Hybrid Probabilistic Inference with Logical Constraints: Tractability and Message Passing

Zhe Zeng[*]
Computer Science Department
University of California, LA

Fanqi Yan[*]
AMSS
Chinese Academy of Sciences

Paolo Morettin[*]
DISI
University of Trento, Italy

Antonio Vergari
Computer Science Department
University of California, LA

Guy Van den Broeck
Computer Science Department
University of California, LA

Abstract

Weighted model integration (WMI) is a very appealing framework for probabilistic inference: it allows to express the complex dependencies of real-world hybrid scenarios where variables are heterogeneous in nature (both continuous and discrete) via the language of Satisfiability Modulo Theories (SMT), as well as computing probabilistic queries with complex logical constraints. Recent work has shown WMI inference to be reducible to a model integration (MI) problem, under some assumptions, thus effectively allowing hybrid probabilistic reasoning by volume computations. In this paper, our first contribution is that we theoretically trace the tractability boundaries of exact MI. Indeed, we prove that in terms of the structural requirements on the primal graphs representing formula structures, bounding graph diameter and treewidth is not only sufficient, but also necessary for tractable exact inference via MI. Our second contribution is that we we introduce a novel formulation of MI via an exact message passing scheme on the tractable MI problems. It allows to efficiently compute the marginal densities and statistical moments of all the variables in linear time. As such, we are able to amortize inference for rich MI queries when they conform to the problem structure, i.e., the primal graph.

1 Introduction

In many real-world scenarios, performing probabilistic inference requires reasoning over domains with complex logical constraints while dealing with variables that are heterogeneous in nature, i.e., both continuous and discrete. Consider for instance an autonomous agent such as a self-driving vehicle. It would have to model continuous variables like the speed and position of other vehicles, which are constrained by the geometry of vehicles and roads and the laws of physics. It should also be able to reason over discrete attributes like the color of traffic lights and the number of pedestrians.

These scenarios are beyond the reach of probabilistic models like variational autoencoders [20] and generative adversarial networks [17], whose inference capabilities, despite their recent success, are severely limited. Classical probabilistic graphical models [23], while providing more flexible inference routines, are generally incapacitated when dealing with continuous and discrete variables at once [32], or make simplistic [18, 24] or overly strong assumptions about their parametric forms [37]. Even recent efforts in modeling these heterogeneous scenarios flexibly [27, 34], while delivering tractable inference, are not able to perform inference in the presence of logical constraints.

[*]Authors contributed equally. This research was performed while F.Y. and P.M. were visiting UCLA.

Preprint. Under review.

Weighted Model Integration (WMI) [3, 28] is a recently introduced framework for probabilistic inference that offers all the aforementioned "ingredients" needed for hybrid probabilistic reasoning with logical constraints, *by design*. First, WMI leverages the expressive representation language of Satisfiability Modulo Theories (SMT) [2] for describing both a problem (theory) over continuous and discrete variables, and complex logical formulas to query it. Second, analogous to how Weighted Model Counting (WMC) [7] enables state-of-the-art probabilistic inference over discrete variables [14, 33, 19, 35], probabilistic inference over hybrid domains can be carried in a principled way by WMI. Indeed, parameterizing a WMI problem by the choice of some simple weight functions (e.g., per-literal polynomials [4]) induces a probability distribution over the models of the formula.

These appealing properties motivated several recent works on WMI [28, 29, 21, 39], pushing the boundaries of state-of-the-art solvers over SMT formulas. Recently, a polytime reduction of WMI problems to unweighted Model Integration (MI) problems over real variables has been proposed [38], opening a new perspective on building such algorithms. Solving an MI problem effectively reduces probabilistic reasoning to computing volumes over constrained regions. In fact, as we will prove in this paper, computing MI is inherently hard, whenever the problem structure, here represented by the primal graph associated to the SMT formula, does not abide by some requirements.

The contribution we make in this work is twofold: we theoretically trace the requirements for tractable exact MI inference and we propose an efficient algorithm that satisfies them. Indeed, we first prove that performing MI is #P-hard unless the primal graph of the associated SMT formula is a tree and has a balanced diameter. Under these structural assumptions, we devise a novel MI inference scheme via *message passing* which is able to exactly compute all the variable marginal densities – as well as statistical moments – at once. As such, we are able to amortize inference *inter-queries* for rich univariate and bivariate MI queries when they conform to the formula structure. By doing so, we are able to scale inference beyond the capabilities of all current exact WMI solvers.

The paper is organized as follows. We start by reviewing the necessary WMI and SMT background. Then we introduce MI while presenting our theoretical results on the hardness of MI in the following section. Next, we present our exact message passing scheme, after which we perform experiments.

2 Background

Notation. We use uppercase letters for random variables, e.g., X, B, and lowercase letters for their assignments e.g., x, b. Bold uppercase letters denote sets of variables, e.g., $\mathbf{X}, \mathbf{B}$, and their lowercase version their assignments, e.g., $\boldsymbol{x}, \boldsymbol{b}$. We denote with capital greek letters, e.g., Λ, Φ, Δ, (quantifier free) logical formulas and literals (i.e., atomic formulas or their negation) with lowercase ones, e.g., ℓ, ϕ, δ. We denote satisfaction of a formula Φ by one assignment $\boldsymbol{x}$ by $\boldsymbol{x} \models \Phi$ and we denote its corresponding indicator function as $[\![\boldsymbol{x} \models \Phi]\!]$. We use $[n]$ to denote the set of integers from 1 to n.

Satisfiability Modulo Theories (SMT). SMT [1] generalizes the well-known SAT problem [6] to determining the satisfiability of a logical formula w.r.t. a decidable background theory. Rich mixed logical/algebraic constraints can be expressed in SMT for hybrid domains. In particular, we consider quantifier-free SMT formulas in the theory of linear arithmetic over the reals, or SMT($\mathcal{LRA}$). Here, formulas are Boolean combinations of atomic propositions (e.g., α, β), and of atomic $\mathcal{LRA}$ formulas over real variables (e.g., $\ell : X_i < X_j + 5$), for which satisfaction is defined in an obvious way. W.l.o.g. we assume SMT formulas to be in conjunctive normal form (CNF) (see Figure 4 for some examples).

In order to characterize the dependency structure of an SMT($\mathcal{LRA}$) formula Δ as well as the hardness of inference, we denote the *primal graph* [13] of formula Δ by $\mathcal{G}_\Delta$, as the undirected graph whose vertices are variables in Δ and whose edges connect any two variables that appear together in at least one clause in Δ. In the next sections, we will extensively refer to the *diameter* and *treewidth* of a primal graph defined as usual for undirected graphs [23]. Recall that trees have treewidth one.

Weighted Model Integration (WMI). Weighted Model Integration (WMI) [3, 28] provides a framework for probabilistic inference over models defined over the logical constraints given by SMT($\mathcal{LRA}$) formulas. Formally, let $\mathbf{X}$ be a set of continuous random variables defined over $\mathbb{R}$, and $\mathbf{B}$ a set of Boolean random variables defined over $\mathbb{B} = \{\top, \bot\}$. Given an SMT formula Δ over (subsets of) $\mathbf{X}$ and $\mathbf{B}$, a *weight function* $w : (\boldsymbol{x},\boldsymbol{b}) \mapsto \mathbb{R}^+$, the task of computing the WMI over formula Δ, w.r.t. weight function w, and variables $\mathbf{X}$ and $\mathbf{B}$ is defined as:

$$\mathsf{WMI}(\Delta, w; \mathbf{X}, \mathbf{B}) \triangleq \sum_{\boldsymbol{b} \in \mathbb{B}^{|\mathbf{B}|}} \int_{\Delta(\boldsymbol{x}, \boldsymbol{b})} w(\boldsymbol{x}, \boldsymbol{b})\, d\boldsymbol{x}, \tag{1}$$

that is, summing over all possible Boolean assignments $\boldsymbol{b} \in \mathbb{B}^{|\mathbf{B}|}$ while integrating over those assignments of $\mathbf{X}$ such that the evaluation of the formula $\Delta(\boldsymbol{x}, \boldsymbol{b})$ is SAT. Intuitively, $\mathsf{WMI}(\Delta, w; \mathbf{X}, \mathbf{B})$ equals the partition function of the unnormalized probability distribution induced by weight w on formula Δ. As such, the weight function w acts as an unnormalized probability density while the formula Δ represents logical constraints defining its structure. In the following, we will adopt the shorthand $\mathsf{WMI}(\Delta, w)$ for computing the WMI of all the variables in Δ. More generally, the choice of the weight function w can by guided by some domain-specific knowledge or efficiency reasons. We follow the common assumption that the weight function w factorizes over the literals ℓ in Δ that are satisfied by one joint assignment $(\boldsymbol{x}, \boldsymbol{b})$, i.e., $w(\boldsymbol{x}, \boldsymbol{b}) = \prod_{\ell:(\boldsymbol{x},\boldsymbol{b})\models\ell} w_\ell(\boldsymbol{x}, \boldsymbol{b})$. Moreover, we adopt *polynomial functions* [3, 4, 28] for the per-literal weight w_ℓ. Note that this induces a global *piecewise polynomial* parametric form for weight w, where each piece is defined as the polynomial associated to a region induced by the truth assignments to formula Δ [28]. Furthermore, multivariate piecewise polynomials can be integrated efficiently over given bounds [11].

For example, consider the WMI problem over formula $\Delta = (0 < X_1 < 2) \wedge (0 < X_2 < 2) \wedge (X_1 + X_2 < 2) \wedge (B \vee (X_1 > 1))$ on variables $\mathbf{X} = \{X_1, X_2\}$, $\mathbf{B} = \{B\}$. Let the weight function w decompose on per-literal weights as follows: $w_{\ell_1}(X_1, X_2) = X_1 X_2$, $w_{\ell_2}(X_1) = 2$ and $w_{\ell_3}(B) = 3$, where $\ell_1 = X_1 + X_2 < 2$, $\ell_2 = X_1 > 1$ and $\ell_3 = B$. Note that literals not mentioned are assumed to have weight one. Then, the WMI can be computed as:

$$\begin{aligned}\mathsf{WMI}(\Delta, w; \mathbf{X}, B) = {} & 3 \int_0^1 dx_1 \int_0^{2-x_1} x_1 x_2\, dx_2 + 3 \int_1^2 dx_1 \int_0^{2-x_1} 2x_1 x_2\, dx_2 \\ & + \int_1^2 dx_1 \int_0^{2-x_1} 2x_1 x_2\, dx_2 = \frac{73}{24}.\end{aligned} \tag{2}$$

Model Integration is all you need. Recently, Zeng and Van den Broeck (2019) showed that a WMI problem can be reduced in poly-time to an Model Integration (MI) problem over continuous variables only. This reduction is appealing because it allows to perform hybrid probabilistic reasoning with logical constraints in terms of volume computations over convex polytopes. We now briefly review this poly-time reduction. We refer the readers to [38] for a detailed exposition.

First, w.l.o.g., a WMI problem on continuous and Boolean variables of the form $\mathsf{WMI}(\Delta, w; \mathbf{X}, \mathbf{B})$ can always be reduced to new WMI problem $\mathsf{WMI}_{\mathbb{R}}(\Delta', w'; \mathbf{X}')$ on continuous variables only. To do so, we substitute the Boolean variables $\mathbf{B}$ in formula Δ with fresh continuous variables in $\mathbf{X}'$ and replace each Boolean atom and its negation in formula Δ by two exclusive $\mathcal{LRA}$ atoms over the new real variables in formula Δ', and distilling a new weight function w accordingly. Note that the primal graph of formula Δ' retains its treewidth, i.e., if primal graph $\mathcal{G}_\Delta$ is a tree, then so is the graph $\mathcal{G}_{\Delta'}$.

Furthermore, WMI problems on continuous variables with polynomial weights can be reduced to equivalent MI problems whose definition will be formally presented in the next paragraph. Specifically, $\mathsf{WMI}_{\mathbb{R}}$ with polynomial weights w' have equivalent MI problems $\mathsf{WMI}_{\mathbb{R}}(\Delta', w'; \mathbf{X}') = \mathsf{MI}(\Delta''; \mathbf{X}'')$, with $\mathbf{X}''$ containing auxiliary continuous variables whose extrema of integration are chosen such that their integration is precisely the value of weights w'. In the case of monomial weights, the treewidth of $\mathcal{G}''_\Delta$ will not increase w.r.t. $\mathcal{G}'_\Delta$. This is not guaranteed for generic polynomial weights. A detailed description of these reduction processes is included in Appendix, where we also show the $\mathsf{WMI}_{\mathbb{R}}$ and MI problems equivalent to the WMI one in Equation 2.

Computing MI. Given a set $\mathbf{X}$ of continuous random variables over $\mathbb{R}$, and an SMT($\mathcal{LRA}$) formula $\Delta = \bigwedge_i \Gamma_i$ over $\mathbf{X}$, the task of MI over formula Δ, w.r.t. variables $\mathbf{X}$ is defined as computing the following integral [38]:

$$\mathsf{MI}(\Delta; \mathbf{X}) \triangleq \int_{\boldsymbol{x}\models\Delta} 1\, d\boldsymbol{x} = \int_{\mathbb{R}^{|\mathbf{X}|}} [\![\boldsymbol{x} \models \Delta]\!]\, d\boldsymbol{x} = \int_{\mathbb{R}^{|\mathbf{X}|}} \prod_{\Gamma\in\Delta} [\![\boldsymbol{x} \models \Gamma]\!]\, d\boldsymbol{x}. \tag{3}$$

The first equality can be seen as computing the volume of the constrained regions defined by formula Δ, and the last one is obtained by eliciting the "pieces" associated to each clause $\Gamma \in \Delta$. Again, in the following we will use the shorthand $\mathsf{MI}(\Delta)$ when integrating over all variables in formula Δ. Moreover, the MI problem $\mathsf{MI}(\Delta)$ can be rewritten in an iterated integral form as follows.

$$\mathsf{MI}(\Delta) = \int_{\mathbb{R}} dx_1 \cdots \int_{\mathbb{R}} dx_{i-1} \int_{\mathbb{R}} f_i(x_i)\, dx_i, \quad i = 2, \cdots, n. \tag{4}$$

In a general way, we can always define a univariate piecewise polynomial f_i as a function of the MI over the remaining variables in a recursive way as follow:

$$f_i(x_i) := \int_{\mathbb{R}} [\![x_i, x_{i+1} \models \tilde{\Delta}_i]\!] \cdot f_{i+1}(x_{i+1})\, dx_{i+1},\ i \in [1, n-1] \qquad f_n(x_n) := [\![x_n \models \tilde{\Delta}_n]\!]$$

where the formula $\tilde{\Delta}_i := \exists \boldsymbol{x}_{1:i-1}.\Delta$ is defined by the forgetting operation [25]. Recall that the formula Δ is defined by SMT($\mathcal{LRA}$) which means that the integration bounds are linear arithmetic over the real variables. Thus the MI can be expressed as the integration over an arbitrary variable $X_r \in \mathbf{X}$ where the integrand f_r is a univariate piecewise polynomial and the pieces are the collection I of intervals of the form $[l, u]$:

$$\mathsf{MI}(\Delta) = \int_{\mathbb{R}} f_r(x_r)\, dx_r = \sum\nolimits_{[l,u]\in I} \int_l^u f_{l,u}(x_r) dx_r. \tag{5}$$

Hybrid inference via MI. Before moving to our theoretical and algorithmic contributions, we review the kind of probabilistic queries computable via MI.[2] Analogously to $\mathsf{WMI}(\Delta, w)$, $\mathsf{MI}(\Delta)$ computes the partition function of the unnormalized distribution induced over the models of formula Δ. Therefore, it is possible to compute the (now normalized) probability of any logical *query* Φ expressable as an SMT($\mathcal{LRA}$) formula involving complex logical and numerical constraints as

$$\mathsf{Pr}_{\Delta}(\Phi) = \mathsf{MI}(\Delta \wedge \Phi)\ /\ \mathsf{MI}(\Delta).$$

In the next section, we will show how to compute the probabilities of a collection of rich queries $\{\Phi_t\}_t$ in a single message-passing evaluation if all Φ_t are univariate formulas, i.e., contain only one variable $X_i \in \mathbf{X}$, or bivariate ones conforming to graph $\mathcal{G}_{\Delta}$, i.e., Φ_t contains only $X_i, X_j \in \mathbf{X}$ and they are connected by at least one edge in $\mathcal{G}_{\Delta}$. Moreover, one might want to statistically reason about the marginal distribution of the variables in $\mathbf{X}$, i.e., $p_{\Delta}(x_i)$ which is defined as:

$$p_{\Delta}(x_i) \triangleq \frac{1}{\mathsf{MI}(\Delta)} f_i(x_i) = \frac{1}{\mathsf{MI}(\Delta)} \int_{\mathbb{R}^{|\mathbf{X}|-1}} [\![\boldsymbol{x} \models \Delta]\!]\, d\boldsymbol{x} \setminus \{x_i\}. \tag{6}$$

3 On the inherent hardness of MI

It is well-known that for discrete probabilistic graphical models, the simplest structural requirement to guarantee tractable inference is to bound their treewidth [23]. For instance, for tree-shaped Bayesian Networks, all exact marginals can be computed at once in polynomial time [30]. However, existing WMI solvers show exponential blow-up in their runtime even when the WMI problems have primal graphs with simple tree structures [38]. This observation motivates us to trace the theoretical boundaries for tractable probabilistic inference via MI. As we will show in this section, we find out that requiring a MI problem to only have a tree-shaped structure is not sufficient to ensure tractability. Therefore, inference on MI problems is inherently harder than its discrete-only counterpart.

Specifically, we will show how the hardness of MI depends on two structural properties: the treewidth of the primal graph and the length of its diameter. To begin, we prove that even for SMT($\mathcal{LRA}$) formulas Δ whose primal graphs $\mathcal{G}_{\Delta}$ are trees but have unbounded diameters (i.e., they are unbalanced trees, like paths), computing $\mathsf{MI}(\Delta)$ is hard. This is surprising since for its discrete counterpart, the complexity of model counting problem is exponential in the treewidth but not in the diameter.

Theorem 1. *Computing $\mathsf{MI}(\Delta)$ of an SMT($\mathcal{LRA}$) formula Δ whose primal graph is a tree with diameter $O(n)$ is #P-hard, with n being the number of variables.*

Proof. We prove our complexity result by reducing a #P-complete variant of the subset sum problem [15] to an MI problem over an SMT($\mathcal{LRA}$) theory Δ with a tree primal graph whose diameter is $O(n)$. This problem is a counting version of the subset sum problem saying that given a set of

[2] Note that equivalent queries can be defined for WMI and $\mathsf{WMI}_{\mathbb{R}}$ problem formulations.

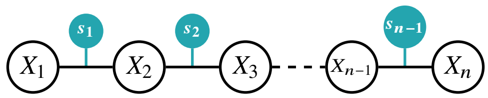

Figure 1: Primal graph used for #P-hardness reduction

positive integers $S = \{s_0, s_1, \cdots, s_{n-1}\}$, and a positive integer L, the goal is to count the number of subsets $S' \subseteq S$ such that the sum of all the integers in the subset S' equals to L.

In a nutshell, one can always construct in polynomial time a theory Δ such that its primal graph $\mathcal{G}_\Delta$ is a chain (hence has $O(n)$ diameter) and where computing $\mathsf{MI}(\Delta)$ equals solving (up to a constant) the aforementioned subset sum problem variant, which is known to be #P-hard [10, 8]. Firstly, we reduce the counting subset sum problem in polynomial time to a model integration problem with the following SMT($\mathcal{LRA}$) theory whose primal graph is shown in Figure 1.

$$\Delta = \begin{cases} s_1 - \frac{1}{2n} < X_1 < s_1 + \frac{1}{2n} \vee -\frac{1}{2n} < X_1 < \frac{1}{2n} \\ X_{i-1} + s_i - \frac{1}{2n} < X_i < X_{i-1} + s_i + \frac{1}{2n} \ \vee \ X_{i-1} - \frac{1}{2n} < X_i < X_{i-1} + \frac{1}{2n}, \quad i = 2, \cdots n \end{cases}$$

For brevity, denote the first and the second literal in the i-th clause by $\ell(i, 0)$ and $\ell(i, 1)$, respectively. Also we choose two constants $l = L - \frac{1}{2}$ and $u = L + \frac{1}{2}$. In the following, we prove that $n^n \mathsf{MI}(\Delta \wedge (l < X_n < u))$ equals the number of subset $S' \subseteq S$ whose element sum is L.

Let $\boldsymbol{a}^k = (a_1, a_2, \cdots, a_k)$ be some assignment to Boolean variables $\{A_1, A_2, \cdots, A_k\}$ with $a_i \in \{0, 1\}$, $i \in [k]$. Define $S(\boldsymbol{a}^k) \triangleq \sum_{i=1}^k a_i s_{i-1}$ as subset sums, and formulas $\Delta_{\boldsymbol{a}^k} \triangleq \bigwedge_{i=1}^k \ell(i, a_i)$ for each $\boldsymbol{a}^k$.

We claim that given an assignment $\boldsymbol{a}^k \in \{0, 1\}^k$, the model integration for theory $\Delta_{\boldsymbol{a}^k}$ is $(\frac{1}{n})^k$. Moreover, using induction we conclude that for each variable x_i in $\Delta_{\boldsymbol{a}^k}$, its satisfying assignments form the interval $[\sum_{j=1}^i a_j s_{j-1} - \frac{i}{2n}, \sum_{j=1}^i a_j s_{j-1} + \frac{i}{2n}]$. Specifically, the satisfying assignments for variable x_n in theory $\Delta_{\boldsymbol{a}^n}$ is the interval $[S(\boldsymbol{a}^n) - \frac{1}{2}, S(\boldsymbol{a}^n) + \frac{1}{2}]$. For any subset $S' \subseteq S$, we can have one-to-one correspondence to assignments in $\{0, 1\}^n$ by defining $\boldsymbol{a}^n$ as $a_i = 1$ if and only if $s_i \in S'$. Furthermore, each assignments $\boldsymbol{a}^n \in \{0, 1\}^n$, the model integration of $\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)$ falls into one of the following two cases: 1) if $S(\boldsymbol{a}^n) < L$ or $S(\boldsymbol{a}^n) > L$, then $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = 0$; 2) otherwise if $S(\boldsymbol{a}^n) = L$, then $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = (\frac{1}{n})^n$. Indeed, we have shown that variable X_n has its satisfying assignments in interval $[S(\boldsymbol{a}^n) - \frac{1}{2}, S(\boldsymbol{a}^n) + \frac{1}{2}]$ in theory $\Delta_{\boldsymbol{a}^n}$ for each $\boldsymbol{a}^n \in \{0, 1\}^n$. If $S(\boldsymbol{a}^n) < L$, given that $S(\boldsymbol{a}^n)$ is a sum of positive integers, then it holds that $S(\boldsymbol{a}^n) + \frac{1}{2} \leq (L - 1) + \frac{1}{2} = L - \frac{1}{2} = l$ and therefore, $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = 0$; similarly, if $S(\boldsymbol{a}^n) > L$, then it holds that $S(\boldsymbol{a}^n) - \frac{1}{2} \geq u$ and therefore, $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = 0$. If $S(\boldsymbol{a}^n) = L$, we have $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = \mathsf{MI}(\Delta_{\boldsymbol{a}^n}) = (\frac{1}{n})^n$.

Observe that for each clause in SMT($\mathcal{LRA}$) Δ, literals are mutually exclusive since each s_i is a positive integer. Then we have that formulas $\Delta_{\boldsymbol{a}^n}$ are mutually exclusive and meanwhile $\Delta = \bigvee_{\boldsymbol{a}^n} \Delta_{\boldsymbol{a}^n}$. Thus it holds that $\mathsf{MI}(\Delta) = \sum_{\boldsymbol{a}^n} \mathsf{MI}(\Delta_{\boldsymbol{a}^n})$. Similarly, we have formulas $(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u))$'s are mutually exclusive and meanwhile $\mathsf{MI}(\Delta \wedge (l < X_n < u)) = \sum_{\boldsymbol{a}^n} \mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u))$. From the above results, we can conclude that $\mathsf{MI}(\Delta \wedge (l < X_n < u)) = t(\frac{1}{n})^n$ where t is the number of assignments $\boldsymbol{a}^n$ such that $S(\boldsymbol{a}^n) = L$. Notice that for each $\boldsymbol{a}^n \in \{0, 1\}^n$, it one-to-one corresponds to a subset $S' \subseteq S$ and $S(\boldsymbol{a}^n)$ equals to L if and only if the sum of elements in S' is L. This finishes our proof of the statement that $n^n \mathsf{MI}(\Delta \wedge (l < X_n < u))$ equals to the number of subset $S' \subseteq S$ whose element sum equals to L. Therefore, model integration problems with tree primal graphs whose diameter is $O(n)$ is #P-hard. A more complete proof is provided in Appendix B.1. □

Furthermore, when the primal graphs are balanced trees, i.e., they have $O(\log(n))$ diameters, increasing their treewidth from one to two is sufficient to turn MI problems from tractable to #P-hard.

Theorem 2. *Computing $\mathsf{MI}(\Delta)$ of an SMT($\mathcal{LRA}$) formula Δ whose primal graph $\mathcal{G}_\Delta$ has treewidth two and diameter of length $O(\log(n))$ is #P-hard, with n being the number of variables.*

Sketch of proof. As before, we construct a poly-time reduction from the #P-complete variant of the subset sum problem to an MI problem. This time, the SMT($\mathcal{LRA}$) formula Δ is built such that the graph $\mathcal{G}_\Delta$ has treewidth two with cliques (hence not a tree). Meanwhile the primal graph has

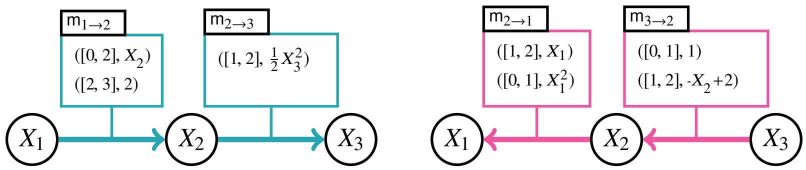

Figure 2: **An example of a run of MP-MI**. Upward (blue) and downward (pink) messages are shown as piecewise polynomials in boxes for a tree primal graph rooted at X_3. The final belief for X_2 can then be computed as the piecewise polynomial $\mathsf{b}_2 = \mathsf{m}_{1\to2} \cdot \mathsf{m}_{3\to2} = \{([0,1], X_2), ([1,2], -X_2^2 + 2X_2)\}$.

diameter to be at most $\log(n)$ by putting the cliques in a balanced way. Then the MI over a subtree could potentially be a subset sum over the integers that appear in the formulas associated with the subtree. Then computing the MI of formula Δ equals to solving the subset sum problem. Complete proof is provided in Appendix B.2. □

From Theorems 1 and 2 we can deduce that having a tree-shaped *and* balanced primal graph is a necessary condition for tractability. Together with the fact that our proposed exact MI solver, which will be introduce in the next section, achieves quasi-polynomial complexity on MI with balanced tree-shaped primal graphs, we shows that the balanced tree-shape is both necessary and sufficient. This sets the standard for the solver complexity: every exact MI solver that aims to be efficient, need to operate in the aforementioned regime.

In next section we introduce a novel and efficient exact MI solver on MI with balanced tree-shaped primal graphs together with its complexity analysis. It computes exact MI by exchanging messages among the nodes of the primal graph of an SMT($\mathcal{LRA}$) formula. As the reader might intuit at this point, devising a message passing inference scheme for MI will be inherently more challenging than for discrete domains.

4 MP-MI: exact MI inference via message passing

Deriving an equivalent message passing scheme for MI to what belief propagation is for the discrete case [30] poses unique and considerable challenges. First, by allowing complex logical constraints such as those creating disjoint feasible regions, one might have to integrate over exponentially many feasible regions, i.e., polytopes [28]. This is computationally expensive even though numerical integration is a consolidated field. Additionally, different from discrete domains, in real or hybrid domains one generally does not have universal and compact representations for distributions [23, 36]. And when these are available, e.g. in the case of Gaussians, the corresponding density models might have restricted expressiveness and not allow for efficient integration over arbitrary constraints. In fact, in the case of exponentiated polynomials, exact integration of is limited to polynomials of low degree (usually, two).

General propagation scheme. Let Δ be an SMT($\mathcal{LRA}$) formula and $\mathcal{G}_\Delta$ its tree primal graph, rooted at node r corresponding to variable $X_r \in \mathbf{X}$. This can always be done by choosing an arbitrary node r as root and then orienting all edges away from node r. Also let $\mathcal{V}$ be the set of indexes of variables $\mathbf{X}$ in formula Δ and let $\mathcal{E}$ be the set of edges $i - j$ in graph $\mathcal{G}_\Delta$ connecting variables X_i and X_j. Then the formula can be rewritten as $\Delta = \bigwedge_{i\in\mathcal{V}} \Delta_i \wedge \bigwedge_{(i,j)\in\mathcal{E}} \Delta_{i,j}$, with Δ_i being formulas involving only variable X_i and, analogously, formula $\Delta_{i,j}$ involving only variables X_i and X_j. Notice that the tree primal graph assumption implies that each clause in the formula has at most two variables; otherwise if there are more than two variables in one clause, these variables would be pairwise connected by one edge in the primal graph and create loops.

Our inference scheme, which we name MP-MI, will be able to compute the unnormalized marginals of each node in $\mathcal{G}_\Delta$, defining the *belief* associated to that node.

Definition 3. *Let Δ be an SMT($\mathcal{LRA}$) formula with tree primal graph $\mathcal{G}_\Delta$. The belief b_i of node i in graph $\mathcal{G}_\Delta$ is the unnormalized marginal $p_i(x_i)$ of variable $X_i \in \mathbf{X}$.*

As we will show, in MP-MI beliefs can be computed by exchanging *messages* between nodes in $\mathcal{G}_\Delta$.

Algorithm 1 MP-MI(Δ) – Message Passing Model Integration

$\mathbf{V}_{\mathsf{up}} \leftarrow$ sort nodes in $\mathcal{G}_\Delta$, children before parents
for each $X_i \in \mathbf{V}_{\mathsf{up}}$ **do** $\mathsf{send\text{-}message}(X_i, X_{\mathsf{parent}(X_i)})$ **end** ▹ upward pass
$\mathbf{V}_{\mathsf{down}} \leftarrow$ sort nodes in $\mathcal{G}_\Delta$, parents before children
for each $X_i \in \mathbf{V}_{\mathsf{down}}$ **do** ▹ downward pass
  for each $X_c \in \mathsf{ch}(X_i)$ **do** $\mathsf{send\text{-}message}(X_i, X_c)$ **end**
Return $\{\mathsf{b}_i\}_{i:X_i \in \mathcal{G}_\Delta}$

send-message(X_i, X_j)

$\mathsf{b}_i \leftarrow \mathsf{compute\text{-}beliefs}$ ▹ cf. Equation 8
$P \leftarrow \mathsf{critical\text{-}points}(\mathsf{b}_i, \Delta_i, \Delta_{i,j})$, $\quad I \leftarrow \mathsf{intervals\text{-}from\text{-}points}(P)$ ▹ cf. SMI in [38]
for interval $[l, u] \in I$ consistent with formula $\Delta_i \wedge \Delta_{i,j}$ **do**
  $\langle l_s, u_s, f\rangle \leftarrow \mathsf{symbolic\text{-}bounds}(\mathsf{b}_i, [l, u], \Delta_{i,j})$
  $f' \leftarrow \int_{l_s}^{u_s} f(x_i)\, dx_i, \quad \mathsf{m}_{i\to j} \leftarrow \mathsf{m}_{i\to j} \cup \langle l, u, f'\rangle$
Return $\mathsf{m}_{i\to j}$

Definition 4. *The message from a node i (corresponding to variable $X_i \in \mathbf{X}$) in primal graph $\mathcal{G}_\Delta$ to one of its neighbor nodes $j \in \mathsf{neigh}(i)$ is computed recursively as follows,*

$$\mathsf{m}_{i\to j}(x_j) = \int_{\mathbb{R}} [\![x_i, x_j \models \Delta_{i,j}]\!] [\![x_i \models \Delta_i]\!] \times \prod_{c \in \mathsf{neigh}(i)\setminus\{j\}} \mathsf{m}_{c\to i}(x_i)\, dx_i \tag{7}$$

MP-MI operates in two phases: an *upward pass* and a *downward* one. First we send messages up from the leaves to the root (upward pass) such that each node has all information from its children and then we incorporate messages from the root down to the leaves (downward pass) such that each node also has information from its parent node. When the message passing process finishes, each node in graph $\mathcal{G}_\Delta$ is able to compute its belief by aggregating the messages received neighbors.

Proposition 5. *Let $\mathsf{ch}(i)$ be the set of children nodes for node i in $\mathcal{G}_\Delta$. The belief of node i in the upward pass, b_i^+, and the downward belief b_i^-, can be computed as:*

$$\mathsf{b}_i^+(x_i) = \prod_{c \in \mathsf{ch}(i)} \mathsf{m}_{c\to i}(x_i), \quad \mathsf{b}_i(x_i) = \mathsf{b}_i^-(x_i) = \prod_{c \in \mathsf{neigh}(i)} \mathsf{m}_{c\to i}(x_i) \tag{8}$$

where $\mathsf{m}_{c\to i}$ denotes the message sent from a node c to its neighbor node i. The final belief of node i is its downward belief which is the unnormalized marginal, i.e. $\mathsf{MI}(\Delta) = \int_{\mathbb{R}} [\![x_i \models \Delta_i]\!] \cdot \mathsf{b}_i(x_i)\, dx_i$.

Notice that even though the integration is symbolically defined over the whole real domain, the SMT($\mathcal{LRA}$) logical constraints in formulas $\Delta_{i,j}$ and Δ_i would give integration bounds that are linear in the variables. This guarantees that our messages will be univariate piecewise polynomials.

Proposition 6. *Let Δ be an SMT($\mathcal{LRA}$) formula with tree primal graph, then the messages as defined in Equation 7 and beliefs as defined in Equation 8 are univariate piecewise polynomials.*

Remark. The multiplication of two piecewise polynomial functions $f_1(x)$ and $f_2(x)$ is defined as a piecewise polynomial function $f(x)$ whose domain is the intersection of the domains of these two functions and for each x in its domain, the value is defined as $f(x) := f_1(x) \cdot f_2(x)$.

In figure 2 we show an example of the two passes in MP-MI and we summarize the whole MP-MI scheme in Algorithm 1. There, two functions *critical-points* and *symbolic-bounds* are subroutines used to compute the numeric and symbolic bounds of integration for our pieces of univariate polynomials. Both of them can be efficiently implemented, see [38] for details. Concerning the actual integration of the polynomial pieces, this can be done efficiently symbolically, a task supported by many scientific computing packages. Next we will show how the beliefs and messages obtained from MP-MI can be leveraged for inference tasks.

Amortizing Queries. Given a SMT($\mathcal{LRA}$) formula Δ, in the next Propositions, we show that we can leverage beliefs and messages computed by MP-MI to speed up (amortize) inference time over multiple queries on formula Δ. More specifically, when given queries that conform to the structure of formula Δ, i.e. queries on a node variable or queries over variables that are connected by an edge in graph $\mathcal{G}_\Delta$, we can reuse the local information encoded in beliefs.

From the *"MI is all you need"* perspective, we can compute the probability of a logical query as a ratio of two MI computations. Expectations and moments can also be computed efficiently by leveraging beliefs and taking ratios. They are pivotal in several scenarios including inference and learning.

Proposition 7. *Let Δ be an SMT($\mathcal{LRA}$) formula with a tree primal graph, and let Φ be an SMT($\mathcal{LRA}$) query over variable $X_i \in \mathbf{X}$. It holds that $\mathsf{MI}(\Delta \wedge \Phi) = \int_{\mathbb{R}} [\![x_i \models \Phi]\!][\![x_i \models \Delta_i]\!]\mathsf{b}_i(x_i)dx_i$.*

Proposition 8. *Let Δ be an SMT($\mathcal{LRA}$) formula and let Φ be an SMT($\mathcal{LRA}$) query over $X_i, X_j \in \mathbf{X}$ that are connected in tree primal graph $\mathcal{G}_\Delta$. The updated message from node j to node i is as follows.*

$$\mathsf{m}^*_{j\to i}(x_i) = \int_{\mathbb{R}} \mathsf{b}_j(x_j)/\mathsf{m}_{i\to j}(x_j) \times [\![x_i, x_j \models \Delta_{i,j} \wedge \Phi]\!][\![x_j \models \Delta_j]\!]\, dx_j$$

It holds that $\mathsf{MI}(\Delta \wedge \Phi) = \int_{\mathbb{R}} [\![x_i \models \Delta_i]\!] \cdot \mathsf{b}^*_i(x_i)dx_i$ with b^*_i obtained from the updated message $\mathsf{m}^*_{j\to i}$.

Proposition 9. *Let Δ be an SMT($\mathcal{LRA}$) formula with tree primal graph, then the k-th moment of variable $X_i \in \mathbf{X}$ can be obtained by $\mathbb{E}[X_i^k] = \frac{1}{\mathsf{MI}(\Delta)} \int_{\mathbb{R}} [\![x_i \models \Delta_i]\!] \times x_i^k \mathsf{b}_i(x_i)\, dx_i$.*

Pre-computing beliefs and messages can dramatically speed up inference by amortization, as we will show in our experiments. This is especially important when the primal graphs have large diameter. In fact, recall from Section 3 that even when the formula Δ has a tree-shaped primal graph, but unbounded diameter, computing MI is still hard.

Complexity of MP-MI. As we mention in our analysis on the inherent hardness of MI problems in Section 3, our proposed MP-MI scheme runs efficiently on MI problems with tree-shaped and balanced tree primal graphs. Here we formally derive its algorithmic complexity. To do so, we leverage the concept of a pseudo tree. The pseudo tree is a directed tree with the shortest diameter among all the spanning trees of an undirected primal graph. In MP-MI this is equivalent to select a root r in the primal graph such that it is the root of the pseudo tree and its child-parent relationships guide the execution of the upward and downward passes.

Theorem 10. *Consider an SMT($\mathcal{LRA}$) formula Δ with a tree primal graph with diameter h_p, and a pseudo tree with l leaves and diameter h_t. Let m be the number of literals in formula Δ, and n be the number of variables. Then MI(Δ) can be computed in $O(l \cdot (n^3 \cdot m^{h_p})^{h_t})$ by the MP-MI algorithm.*

This result comes from the fact that when choosing the same node as root, the *upward pass* of MP-MI essentially corresponds to the SMI algorithm in Zeng and Van den Broeck when symbolic integration is applied. While SMI can only compute the unnormalized marginal of the root node, MP-MI can obtain all unnormalized marginals for all nodes. Therefore, the complexity of MP-MI is linear in the complexity of one run of SMI. Based on the complexity results in Theorem 10, MP-MI is potentially exponential in the diameter of $\mathcal{G}_\Delta$. Note that this results from the fact that the size of a single message, i.e. the number of pieces in its piecewise polynomial representation (cf. Eq. 7), is not bounded by the treewidth of primal graph $\mathcal{G}_\Delta$ but increases exponentially in the diameter of $\mathcal{G}_\Delta$.

This, together with the fact that belief propagation is polynomial for discrete domains with tree primal graphs, indicates that performing inference over hybrid or continuous domains with logical constrains in SMT($\mathcal{LRA}$)is inherently more difficult than that in discrete domains. The increase in complexity from discrete domains to continuous domains is not simply a matter of our inability to find good algorithms but the inherent hardness of the problem.

5 Related Work

WMI generalizes weighted model counting (WMC) [31] to hybrid domains [3]. WMC is one of the state-of-the-art approaches for inference in many discrete probabilistic models. Existing general techniques for exact WMI include DPLL-based search with numerical [3, 28, 29] or symbolic integration [12] and compilation-based algorithms [21, 39]. Motivated by its success in WMC, Belle et al. [5] presented a component caching scheme for WMI that allows to reuse cached computations at the cost of not supporting algebraic constraints between variables. Differently from usual, Merrell et al. [26] adopt Gaussian distributions, while Zuidberg Dos Martires et al. [39] fixed univariate parametric assumptions for weight functions.

Closest to our MP-MI, Search-based MI (SMI) [38] is an exact solver which leverages context-specific independence to perform efficient search. SMI recovers univariate piecewise polynomials

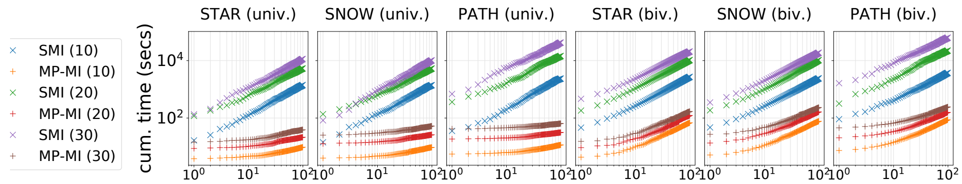

Figure 3: Log-log plot of cumulative time (seconds, y-axis) for MP-MI (orange, red) and SMI (blue, green) over STAR, SNOW and PATH primal graphs (see text) with 10, 20 and 30 variables for 100 univariate and bivariate queries (x-axis).

by interpolation while we adopt symbolic integration. As already discussed, MP-MI shares the same complexity as SMI in that its worst-case complexity is exponential in the primal graph diameter. Many recent efforts in WMI converged in the *pywmi* [22] python framework.

6 Experiments

In this section, we present a preliminary empirical evaluation to answer the following research questions: i) how does our MP-MI compare with SMI, the search-based approach to MI [38]? ii) how beneficial is amortizing multiple queries with MP-MI? We implemented MP-MI in Python 3, using the scientific computing python package *sympy* for symbolic integration, the *MathSAT5* SMT solver [9] and the *pysmt* package [16] for manipulating and representing SMT($\mathcal{LRA}$) formulas. We compare MP-MI with SMI on synthetic SMT($\mathcal{LRA}$) formulas over $n \in \{10, 20, 30\}$ variables and comprising both univariate and bivariate literals. In order to investigate the effect of adopting tree primal graphs with different diameters we considered: star-shaped graphs (STAR) with diameters two and complete ternary trees (SNOW) with diameters being $\log(n)$ and linear chains (PATH) with diameters of length n. These synthetic structures were originally investigated by the authors of SMI and are prototypical tree structures in the real world while easy to interpret due to their regularity.

Figure 3 shows the cumulative runtime of random queries that involve both univariate and bivariate literals. As expected, MP-MI takes a fraction time than SMI (up to two order of magnitudes) to answer 100 univariate or bivariate queries in all experimental scenarios, since it is able to amortize inference *inter-query*. More surprisingly, MP-MI is even faster than SMI to compute a single query. This is due to the fact that SMI solves polynomial integration numerically, by first reconstructing the univariate polynomials using interpolation e.g. Lagrange interpolation, while in MP-MI we adopt symbolical integration. Hence the complexity of the former is always quadratic in the degree of the polynomial, while for the latter the average case is linear in the number of monomials in the polynomial to integrate, which in practice might be much less than the degree of the polynomial.

7 Conclusions

In this paper, we theoretically traced the exact boundaries of tractability for MI problems. Specifically, we proved that balanced tree-shaped primal graphs are not only a sufficient condition for tractability in MI, but also a necessary one. Then we presented MP-MI, the first exact message passing algorithm for MI, which works efficiently on the aforementioned class of tractable MI problems with balanced-tree-shaped primal graphs. MP-MI also dramatically reduces the answering time of several queries including expectations and moments by amortizing computations. All these advancements suggest interesting future research venues. For instance, the efficient computation of the moments could enable the development of moment matching algorithms for approximate probabilistic inference over more challenging problems that do not admit tractable computations. Another promising direction is to perform exact inference over approximate (tree-shaped and diameter-bounded) primal graphs. Therefore, we have laid the foundations to scale hybrid probabilistic inference with logical constraints.

Acknowledgements

This work is partially supported by NSF grants #IIS-1633857, #CCF-1837129, DARPA XAI grant #N66001-17-2-4032, NEC Research, and gifts from Intel and Facebook Research.

References

[1] Clark Barrett and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of Model Checking*, pages 305–343. Springer, 2018.

[2] Clark Barrett, Leonardo de Moura, Silvio Ranise, Aaron Stump, and Cesare Tinelli. The smt-lib initiative and the rise of smt (hvc 2010 award talk). In *Proceedings of the 6th international conference on Hardware and software: verification and testing*, pages 3–3. Springer-Verlag, 2010.

[3] Vaishak Belle, Andrea Passerini, and Guy Van den Broeck. Probabilistic inference in hybrid domains by weighted model integration. In *Proceedings of 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2770–2776, 2015.

[4] Vaishak Belle, Guy Van den Broeck, and Andrea Passerini. Hashing-based approximate probabilistic inference in hybrid domains. In *UAI*, pages 141–150, 2015.

[5] Vaishak Belle, Guy Van den Broeck, and Andrea Passerini. Component caching in hybrid domains with piecewise polynomial densities. In *AAAI*, pages 3369–3375, 2016.

[6] Armin Biere, Marijn Heule, and Hans van Maaren. *Handbook of satisfiability*, volume 185. IOS press, 2009.

[7] Mark Chavira and Adnan Darwiche. On probabilistic inference by weighted model counting. 2008.

[8] Qi Cheng, Joshua Hill, and Daqing Wan. Counting value sets: algorithm and complexity. *The Open Book Series*, 1(1):235–248, 2013.

[9] Alessandro Cimatti, Alberto Griggio, Bastiaan Joost Schaafsma, and Roberto Sebastiani. The mathsat5 smt solver. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 93–107. Springer, 2013.

[10] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.

[11] Jesús A De Loera, Brandon Dutra, Matthias Koeppe, Stanislav Moreinis, Gregory Pinto, and Jianqiu Wu. Software for exact integration of polynomials over polyhedra. *Computational Geometry*, 46(3):232–252, 2013.

[12] Rodrigo de Salvo Braz, Ciaran O'Reilly, Vibhav Gogate, and Rina Dechter. Probabilistic inference modulo theories. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 3591–3599. AAAI Press, 2016.

[13] Rina Dechter and Robert Mateescu. And/or search spaces for graphical models. *Artificial intelligence*, 171(2-3):73–106, 2007.

[14] Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens, and Luc De Raedt. Inference and learning in probabilistic logic programs using weighted boolean formulas. *Theory and Practice of Logic Programming*, 15(3):358–401, 2015.

[15] Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002.

[16] Marco Gario and Andrea Micheli. Pysmt: a solver-agnostic library for fast prototyping of smt-based algorithms. In *SMT Workshop 2015*, 2015.

[17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[18] David Heckerman and Dan Geiger. Learning bayesian networks: a unification for discrete and gaussian domains. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pages 274–284. Morgan Kaufmann Publishers Inc., 1995.

[19] Steven Holtzen, Todd Millstein, and Guy Van den Broeck. Symbolic exact inference for discrete probabilistic programs. In *Proceedings of the ICML Workshop on Tractable Probabilistic Modeling (TPM)*, 2019.

[20] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[21] Samuel Kolb, Martin Mladenov, Scott Sanner, Vaishak Belle, and Kristian Kersting. Efficient symbolic integration for probabilistic inference. In *IJCAI*, pages 5031–5037, 2018.

[22] Samuel Kolb, Paolo Morettin, Pedro Zuidberg Dos Martires, Francesco Sommavilla, Andrea Passerini, Roberto Sebastiani, and Luc De Raedt. The pywmi framework and toolbox for probabilistic inference using weighted model integration. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6530–6532. International Joint Conferences on Artificial Intelligence Organization, 2019.

[23] Daphne Koller and Nir Friedman. Probabilistic graphical models. 2009.

[24] Steffen Lilholt Lauritzen and Nanny Wermuth. Graphical models for associations between variables, some of which are qualitative and some quantitative. *The annals of Statistics*, pages 31–57, 1989.

[25] Fangzhen Lin and Ray Reiter. Forget it. In *Working Notes of AAAI Fall Symposium on Relevance*, pages 154–159, 1994.

[26] David Merrell, Aws Albarghouthi, and Loris D'Antoni. Weighted model integration with orthogonal transformations. *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 2017. doi: 10.24963/ijcai.2017/643. URL http://par.nsf.gov/biblio/10061762.

[27] Alejandro Molina, Antonio Vergari, Nicola Di Mauro, Sriraam Natarajan, Floriana Esposito, and Kristian Kersting. Mixed sum-product networks: A deep architecture for hybrid domains. In *Thirty-second AAAI conference on artificial intelligence*, 2018.

[28] Paolo Morettin, Andrea Passerini, and Roberto Sebastiani. Efficient weighted model integration via smt-based predicate abstraction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 720–728. AAAI Press, 2017.

[29] Paolo Morettin, Andrea Passerini, and Roberto Sebastiani. Advanced smt techniques for weighted model integration. *Artificial Intelligence*, 275:1–27, 2019.

[30] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier, 2014.

[31] Tian Sang, Paul Beame, and Henry A Kautz. Performing bayesian inference by weighted model counting. In *AAAI*, volume 5, pages 475–481, 2005.

[32] Prakash P Shenoy and James C West. Inference in hybrid bayesian networks using mixtures of polynomials. *International Journal of Approximate Reasoning*, 52(5):641–657, 2011.

[33] Guy Van den Broeck and Dan Suciu. *Query Processing on Probabilistic Data: A Survey*. Foundations and Trends in Databases. Now Publishers, August 2017.

[34] Antonio Vergari, Alejandro Molina, Robert Peharz, Zoubin Ghahramani, Kristian Kersting, and Isabel Valera. Automatic bayesian density analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5207–5215, 2019.

[35] Jonas Vlasselaer, Wannes Meert, Guy Van den Broeck, and Luc De Raedt. Exploiting local and repeated structure in dynamic bayesian networks. *Artificial Intelligence*, 232:43 – 53, March 2016. ISSN 0004-3702.

[36] Dilin Wang, Zhe Zeng, and Qiang Liu. Stein variational message passing for continuous graphical models. In *Proceedings of the International Conference of Machine Learning*, 2018.

[37] Eunho Yang, Yulia Baker, Pradeep Ravikumar, Genevera Allen, and Zhandong Liu. Mixed graphical models via exponential families. In *Artificial Intelligence and Statistics*, pages 1042–1050, 2014.

[38] Zhe Zeng and Guy Van den Broeck. Efficient search-based weighted model integration. *Proceedings of UAI*, 2019.

[39] Pedro Miguel Zuidberg Dos Martires, Anton Dries, and Luc De Raedt. Exact and approximate weighted model integration withprobability density functions using knowledge compilation. In *Proceedings of the 30th Conference on Artificial Intelligence*. AAAI Press, 2019.

A Reduction From WMI to MI

$$\Delta = \begin{cases} \Gamma_1 : 0 < X_1 < 2 \\ \Gamma_2 : 0 < X_2 < 2 \\ \Gamma_3 : X_1 + X_2 < 2 \\ \Gamma_4 : B \vee (X_1 > 1) \end{cases} \qquad \Delta' = \begin{cases} \Gamma_1 : 0 < X_1 < 2 \\ \Gamma_2 : 0 < X_2 < 2 \\ \Gamma_3 : X_1 + X_2 < 2 \\ \Gamma_4' : (0 < Z_B < 1) \\ \quad \vee (X_1 > 1) \end{cases} \qquad \Delta'' = \Delta' \wedge \begin{cases} \Gamma_5 : 0 < Z_{X_1}' < X_1 \\ \Gamma_6 : 0 < Z_{X_2} < X_2 \\ \Gamma_7 : 0 < Z_{X_1}'' < 2 \end{cases}$$

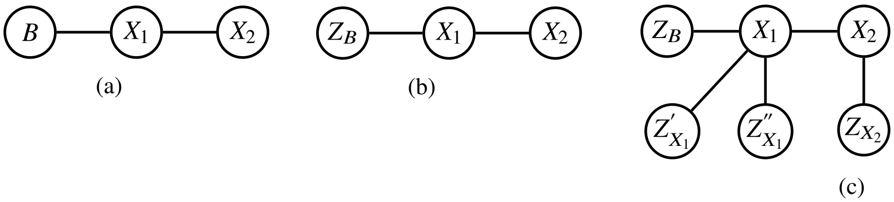

Figure 4: **From** WMI **to** MI**, passing by** $\mathsf{WMI}_{\mathbb{R}}$**.** An example of a WMI problem with an SMT($\mathcal{LRA}$) CNF formula Δ over real variables $\mathbf{X}$ and Boolean variables B and corresponding primal graph $\mathcal{G}_\Delta$ in (a). Their reductions to Δ' and $\mathcal{G}_{\Delta'}$ as an $\mathsf{WMI}_{\mathbb{R}}$ problem in (b). The equivalent MI problem with formula Δ'' and primal graph $\mathcal{G}_{\Delta''}$ over only real variables $\mathbf{X}'' = \mathbf{X} \cup \{Z_B, Z_{X_1}, Z_{X_2}\}$ after the introduction of auxiliary variables Z_B, Z_{X_1}, Z_{X_2}. Note that $\mathcal{G}_\Delta$ and $\mathcal{G}_\Delta''$ have the same treewidth one.

Figure 4 illustrates one example of a reduction of a WMI problem to one $\mathsf{WMI}_{\mathbb{R}}$ one to a MI problem. Consider the WMI problem over formula $\Delta = (0 < X_1 < 2) \wedge (0 < X_2 < 2) \wedge (X_1 + X_2 < 1) \wedge (B \vee (X_1 > 1))$ on variables $\mathbf{X} = \{X_1, X_2\}, \mathbf{B} = \{B\}$ whose primal graph $\mathcal{G}_\Delta$ is also shown in Figure 4a. Assume a weight function which decomposes as $w(X_1, X_2, B) = w_{\Gamma_3}(X_1, X_2) w_{\Gamma_4}(X_1, B) = w_{\Gamma_3}(X_1, X_2) w_{\Gamma_4}(X_1) w_{\Gamma_4}(B)$ and whose values are $w_{\Gamma_3}(X_1, X_2) = X_1 X_2$, $w_{\Gamma_4}(X_1) = 2$ and $w_{\Gamma_4}(B) = 3$ when B is true and $w(B) = 1$ otherwise. The WMI of formula Δ is:

$$\begin{aligned} \mathsf{WMI}(\Delta, w; \mathbf{X}, B) = &\int_0^1 dx_1 \int_0^{2-x_1} 1 \times 3 x_1 x_2 \, dx_2 \\ + &\int_1^2 dx_1 \int_0^{2-x_1} 2 \times 3 x_1 x_2 \, dx_2 \\ + &\int_1^2 dx_1 \int_0^{2-x_1} 2 \times 1 x_1 x_2 \, dx_2 \, . \end{aligned} \tag{9}$$

In Figure 4b, we show the reduction to the above example problem to a $\mathsf{WMI}_{\mathbb{R}}$ one. A free real variable Z_B is introduced to replace Boolean variable B. Then, the equivalent problem to the WMI one in Equation 9, can be computed as:

$$\begin{aligned} \mathsf{WMI}_{\mathbb{R}}(\Delta', w') = &\int_0^1 dz_B \int_0^1 dx_1 \int_0^{2-x_1} 1 \times 3 x_1 x_2 \, dx_2 \\ + &\int_0^1 dz_B \int_1^2 dx_1 \int_0^{2-x_1} 2 \times 3 x_1 x_2 \, dx_2 \\ + &\int_{-1}^0 dz_B \int_1^2 dx_1 \int_0^{2-x_1} 2 \times 1 x_1 x_2 \, dx_2 \, . \end{aligned} \tag{10}$$

Figure 4c illustrates the additional reduction from the above $\mathsf{WMI}_{\mathbb{R}}$ problem to a MI one. There, additional real variables Z_{X_1}', Z_{X_2} and Z_{X_1}'' are added to formula Δ'' in substitution of the monomial weights attached to literal Γ_3 and Γ_4, respectively. Therefore, the same result as Equation 9 and Equation 10 can be obtained as

$$\mathsf{MI}(\Delta'') = \int_0^3 dz_B \int_0^1 dx_1 \int_0^{2-x_1} dx_2 \int_0^{x_1} dz'_{X_1} \int_0^{x_2} dz_{X_2} \tag{11}$$
$$+ \int_0^2 dz''_{X_1} \int_0^3 dz_B \int_1^2 dx_1 \int_0^{2-x_1} dx_2 \int_0^{x_1} dz'_{X_1} \int_0^{x_2} dz_{X_2}$$
$$+ \int_0^2 dz''_{X_1} \int_1^2 dx_1 \int_0^{2-x_1} dx_2 \int_0^{x_1} dz'_{X_1} \int_0^{x_2} dz_{X_2} \; .$$

B Proofs

B.1 THEOREM 1 (MI of a formula with tree primal graph with unbounded diameter is #P-Hard)

Proof. (Theorem 1) We prove our complexity result by reducing a #P-complete variant of the subset sum problem [15] to an MI problem over an SMT($\mathcal{LRA}$) formula Δ with tree primal graph whose diameter is $O(n)$. This problem is a counting version of subset sum problem saying that given a set of positive integers $S = \{s_1, s_2, \cdots, s_n\}$, and a positive integer L, the goal is to count the number of subsets $S' \subseteq S$ such that the sum of all the integers in the subset S' equals to L.

First, we reduce the counting subset sum problem in polynomial time to a model integration problem by constructing the following SMT($\mathcal{LRA}$) formula Δ on real variables $\mathbf{X}$ whose primal graph is shown in Figure 1:

$$\Delta = \begin{cases} s_1 - \frac{1}{2n} < X_1 < s_1 + \frac{1}{2n} \vee -\frac{1}{2n} < X_1 < \frac{1}{2n} \\ X_{i-1} + s_i - \frac{1}{2n} < X_i < X_{i-1} + s_i + \frac{1}{2n} \;\vee\; X_{i-1} - \frac{1}{2n} < X_i < X_{i-1} + \frac{1}{2n}, \quad i = 2, \cdots n \end{cases}$$

For brevity, we denote the first and the second literal in the i-th clause by $\ell(i, 0)$ and $\ell(i, 1)$ respectively. Also We choose two constants $l = L - \frac{1}{2}$ and $u = L + \frac{1}{2}$.

In the following, we prove that $n^n \mathsf{MI}(\Delta \wedge (l < X_n < u))$ equals to the number of subset $S' \subseteq S$ whose element sum equals to L, which indicates that model integration problem whose tree primal graph has diameter $O(n)$ is #P-hard.

Let $\boldsymbol{a}^k = (a_1, a_2, \cdots, a_k)$ be some assignment to Boolean variables $(A_1, A_2, \cdots, A_k)$ with $a_i \in \{0, 1\}$, $i \in [k]$. Given an assignment $\boldsymbol{a}^k$, we define subset sums to be $S(\boldsymbol{a}^k) \triangleq \sum_{i=1}^k a_i s_i$, and formulas $\Delta_{\boldsymbol{a}^k} \triangleq \bigwedge_{i=1}^k \ell(i, a_i)$.

Claim 11. *The model integration for formula $\Delta_{\boldsymbol{a}^k}$ with an given assignment $\boldsymbol{a}^k \in \{0,1\}^k$ is $\mathsf{MI}(\Delta_{\boldsymbol{a}^k}) = (\frac{1}{n})^k$. Moreover, for each variable X_i in $\Delta_{\boldsymbol{a}^k}$, its satisfying assignments consist of the interval $[\sum_{j=1}^i a_j s_j - \frac{i}{2n}, \sum_{j=1}^i a_j s_j + \frac{i}{2n}]$. Specifically, the satisfying assignments for variable X_n in formula $\Delta_{\boldsymbol{a}^n}$ can be denoted by the interval $[S(\boldsymbol{a}^n) - \frac{1}{2}, S(\boldsymbol{a}^n) + \frac{1}{2}]$.*

Proof. (Claim 11) First we prove that $\mathsf{MI}(\Delta_{\boldsymbol{a}^k}) = (\frac{1}{n})^k$. For brevity, denote $a_i s_i$ by $\hat{s}_i$. By definition of model integration and the fact that the integral is absolutely convergent (since we are integrating a constant function, i.e., one, over finite volume regions), we have the following equation.

$$\mathsf{MI}(\Delta_{\boldsymbol{a}^k}) = \int_{(x_1,\cdots,x_k)\models\Delta_{\boldsymbol{a}^k}} 1 \; dx_1 \cdots dx_k = \int_{\hat{s}_1 - \frac{1}{2n}}^{\hat{s}_1 + \frac{1}{2n}} dx_1 \cdots \int_{x_{k-2}+\hat{s}_{k-1}-\frac{1}{2n}}^{x_{k-2}+\hat{s}_{k-1}+\frac{1}{2n}} dx_{k-1} \int_{x_{k-1}+\hat{s}_k-\frac{1}{2n}}^{x_{k-1}+\hat{s}_k+\frac{1}{2n}} 1 \; dx_k$$

Observe that for the most inner integration over variable x_k, the integration result is $\frac{1}{n}$. By doing this iteratively, we have that $\mathsf{MI}(\Delta_{\boldsymbol{a}^k}) = (\frac{1}{n})^k$.

Next we prove that satisfying assignments for variable X_i in formula $\Delta_{\boldsymbol{a}^k}$ is the interval $[\sum_{j=1}^i a_j s_j - \frac{i}{2n}, \sum_{j=1}^i a_j s_j + \frac{i}{2n}]$ by mathematical induction. For $i = 1$, since X_1 is in interval $[a_1 s_1 - \frac{1}{2n}, a_1 s_1 + \frac{1}{2n}]$, the statement holds in this case. Suppose that the statement holds for $i = m$, i.e. variable X_m has its satisfying assignments in interval $[\sum_{j=1}^m a_j s_j - \frac{m}{2n}, \sum_{j=1}^m a_j s_j + \frac{m}{2n}]$. Since variable X_{m+1} has its satisfying assignments in interval $[X_m + a_{m+1}s_{m+1} - \frac{1}{2n}, X_m + a_{m+1}s_{m+1} + \frac{1}{2n}]$, then its satisfying assignments consist interval $[\sum_{j=1}^{m+1} a_j s_j - \frac{m+1}{2n}, \sum_{j=1}^{m+1} a_j s_j + \frac{m+1}{2n}]$, that is, the statement also holds for $i = m + 1$. Thus the statement holds.

□

The above claim shows how to compute the model integration of formula $\Delta_{\boldsymbol{a}^k}$. We will show in the next claim how to compute the model integration of formula $\Delta_{\boldsymbol{a}^n}$ conjoined with a query $l < X_n < u$.

Claim 12. *For each assignment $\boldsymbol{a}^n \in \{0, 1\}^n$, the model integration of formula $\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)$ falls into one of the following cases:*

- *If $S(\boldsymbol{a}^n) < L$ or $S(\boldsymbol{a}^n) > L$, it holds that $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = 0$.*
- *If $S(\boldsymbol{a}^n) = L$, it holds that $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = (\frac{1}{n})^n$.*

Proof. (Claim 12) From the previous Claim 11, it is shown that variable X_n has its satisfying assignments in interval $[S(\boldsymbol{a}^n) - \frac{1}{2}, S(\boldsymbol{a}^n) + \frac{1}{2}]$ in formula $\Delta_{\boldsymbol{a}^n}$ for each $\boldsymbol{a}^n \in \{0, 1\}^n$. If $S(\boldsymbol{a}^n) < L$, given that $S(\boldsymbol{a}^n)$ is a sum of positive integers, then it holds that $S(\boldsymbol{a}^n) + \frac{1}{2} \leq (L - 1) + \frac{1}{2} = L - \frac{1}{2} = l$ and therefore, $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = 0$; similarly, if $S(\boldsymbol{a}^n) > L$, then it holds that $S(\boldsymbol{a}^n) - \frac{1}{2} \geq u$ and therefore, $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = 0$. If $S(\boldsymbol{a}^n) = L$, by Claim 11 we have that the satisfying assignment interval is inside the interval $[l, u]$ and thus it holds that $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)) = \mathsf{MI}(\Delta_{\boldsymbol{a}^n}) = (\frac{1}{n})^n$. □

In the next claim, we show how to compute the model integration of formula Δ as well as for formula Δ conjoined with query $l < X_n < u$ based on the already proven Claim 11 and Claim 12.

Claim 13. *The following two equations hold:*

1. $\mathsf{MI}(\Delta) = \sum_{\boldsymbol{a}^n} \mathsf{MI}(\Delta_{\boldsymbol{a}^n})$.
2. $\mathsf{MI}(\Delta \wedge (l < X_n < u)) = \sum_{\boldsymbol{a}^n} \mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u))$.

Proof. (Claim 13) Observe that for each clause in Δ, literals are mutually exclusive since each s_i is a positive integer. Then we have that formulas $\Delta_{\boldsymbol{a}^n}$ are mutually exclusive and meanwhile $\Delta = \bigvee_{\boldsymbol{a}^n} \Delta_{\boldsymbol{a}^n}$. Thus it holds that $\mathsf{MI}(\Delta) = \sum_{\boldsymbol{a}^n} \mathsf{MI}(\Delta_{\boldsymbol{a}^n})$. Similarly, we have formulas $(\Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u))$'s are mutually exclusive and meanwhile $\Delta \wedge (l < X_n < u) = \bigvee_{\boldsymbol{a}^n} \Delta_{\boldsymbol{a}^n} \wedge (l < X_n < u)$. Thus the second equation holds. □

From the above claims, we can conclude that $\mathsf{MI}(\Delta \wedge (l < X_n < u)) = t(\frac{1}{n})^n$ where t is the number of assignments $\boldsymbol{a}^n$ s.t. $S(\boldsymbol{a}^n) = L$. Notice that for each $\boldsymbol{a}^n \in \{0, 1\}^n$, there is a one-to-one correspondance to a subset $S' \subseteq S$ by defining $\boldsymbol{a}^n$ as $a_i = 1$ if and only if $s_i \in S'$; and $S(\boldsymbol{a}^n)$ equals to L if and only if the sum of elements in $S'$is L. Therefore $n^n \mathsf{MI}(\Delta \wedge (l < X_n < u))$ equals to the number of subset $S' \subseteq S$ whose element sum equals to L.

This finishes the proof for the statement that a model integration problem whose tree primal graph has diameter $O(n)$ is #P-hard.

□

B.2 THEOREM 2 (MI of a formula with primal graph with logarithmic diameter and treewidth two is #P-Hard)

Proof. (Theorem 2) Again we prove our complexity result by reducing the #P-complete variant of the subset sum problem [15] to an MI problem over an SMT($\mathcal{LRA}$) formula Δ with primal graph whose diameter is $O(\log n)$ and treewidth two. In the #P-complete subset sum problem, we are given a set of positive integers $S = \{s_1, s_2, \cdots, s_n\}$, and a positive integer L. The goal is to count the number of subsets $S' \subseteq S$ such that the sum of all the integers in S' equals L.

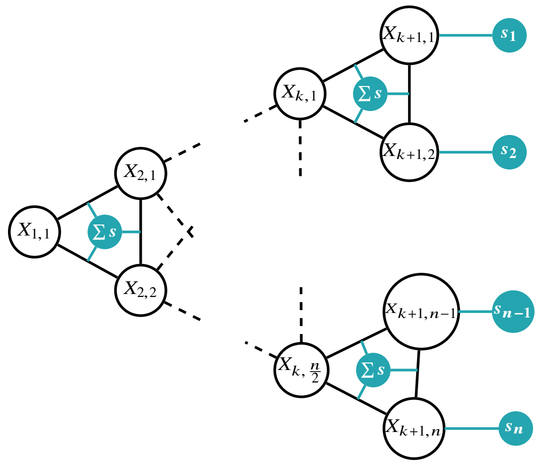

Figure 5: Primal graph used for #P-hardness reduction in Theorem 7. We also put blue nodes to indicate that integer s_i's in set S are contained in some clauses and that model integration over some cliques is the sum of some s_i's.

First, we reduce this problem in polynomial time to a model integration problem with the following SMT($\mathcal{LRA}$) formula Δ where variables are real and u and l are two constants. Its primal graph is shown in Figure 5. Consider $n = 2^k$, $n, k \in \mathbb{N}$.

$$\Delta = \bigwedge_{i \in [n]} (-\frac{1}{4n} < X_{k+1,i} < \frac{1}{4n} \vee -\frac{1}{4n} + s_i < X_{k+1,i} < \frac{1}{4n} + s_i) \bigwedge \Delta_t$$

$$\text{where } \Delta_t = \bigwedge_{j \in [k], i \in [2^j]} -\frac{1}{4n} + X_{j+1,2i-1} + X_{j+1,2i} < X_{j,i} < \frac{1}{4n} + X_{j+1,2i-1} + X_{j+1,2i}$$

For brevity, we denote all the variables by $\mathbf{X}$ and denote the literal $-\frac{1}{4n} < X_{k+1,i} < \frac{1}{4n}$ by $\ell(i,0)$ and literal $-\frac{1}{4n} + s_i < X_{k+1,i} < \frac{1}{4n} + s_i$ by $\ell(i,1)$ respectively. Also We choose two constants $l = L - \frac{1}{2}$ and $u = L + \frac{1}{2}$. In the following, we prove that $(2n)^{2n-1}\mathsf{MI}(\Delta \wedge (l < X_{1,1} < u))$ equals to the number of subset $S' \subseteq S$ whose element sum equals to L, which indicates that model integration problem with primal graph whose diameter is $O(\log n)$ and treewidth two is #P-hard.

Let $\boldsymbol{a}^n = (a_1, a_2, \cdots, a_n) \in \{0,1\}^n$ be some assignment to Boolean variables $(A_1, A_2, \cdots, A_n)$. Given an assignment $\boldsymbol{a}^n$, define the sum as $S(\boldsymbol{a}^n) \triangleq \sum_{i=1}^n a_i s_i$, and formula as $\Delta_{\boldsymbol{a}^n} \triangleq \bigwedge_{i=1}^n \ell(i, a_i) \wedge \Delta_t$.

Claim 14. *The model integration for formula $\Delta_{\boldsymbol{a}^n}$ with given $\boldsymbol{a}^n \in \{0,1\}^n$ is $\mathsf{MI}(\Delta_{\boldsymbol{a}^n}) = (\frac{1}{2n})^{2n-1}$. Moreover, for each variable $X_{j,i}$ in formula $\Delta_{\boldsymbol{a}^n}$, its satisfying assignments consist of the interval $[\sum_l a_l s_l - \frac{2^{k-j+2}-1}{4n}, \sum_l a_l s_l + \frac{2^{k-j+2}-1}{4n}]$ where $l \in \{l \mid X_{k+1,l}$ is a descendant of $X_{j,i}\}$. Specifically, the satisfying assignments for the root variable $X_{1,1}$ can be denoted the interval $[S(\boldsymbol{a}^n) - \frac{2n-1}{4n}, S(\boldsymbol{a}^n) + \frac{2n-1}{4n}] \subset [S(\boldsymbol{a}^n) - \frac{1}{2}, S(\boldsymbol{a}^n) + \frac{1}{2}]$.*

Proof. (Claim 14)

First we prove that $\mathsf{MI}(\Delta_{\boldsymbol{a}^n}) = (\frac{1}{2n})^{2n-1}$. For brevity, denote $a_i s_i$ by $\hat{s}_i$. By definition of model integration and the fact that the integral is absolutely convergent (since we are integrating a constant function, i.e., one, over finite volume regions), we have the following equations

$$\begin{aligned}\mathsf{MI}(\Delta_{\boldsymbol{a}^n}) &= \int_{\boldsymbol{x} \models \Delta_{\boldsymbol{a}^n}} 1 \, d\mathbf{X} \\ &= \int_{-\frac{1}{4n}+\hat{s}_n}^{\frac{1}{4n}+\hat{s}_n} dx_{k+1,n} \cdots \int_{-\frac{1}{4n}+\hat{s}_1}^{\frac{1}{4n}+\hat{s}_1} dx_{k+1,1} \int_{-\frac{1}{4n}+x_{k+1,n-1}+x_{k+1,n}}^{\frac{1}{4n}+x_{k+1,n-1}+x_{k+1,n}} dx_{k,2^{k-1}} \cdots \int_{-\frac{1}{4n}+x_{2,1}+x_{2,2}}^{\frac{1}{4n}+x_{2,1}+x_{2,2}} 1 \, dx_{1,1} \ .\end{aligned}$$

Observe that for the most inner integration over variable $x_{1,1}$, the integration result is $\frac{1}{2n}$. By doing this iteratively, we have that $\mathsf{MI}(\Delta_{\boldsymbol{a}^k}) = (\frac{1}{2n})^{2n-1}$ where the $2n-1$ comes from the number of variables.

Then we prove that satisfying assignments for variable $X_{j,i}$ in formula $\Delta_{\boldsymbol{a}^n}$ lie in the interval $[\sum_l a_l s_l - \frac{2^{k-j+2}-1}{4n}, \sum_l a_l s_l + \frac{2^{k-j+2}-1}{4n}]$ where $l \in \{l \mid X_{k+1,l}$ is a descendant of $X_{j,i}\}$ by performing mathematical induction in a bottom-up way.

For $j = 1$, any variable $X_{k+2-j,i}$ with $i \in [2^{k+2-j}]$ has satisfying assignments consisting of the interval $[a_i s_i - \frac{1}{4n}, a_i s_i + \frac{1}{4n}]$. Thus the statement holds for this case.

Suppose that the statement holds for $j = m$, that is, for any $i \in [2^{k+2-m}]$, any variable $X_{k+2-m,i}$ has satisfying assignments consisting interval $[\sum_l a_l s_l - \frac{2^m-1}{4n}, \sum_l a_l s_l + \frac{2^m-1}{4n}]$ where $l \in \{l \mid X_{k+1,l}$ is a descendant of $X_{k+2-m,i}\}$.

Then for $j = m+1$ and any $i \in [2^{k+1-m}]$, the variable $X_{k+1-m,i}$ has two descendants, variable $X_{k+2-m,2i-1}$ and variable $X_{k+2-m,2i}$. Moreover, we have that $-\frac{1}{4n} + X_{k+2-m,2i-1} + X_{k+2-m,2i} < X_{k+1-m,i} < \frac{1}{4n} + X_{k+2-m,2i-1} + X_{k+2-m,2i}$. Then the lower bound of the interval for variable $X_{k+1-m,i}$ is $-\frac{1}{4n} + \sum_l a_l s_l - 2\frac{2^m-1}{4n} = \sum_l a_l s_l - \frac{2^{m+1}-1}{4n}$; similarly the upper bound of the interval is $\sum_l a_l s_l + \frac{2^{m+1}-1}{4n}$, where $l \in \{l \mid X_{k+1,l}$ is a descendant of $X_{k+1-m,i}\}$. That is, the statement also holds for $j = m+1$ which finishes our proof. □

The above claim shows what the model integration of formula $\Delta_{\boldsymbol{a}^k}$ is like. We'll show in the next claim what the model integration of formula $\Delta_{\boldsymbol{a}^n}$ conjoined with a query $l < X_{1,1} < u$ is like.

Claim 15. *For each assignments $\boldsymbol{a}^n \in \{0,1\}^n$, the model integration of $\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u)$ falls into one of the following cases:*

- *If $S(\boldsymbol{a}^n) < L$ or $S(\boldsymbol{a}^n) > L$, then* $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u)) = 0$.
- *If $S(\boldsymbol{a}^n) = L$, then* $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u)) = (\frac{1}{2n})^{2n-1}$.

Proof. (Claim 15) From previous Claim 14, it is shown that variable $X_{1,1}$ has its satisfying assignments in the interval $[S(\boldsymbol{a}^n) - \frac{2n-1}{4n}, S(\boldsymbol{a}^n) + \frac{2n-1}{4n}]$ in formula $\Delta_{\boldsymbol{a}^n}$ for each $\boldsymbol{a}^n \in \{0,1\}^n$.

If $S(\boldsymbol{a}^n) < L$, given that $S(\boldsymbol{a}^n)$ is a sum of positive integers, then it holds that $S(\boldsymbol{a}^n) + \frac{1}{2} \leq (L-1) + \frac{2n-1}{4n} < L - \frac{1}{2} = l$ and therefore, $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u)) = 0$; similarly, if $S(\boldsymbol{a}^n) > L$, then it holds that $S(\boldsymbol{a}^n) - \frac{1}{2} > u$ and therefore, $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u)) = 0$. If $S(\boldsymbol{a}^n) = L$, then by Claim 14 we have that the satisfying assignment interval is inside the interval $[l, u]$ and thus it holds that $\mathsf{MI}(\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u)) = \mathsf{MI}(\Delta_{\boldsymbol{a}^n}) = (\frac{1}{2n})^{2n-1}$. □

Claim 16. *The following two equations hold:*

1. $\mathsf{MI}(\Delta) = \sum_{\boldsymbol{a}^n} \mathsf{MI}(\Delta_{\boldsymbol{a}^n})$.
2. $\mathsf{MI}(\Delta \wedge (l < X_{1,1} < u)) = \sum_{\boldsymbol{a}^n} MI(\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u))$.

Proof. (Claim 16) Observe that for each pair of literals $\ell(i,0)$ and $\ell(i,1), i \in [n]$, literals are mutually exclusive since each s_i is a positive integer. Then we have that formulas $\Delta_{\boldsymbol{a}^n}$ are mutually exclusive and meanwhile formula $\Delta = \bigvee_{\boldsymbol{a}^n} \Delta_{\boldsymbol{a}^n}$. Thus it holds that $\mathsf{MI}(\Delta) = \sum_{\boldsymbol{a}^n} \mathsf{MI}(\Delta_{\boldsymbol{a}^n})$. Similarly, we have formulas $(\Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u))$'s are mutually exclusive and meanwhile $\Delta \wedge (l < X_{1,1} < u) = \bigvee_{\boldsymbol{a}^n} \Delta_{\boldsymbol{a}^n} \wedge (l < X_{1,1} < u)$. Thus the second equation holds. □

From the above claims, we can conclude that $\mathsf{MI}(\Delta \wedge (l < X_{1,1} < u)) = t(\frac{1}{2n})^{2n-1}$ where t is the number of assignments $\boldsymbol{a}^n$ s.t. $S(\boldsymbol{a}^n) = L$. Notice that for each $\boldsymbol{a}^n \in \{0,1\}^n$, there is a one-to-one correspondence to a subset $S' \subseteq S$ by defining $\boldsymbol{a}^n$ as $a_i = 1$ if and only if $s_i \in S'$; and $S(\boldsymbol{a}^n)$ equals to L if and only if the sum of elements in S' is L. Therefore $(2n)^{2n-1}\mathsf{MI}(\Delta \wedge (l < X_{1,1} < u))$ equals to the number of subset $S' \subseteq S$ whose element sum equals to L.

This finishes the proof for the statement that a model integration problem with primal graph whose diameter is $O(\log n)$ and treewidth two is #P-hard. □

B.3 PROPOSITION 3 (MI via message passing)

Proof. (Proposition 3)

By the definition of downward pass beliefs and messages, we have that the downward pass belief b_{i^*} of a node i^* can be written as follows

$$\begin{aligned}\mathsf{b}_{i^*}(x_{i^*}) &= \prod_{j\in\mathsf{neigh}(i^*)} \mathsf{m}_{j\to i^*}(x_{i^*}) = \prod_{j\in\mathsf{neigh}(i^*)} \int_{\mathbb{R}} [\![x_{i^*}, x_j \models \Delta_{i^*,j}]\!][\![x_j \models \Delta_j]\!] \prod_{c\in\mathsf{neigh}(j)\setminus\{i^*\}} \mathsf{m}_{c\to j}(x_j)\, dx_j \\ &= \int_{\mathbb{R}^{|\mathsf{neigh}(i^*)|}} \prod_{(i^*,j)\in\mathcal{E}} [\![x_{i^*}, x_j \models \Delta_{i^*,j}]\!][\![x_j \models \Delta_j]\!] \prod_{(j,c)\in\mathcal{E}, c\neq i^*} \mathsf{m}_{c\to j}(x_j)\, d\boldsymbol{x}_{i^*} ,\end{aligned}$$

where the last equality comes from interchanging integration with product, and $\boldsymbol{x}_{i^*}$ is defined as $\boldsymbol{x}_{i^*} = \{x_j \mid (i^*, j) \in \mathcal{E}\}$. By doing this recursively, i.e. plugging in the messages as defined in Equation 7, the belief of node i^* can be expressed as follows

$$\mathsf{b}_{i^*}(x_{i^*}) = \int_{\mathbb{R}^{|\mathbf{X}|-1}} \prod_{(i^*,j)\in\mathcal{E}} [\![x_{i^*}, x_j \models \Delta_{i^*,j}]\!] \prod_{j\in\mathcal{V}\setminus\{i^*\}} [\![x_j \models \Delta_j]\!]\, d\boldsymbol{x} \setminus \{x_{i^*}\} = \int_{\mathbb{R}^{|\mathbf{X}|-1}} [\![\boldsymbol{x} \models \Delta]\!]\, d\boldsymbol{x} \setminus \{x_{i^*}\} .$$

The last equality comes from the fact that the formula Δ has a tree primal graph $\mathcal{G}_\Delta$, i.e. $\Delta = \wedge_{(i,j)\in\mathcal{E}}\Delta_{i,j} \wedge_{i\in\mathcal{V}} \Delta_i$. Recall the definition of MI as defined in Equation 6, we have that the final belief b_{i^*} is the unnormalized marginal of variable $X_{i^*} \in \mathbf{X}$, i.e. $\mathsf{b}_{i^*}(x_{i^*}) = p_\Delta(x_{i^*}) \cdot \mathsf{MI}(\Delta)$. Besides, this also indicates that the integration over the belief of X_{i^*} is equal to the MI of formula Δ. □

B.4 PROPOSITION 4 (Messages and beliefs)

Proof. (Proposition 4)

This follows by induction on both the level of the node and the number of its neighbors. Consider the base case of a node i with only one neighbor j being the leaf node. Then the message sent from node j to node i would be $\mathsf{m}_{i\to j}(x_j) = \int_{\mathbb{R}} [\![x_i, x_j \models \Delta_{i,j}]\!][\![x_i \models \Delta_i]\!]\, dx_i$. This integral has one as an integrand over pieces that satisfy the logical constraints $[\![x_i, x_j \models \Delta_{i,j}]\!][\![x_i \models \Delta_i]\!]$ with integration bounds linear in variable x_j. Therefore the resulting message from node j to node i is a piecewise linear function in variable x_j. Since node i has only one child by assumption, its upward-pass belief is also piecewise univariate polynomial.

From here, the proof follows for any message and belief for more complex tree structures by considering that the piecewise polynomial family is closed under multiplication and integration. □

B.5 PROPOSITION 5 (Univariate queries via message passing)

Proof. (Proposition 5)

For an SMT($\mathcal{LRA}$) query Φ over a variable $X_i \in \mathbf{X}$, the MI over formula Δ conjoined with query Φ can be expressed as follows by the definition of model integration.

$$\mathsf{MI}(\Delta \wedge \Phi) = \int_{\mathbb{R}^{|\mathbf{X}|}} [\![\boldsymbol{x} \models \Delta \wedge \Phi]\!]\, d\boldsymbol{x} = \int_{\mathbb{R}^{|\mathbf{X}|}} [\![x_i \models \Phi]\!][\![\boldsymbol{x} \models \Delta]\!]\, d\boldsymbol{x} \setminus \{x_i\} dx_i .$$

Notice that by the proof of Proposition 2, we have that the downward pass belief of node i is $\mathsf{b}_i(x_i) = \int_{\mathbb{R}^{|\mathbf{X}|-1}} [\![\boldsymbol{x} \models \Delta]\!]\, d\boldsymbol{x} \setminus \{x_i\}$. By plugging the belief b_i in the above equation of MI over formula $\Delta \wedge \Phi$, we have that

$$\mathsf{MI}(\Delta \wedge \Phi) = \int_{\mathbb{R}} [\![x_i \models \Phi]\!][\![x_i \models \Delta_i]\!]\mathsf{b}_i(x_i) dx_i .$$

□

B.6 PROPOSITION 6 (Bivariate queries via message passing)

Proof. (Proposition 6)

Denote the SMT($\mathcal{LRA}$) formula $\Delta \wedge \Phi$ by Δ^* where Φ is an SMT($\mathcal{LRA}$) query over variables $X_{i^*}, X_{j^*} \in \mathbf{X}$. We also denote the belief and messages in formula Δ^* by b_i^* and $\mathsf{m}_{i\to j}^*$ respectively.

Notice that since query Φ is defined over variables X_{i^*}, X_{j^*}, then it holds that for any $(i,j) \in \mathcal{E}$, $\Delta_{i,j} = \Delta_{i,j}^*$ if $(i,j) \neq (i^*, j^*)$; else $\Delta_{i^*,j^*}^* = \Delta_{i^*,j^*} \wedge \Phi$. Also for any $i \in \mathcal{V}$, it holds that $\Delta_i = \Delta_i^*$. Therefore, we have that $\mathsf{b}_{j^*}^*(x_{j^*})/\mathsf{m}_{i^*\to j^*}^* = \mathsf{b}_{j^*}(x_{j^*})/\mathsf{m}_{i^*\to j^*}$ by the definition of beliefs and messages. Moreover, we can compute the message sent from node j^* to node i^* in formula Δ^* as follows:

$$\begin{aligned}
\mathsf{m}_{j^*\to^* i}^*(x_{i^*}) &= \int_{\mathbb{R}} \mathsf{b}_{j^*}^*(x_{j^*})/\mathsf{m}_{i^*\to j^*}^*(x_{j^*}) \times [\![x_{i^*}, x_{j^*} \models \Delta_{i^*,j^*}^*]\!][\![x_{j^*} \models \Delta_{j^*}]\!]\, dx_{j^*} \\
&= \int_{\mathbb{R}} \mathsf{b}_{j^*}(x_{j^*})/\mathsf{m}_{i^*\to j^*}(x_{j^*}) \times [\![x_{i^*}, x_{j^*} \models \Delta_{i^*,j^*} \wedge \Phi]\!][\![x_{j^*} \models \Delta_j]\!]\, dx_{j^*}.
\end{aligned}$$

Similarly, we have that the final belief on node i^* is as follows:

$$\mathsf{b}_{i^*}^*(x_{i^*}) = \prod_{j \in \mathsf{neigh}(i^*)} \mathsf{m}_{j\to i^*}^*(x_{i^*}) = \mathsf{m}_{i^*\to j^*}^*(x_{i^*}) \prod_{j \in \mathsf{neigh}(i^*), j \neq j^*} \mathsf{m}_{j\to i^*}(x_{i^*}).$$

Then the MI over formula $\Delta \wedge \Phi$ can be computed by doing $\mathsf{MI}(\Delta \wedge \Phi) = \int_{\mathbb{R}} [\![x_{i^*} \models \Delta_{i^*}]\!] \cdot \mathsf{b}_{i^*}^*(x_{i^*})\, dx_{i^*}$ where messages except $\mathsf{m}_{i^*\to j^*}^*$ are pre-computed and the computation of the message $\mathsf{m}_{i^*\to j^*}^*$ can reuse the pre-computed beliefs as shown above.

□

B.7 PROPOSITION 7 (Statistical moments via message passing)

Proof. (Proposition 7)

By the definition of the k-th moment of the random variables and Proposition 2 that belief b_i of node i is the unnormalized marginal $p_i(x_i)$ of variable $X_i \in \mathbf{X}$, we have that

$$\mathbb{E}[X_i^k] = \int_{\mathbb{R}} [\![x_i \models \Delta_i]\!] \times x_i^k p_\Delta(x_i)\, dx_i = \frac{1}{\mathsf{MI}(\Delta)} \int_{\mathbb{R}} [\![x_i \models \Delta_i]\!] \times x_i^k \mathsf{b}_i(x_i)\, dx_i \ .$$

□