nOBDD). We call an nFBDD *1-complete* when along every path from the source to the 1-sink, every variable occurs exactly once. We call an nFBDD *0-reduced* when it contains no decision nodes $ite(x, 0\text{-sink}, 0\text{-sink})$ and no $\vee$ -nodes that have the 0-sink among their children. Technically, 0-reduced nFBDD cannot represent functions with no models, but these functions are not considered in this paper. An nFBDD is *alternating* when its source node is a $\vee$ -node, when every $\vee$ -node only has decision nodes for children, and when every decision node has only $\vee$ -nodes and sinks for children.

► **Lemma 2.** *Every nFBDD B over n variables can be made 1-complete, 0-reduced, and alternating in time $O(n|B|^2)$.*

Proof sketch. First, we make B 0-reduced by repeating the following until reaching a fixed point: replace all nodes $ite(x, 0\text{-sink}, 0\text{-sink})$ by the 0-sink, remove the 0-sink from $\vee$ -nodes' children, and replace all $\vee$ -nodes with no child by the 0-sink. Doing these replacements bottom-up in B takes time $O(|B|)$ and results in a 0-reduced nFBDD B' with $|B'| \leq |B|$.

Second, we make B' alternating w.r.t. the $\vee$ -nodes. Replace every $\vee$ -node that have the 1-sink as a child by the 1-sink. At that point, no sink is a child of any $\vee$ -node. Next, for every $\vee$ -node q in B' , if q has a parent q' that is a $\vee$ -node, then remove q from q' 's children and add all of q 's children to q' 's children. Doing this replacement bottom-up yields an nFBDD whose $\vee$ -nodes all have only decision nodes children. The number of children of each $\vee$ -node is increased by at most $|B'|$ so the running time is $O(|B'|^2) = O(|B|^2)$. Let B'' be the resulting nFBDD. B'' is still 0-reduced.

Third, we make B' 1-complete. For every $q \in B''$, let $\text{children}(q) = (q_1, \dots, q_k)$. While there exists $i \in [k]$ such that $\text{var}(q_i) \neq \text{var}(q)$, choose $x \in \text{var}(q) \setminus \text{var}(q_i)$ and replace q_i by the node $ite(x, q_i, q_i)$ in the children of q . This adds at most n decision nodes per original child of q . So doing this for all decision nodes takes time $O(n|B''|) = O(n|B|^2)$ and gives a 1-complete nFBDD B''' which is still 0-reduced and alternating w.r.t. the $\vee$ nodes.

Finally, to make B''' alternating, we just have to consider every $ite(x, q_1, q_0)$ and, if q_b is not a $\vee$ -node or a sink, to replace it by a $\vee$ -node whose unique child is q_b . This takes time $O(B''')$. One $\vee$ -node is added for the source of the nFBDD if needed. ◀

The nodes of 1-complete 0-reduced alternating nFBDD are organized in layers $L_0, L_1, \dots, L_{2n}$. L_0 contains the sinks and, for $1 \leq i \leq 2n$ the layer L_i contains all nodes whose children (except the 0-sink) are in L_{i-1} . We write $L_{\leq i} = L_0 \cup \dots \cup L_i$, and similarly for $L_{\geq i}$, $L_{< i}$ and $L_{> i}$. Note that for all $1 \leq i \leq n$, L_{2i-1} contains only decision nodes whereas L_{2i} contains only $\vee$ -nodes. Importantly, we assume an arbitrary ordering on the children of the nodes; for every node q we have a sequence (not a set) $\text{children}(q)$ of its children.

FPRAS. For a counting problem that, given an input x of size $|x|$, aims at computing some integer value $N(x)$, a fully polynomial-time randomized approximation scheme (FPRAS) is an algorithm that, given x , $\varepsilon > 0$, and $0 < \delta < 1$, runs in time polynomial in $|x|$, $1/\varepsilon$, and $\log(1/\delta)$, and returns $\tilde{N}$ with the guarantee that $\Pr[\tilde{N} \in (1 \pm \varepsilon)N(x)] \geq 1 - \delta$. In this paper we give an #FPRAS for the problem #nFBDD.

#nFBDD

Input: an nFBDD B

Output: its number of models $|B^{-1}(1)|$

2.1 Related Work

As noted in Section 1, the literature on binary decision diagrams is extensive; therefore, we will focus solely on related results in the context of the model counting problem. The problem of #nFBDD is #P-complete: membership in #P is immediate as every assignment can be evaluated in PTIME, and the #P-hardness follows from the observation that the problem of #DNF, i.e., counting the number of satisfying assignments of Boolean formulas in Disjunctive Normal Form, is #P-hard [22]. Moreover, every DNF can be represented as an nFBDD such that the size of the resulting nFBDD is polynomial in the size of the DNF. Furthermore, it is also known that the problem of #nOBDD is SpanL-complete [6].

Given the #P-hardness, a natural direction of research is to understand the complexity of approximation. The discovery of polynomial-time randomized approximation schemes for #P-hard problems has been of long-standing interest and has led to several theoretically deep and elegant results. One such result was that of Karp and Luby [17] in the context of #DNF, relying on Monte Carlo sampling. Building on Monte Carlo sampling, Kannan, Sampath, and Mahaney [16] proposed a quasi-polynomial running time approximation scheme for #nOBDD.² In a follow-up work [12], this result was extended to handle context-free grammars by reducing the problem of counting words of a context-free grammar to estimating the support size of multilinear $(+, \times)$ -programs. It is straightforward to see that the same reduction can be applied in the context of #DNNF, implying a quasi-polynomial runtime approximation for #nFBDD. Since then, the question of the existence of a fully polynomial-time randomized approximation scheme for #nFBDD and its generalizations has remained open.

In a major breakthrough, Arenas et al. [6] provided an FPRAS for #nOBDD. Their technique relied on union of sets estimation à la Karp-Luby and the generation of independent samples via the self-reducibility union property.³ The self-reducibility union property can be informally stated as follows: The set of models conditioned on a given partial assignment (according to the variable ordering of the given nOBDD) can be expressed as the union of models of the states of the given nOBDD. In a follow-up work [7], Arenas et al. observed that the problem of model counting over structured DNNF (st-DNNF) circuits also admits FPRAS. In this context, it is worth highlighting that the self-reducibility union property does not hold for nFBDD and there exists exponential separation between nFBDD and st-DNNF, i.e., there is a family of functions for which the smallest FBDD are exponentially smaller than the smallest st-DNNF, and therefore, the problem of whether there exists an FPRAS for #nFBDD remains open.

3 Technical Overview

Our algorithm proceeds in a bottom-up manner and for every node q of given nFBDD B , we keep: (1) a number $p(q) \in (0, 1]$ which seeks to approximate $\frac{1}{|\text{mod}(q)|}$, and therefore, $\frac{1}{p(q)}$ can be used to estimate $|\text{mod}(q)|$, and (2) $n_s n_t$ sets of samples $S^1(q), \dots, S^{n_s n_t}(q)$, each a subset of $\text{mod}(q)$, where n_s and n_t are polynomial in $n = |\text{var}(B)|$, ε^{-1} and $\log |B|$. Few

² The result of [16] was stated for regular languages, but the underlying techniques can handle #nOBDD.

³ The term *self-reducibility union property* was coined by Meel, Chakraborty, and Mathur [18] to explain the high-level idea of [6].

comments are in order: we keep many $(n_s \cdot n_t)$ independent sets of samples so as to rely on the median of means estimator. As mentioned earlier, our algorithm works in bottom-up manner, in particular, for a node q , we will compute $(p(q), \{S^r(q)\}_r)$ using the values of $p(\cdot)$ and $\{S^r(\cdot)\}_r$ of its children.

Ideally, we want every model of q to be in $S^r(q)$ identically and independently with probability $p(q)$, and thus that the expected size of $S^r(q)$ is small. However, obtaining iid samples is computationally expensive, which resulted in quasi-polynomial runtimes in earlier studies [12]. Recent works on FPRAS for nOBDD achieved independence by leveraging self-reducibility union property [6, 7], but, as remarked in Section 2.1, the self-reducibility union property does not hold for nFBDD and therefore, it is not known how to accomplish independence among samples without giving up on the desiderata of polynomial time.

The key insight in our approach is to give up the desire for independence altogether, in particular, we do not even ensure pairwise independence, i.e., even for $\alpha, \alpha' \in \text{mod}(q)$, it may not be the case that $\Pr[\alpha \in S^r(q) | \alpha' \in S^r(q)] = \Pr[\alpha \in S^r(q)]$. Of course, we do need to quantify the dependence. In order to discuss the dependence, we first discuss how we update $p(q)$ and $S^r(q)$ for decision nodes and $\vee$ -nodes.

- Let $q = \text{ite}(x, q_1, q_0)$. Then we compute $p(q)$ and $S^r(q)$ as $p(q) = \left(\frac{1}{p(q_0)} + \frac{1}{p(q_1)}\right)^{-1}$ and $S^r(q) = \left(\text{reduce}\left(S^r(q_0), \frac{p(q)}{p(q_0)}\right) \otimes \{x \mapsto 0\}\right) \cup \left(\text{reduce}\left(S^r(q_1), \frac{p(q)}{p(q_1)}\right) \otimes \{x \mapsto 1\}\right)$, where $\text{reduce}(S, p)$ is the operation that keeps each element of a set S with probability p .
- Let q be a $\vee$ -node such that $q = q_1 \vee q_2$ (assuming two children for simplicity). Furthermore, for simplicity of exposition and for the sake of high-level intuition, assume $p(q_1) = p(q_2)$. The technical difficulty for $\vee$ -nodes arises from the fact that it is possible that for a given $\alpha \in \text{mod}(q)$, we have $\alpha \in \text{mod}(q_1) \cap \text{mod}(q_2)$. Therefore, in order to ensure no α has undue advantage even if it lies in the models set of multiple children, we update $p(q)$ and $S^r(q)$ as follows: we first compute $\rho = \min(p(q_1), p(q_2))$ and $\hat{S}^r(q) = S^r(q_1) \cup (S^r(q_2) \setminus \text{mod}(q_1))$ and then $p(q) = \text{median}_{0 \leq j < n_t} \left(\frac{1}{\rho \cdot n_s} \sum_{r=j \cdot n_s + 1}^{(j+1)n_s} |\hat{S}^r(q)|\right)^{-1}$ followed by $S^r(q) = \text{reduce}\left(\hat{S}^r(q), \frac{p(q)}{\rho}\right)$.

Observe that the usage of $S^r(q_2) \setminus \text{mod}(q_1)$ ensures that for every $\alpha \in \text{mod}(q)$, there is exactly one child of $q' \in \text{children}(q)$ such that if $\alpha \in \hat{S}^r(q)$ then $\alpha \in S^r(q')$, and therefore, no α has *undue advantage*.

It is worth re-emphasizing the crucial departure in our work from earlier efforts is embrace of dependence. For instance, consider $q = q_1 \vee q_2$ and $\hat{q} = q_1 \vee q_3$, then $S^r(q)$ and $S^r(\hat{q})$ will of course be reusing samples $S^r(q_1)$, and therefore, do not even have pairwise independence. Now, of course, we need to bound dependence so as to retain any hope of computing $p(q)$ from $S^r(q)$. To this end, we focus on the following quantity of interest: $\Pr[\alpha \in S^r(q) | \alpha' \in S^r(q)]$ for $\alpha, \alpha' \in \text{mod}(q)$, which determines the variance for the estimator. In this regard, for every (α, q) , we can define a derivation path, denoted as $\text{path}(\alpha, q)$, where for every $\vee$ -node, $\text{path}(\alpha, q)$ is $\text{path}(\alpha, q')$ appended with q , where q' is the first child of q such that $\alpha \in \text{mod}(q')$. The key observation is that our computations ensure $\Pr[\alpha \in S^r(q) | \alpha' \in S^r(q)]$ depends on the first node (starting from the 1-sink) q^* where the derivation paths $\text{path}(\alpha, q)$ and $\text{path}(\alpha, q')$ diverge. In particular, it turns out:

$$\Pr[\alpha \in S^r(q) | \alpha' \in S^r(q)] \leq \frac{p(q)}{p(q^*)}.$$

One might wonder whether the above expression suffices: it turns out it does, because the number of pairs (α, α') whose derivation paths diverge for the first time at q^* can be shown to be bounded by $\frac{|\text{mod}(q)|^2}{|\text{mod}(q^*)|}$, which suffices to show that the variance of the estimator can be bounded by constant factor of square of its mean, thus allowing us to use median of means estimator.

We close off by remarking that the situation is more nuanced than previously described, as $p(q)$ is itself a random variable. Although the high-level intuition remains consistent, the technical analysis requires coupling, based on a carefully defined random process, detailed in Section 6. Simplifying the rigorous technical argument would be an interesting direction of future research.

4 Algorithm

The core of our FPRAS, `approxMCnFBDD_core`, takes in a 1-complete, 0-reduced and alternating nFBDD B . B 's nodes are processed in bottom-up order, so from the sinks to the source. For each node q , the algorithm computes $p(q)$ which seeks to estimate $|\text{mod}(q)|^{-1}$, and polynomially-many subsets of $\text{mod}(q)$ called *sample sets*: $S^1(q), \dots, S^N(q)$. The algorithm stops after processing the source node q_{source} and returns $p(q_{\text{source}})^{-1}$. The procedure that computes the content of $S^r(q)$ and the value for $p(q)$ is `estimateAndSample(q)`. Since this procedure uses randomness, the $\{S^r(q)\}_r$ and $p(q)$ are random variables. Our FPRAS works towards ensuring that “ $\Pr[\alpha \in S^r(q)] = p(q)$ ” holds true for every $q \in B$ and $\alpha \in \text{mod}(q)$. Thus, if $p(q)$ is a good estimate of $|\text{mod}(q)|^{-1}$, then $S^r(q)$ should be small in expectation. We put the equality between quotes because it does not make much sense as the left-hand side is a probability, so a fixed number, whereas the right-hand side is a random variable.

■ **Algorithm 1** `estimateAndSample(q)` with $q = \text{ite}(x, q_1, q_0)$.

```

1  $p(q) = \left( \frac{1}{p(q_0)} + \frac{1}{p(q_1)} \right)^{-1}$ 
2 for  $1 \leq r \leq n_s n_t$  do
3    $S^r(q) = \left( \text{reduce} \left( S^r(q_0), \frac{p(q)}{p(q_0)} \right) \otimes \{x \mapsto 0\} \right) \cup \left( \text{reduce} \left( S^r(q_1), \frac{p(q)}{p(q_1)} \right) \otimes \{x \mapsto 1\} \right)$ 

```

Decision nodes and $\vee$ -nodes are handled differently by `estimateAndSample`. When q is a decision node $\text{ite}(x, q_1, q_0)$, $p(q)$ is computed deterministically from $p(q_0)$ and $p(q_1)$, and $S^r(q)$ is *reduced from* $(S^r(q_0) \otimes \{x \mapsto 0\}) \cup (S^r(q_1) \otimes \{x \mapsto 1\})$ using the `reduce` procedure.

■ **Algorithm 2** `reduce(S, p)` with $p \in [0, 1]$.

```

1  $S' \leftarrow \emptyset$ 
2 for  $s \in S$  do
3    $\lfloor$  add  $s$  to  $S'$  with probability  $p$ 
4 return  $S'$ 

```

`estimateAndSample(q)` is more complicated when q is a $\vee$ -node. We explain it gradually. For starter, let $q = q_1 \vee \dots \vee q_k$ with its children ordered as follows: $(q_1, \dots, q_k)$, and consider the problem of approximating $|\text{mod}(q)|$ when a sample set $S(q_i) \subseteq \text{mod}(q_i)$ is available for every $i \in [k]$ and $b \in \{0, 1\}$ with the guarantee that $\Pr[\alpha \in S(q_i)] = \rho$ holds for every $\alpha \in \text{mod}(q_i)$. Every $S(q_i)$ is a subset of $\text{mod}(q_i)$. We compute $\hat{S}(q) = \text{union}(q, S(q_1), \dots, S(q_k))$ as follows: for every $\alpha \in S(q_i)$, α is added to $\hat{S}(q)$ if and only if q_i is the first child of q for which α is a model. Simple computations show that $\Pr[\alpha \in \hat{S}(q)] = \rho$ holds for every $\alpha \in \text{mod}(q)$, and therefore $\rho^{-1}|\hat{S}(q)|$ is an unbiased estimate of $|\text{mod}(q)|$ (i.e., the expected value of the estimate is $|\text{mod}(q)|$).

■ **Algorithm 3** $\text{union}(q, S_1, \dots, S_k)$ with $\text{children}(q) = (q_1, \dots, q_k)$.

```

1  $S' = \emptyset$ 
2 for  $1 \leq i \leq k$  do
3   for  $\alpha \in S_i$  do
4     if  $\alpha \notin \text{mod}(q_j)$  for every  $j < i$  then add  $\alpha$  to  $S'$ 
5 return  $S'$ 

```

Now suppose that for every i , we only know that $\Pr[\alpha \in S(q_i)] = p(q_i)$ for every $\alpha \in \text{mod}(q_i)$ for some number $p(q_i)$ independent of α . Then we normalize the probabilities before computing the union. This is done using the `reduce` procedure. Let $\rho = \min(p(q_1), \dots, p(q_k))$ and $\bar{S}(q_i) = \text{reduce}(S(q_i), \frac{\rho}{p(q_i)})$. We have that $\Pr[\alpha \in \bar{S}(q_i)] = \rho$ holds for every $\alpha \in \text{mod}(q)$. So $\hat{S}(q) = \text{union}(q, \bar{S}(q_1), \dots, \bar{S}(q_k))$, and $\rho^{-1}|\hat{S}(q)|$ is an unbiased estimate of $|\text{mod}(q)|$.

■ **Algorithm 4** $\text{estimateAndSample}(q)$ with $q = q_1 \vee \dots \vee q_k$ and $\text{children}(q) = (q_1, \dots, q_k)$.

```

1  $\rho = \min(p(q_1), \dots, p(q_k))$ 
2 for  $1 \leq r \leq n_s n_t$  do
3    $\hat{S}^r(q) = \text{union}\left(q, \text{reduce}(S^r(q_1), \frac{\rho}{p(q_1)}), \dots, \text{reduce}(S^r(q_k), \frac{\rho}{p(q_k)})\right)$ 
4 for  $0 \leq j < n_t$  do
5    $M_j = \frac{1}{\rho \cdot n_s} \sum_{r=j \cdot n_s + 1}^{(j+1)n_s} |\hat{S}^r(q)|$ 
6  $\hat{\rho} = \text{median}_{0 \leq j < n_t}(M_j)^{-1}$ 
7  $p(q) = \min(\rho, \hat{\rho})$ 
8 for  $1 \leq r \leq n_s n_t$  do
9    $S^r(q) = \text{reduce}(\hat{S}^r(q), \frac{p(q)}{\rho})$ 

```

To find an estimate that is concentrated around $|\text{mod}(q)|$, we use the “median of means” technique. Suppose that instead of one sample set $S(q_i)$ we have several sample sets $S^1(q_i), S^2(q_i), \dots, S^N(q_i)$ all verifying $\Pr[\alpha \in S^r(q_i)] = p(q_i)$. Then define $\bar{S}^r(q_i)$ similarly to $\bar{S}(q_i)$ and $\hat{S}^r(q)$ similarly to $\hat{S}(q)$. Each $\rho^{-1}|\hat{S}^r(q)|$ is an estimate of $|\text{mod}(q)|$. Say $N = n_s n_t$ and partition $\hat{S}^1(q), \hat{S}^2(q), \dots, \hat{S}^N(q)$ into n_t batches of n_s sets. The median of means technique computes the average size $M_j = \frac{1}{n_s} \sum_{r=j \cdot n_s + 1}^{(j+1)n_s} |\hat{S}^r(q)|$ over each batch and uses $\rho^{-1} \text{median}(M_1, \dots, M_{n_t})$ to estimate $|\text{mod}(q)|$. The mean computation aims at reducing the variance of the estimate. The parameter n_s can be chosen so that $\Pr[\rho^{-1} M_j \in (1 \pm \varepsilon)|\text{mod}(q)|] > \frac{1}{2}$ holds true. With the appropriate value for n_t , $\rho^{-1} \text{median}(M_1, \dots, M_{n_t})$ lies in $(1 \pm \varepsilon)|\text{mod}(q)|$ with high probability (though the estimate is not unbiased anymore).

So this is how `estimateAndSample` works for $\vee$ -nodes. The median of means serves to compute $\hat{\rho}$, the inverse of the estimate of $|\text{mod}(q)|$ which in turn is used to compute $p(q)$. When q is not the source node q_{source} , we compute sample sets $S^r(q)$ in preparation for processing of q 's ancestors. For this we reuse $\hat{S}^r(q)$ and compute $S^r(q) = \text{reduce}(\hat{S}^r(q), \frac{p(q)}{\rho})$.

Algorithm 5 $\text{approxMCnFBDD_core}(B, n, n_s, n_t, \theta)$.

```

1  $p(1\text{-sink}) = 1, p(0\text{-sink}) = \infty$ 
2 for  $1 \leq r \leq n_s n_t$  do
3    $S^r(1\text{-sink}) = \{\alpha_\emptyset\}, S^r(0\text{-sink}) = \emptyset$ 
4 for  $1 \leq i \leq 2n$  do
5   for  $q \in L_i$  do
6      $\text{estimateAndSample}(q)$ 
7     if  $|S^r(q)| \geq \theta$  then return 0
8 return  $\frac{1}{p(q_{\text{source}})}$ 

```

To ensure a polynomial running time, $\text{approxMCnFBDD_core}$ terminates as soon as the number of samples grows too large (Line 7) and returns 0. This output is erroneous but we will show that the probability of terminating this way is negligible. For parameters carefully chosen, $\text{approxMCnFBDD_core}$ returns a good estimate of $|B^{-1}(1)|$ with probability larger than $1/2$. The full FPRAS approxMCnFBDD amplifies this probability to $1 - \delta$ by returning the median output of independent runs of approxMCnFBDD .

Algorithm 6 $\text{approxMCnFBDD}(B, \varepsilon, \delta)$.

```

1 make  $B$  alternating, 1-complete and 0-reduced
2  $n = |\text{var}(B)|, m = \lceil 8 \ln(1/\delta) \rceil$ 
3  $\kappa = \varepsilon/(1 + \varepsilon), n_s = \lceil 4n/\kappa^2 \rceil, n_t = \lceil 8 \ln(16|B|) \rceil, \theta = 16n_s n_t (1 + \kappa)|B|$ 
4 for  $1 \leq j \leq m$  do
5    $\text{est}_j = \text{approxMCnFBDD\_core}(B, n, n_s, n_t, \theta)$ 
6 return  $\text{median}(\text{est}_1, \dots, \text{est}_m)$ 

```

5 Derivation paths

Models can have several accepting paths in an nFBDD. For q a node of B and $\alpha \in \text{mod}(q)$, we map (α, q) to a canonical accepting path, called the *derivation path* of α for q , denoted by $\text{path}(\alpha, q)$. A path $\mathcal{P}$ is formally represented with a tuple $(V(\mathcal{P}), E(\mathcal{P}))$, with $V(\mathcal{P})$ a sequence of vertices and $E(\mathcal{P})$ a sequence of edges.

► **Definition 3.** For $q \in B$ and $\alpha \in \text{mod}(q)$, the derivation path $\text{path}(\alpha, q)$ is defined as follows:

- If q is the 1-sink then $\alpha = \alpha_\emptyset$ and the only derivation path is $\text{path}(\alpha_\emptyset, q) = (\{q\}, \emptyset)$.
- If $q = \text{ite}(x, q_1, q_0)$, let α' be the restriction of α to $\text{var}(\alpha) \setminus \{x\}$, then $V(\text{path}(\alpha, q)) = V(\text{path}(\alpha', q_{\alpha(x)})) \cdot q$ and $E(\text{path}(\alpha, q))$ is $E(\text{path}(\alpha', q_{\alpha(x)}))$ plus the $\alpha(x)$ -edge of q .
- If $q = q_1 \vee \dots \vee q_k$ with the children ordering $\text{children}(q) = (q_1, \dots, q_k)$, let i be the smallest integer between 1 and k such that $\alpha \in \text{mod}(q_i)$ then $V(\text{path}(\alpha, q)) = V(\text{path}(\alpha, q_i)) \cdot q$ and $E(\text{path}(\alpha, q))$ is $E(\text{path}(\alpha, q_i))$ plus the edge between q_i and q .

Our algorithm constructs sample sets in a way that respect derivation paths. That is, an assignment $\alpha \in \text{mod}(q)$ may end up in $S^r(q)$ only if it is derived through $\text{path}(\alpha, q)$.

► **Lemma 4.** *Let $q \in B$ and $\alpha \in \text{mod}(q)$, let $V(\text{path}(\alpha, q)) = (q_0, q_1, \dots, q_{i-1}, q_i)$ with q_0 the 1-sink and $q_i = q$. For every $j \in [0, i-1]$, let α_j be the restriction of α to $\text{var}(q_j)$. In a run of `approxMCnFBDD_core`, $\alpha \in S^r(q)$ holds only if $\alpha_j \in S^r(q_j)$ holds for every $j \in [0, i-1]$.*

Proof. $\alpha_0 = \alpha_\emptyset \in S^r(1\text{-sink}) = S^r(q_0)$ holds by construction. Now consider $j > 0$, it is sufficient to show that $\alpha_j \in S^r(q_j)$ only if $\alpha_{j-1} \in S^r(q_{j-1})$.

- If $q_j = \text{ite}(x, q_{j,1}, q_{j,0})$ then $\alpha_j = \alpha_{j-1} \cup \{x \mapsto \alpha(x)\}$ and $q_{j-1} = q_{j,\alpha(x)}$. Looking at `estimateAndSample` for decision nodes, one sees that $\alpha_j \in S^r(q_j)$ only if $\alpha_{j-1} \in \text{reduce}(S^r(q_{j,\alpha(x)}), p(q_j)/p(q_{j-1}))$, so only if $\alpha_{j-1} \in S^r(q_{j,\alpha(x)}) = S^r(q_{j-1})$.
- If q_j is a $\vee$ -node with $\text{children}(q_j) = (q_j^0, \dots, q_j^k)$ then $\alpha_j = \alpha_{j-1}$ and there is an i such that $q_{j-1} = q_j^i$. Observe that $\alpha_j \in S^r(q_j)$ only if $\alpha_j \in S^r(q_j^i)$ – so only if $\alpha_j \in \text{reduce}(S^r(q_j^i), \rho/p(q_j^i))$ – for the smallest ℓ such that $\alpha_j \in \text{mod}(q_j^\ell)$; so only if $\alpha_j = \alpha_{j-1} \in S^r(q_j^\ell)$. The definition of $\text{path}(\alpha, q)$ implies that $\ell = i$. ◀

Given two derivation paths $\mathcal{P}$ and $\mathcal{P}'$. We call their *last common prefix nodes* denoted by $\text{lcpn}(\mathcal{P}, \mathcal{P}')$, the deepest node where the two paths diverge, that is, the first node contained in both paths from which they follow different edges. Note that if $\mathcal{P}$ and $\mathcal{P}'$ are consistent up to node q' , and $q = \text{ite}(x, q', q')$, and $\mathcal{P}$ follows the 0-edge while $\mathcal{P}'$ follows the 1-edge, then the two paths diverge at q' even though they both contain q .

► **Definition 5.** *Let $\mathcal{P} = ((q_0, \dots, q_k), (e_1, \dots, e_k))$ and $\mathcal{P}' = ((q'_0, \dots, q'_\ell), (e'_1, \dots, e'_\ell))$ be two derivation paths. The last common prefix node, denoted by $\text{lcpn}(\mathcal{P}, \mathcal{P}')$, is the node q_i for the biggest i such that $(q_0, \dots, q_i) = (q'_0, \dots, q'_i)$ and $(e_1, \dots, e_i) = (e'_1, \dots, e'_i)$.*

Note that every derivation path contains the 1-sink for first node, so the last common prefix node is well-defined. Let $V(\text{path}(\alpha, q)) = (q_0, \dots, q_i)$ with q_0 the 1-sink and $q_i = q$. For every $0 \leq \ell \leq i$, we define

$$I(\alpha, q, \ell) := \{\alpha' \in \text{mod}(q) \mid \text{lcpn}(\text{path}(\alpha, q), \text{path}(\alpha', q)) = q_\ell\}.$$

The following result will play a key role in bounding the variance of estimators in the analysis.

► **Lemma 6.** *Let $\alpha \in \text{mod}(q)$ and $V(\text{path}(\alpha, q)) = (q_0, \dots, q_i)$ with q_0 the 1-sink and $q_i = q$. For every $0 \leq \ell \leq i$, $|I(\alpha, q, \ell)| \leq \frac{|\text{mod}(q)|}{|\text{mod}(q_\ell)|}$.*

Proof. Let $I(\alpha, q, \ell) = \{\alpha^1, \alpha^2, \dots\}$. Let α_ℓ be the restriction of α to $\text{var}(q_\ell)$. By definition, every α^i is of the form $\alpha_\ell \cup \beta^i$ for some assignment β^i to $\text{var}(q) \setminus \text{var}(q_\ell)$. Consider $\alpha'_\ell \in \text{mod}(q_\ell)$, then every $\alpha'_\ell \cup \beta^i$ is in $\text{mod}(q)$. Since the α^i s differ on $\text{var}(q) \setminus \text{var}(q_\ell)$, the β^i s are pairwise distinct, and therefore $\{\alpha'_\ell \cup \beta^i\}_i$ is a set of $|I(\alpha, q, \ell)|$ distinct assignments. Considering all $|\text{mod}(q_\ell)|$ possible choices for α'_ℓ , we find that $\{\alpha'_\ell \cup \beta^i\}_{\alpha'_\ell, i}$ is a set of $|\text{mod}(q_\ell)| \cdot |I(\alpha, q, \ell)|$ distinct assignments in $\text{mod}(q)$. Hence $|\text{mod}(q_\ell)| \cdot |I(\alpha, q, \ell)| \leq |\text{mod}(q)|$. ◀

6 The Framework for the Analysis

We introduce a random process that simulates `approxMCnFBDD_core`. Our intuition is that, for every $\alpha \in \text{mod}(q)$, a statement in the veins of “ $\Pr[\alpha \in S^r(q)] = p(q)$ ” should hold. The problem is that this equality makes no sense because $\Pr[\alpha \in S^r(q)]$ is a fixed real value whereas $p(q)$ is a random variable. The variables $S^r(q)$ and $p(q)$ for different q are too dependent of each other so we use a random process to work with new variables that behave more nicely. The random process simulates several runs of the algorithm for all possible values of the $p(q)$ s. There, $S^r(q)$ is simulated by a different variable for each possible run. A coupling argument then allows us to replace $S^r(q)$ by one of these variables assuming enough knowledge on the algorithm run up to q , encoded in what we call a *history* for q .

6.1 History

A *history* h for a set of nodes Q is a mapping $h : Q \rightarrow \mathbb{Q} \cup \{\infty\}$. h is *realizable* when there exists a run of `approxMCnFBDD_core*` that gives the value $h(q)$ to $p(q)$ for every $q \in Q$. Such a run is said *compatible* with h . Two histories h for Q and h' for Q' are *compatible* when $h(q) = h'(q)$ for all $q \in Q \cap Q'$. Compatible histories can be merged into an history $h \cup h'$ for $Q \cup Q'$. For $q \in Q$ and $t \in \mathbb{Q}$, we write $h \cup (q \mapsto t)$ to refer to the history h augmented with $h(q) = t$. For $q \in B$, we define the set $\text{desc}(q)$ of its descendants by $\text{desc}(1\text{-sink}) = \emptyset$ and $\text{desc}(q) = \text{children}(q) \cup \bigcup_{q' \in \text{children}(q)} \text{desc}(q')$. Note that $q \notin \text{desc}(q)$. We only study histories realizable for sets Q that are closed for desc , that is, if $q \in Q$ and q' is a descendant of q , then $q' \in Q$. Thus we abuse terminology and refer to a history for $\text{desc}(q)$ as a history for q . The only history for the sinks is the vacuous history $h_\emptyset$ for $Q = \emptyset$ (because no descendants).

6.2 Random Process

The random process comprises $n_s n_t$ independent copies identified by the superscript r . For $q \in B$, $t \in \mathbb{Q} \cup \{\infty\}$ and h a realizable history for q , we have a random variable $\mathfrak{S}_{h,t}^r(q)$ whose domain is all possible subsets of $\text{mod}(q)$. $\mathfrak{S}_{h,t}^r(q)$ simulates $S^r(q)$ in runs of compatible with h and where the value t is assigned to $p(q)$.

- If q is the 0-sink, then $\text{mod}(q) = \emptyset$ and only $\mathfrak{S}_{h_\emptyset, \infty}^r(q) = \emptyset$ is defined.
- If q is the 1-sink, then $\text{mod}(q) = \{\alpha_\emptyset\}$ and only $\mathfrak{S}_{h_\emptyset, 1}^r(q) = \{\alpha_\emptyset\}$ is defined.
- If $q = \text{ite}(x, q_1, q_0)$, then for every $\mathfrak{S}_{h_0, t_0}^r(q_0)$, $\mathfrak{S}_{h_1, t_1}^r(q_1)$ with h_1 and h_0 realizable and compatible histories for q_1 and q_0 , let $h = h_1 \cup h_0 \cup (q_0 \mapsto t_0, q_1 \mapsto t_1)$ and $t = \left(\frac{1}{t_0} + \frac{1}{t_1}\right)^{-1}$

$$\mathfrak{S}_{h,t}^r(q) = \text{reduce} \left(\mathfrak{S}_{h_0, t_0}^r(q_0) \otimes \{x \mapsto 0\}, \frac{t}{t_0} \right) \cup \text{reduce} \left(\mathfrak{S}_{h_1, t_1}^r(q_1) \otimes \{x \mapsto 1\}, \frac{t}{t_1} \right).$$

- If $q = q_1 \vee \dots \vee q_k$ then for every $\mathfrak{S}_{h_1, t_1}^r(q_1), \dots, \mathfrak{S}_{h_k, t_k}^r(q_k)$ with realizable and pairwise compatible histories, we define $h = h_1 \cup \dots \cup h_k \cup (q_1 \mapsto t_1, \dots, q_k \mapsto t_k)$, $t_{\min} = \min(t_1, \dots, t_k)$ and the variable $\hat{\mathfrak{S}}_h^r(q)$ that simulates $\hat{S}^r(q)$ when the history for q is h

$$\hat{\mathfrak{S}}_h^r(q) = \text{union} \left(q, \text{reduce} \left(\mathfrak{S}_{h_1, t_1}^r(q_1), \frac{t_{\min}}{t_1} \right), \dots, \text{reduce} \left(\mathfrak{S}_{h_k, t_k}^r(q_k), \frac{t_{\min}}{t_k} \right) \right).$$

For all $t \leq t_{\min}$ we define $\mathfrak{S}_{h,t}^r(q) = \text{reduce} \left(\hat{\mathfrak{S}}_h^r(q), \frac{t}{t_{\min}} \right)$.

We make important observations to motivate the interest of the random process. Let $H(q)$ be the random variable on the history for q obtained when running `approxMCnFBDD_core*`.

► **Fact 7.** *It holds that $(H(q), (\hat{S}^r(q) \mid r \in R)) = (H(q), (\hat{\mathfrak{S}}_{H(q)}^r(q) \mid r \in R))$ and $(H(q), p(q), (S^r(q) \mid r \in R)) = (H(q), p(q), (\mathfrak{S}_{H(q), p(q)}^r(q) \mid r \in R))$*

Fact 7 is explained in more details in the full version of the paper. The equalities should be interpreted as follows: for every $(A^r \subseteq \text{mod}(q) \mid r \in R)$, the following holds:

$$\begin{aligned} \Pr \left[H(q) = h, p(q) = t, \bigcap_{r \in R} S^r(q) = A^r \right] &= \Pr \left[H(q) = h, p(q) = t, \bigcap_{r \in R} \mathfrak{S}_{h,t}^r(q) = A^r \right] \\ \Pr \left[H(q) = h \text{ and } \bigcap_{r \in R} \hat{S}^r(q) = A^r \right] &= \Pr \left[H(q) = h \text{ and } \bigcap_{r \in R} \hat{\mathfrak{S}}_h^r(q) = A^r \right]. \end{aligned}$$

A second key observation is that the variables $\mathfrak{S}_{h,t}^r(q)$ and $\hat{\mathfrak{S}}_h^r(q)$ are independent of $H(q)$, $p(q)$, $S^r(q)$ and $\hat{S}^r(q)$. This is because the latter variables comes from the algorithm while the former are defined within the random process, and the two do not interact in any way.

► **Fact 8.** The variables $\mathfrak{S}_{h,t}^r(q)$ and $\hat{\mathfrak{S}}_h^r(q)$ are independent of any combination of $H(q')$, $p(q')$, $S^r(q')$ and $\hat{S}^r(q')$ for every q' (including $q' = q$).

In the random process we have the correct variant of the equality “ $\Pr[\alpha \in S^r(q)] = p(q)$ ”.

► **Lemma 9.** For every $\mathfrak{S}_{h,t}^r(q)$ and $\alpha \in \text{mod}(q)$, it holds that $\Pr[\alpha \in \mathfrak{S}_{h,t}^r(q)] = t$. In addition, if q is a $\vee$ -node with $\text{children}(q) = (q_1, \dots, q_k)$ then $\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q)] = \min(h(q_1), \dots, h(q_k))$.

► **Lemma 10.** For every $\mathfrak{S}_{h,t}^r(q)$ with $\text{children}(q) = (q_1, \dots, q_k)$ and $\alpha, \alpha' \in \text{mod}(q)$ with $\alpha \neq \alpha'$, let $t^* = h(\text{lcpn}(\text{path}(\alpha, q), \text{path}(\alpha', q)))$, then we have that $\Pr[\alpha \in \mathfrak{S}_{h,t}^r(q) \mid \alpha' \in \mathfrak{S}_{h,t}^r(q)] \leq \frac{t}{t^*}$ and if q is $\vee$ -node then $\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q) \mid \alpha' \in \hat{\mathfrak{S}}_h^r(q)] \leq \frac{\min(h(q_1), \dots, h(q_k))}{t^*}$.

Lemmas 9 and 10 are proved in appendix. Compared to “ $\Pr[\alpha \in S^r(q)] = p(q)$ ”, there is nothing wrong with the lemmas as t , t' and the $h(q_i)$ s are fixed real numbers.

7 Analysis

We now conduct the analysis of `approxMCnFBDD`. The hardest part to analyze is the core algorithm `approxMCnFBDD_core`, for which we will prove the following.

► **Lemma 11.** Let B be a 1-complete 0-reduced and alternating n FBDD over n variables. Let $m = \max_i |L^i|$, $\varepsilon > 0$, and $\kappa = \frac{\varepsilon}{1+\varepsilon}$. If $n_s \geq \frac{4n}{\kappa^2}$, $n_t \geq 8 \ln(16|B|)$ and $\theta = 16n_s n_t (1 + \kappa)|B|$ then `approxMCnFBDD_core`(B, n, n_s, n_t, θ) runs in time $O(n_s n_t \log(n_t) \cdot \theta \cdot |B|^2)$ and returns `est` with the guarantee $\Pr[\text{est} \notin (1 \pm \varepsilon)|B^{-1}(1)|] \leq \frac{1}{4}$.

Our main result is obtained decreasing this $1/4$ down to any $\delta > 0$ with the median technique.

► **Theorem 1.** Let B be an n FBDD over n variables, $\varepsilon > 0$ and $\delta > 0$. Algorithm `approxMCnFBDD`(B, ε, δ) runs in time $O(n^5 \varepsilon^{-4} \log(\delta^{-1})|B|^6 \log|B|)$ and returns `est` with the guarantee that $\Pr[\text{est} \in (1 \pm \varepsilon)|B^{-1}(1)|] \geq 1 - \delta$.

Proof. Let $\text{est}_1, \dots, \text{est}_m$ be the estimates from m independent calls to `approxMCnFBDD_core`. Let X_i be the indicator variable that takes value 1 if and only if $\text{est}_i \notin (1 \pm \varepsilon)|B^{-1}(1)|$, and define $\bar{X} = X_1 + \dots + X_m$. By Lemma 11, $\mathbb{E}[\bar{X}] \leq m/4$. Hoeffding bound gives

$$\Pr\left[\text{median}_{1 \leq i \leq m}(X_i) \notin (1 \pm \varepsilon)|B^{-1}(1)|\right] = \Pr\left[\bar{X} > \frac{m}{2}\right] \leq \Pr\left[\bar{X} - \mathbb{E}[\bar{X}] > \frac{m}{4}\right] \leq e^{-m/8} \leq \delta.$$

The running time is $O(|\log(\frac{1}{\delta})|)$ times that of `approxMCnFBDD_core`(B', n, n_s, n_t) where B' is B after it has been made 1-complete, 0-reduced, and alternating. By Lemma 2, $|B'| = O(n|B|^2)$ and B' is constructed in time $O(n|B|^2)$. So each call to `approxMCnFBDD_core`(B', n, n_s, n_t) takes time $O(n^5|B|^6 \log|B|\varepsilon^{-4})$ ◀

Recall that `approxMCnFBDD_core*` is `approxMCnFBDD_core` without the terminating condition of Line 7, that is, where the sets $S^r(q)$ can grow big. Analyzing `approxMCnFBDD_core*` is enough to prove Lemma 11 without running time requirements. In particular, it is enough to prove Lemmas 12 and 13. For these lemmas the settings described in Lemma 11 are assumed.

► **Lemma 12.** The probability that `approxMCnFBDD_core*`(B, n, n_s, n_t, θ) computes $p(q) \notin (1 \pm \kappa)|\text{mod}(q)|^{-1}$ for some $q \in B$ is at most $1/8$.

► **Lemma 13.** *The probability that $\text{approxMCnFBDD_core}^*(B, n, n_s, n_t, \theta)$ constructs sets $S^r(q)$ such that $|S^r(q)| \geq \theta$ for some $q \in B$ is at most $1/8$.*

Proof of Lemma 11. Let $A^{(*)} = \text{approxMCnFBDD_core}^{(*)}(B, n, n_s, n_t, \theta)$.

$$\begin{aligned}
& \Pr_A [\text{est} \notin (1 \pm \varepsilon)|B^{-1}(1)|] \\
&= \Pr_{A^*} \left[\bigcup_{r,q} |S^r(q)| \geq \theta \right] + \Pr_{A^*} \left[\bigcap_{r,q} |S^r(q)| < \theta \text{ and } \text{est} \notin (1 \pm \varepsilon)|B^{-1}(1)| \right] \\
&\leq \Pr_{A^*} \left[\bigcup_{r,q} |S^r(q)| \geq \theta \right] + \Pr_{A^*} [\text{est} \notin (1 \pm \varepsilon)|B^{-1}(1)|] \\
&\leq \Pr_{A^*} \left[\bigcup_{r,q} |S^r(q)| \geq \theta \right] + \Pr_{A^*} \left[\bigcup_q p(q) \notin \frac{1}{(1 \pm \varepsilon)|\text{mod}(q)|} \right]
\end{aligned}$$

where q ranges over B 's nodes and r ranges in $[n_s n_t]$. The parameter κ has been set so that $p(q) \notin \frac{1 \pm \kappa}{|\text{mod}(q)|}$ implies $p(q) \notin \frac{1}{(1 \pm \varepsilon)|\text{mod}(q)|}$ so, using Lemmas 12 and 13:

$$\Pr_A [\text{est} \notin (1 \pm \varepsilon)|B^{-1}(1)|] \leq \Pr_{A^*} \left[\bigcup_{r,q} |S^r(q)| \geq \theta \right] + \Pr_{A^*} \left[\bigcup_q p(q) \notin \frac{1 \pm \kappa}{|\text{mod}(q)|} \right] \leq \frac{1}{4}$$

The algorithm stops whenever the number of samples grows beyond θ so, for the worst-case running time, the number of samples is less than θ . Each node goes once through `estimateAndSample`. For a decision node, `estimateAndSample` takes time $O(n_s n_t \theta)$. For a $\vee$ -node q , `estimateAndSample` calls `union` $n_s n_t$ times, does a median of means where it computes the median of n_t means of n_s integers, and updates the sample sets. Updating the sample sets takes time $O(n_s n_t \theta)$. Each mean costs $O(n_s)$ and the median requires sorting so the whole thing is done in $O(n_s n_t \log(n_t))$ time. For each sample, the `union` tests whether it is a model of the children of q , model checking is a linear-time operation on nFBDD so the total cost of one union is $O(|\text{children}(q)| \cdot |B| \cdot \theta)$. So the total cost of `estimateAndSample` for all $\vee$ -nodes is at most $O(n_s n_t \log(n_t) \cdot \theta \cdot |B| \cdot \sum_q |\text{children}(q)|) = O(n_s n_t \log(n_t) \cdot \theta \cdot |B|^2)$. ◀

It remains to prove Lemmas 12 and 13.

7.1 Proof of Lemma 12

Let $\Delta(q)$ be the interval $\frac{1 \pm \kappa}{|\text{mod}(q)|}$ and $\nabla(q)$ be the interval $\frac{|\text{mod}(q)|}{1 \pm \kappa}$.

► **Claim 14.** The event $\bigcup_{q \in B} (p(q) \notin \Delta(q))$ occurs if and only if the event $\bigcup_{q \in L_{>0}} (p(q) \notin \Delta(q))$ and for all $q' \in \text{desc}(q)$, $p(q') \in \Delta(q')$ occurs.

Proof. The “if” direction is trivial. For the other direction, suppose that $p(q) \notin \Delta(q)$ holds for some q . Let i be the smallest integer such that there is $q \in L_i$ and $p(q) \notin \Delta(q)$. i cannot be 0 because the only node in L_0 is the 1-sink and $p(1\text{-sink}) = 1 = |\text{mod}(1\text{-sink})|^{-1} \in \Delta(1\text{-sink})$. So $q \in L_{>0}$ and, by minimality of i , we have that $p(q') \in \Delta(q')$ for all $q' \in \text{desc}(q)$. ◀

$$\begin{aligned}
\Pr \left[\bigcup_{q \in B} p(q) \notin \Delta(q) \right] &= \Pr \left[\bigcup_{q \in L_{>0}} p(q) \notin \Delta(q) \text{ and } \forall q' \in \text{desc}(q), p(q') \in \Delta(q') \right] \\
&\leq \sum_{q \in L_{>0}} \underbrace{\Pr [p(q) \notin \Delta(q) \text{ and } \forall q' \in \text{desc}(q), p(q') \in \Delta(q')]}_{P(q)}. \tag{1}
\end{aligned}$$

We bound $P(q)$ from above. If $q = \text{ite}(x, q_1, q_0)$ is a decision node, then by construction, $p(q_0) \in \Delta(q_0)$ and $p(q_1) \in \Delta(q_1)$ implies $p(q)^{-1} \in \frac{|\text{mod}(q_0)| + |\text{mod}(q_1)|}{1 \pm \kappa} = \frac{|\text{mod}(q)|}{1 \pm \kappa} = \nabla(q)$ with probability 1. Thus $P(q) = 0$ for decision nodes and only the case of $\vee$ -nodes remains.

Going to the random process

To bound $P(q)$ when q is a $\vee$ -node, we move the analysis to the random process, whose variables we can analyze using Lemmas 9 and 10. Consider the set $\mathcal{H}_q$ of realizable histories for q and denote by $H(q) = h$ the event that the algorithm sets $p(q')$ to $h(q')$ for all $q' \in \text{desc}(q)$.

$$\begin{aligned} P(q) &= \Pr \left[\bigcup_{h \in \mathcal{H}_q} H(q) = h \text{ and } p(q) \notin \Delta(q) \text{ and for all } q' \in \text{desc}(q), p(q') \in \Delta(q') \right] \\ &\leq \sum_{h \in \mathcal{H}_q} \Pr [H(q) = h \text{ and } p(q) \notin \Delta(q) \text{ and for all } q' \in \text{desc}(q), p(q') \in \Delta(q')] \end{aligned}$$

Let $\mathcal{H}_q^*$ be the subset of $\mathcal{H}_q$ where $h(q') \in \Delta(q')$ holds for every $q' \in \text{desc}(q)$, then

$$\begin{aligned} P(q) &\leq \sum_{h \in \mathcal{H}_q^*} \Pr [H(q) = h \text{ and } p(q) \notin \Delta(q) \text{ and for all } q' \in \text{desc}(q), p(q') \in \Delta(q')] \\ &= \sum_{h \in \mathcal{H}_q^*} \Pr [H(q) = h \text{ and } p(q) \notin \Delta(q)]. \end{aligned}$$

Let $\rho = \min(p(q_1), \dots, p(q_k))$, $\rho_h = \min(h(q_1), \dots, h(q_k))$ and $M_j = \frac{1}{\rho \cdot n_s} \sum_{r=j \cdot n_s + 1}^{(j+1)n_s} |\hat{S}^r(q)|$. Note that ρ is a random variable whereas ρ_h is a constant. If $H(q) = h$ then the events $\rho = \rho_h$ and $p(q) = \min(\rho_h, \hat{\rho})$ both hold, where $\hat{\rho} = \text{median}(M_0, \dots, M_{n_t-1})^{-1}$.

▷ **Claim 15.** If $H(q) = h$ then $\hat{\rho} \in \Delta(q)$ implies that $p(q) \in \Delta(q)$.

Proof. $H(q) = h$ implies $p(q) = \min(\rho_h, \hat{\rho})$. If $p(q) = \hat{\rho}$ holds then, trivially, $\hat{\rho} \in \Delta(q)$ implies that $p(q) \in \Delta(q)$. If $p(q) = \rho_h$ occurs then $\hat{\rho} \in \Delta(q)$ implies that $p(q) \leq \hat{\rho} \leq \frac{1+\kappa}{|\text{mod}(q)|}$ and, since $h \in \mathcal{H}_q^*$ guarantees that $\rho_h \geq \frac{1-\kappa}{\max_{j \in [k]} |\text{mod}(q_j)|} \geq \frac{1-\kappa}{|\text{mod}(q)|}$, we have that $p(q) \in \Delta(q)$. $\triangleleft$

Let $\mathfrak{M}_{j,h} = \frac{1}{\rho_h n_s} \sum_{r=j \cdot n_s + 1}^{(j+1)n_s} |\hat{\mathfrak{S}}_h^r(q)|$. Using that $H(q) = h$ implies $\rho = \rho_h$ and Fact 7, we find that $(H(q), \text{median}_j(M_j)) = (H(q), \text{median}_j(\mathfrak{M}_{j,H(q)}))$.

$$\begin{aligned} \Pr [H(q) = h \text{ and } p(q) \notin \Delta(q)] &\leq \Pr [H(q) = h \text{ and } \hat{\rho} \notin \Delta(q)] && \text{(Claim 15)} \\ &= \Pr [H(q) = h \text{ and } \text{median}_{0 \leq j < n_t}(M_j) \notin \nabla(q)] && \text{(by definition of } \hat{\rho}) \\ &\leq \Pr [H(q) = h \text{ and } \text{median}_{0 \leq j < n_t}(\mathfrak{M}_{j,h}) \notin \nabla(q)] && \text{(Fact 7)} \\ &= \Pr [H(q) = h] \Pr [\text{median}_{0 \leq j < n_t}(\mathfrak{M}_{j,h}) \notin \nabla(q)] && \text{(Fact 8)} \end{aligned}$$

We have reached our goal to replace variables by their counterpart in the random process. Now we bound $\Pr [\text{median}_j(\mathfrak{M}_{j,h}) \notin \nabla(q)]$ using Chebyshev's inequality and Hoeffding bound.

Variance upper bound

By Lemma 9, the expected value of $|\hat{\mathfrak{S}}_h^r(q)|$ is $\mu = \rho_h |\text{mod}(q)|$. Now for the variance,

$$\begin{aligned} \text{Var} \left[|\hat{\mathfrak{S}}_h^r(q)| \right] &\leq \mathbb{E} \left[|\hat{\mathfrak{S}}_h^r(q)|^2 \right] = \mu + \sum_{\substack{\alpha, \alpha' \in \text{mod}(q) \\ \alpha \neq \alpha'}} \Pr \left[\alpha \in \hat{\mathfrak{S}}_h^r(q) \text{ and } \alpha' \in \hat{\mathfrak{S}}_h^r(q) \right] \\ &= \mu + \sum_{\substack{\alpha, \alpha' \in \text{mod}(q) \\ \alpha \neq \alpha'}} \Pr \left[\alpha \in \hat{\mathfrak{S}}_h^r(q) \mid \alpha' \in \hat{\mathfrak{S}}_h^r(q) \right] \Pr \left[\alpha' \in \hat{\mathfrak{S}}_h^r(q) \right] \\ &\leq \mu + \sum_{\substack{\alpha, \alpha' \in \text{mod}(q) \\ \alpha \neq \alpha'}} \frac{\rho_h^2}{h(\text{lcpn}(\text{path}(q, \alpha), \text{path}(q, \alpha')))} \quad (\text{Lemmas 9 and 10}) \end{aligned}$$

Let $\mathcal{P} = \text{path}(\alpha, q)$ and $V(\mathcal{P}) = (q_\alpha^0, q_\alpha^1, q_\alpha^2, \dots, q_\alpha^{i-1}, q)$, with $q_\alpha^0 = 1\text{-sink}$. Let $\mathcal{P}' = \text{path}(\alpha', q)$ for any $\alpha' \in \text{mod}(q)$ distinct from α . Then $\text{lcpn}(\mathcal{P}, \mathcal{P}')$ is one of the q_α^j . Recall that $I(\alpha, q, j)$ is the set of $\alpha' \in \text{mod}(q)$ such that $\text{lcpn}(\mathcal{P}, \mathcal{P}') = q_\alpha^j$.

$$\begin{aligned} \sum_{\substack{\alpha, \alpha' \in \text{mod}(q) \\ \alpha \neq \alpha'}} \frac{\rho_h^2}{h(\text{lcpn}(\text{path}(q, \alpha), \text{path}(q, \alpha')))} &= \sum_{\alpha \in \text{mod}(q)} \sum_{j \in [0, i-1]} |I(\alpha, q, j)| \frac{\rho_h^2}{h(q_\alpha^j)} \\ &\leq \sum_{\alpha \in \text{mod}(q)} \sum_{j \in [0, i-1]} \frac{\rho_h^2 |\text{mod}(q)|}{|\text{mod}(q_\alpha^j)| \cdot h(q_\alpha^j)} \quad (\text{Lemma 6}) \end{aligned}$$

Because $h \in \mathcal{H}_q^*$ and $h(q_\alpha^j)$ is in $\Delta(q_\alpha^j)$ we have $|\text{mod}(q_\alpha^j)| \cdot h(q_\alpha^j) \geq 1 - \kappa$.

$$\sum_{\alpha \in \text{mod}(q)} \sum_{j \in [0, i-1]} \frac{\rho_h^2 |\text{mod}(q)|}{|\text{mod}(q_\alpha^j)| \cdot h(q_\alpha^j)} \leq \frac{\rho_h^2}{1 - \kappa} \sum_{\alpha \in \text{mod}(q)} \sum_{j \in [0, i-1]} |\text{mod}(q)| \leq \frac{\mu^2 n}{1 - \kappa}$$

Putting everything together, we conclude that $\text{Var}[|\hat{\mathfrak{S}}_h^r(q)|] \leq \mu + \frac{\mu^2 n}{1 - \kappa}$.

Median of means

We have that $\mathbb{E}[\mathfrak{M}_{j,h}] = \frac{\mu}{\rho_h} = |\text{mod}(q)|$ and, by independence of the variables $\{\hat{\mathfrak{S}}_h^r(q)\}_r$

$$\text{Var}[\mathfrak{M}_{j,h}] = \sum_{r=j \cdot n_s + 1}^{(j+1)n_s} \frac{\text{Var}[|\hat{\mathfrak{S}}_h^r(q)|]}{\rho_h^2 n_s^2} \leq \frac{1}{\rho_h^2 n_s} \left(\mu + \frac{\mu^2 n}{1 - \kappa} \right) = \frac{1}{n_s} \left(\frac{|\text{mod}(q)|}{\rho_h} + \frac{n |\text{mod}(q)|^2}{1 - \kappa} \right).$$

$\mathfrak{M}_{j,h} \in \frac{|\text{mod}(q)|}{1 \pm \kappa}$ occurs if and only if $\frac{-\kappa |\text{mod}(q)|}{1 + \kappa} \leq \mathfrak{M}_{j,h} - |\text{mod}(q)| \leq \frac{\kappa |\text{mod}(q)|}{1 - \kappa}$, which is subsumed by $|\mathfrak{M}_{j,h} - |\text{mod}(q)|| \leq \frac{\kappa |\text{mod}(q)|}{1 + \kappa}$. So Chebyshev's inequality gives

$$\begin{aligned} \Pr \left[\mathfrak{M}_{j,h} \notin \frac{|\text{mod}(q)|}{1 \pm \kappa} \right] &\leq \Pr \left[|\mathfrak{M}_{j,h} - |\text{mod}(q)|| > \frac{\kappa |\text{mod}(q)|}{1 + \kappa} \right] \leq \frac{(1 + \kappa)^2}{\kappa^2 |\text{mod}(q)|^2} \text{Var}[\mathfrak{M}_{j,h}] \\ &\leq \frac{(1 + \kappa)^2}{\kappa^2 n_s} \left(\frac{1}{|\text{mod}(q)| \rho_h} + \frac{n}{1 - \kappa} \right) \\ &\leq \frac{(1 + \kappa)^2}{\kappa^2 n_s} \left(\frac{1}{1 - \kappa} + \frac{n}{1 - \kappa} \right) \quad (\rho_h \geq \frac{1 - \kappa}{\max_j |\text{mod}(q_j)|} \geq \frac{1 - \kappa}{|\text{mod}(q)|}) \\ &\leq \frac{2n}{\kappa^2 n_s} \leq \frac{1}{4} \quad (\frac{(1 + \kappa)^2}{1 - \kappa} \text{ decreases to 1 and } n_s \geq \frac{4n}{\kappa^2}) \end{aligned}$$

By taking the median, we decrease the $\frac{1}{4}$ upper bound to a much smaller value. Let E_j be the indicator variable taking value 1 if and only if $\mathfrak{M}_{j,h} \notin \frac{|\text{mod}(q)|}{1 \pm \kappa}$ and let $\bar{E} = \sum_{j=0}^{n_t-1} E_j$. We have $\mathbb{E}[\bar{E}] \leq \frac{n_t}{4}$ so Hoeffding bound gives

$$\Pr \left[\text{median}_{0 \leq j < n_t}(\mathfrak{M}_{j,h}) \notin \frac{|\text{mod}(q)|}{1 \pm \kappa} \right] = \Pr \left[\bar{E} > \frac{n_t}{2} \right] \leq \Pr \left[\bar{E} - \mathbb{E}(\bar{E}) \geq \frac{n_t}{4} \right] \leq e^{-n_t/8} \leq \frac{1}{16|B|}$$

where the last inequality comes from $n_t \geq 8 \ln(16|B|)$. Putting everything together we have

$$P(q) \leq \sum_{h \in \mathcal{H}_q^*} \frac{1}{16|B|} \Pr[H(q) = h] \leq \frac{1}{16|B|}.$$

Used in (1), this gives $\Pr \left[\bigcup_{q \in B} p(q) \notin \Delta(q) \right] \leq \frac{1}{16}$, thus finishing the proof of Lemma 12. We have shown a $\frac{1}{16}$ bound instead of a $\frac{1}{8}$ bound in preparation for the proof of Lemma 13.

7.2 Proof of Lemma 13

We first bound $\Pr \left[\bigcup_{r,q} |S^r(q)| \geq \theta \right]$ from above by

$$\underbrace{\Pr \left[\bigcup_{r,q} |S^r(q)| \geq \theta \text{ and } \bigcap_{q' \in B} p(q') \in \frac{1 \pm \kappa}{|\text{mod}(q')|} \right]}_{P_1} + \underbrace{\Pr \left[\bigcup_{q' \in B} p(q') \notin \frac{1 \pm \kappa}{|\text{mod}(q')|} \right]}_{P_2}.$$

We have already a $\frac{1}{16}$ upper bound on P_2 , so we focus on P_1 .

$$\begin{aligned} P_1 &\leq \sum_{r,q} \Pr \left[|S^r(q)| \geq \theta \text{ and } \bigcap_{q' \in B} p(q') \in \Delta(q') \right] && \text{(Union bound)} \\ &\leq \frac{1}{\theta} \cdot \sum_{r,q} \mathbb{E} \left[|S^r(q)| \cdot \prod_{q' \in B} \mathbb{1}(p(q') \in \Delta(q')) \right] && \text{(Markov's inequality)} \\ &\leq \frac{1}{\theta} \cdot \sum_{r,q} \underbrace{\mathbb{E} [|S^r(q)| \cdot \mathbb{1}(p(q) \in \Delta(q))]}_{E(r,q)} \end{aligned}$$

To bound $E(r,q)$ we introduce the history of q and move to the variables of the random process. Recall that $\mathcal{H}_q$ is the set of all realizable histories for q .

$$\begin{aligned} E(r,q) &= \sum_{t \in \Delta(q)} \mathbb{E} [|S^r(q)| \cdot \mathbb{1}(p(q) = t)] = \sum_{h \in \mathcal{H}_q} \sum_{t \in \Delta(q)} \mathbb{E} [|S^r(q)| \cdot \mathbb{1}(p(q) = t \text{ and } H(q) = h)] \\ &= \sum_{h \in \mathcal{H}_q} \sum_{t \in \Delta(q)} \mathbb{E} [|\mathfrak{S}_{h,t}^r(q)| \cdot \mathbb{1}(p(q) = t \text{ and } H(q) = h)] && \text{(Fact 7)} \\ &= \sum_{h \in \mathcal{H}_q} \sum_{t \in \Delta(q)} \mathbb{E} [|\mathfrak{S}_{h,t}^r(q)|] \cdot \Pr[p(q) = t \text{ and } H(q) = h] && \text{(Fact 8)} \\ &= \sum_{h \in \mathcal{H}_q} \sum_{t \in \Delta(q)} t \cdot |\text{mod}(q)| \cdot \Pr[p(q) = t \text{ and } H(q) = h] && \text{(Lemma 9)} \\ &\leq \sum_{h \in \mathcal{H}_q} \sum_{t \in \Delta(q)} (1 + \kappa) \cdot \Pr[p(q) = t \text{ and } H(q) = h] \leq (1 + \kappa) && (t \leq \frac{1+\kappa}{|\text{mod}(q)|}) \end{aligned}$$

It follows that $P_1 \leq \frac{(1+\kappa)n_s n_t |B|}{\theta} \leq \frac{1}{16}$ (because $\theta \geq 16(1 + \kappa)n_s n_t |B|$) and therefore $\Pr \left[\bigcup_{r,q} |S^r(q)| \geq \theta \right] \leq \frac{1}{16} + \frac{1}{16} = \frac{1}{8}$. This finishes the proof of Lemma 13.

8 Conclusion

In this paper, we resolved the open problem of designing an FPRAS for $\#n\text{FBDD}$. Our work also introduces a new technique to quantify dependence, and it would be interesting to extend this technique to other languages that generalize $n\text{FBDD}$. Another promising direction for future work would be to improve the complexity of the proposed FPRAS to enable practical adoption.

References

- 1 Antoine Amarilli, Marcelo Arenas, YooJung Choi, Mikaël Monet, Guy Van den Broeck, and Benjie Wang. A circus of circuits: Connections between decision diagrams, circuits, and automata. *arXiv preprint arXiv:2404.09674*, 2024. doi:10.48550/arXiv.2404.09674.
- 2 Antoine Amarilli and Florent Capelli. Tractable circuits in database theory. *SIGMOD Rec.*, 53(2):6–20, 2024. doi:10.1145/3685980.3685982.
- 3 Antoine Amarilli and Florent Capelli. Tractable circuits in database theory. *ACM SIGMOD Record*, 53(2):6–20, 2024. doi:10.1145/3685980.3685982.
- 4 Antoine Amarilli, Florent Capelli, Mikaël Monet, and Pierre Senellart. Connecting knowledge compilation classes and width parameters. *Theory Comput. Syst.*, 64(5):861–914, 2020. doi:10.1007/S00224-019-09930-2.
- 5 Antoine Amarilli, Timothy van Bremen, and Kuldeep S Meel. Conjunctive queries on probabilistic graphs: The limits of approximability. In *27th International Conference on Database Theory*, 2024.
- 6 Marcelo Arenas, Luis Alberto Croquevielle, Rajesh Jayaram, and Cristian Riveros. Efficient logspace classes for enumeration, counting, and uniform generation. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2019*, pages 59–73. ACM, 2019. doi:10.1145/3294052.3319704.
- 7 Marcelo Arenas, Luis Alberto Croquevielle, Rajesh Jayaram, and Cristian Riveros. When is approximate counting for conjunctive queries tractable? In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1015–1027. ACM, 2021. doi:10.1145/3406325.3451014.
- 8 Paul Beame and Vincent Liew. New limits for knowledge compilation and applications to exact model counting. *arXiv preprint arXiv:1506.02639*, 2015. arXiv:1506.02639.
- 9 Randal E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. Computers*, 35(8):677–691, 1986. doi:10.1109/TC.1986.1676819.
- 10 Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *J. Artif. Intell. Res.*, 17:229–264, 2002. doi:10.1613/JAIR.989.
- 11 Daniel Deutch, Nave Frost, Benny Kimelfeld, and Mikaël Monet. Computing the shapley value of facts in query answering. In *SIGMOD '22: International Conference on Management of Data*, pages 1570–1583. ACM, 2022. doi:10.1145/3514221.3517912.
- 12 Vivek Gore, Mark Jerrum, Sampath Kannan, Z. Sweedyk, and Stephen R. Mahaney. A quasi-polynomial-time algorithm for sampling words from a context-free language. *Inf. Comput.*, 134(1):59–74, 1997. doi:10.1006/INCO.1997.2621.
- 13 Abhay Kumar Jha and Dan Suciu. On the tractability of query compilation and bounded treewidth. In *15th International Conference on Database Theory, ICDT '12*, pages 249–261. ACM, 2012. doi:10.1145/2274576.2274603.
- 14 Abhay Kumar Jha and Dan Suciu. Knowledge compilation meets database theory: Compiling queries to decision diagrams. *Theory Comput. Syst.*, 52(3):403–440, 2013. doi:10.1007/S00224-012-9392-5.
- 15 Sheldon B. Akers Jr. Binary decision diagrams. *IEEE Trans. Computers*, 27(6):509–516, 1978. doi:10.1109/TC.1978.1675141.

- 16 Sampath Kannan, Z. Sweedyk, and Stephen R. Mahaney. Counting and random generation of strings in regular languages. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 551–557. ACM/SIAM, 1995. URL: <http://dl.acm.org/citation.cfm?id=313651.313803>.
- 17 Richard M. Karp and Michael Luby. Monte-carlo algorithms for enumeration and reliability problems. In *24th Annual Symposium on Foundations of Computer Science, Tucson, Arizona, USA, 7-9 November 1983*, pages 56–64. IEEE Computer Society, 1983. doi:10.1109/SFCS.1983.35.
- 18 Kuldeep S. Meel, Sourav Chakraborty, and Umang Mathur. A faster FPRAS for #nfa. *Proc. ACM Manag. Data*, 2(2):112, 2024. doi:10.1145/3651613.
- 19 Stefan Mengel. *Counting, Knowledge Compilation and Applications*. PhD thesis, Université d’Artois, 2021.
- 20 Mikaël Monet. Solving a special case of the intensional vs extensional conjecture in probabilistic databases. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2020*, pages 149–163. ACM, 2020. doi:10.1145/3375395.3387642.
- 21 Mikaël Monet and Dan Olteanu. Towards deterministic decomposable circuits for safe queries. In *Proceedings of the 12th Alberto Mendelzon International Workshop on Foundations of Data Management*, volume 2100 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2018. URL: <https://ceur-ws.org/Vol-2100/paper19.pdf>.
- 22 Leslie G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comput.*, 8(3):410–421, 1979. doi:10.1137/0208032.
- 23 Ingo Wegener. *Branching Programs and Binary Decision Diagrams*. SIAM, 2000. URL: <http://ls2-www.cs.uni-dortmund.de/monographs/bdd/>.

A Appendix

► **Lemma 9.** *For every $\mathfrak{S}_{h,t}^r(q)$ and $\alpha \in \text{mod}(q)$, it holds that $\Pr[\alpha \in \mathfrak{S}_{h,t}^r(q)] = t$. In addition, if q is a $\vee$ -node with $\text{children}(q) = (q_1, \dots, q_k)$ then $\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q)] = \min(h(q_1), \dots, h(q_k))$.*

Proof. Let $q \in L_i$. We proceed by induction on i . The base case $i = 0$ is immediate since $\mathfrak{S}_{h_0,1}^r(1\text{-sink}) = \{\alpha_\emptyset\}$ and $\mathfrak{S}_{h_0,\emptyset}^r(0\text{-sink}) = \emptyset$ are the only variables for L_0 (and $\frac{1}{\infty} = 0$ by definition). Now let $i > 0$, $q \in L_i$, and suppose that the statement holds for all $\mathfrak{S}_{h',t'}^r(q')$ and $\alpha' \in \text{mod}(q')$ with $q' \in L_{<i}$. If i is odd then q is a decision node $\text{ite}(x, q_1, q_0)$ with q_0 and q_1 in L_{i-1} (because B is alternating). Let $b = \alpha(x)$ and let α' be the restriction of α to $\text{var}(\alpha) \setminus \{x\}$. Then, by induction,

$$\begin{aligned} \Pr[\alpha \in \mathfrak{S}_{h,t}^r(q)] &= \Pr\left[\alpha \in \text{reduce}\left(\mathfrak{S}_{h_b,t_b}^r(q_b), \frac{t}{t_b}\right) \otimes \{x \mapsto b\}\right] \\ &= \Pr[\alpha' \in \mathfrak{S}_{h_b,t_b}^r(q_b)] \Pr\left[\alpha' \in \text{reduce}\left(\mathfrak{S}_{h_b,t_b}^r(q_b), \frac{t}{t_b}\right) \mid \alpha' \in \mathfrak{S}_{h_b,t_b}^r(q_b)\right] = t_b \cdot \frac{t}{t_b} = t \end{aligned}$$

Now if i is even then q is a $\vee$ -node with children $\text{children}(q) = (q_1, \dots, q_k)$ all in L_{i-1} . Let j be the smallest integer such that $\alpha \in \text{mod}(q_j)$ and let $t_{\min} = \min(h(q_1), \dots, h(q_k))$. Then

$$\begin{aligned} \Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q)] &= \Pr\left[\alpha \in \text{reduce}\left(\mathfrak{S}_{h_j,t_j}^r(q_j), \frac{t_{\min}}{t_j}\right)\right] = \Pr[\alpha \in \mathfrak{S}_{h_j,t_j}^r(q_j)] \\ &\quad \cdot \Pr\left[\alpha \in \text{reduce}\left(\mathfrak{S}_{h_j,t_j}^r(q_j), \frac{t_{\min}}{t_j}\right) \mid \alpha \in \mathfrak{S}_{h_j,t_j}^r(q_j)\right] = t_j \cdot \frac{t_{\min}}{t_j} = t_{\min} \end{aligned}$$

And for $t \leq t_{\min}$ we have that $\Pr[\alpha \in \mathfrak{S}_{h,t}^r(q)] = \Pr[\alpha \in \text{reduce}(\hat{\mathfrak{S}}_h^r(q), \frac{t}{t_{\min}})]$

$$= \Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q)] \Pr\left[\alpha \in \text{reduce}\left(\hat{\mathfrak{S}}_h^r(q), \frac{t}{t_{\min}}\right) \mid \alpha \in \hat{\mathfrak{S}}_h^r(q)\right] = t_{\min} \cdot \frac{t}{t_{\min}} = t \quad \blacktriangleleft$$

► **Lemma 10.** *For every $\mathfrak{S}_{h,t}^r(q)$ with $\text{children}(q) = (q_1, \dots, q_k)$ and $\alpha, \alpha' \in \text{mod}(q)$ with $\alpha \neq \alpha'$, let $t^* = h(\text{lcpn}(\text{path}(\alpha, q), \text{path}(\alpha', q)))$, then we have that $\Pr[\alpha \in \mathfrak{S}_{h,t}^r(q) \mid \alpha' \in \mathfrak{S}_{h,t}^r(q)] \leq \frac{t}{t^*}$ and if q is $\vee$ -node then $\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q) \mid \alpha' \in \hat{\mathfrak{S}}_h^r(q)] \leq \frac{\min(h(q_1), \dots, h(q_k))}{t^*}$.*

Proof. We are going to prove a stronger statement, namely, that for every i , for every q and q' (potentially $q = q'$) in L_i and every t and t' such that $t = t'$ when $q = q'$, and every $\alpha \in \text{mod}(q)$ and $\alpha' \in \text{mod}(q')$, and every compatible histories h and h' for q and q' , respectively, we have that

$$\Pr[\alpha \in \mathfrak{S}_{h,t}^r(q) \text{ and } \alpha' \in \mathfrak{S}_{h',t'}^r(q')] \leq \frac{tt'}{t^*}. \quad (2)$$

where $t^* = t$ if $(q, \alpha) = (q', \alpha')$ and $t^* = h(\text{lcpn}(\text{path}(\alpha, q), \text{path}(\alpha', q')))$ otherwise. In addition if i is even, so q and q' are $\vee$ -nodes, then

$$\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q) \text{ and } \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')] \leq \frac{t_{\min} t'_{\min}}{t^*}. \quad (3)$$

where $t_{\min} = \min(h(c) \mid c \in \text{children}(q))$ and $t'_{\min} = \min(h'(c) \mid c \in \text{children}(q'))$ and $t^* = t_{\min}$ if $(q, \alpha) = (q', \alpha')$ and $t^* = h(\text{lcpn}(\text{path}(\alpha, q), \text{path}(\alpha', q')))$ otherwise.

The inequalities (2) and (3) are straightforward when $(\alpha, q) = (\alpha', q')$ because then $h = h'$ (by compatibility) and $t = t' = t^*$ or $t_{\min} = t'_{\min} = t^*$ and we can use Lemma 9. In particular (2) holds when $q = q' = 1$ -sink. Now we assume $(\alpha, q) \neq (\alpha', q')$ and proceed by induction on i . The base case $i = 0$ holds true by the previous remark (note that neither q nor q' can be the 0-sink because $\text{mod}(q)$ and $\text{mod}(q')$ must not be empty).

Case i odd. In this case q and q' are decision nodes. Let $q = \text{ite}(x, q_1, q_0)$ and $q' = \text{ite}(y, q'_1, q'_0)$. Then $h = h_0 \cup h_1 \cup \{q_0 \mapsto t_0, q_1 \mapsto t_1\}$ for some compatible histories h_0 and h_1 for q_0 and q_1 , respectively, and t_0 and t_1 such that $t = (\frac{1}{t_0} + \frac{1}{t_1})^{-1}$. Similarly, $h' = h'_0 \cup h'_1 \cup \{q'_0 \mapsto t'_0, q'_1 \mapsto t'_1\}$. Let $b = \alpha(x)$ and $c = \alpha'(y)$. Let also β be the restriction of α to $\text{var}(\alpha) \setminus \{x\}$ and β' be the restriction of α' to $\text{var}(\alpha') \setminus \{y\}$. Then

$$\begin{aligned} & \Pr[\alpha \in \mathfrak{S}_{h,t}^r(q) \text{ and } \alpha' \in \mathfrak{S}_{h',t'}^r(q')] \\ &= \Pr \left[\beta \in \text{reduce} \left(\mathfrak{S}_{h_b, t_b}^r(q_b), \frac{t}{t_b} \right), \beta' \in \text{reduce} \left(\mathfrak{S}_{h'_c, t'_c}^r(q'_c), \frac{t'}{t'_c} \right) \right. \\ & \quad \left. \mid \beta \in \mathfrak{S}_{h_b, t_b}^r(q_b), \beta' \in \mathfrak{S}_{h'_c, t'_c}^r(q'_c) \right] \Pr \left[\beta \in \mathfrak{S}_{h_b, t_b}^r(q_b) \text{ and } \beta' \in \mathfrak{S}_{h'_c, t'_c}^r(q'_c) \right] \end{aligned}$$

Now, because q and q' are both in L_i , neither is an ancestor of the other and thus the two **reduce** are independent: the output of one **reduce** does not modify the set fed into the second **reduce** nor its output. Thus the probability becomes

$$\begin{aligned} & \Pr \left[\beta \in \mathfrak{S}_{h_b, t_b}^r(q_b), \beta' \in \mathfrak{S}_{h'_c, t'_c}^r(q'_c) \right] \cdot \Pr \left[\beta \in \text{reduce} \left(\mathfrak{S}_{h_b, t_b}^r(q_b), \frac{t}{t_b} \right) \mid \beta \in \mathfrak{S}_{h_b, t_b}^r(q_b) \right] \\ & \quad \cdot \Pr \left[\beta' \in \text{reduce} \left(\mathfrak{S}_{h'_c, t'_c}^r(q'_c), \frac{t'}{t'_c} \right) \mid \beta' \in \mathfrak{S}_{h'_c, t'_c}^r(q'_c) \right] \end{aligned}$$

which is equal to $\frac{tt'}{t_b t'_c} \Pr \left[\beta \in \mathfrak{S}_{h_b, t_b}^r(q_b) \text{ and } \beta' \in \mathfrak{S}_{h'_c, t'_c}^r(q'_c) \right]$. Now, if $(\beta, q_b) = (\beta', q'_c)$ then $t_b = t'_c$ (because h and h' are compatible) and $\Pr \left[\beta \in \mathfrak{S}_{h_b, t_b}^r(q_b) \text{ and } \beta' \in \mathfrak{S}_{h'_c, t'_c}^r(q'_c) \right] = \Pr \left[\beta \in \mathfrak{S}_{h_b, t_b}^r(q_b) \right] = t_b = t'_c$ by Lemma 9. So $\Pr[\alpha \in \mathfrak{S}_{h,t}^r(q) \text{ and } \alpha' \in \mathfrak{S}_{h',t'}^r(q')] \leq \frac{tt'}{t_b} =$

$\frac{tt'}{h(q_b)}$. By assumption, $(\alpha, q) \neq (\alpha', q')$, if $q \neq q'$ then the two derivation paths $\text{path}(\alpha, q)$ and $\text{path}(\alpha', q')$ diverge for the first time at q_b , and if $q = q'$ then $x = y$ and $\alpha(x) = 1 - \alpha'(x)$ (because $(\alpha, q) \neq (\alpha', q')$ by assumption). In this case the derivation paths still diverge for the first time at q_b : one follows the 0-edge and the other follows the 1-edge. So in both cases $\text{lcpn}(\text{path}(\alpha, q), \text{path}(\alpha', q')) = q_b$ and we are done. We still have $(\beta, q_b) \neq (\beta', q'_c)$ to consider. In this case the paths $\text{path}(\beta, q_b)$ and $\text{path}(\beta', q'_c)$ diverge for the first time at some node q^* below q_b and q'_c so by induction $\Pr[\beta \in \mathfrak{S}_{h_b, t_b}^r(q_b) \text{ and } \beta' \in \mathfrak{S}_{h'_c, t'_c}^r(q'_c)] \leq t_b t'_c / h_b(q^*) = t_b t'_c / h(q^*)$. So $\Pr[\alpha \in \mathfrak{S}_{h, t}^r(q) \text{ and } \alpha' \in \mathfrak{S}_{h', t'}^r(q')] \leq tt' / h(q^*)$. But q^* is also the first node where $\text{path}(\alpha, q)$ and $\text{path}(\alpha', q')$ diverge, hence the result.

Case i even. In this case q and q' are both $\vee$ -nodes. Say $q = q_1 \vee \dots \vee q_k$ and $q' = q'_1 \vee \dots \vee q'_m$. Then $h = h_1 \cup \dots \cup h_k \cup \{q_1 \mapsto t_1, \dots, q_k \mapsto t_k\}$ and $h' = h'_1 \cup \dots \cup h'_m \cup \{q'_1 \mapsto t'_1, \dots, q'_m \mapsto t'_m\}$. Let $t_{\min} = \min(t_1, \dots, t_k)$ and $t'_{\min} = \min(t'_1, \dots, t'_m)$.

$$\begin{aligned} & \Pr[\alpha \in \mathfrak{S}_{h, t}^r(q) \text{ and } \alpha' \in \mathfrak{S}_{h', t'}^r(q')] \\ &= \Pr\left[\alpha \in \text{reduce}\left(\hat{\mathfrak{S}}_h^r(q), \frac{t}{t_{\min}}\right) \text{ and } \alpha' \in \text{reduce}\left(\hat{\mathfrak{S}}_{h'}^r(q'), \frac{t'}{t'_{\min}}\right)\right] \\ &= \Pr\left[\alpha \in \text{reduce}\left(\hat{\mathfrak{S}}_h^r(q), \frac{t}{t_{\min}}\right), \alpha' \in \text{reduce}\left(\hat{\mathfrak{S}}_{h'}^r(q'), \frac{t'}{t'_{\min}}\right) \mid \alpha \in \hat{\mathfrak{S}}_h^r(q), \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')\right] \\ & \quad \cdot \Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q) \text{ and } \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')] \end{aligned}$$

The **reduce** events are independent because q and q' both belong to L_i and thus neither in an ancestor of the other: the output of the first **reduce** does not influence the output of the second one, even with the knowledge that $\alpha \in \hat{\mathfrak{S}}_h^r(q)$ and $\alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')$. So the probability becomes

$$\begin{aligned} & \Pr\left[\alpha \in \text{reduce}\left(\hat{\mathfrak{S}}_h^r(q), \frac{t}{t_{\min}}\right) \mid \alpha \in \hat{\mathfrak{S}}_h^r(q)\right] \Pr\left[\alpha' \in \text{reduce}\left(\hat{\mathfrak{S}}_{h'}^r(q'), \frac{t'}{t'_{\min}}\right) \mid \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')\right] \\ & \quad \Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q) \text{ and } \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')] \end{aligned}$$

which is $\frac{tt'}{t_{\min} t'_{\min}} \Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q), \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')]$. Now, there are a unique j and ℓ such that $\alpha \in \hat{\mathfrak{S}}_h^r(q)$ only if $\alpha \in \mathfrak{S}_{h_j, t_j}^r(q_j)$ and $\alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')$ only if $\alpha' \in \mathfrak{S}_{h'_\ell, t'_\ell}^r(q'_\ell)$. Thus $\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q), \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')]$ equals

$$\Pr\left[\alpha \in \text{reduce}\left(\mathfrak{S}_{h_j, t_j}^r(q_j), \frac{t_{\min}}{t_j}\right), \alpha' \in \text{reduce}\left(\mathfrak{S}_{h'_\ell, t'_\ell}^r(q'_\ell), \frac{t'_{\min}}{t'_\ell}\right)\right].$$

q_j and q'_ℓ belong to the same layer so with similar arguments we find that

$$\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q) \text{ and } \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')] = \frac{t_{\min} t'_{\min}}{t_j t'_\ell} \Pr[\alpha \in \mathfrak{S}_{h_j, t_j}^r(q_j) \text{ and } \alpha' \in \mathfrak{S}_{h'_\ell, t'_\ell}^r(q'_\ell)].$$

It is possible that $(\alpha, q_j) = (\alpha', q'_\ell)$ but then $q \neq q'$ for otherwise we would have $(\alpha, q) = (\alpha', q')$, against assumption. In the case $(\alpha, q_j) = (\alpha', q'_\ell)$ we use Lemma 9 and find $\Pr[\alpha \in \mathfrak{S}_{h_j, t_j}^r(q_j) \text{ and } \alpha' \in \mathfrak{S}_{h'_\ell, t'_\ell}^r(q'_\ell)] = \Pr[\alpha \in \mathfrak{S}_{h_j, t_j}^r(q_j)] = t_j = t'_\ell$. So

$$\Pr[\alpha \in \hat{\mathfrak{S}}_h^r(q) \text{ and } \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q')] = \frac{t_{\min} t'_{\min}}{t_j}$$

and $\Pr \left[\alpha \in \mathfrak{S}_{h,t}^r(q) \text{ and } \alpha' \in \mathfrak{S}_{h',t'}^r(q') \right] = \frac{tt'}{t_j}$. When $(\alpha, q_j) = (\alpha', q'_\ell)$, the paths $\text{path}(\alpha, q)$ and $\text{path}(\alpha', q')$ diverge for the first time at $q_j = q'_\ell$. So $t_j = h(\text{lcpn}(\text{path}(\alpha, q), \text{path}(\alpha', q')))$ and we are done. Now let us assume that $(\alpha, q_j) \neq (\alpha', q'_\ell)$, then we use the induction hypothesis and, denoting $q^* = \text{lcpn}(\text{path}(\alpha, q_j), \text{path}(\alpha', q'_\ell))$, we have

$$\Pr \left[\alpha \in \hat{\mathfrak{S}}_h^r(q), \alpha' \in \hat{\mathfrak{S}}_{h'}^r(q') \right] \leq \frac{t_{\min} t'_{\min}}{h(q^*)} \text{ and } \Pr \left[\alpha \in \mathfrak{S}_{h,t}^r(q), \alpha' \in \mathfrak{S}_{h',t'}^r(q') \right] \leq \frac{tt'}{h(q^*)}$$

When $(\alpha, q_j) \neq (\alpha', q'_\ell)$, the first node where $\text{path}(\alpha, q)$ and $\text{path}(\alpha', q')$ diverge is also the first node where $\text{path}(\alpha, q_j)$ and $\text{path}(\alpha', q'_\ell)$ diverge, so q^* . This finishes the proof of the inductive case. $\blacktriangleleft$