


Scalable Precise Computation of Shannon Entropy

Yong Lai 


Key Laboratory of Symbolic Computation and Knowledge Engineering Ministry of Education,
Jilin University, Changchun, China

Haolong Tong 

College of Computer Science and Technology, Jilin University, Changchun, China

Zhenghang Xu 

School of Computer Science and Information Technology, Northeast Normal University,
Changchun, China

Minghao Yin 

School of Computer Science and Information Technology, Northeast Normal University,
Changchun, China

Abstract

Quantitative information flow analyses (QIF) are a class of techniques for measuring the amount of confidential information leaked by a program to its public outputs. Shannon entropy is an important method to quantify the amount of leakage in QIF. This paper focuses on the programs modeled in Boolean constraints and optimizes the two stages of the Shannon entropy computation to implement a scalable precise tool PSE. In the first stage, we design a knowledge compilation language called $\text{ADD}[\wedge]$ that combines Algebraic Decision Diagrams and conjunctive decomposition. $\text{ADD}[\wedge]$ avoids enumerating possible outputs of a program and supports tractable entropy computation. In the second stage, we optimize the model counting queries that are used to compute the probabilities of outputs. We compare PSE with the state-of-the-art probabilistic approximately correct tool EntropyEstimation, which was shown to significantly outperform the previous precise tools. The experimental results demonstrate that PSE solved 56 more benchmarks compared to EntropyEstimation in a total of 459. For 98% of the benchmarks that both PSE and EntropyEstimation solved, PSE is at least $10\times$ as efficient as EntropyEstimation.

2012 ACM Subject Classification Theory of computation \rightarrow Constraint and logic programming

Keywords and phrases Knowledge Compilation, Algebraic Decision Diagrams, Quantitative Information Flow, Shannon Entropy

Digital Object Identifier 10.4230/LIPIcs.SAT.2025.20

Supplementary Material *Software (Source Code)*: <https://github.com/laigroup/PSE>
archived at `swb:1:dir:827b579f6c2ef3a9a79de5d7aaf910102c8ac07c`

Funding This work was supported in part by Jilin Provincial Natural Science Foundation [20240101378JC], Jilin Provincial Education Department Research Project [JJKH20241286KJ], and the National Natural Science Foundation of China [U22A2098, 62172185, and 61976050].

Acknowledgements The authors thank the anonymous reviewers for their constructive feedback.

1 Introduction

Quantitative information flow (QIF) is an important approach to measuring the amount of information leaked about a secret by observing the running of a program [11, 16]. In QIF, we often quantify the leakage using entropy-theoretic notions, such as Shannon entropy [2, 5, 30, 33] or min-entropy [2, 29, 30, 33]. Roughly speaking, a program in QIF can be seen as a

Authors are listed alphabetically by last name. Corresponding author: Minghao Yin



© Yong Lai, Haolong Tong, Zhenghang Xu, and Minghao Yin;
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Theory and Applications of Satisfiability Testing (SAT 2025).

Editors: Jeremias Berg and Jakob Nordström; Article No. 20; pp. 20:1–20:19

Leibniz International Proceedings in Informatics



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

function from a set of secret inputs X to outputs Y observable to an attacker who may try to infer X based on the output Y . Boolean formulas are a basic representation to model programs [14, 15]. In this paper, we focus on precisely computing the Shannon entropy of a program expressed in Boolean formulas.

Let $\varphi(X, Y)$ be a (Boolean) formula that models the relationship between the input variable set X and the output variable set Y in a given program, such that for any assignment of X , at most one assignment of Y satisfies the formula $\varphi(X, Y)$. Let p represent a probability distribution defined over the set $\{\text{false}, \text{true}\}^Y$. For each assignment σ to Y , the probability is defined as $p_\sigma = \frac{|Sol(\varphi(Y \mapsto \sigma))|}{|Sol(\varphi)_{\downarrow X}|}$, where $Sol(\varphi(Y \mapsto \sigma))$ denotes the set of solutions of $\varphi(Y \mapsto \sigma)$ and $Sol(\varphi)_{\downarrow X}$ denotes the set of solutions of φ projected to X . The Shannon entropy of φ is $H(\varphi) = \sum_{\sigma \in 2^Y} -p_\sigma \log p_\sigma$. Then we can immediately obtain a measure of leaked information with the computed entropy and the assumption that X follows a uniform distribution¹ [19].

The workflow of existing precise methods for computing entropy can often be divided into two stages. In the first stage, we enumerate possible outputs, i.e., the satisfying assignments over Y , while in the second stage, we compute the probability of the current output based on the number of inputs mapped to the output [15]. The computation in the second stage often invokes model counting ($\#SAT$), which refers to computing the number of solutions $Sol(\varphi)$ for a given formula φ . Due to the exponential number of possible outputs, the current precise methods are often difficult to scale to programs with a large size of Y . Therefore, researchers have increasingly focused on approximate estimation of Shannon entropy. We remark that Golia et al. [15] proposed the first Shannon entropy estimation tool, EntropyEstimation, which guarantees that the estimate lies within $(1 \pm \epsilon)$ -factor of $H(\varphi)$ with confidence at least $1 - \delta$. EntropyEstimation employs uniform sampling to avoid generating all outputs, and indeed scales much better than the precise methods.

As previously discussed, existing methods for precisely computing Shannon entropy struggle to scale when applied to formulas with a large set of outputs. Theoretically, this requires performing up to $2^{|Y|}$ model counting queries. The primary contribution of this paper is to enhance the scalability of precise Shannon entropy computation by improving both stages of the computation process. For the first stage, we design a knowledge compilation language to guide the search and avoid exhaustive enumeration of possible outputs. This language augments Algebraic Decision Diagrams (ADDs), an influential representation, with conjunctive decomposition. For the second stage, instead of performing model counting queries individually, we leverage shared component caching across successive queries. Moreover, we exploit literal equivalence to pre-process the formula corresponding to a given program. By integrating these techniques, we develop a Precise Shannon Entropy tool PSE. We conducted an extensive experimental evaluation over a comprehensive set of benchmarks (459 in total) and compared PSE with the existing precise Shannon entropy computing methods and the current state-of-the-art Shannon entropy estimation tool, EntropyEstimation. Our experiments indicate that EntropyEstimation is able to solve 276 instances, whereas PSE surpasses this by solving an additional 56 instances. Among the benchmarks that were solved by both PSE and EntropyEstimation, PSE is at least $10\times$ as efficient as EntropyEstimation in 98% of these benchmarks.

The remainder of this paper is organized as follows. Section 2 introduces the notation and provides essential background information. Section 3 introduces Algebraic Decision Diagrams with conjunctive decomposition ($ADD[\wedge]$). Section 4 discusses the application of $ADD[\wedge]$ to QIF and introduces our precise entropy tool, PSE. Section 5 details the experimental setup, results, and analysis. Section 6 reviews related work. Finally, Section 7 concludes the paper.

¹ If X does not follow a uniform distribution, techniques exist for reducing the analysis to a uniform case [1].

2 Notations and Background

In this paper, we focus on the programs modeled by (Boolean) formulas. In the formulas discussed, the symbols x and y denote variables, and literal l refers to either the variable x or its negation $\neg x$, where $\text{var}(l)$ represents the variable underlying the literal l . A formula φ is constructed from the constants *true*, *false* and variables using negation operator \neg , conjunction operator \wedge , disjunction operator \vee , implication operator \rightarrow , and equality operator \leftrightarrow , where $\text{Vars}(\varphi)$ denotes the set of variables appearing in φ . A clause C (resp. term T) is a set of literals representing their disjunction (resp. conjunction). A formula in conjunctive normal form (CNF) is a set of clauses representing their conjunction. Given a formula φ , a variable x , and a constant b , the substitution $\varphi[x \mapsto b]$ refers to the transformed formula obtained by substituting the occurrence of x with b throughout φ .

An assignment σ over variable set V is a mapping from V to $\{\text{false}, \text{true}\}$. The set of all assignments over V is denoted by 2^V . Given a subset $V' \subseteq V$, $\sigma_{\downarrow V'} = \{x \mapsto b \in \sigma \mid x \in V'\}$. Given a formula φ , an assignment over $\text{Vars}(\varphi)$ satisfies φ ($\sigma \models \varphi$) if the substitution $\varphi[\sigma]$ is equivalent to *true*. Given an assignment σ , if all variables are assigned a value in $\{\text{false}, \text{true}\}$, then σ is referred to as a complete assignment. Otherwise it is a partial assignment. A satisfying complete assignment is also called solution or model. We use $\text{Sol}(\varphi)$ to the set of solutions of φ , and model counting is the problem of computing $|\text{Sol}(\varphi)|$. Given two formulas φ and ψ over V , $\varphi \models \psi$ iff $\text{Sol}(\varphi \wedge \neg\psi) = \emptyset$.

2.1 Circuit formula and its Shannon entropy

Given a formula $\varphi(X, Y)$ to represent the relationship between input variables X and output variables Y , if $\sigma_{\downarrow X} = \sigma'_{\downarrow X}$ implies $\sigma = \sigma'$ for each $\sigma, \sigma' \in \text{Sol}(\varphi)$, then φ is referred to as a circuit formula. It is standard in the security community to employ circuit formulas to model programs in QIF [15].

► **Example 1.** The following formula is a circuit formula with input variables $X = \{x_1, \dots, x_{2n}\}$ and output variables $Y = \{y_1, \dots, y_{2n}\}$: $\varphi_n^{\text{sep}} = \bigwedge_{i=1}^n (x_i \wedge x_{n+i} \rightarrow y_i \wedge y_{n+i}) \wedge (\neg x_i \vee \neg x_{n+i} \rightarrow \neg y_i \wedge \neg y_{n+i})$.

In the computation of Shannon entropy, we focus on the probability distribution of outputs. Let p denote a probability distribution defined over the set $\{\text{false}, \text{true}\}^Y$. For each assignment σ to Y , i.e., $\sigma : Y \mapsto \{\text{false}, \text{true}\}$, its weight and probability is defined as $\omega_\sigma = |\text{Sol}(\varphi(Y \mapsto \sigma))|$ and $p_\sigma = \frac{|\text{Sol}(\varphi(Y \mapsto \sigma))|}{|\text{Sol}(\varphi)_{\downarrow X}|}$, respectively, where $\text{Sol}(\varphi(Y \mapsto \sigma))$ denotes the set of solutions of $\varphi(Y \mapsto \sigma)$ and $\text{Sol}(\varphi)_{\downarrow X}$ denotes the set of solutions of φ projected to X . Since φ is a circuit formula, it is easy to prove that $|\text{Sol}(\varphi)_{\downarrow X}| = |\text{Sol}(\varphi)|$. Then, the entropy of φ is $H(\varphi) = \sum_{\sigma \in 2^Y} -p_\sigma \log p_\sigma$. Following the convention in QIF [33], we use base 2 for log, though the base can be chosen freely.

2.2 Knowledge compilation

Knowledge compilation is the approach of compiling CNF formulas into a form to support tractable reasoning tasks such as satisfiability check, equivalence check, and model counting [10]. Ordered binary decision diagram (OBDD) [4] is one of the most influential knowledge compilation forms, which supports many tractable reasoning tasks. Each OBDD is a rooted directed acyclic graph (DAG) defined over a linear ordering of variables \prec . Each internal node v is called decision node and has two outgoing edges, referred to as the low child $lo(v)$ and the high child $hi(v)$, which are typically represented by dashed and solid lines,

respectively. Every node u is labeled with a symbol $\text{sym}(u)$. If u is a terminal node, then $\text{sym}(u) = \perp$ or \top , representing the Boolean constants *false* and *true*, respectively. Otherwise, $\text{sym}(u)$ denotes a variable and u represents $(\neg \text{sym}(u) \wedge \psi) \vee (\text{sym}(u) \wedge \psi')$, where ψ and ψ' are the formulas represented by $\text{lo}(u)$ and $\text{hi}(u)$, respectively. Each decision node v and its parent u have $\text{sym}(u) \prec \text{sym}(v)$. OBDD $[\wedge]$ [24] is an extended form of OBDD with better space efficiency. It augments OBDD with conjunctive decomposition nodes. Each conjunctive decomposition node u has a set of children $\text{Ch}(u)$ representing formulas without shared variables, and u represents a conjunction of the formulas represented by its children. OBDD $[\wedge]$ also supports a set of tractable reasoning tasks, including model counting and equivalence check.

Both OBDD and OBDD $[\wedge]$ can only represent Boolean functions. An Algebraic Decision Diagram (ADD) [3] is an extension of OBDD to represent algebraic functions. ADD is a compact representation of a real-valued function as a directed acyclic graph. While OBDD has two terminal nodes representing *false* and *true*, ADD includes multiple terminal nodes, each assigned a real value. The order in which decision node labels appear in all paths from the root to the terminal nodes of the ADD also align with a given ordering of variables \prec . Given an assignment σ with each variable in \prec , we can obtain a path in a top-down way as follows: for a decision node with x , we pick low child if $\sigma(x) = \text{false}$, and high child otherwise. ADD maps σ to the value on the terminate node of the path. The original design motivation for ADD was to solve matrix multiplication, shortest path algorithms, and direct methods for numerical linear algebra [3]. In subsequent research, ADD has also been used for stochastic model checking [21], stochastic programming [17], and weighted model counting [12, 26].

3 ADD $[\wedge]$: A New Tractable Representation

In order to compute the Shannon entropy of a circuit formula $\varphi(X, Y)$, we need to use the probability distribution over the outputs. Algebraic Decision Diagrams (ADDs) are an influential compact probability representation that can be exponentially smaller than the explicit representation. Macii and Poncino [28] showed that ADD supports efficient exact computation of entropy. However, we observed in the experiments that the sizes of ADDs often exponentially explode with large circuit formulas. We draw inspiration from a Boolean representation known as the Ordered Binary Decision Diagram with conjunctive decomposition (OBDD $[\wedge]$) [24], which reduces its size through recursive component decomposition and divide-and-conquer strategies. This approach enables the representation to be exponentially smaller than the original OBDD. Accordingly, we propose a probabilistic representation called Algebraic Decision Diagrams with conjunctive decomposition (ADD $[\wedge]$) and demonstrate it supports tractable entropy computation. ADD $[\wedge]$ is a general form of ADD and is defined as follows:

► **Definition 2.** An ADD $[\wedge]$ is a rooted DAG, where each node u is labeled with a symbol $\text{sym}(u)$. If u is a terminal node, $\text{sym}(u)$ is a non-negative real weight, also denoted by $\omega(u)$; otherwise, $\text{sym}(u)$ is a variable (called decision node) or operator \wedge (called decomposition node). The children of a decision node u are referred to as the low child $\text{lo}(u)$ and the high child $\text{hi}(u)$, and connected by dashed lines and solid lines, respectively, corresponding to the cases where $\text{var}(u)$ is assigned the value of *false* and *true*. For a decomposition node, its sub-graphs do not share any variables. An ADD $[\wedge]$ is imposed with a linear ordering \prec of variables such that given a node u and its non-terminal child v , $\text{var}(u) \prec \text{var}(v)$.

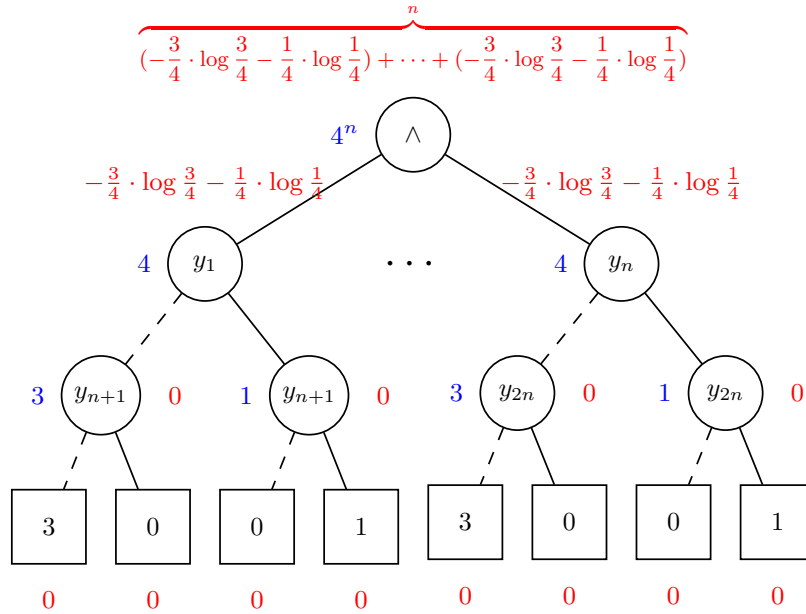
Hereafter, we denote the set of variables that appear in the graph rooted at u as $\text{Vars}(u)$ and the set of child nodes of u as $\text{Ch}(u)$. We now turn to show how an ADD $[\wedge]$ defines a probability distribution:

► **Definition 3.** Let u be an $\text{ADD}[\wedge]$ node over a set of variables Y and let σ be an assignment over Y . The weight of σ is defined as follows:

$$\omega(\sigma, u) = \begin{cases} \omega(u) & \text{terminal} \\ \prod_{v \in \text{Ch}(u)} \omega(\sigma, v) & \text{decomposition} \\ \omega(\sigma, \text{lo}(u)) & \text{decision and } \sigma \models \neg \text{var}(u) \\ \omega(\sigma, \text{hi}(u)) & \text{decision and } \sigma \models \text{var}(u) \end{cases}$$

The weight of an non-terminal $\text{ADD}[\wedge]$ rooted at u is denoted by $\omega(u)$ and defined as $\sum_{\sigma \in 2^{\text{Vars}(u)}} \omega(\sigma, u)$. For nodes with a non-zero weight, the probability of σ over u is defined as $p(\sigma, u) = \frac{\omega(\sigma, u)}{\omega(u)}$.

Figure 1 depicts an $\text{ADD}[\wedge]$ representing the probability distribution of φ_n^{sep} in Example 1 over its outputs with respect to $y_1 \prec y_2 \prec \dots \prec y_{2n}$. The reader can verify that each equivalent ADD with respect to \prec has an exponential number of nodes. In the field of knowledge compilation [10, 13], the concept of succinctness is often used to describe the space efficiency of a representation. Based on the following observations, we can conclude that $\text{ADD}[\wedge]$ is strictly more succinct than ADD. First, OBDD and OBDD $[\wedge]$ are subsets of ADD and $\text{ADD}[\wedge]$, respectively. Second, OBDD $[\wedge]$ is strictly more succinct than OBDD [24]. Finally, each OBDD $[\wedge]$ cannot be transformed into a non-OBDD ADD.



■ **Figure 1** An $\text{ADD}[\wedge]$ representing the probability distribution of φ_n^{sep} in Example 1 over its outputs. According to Proposition 4, the computed weight for each node is marked in blue font. According to Proposition 5, the computed entropy for each node is marked in red font.

3.1 Tractable Computation of Weight and Entropy

The computation of Shannon entropy for an $\text{ADD}[\wedge]$ relies on its weight. We first demonstrate that, for an $\text{ADD}[\wedge]$ node u , its weight $\omega(u)$ can be computed in polynomial time.

► **Proposition 4.** *Given a non-terminal node u in $\text{ADD}[\wedge]$, its weight $\omega(u)$ can be recursively computed as follows in polynomial time:*

$$\omega(u) = \begin{cases} \prod_{v \in Ch(u)} \omega(v) & \text{decomposition} \\ 2^{n_0} \cdot \omega(lo(u)) + 2^{n_1} \cdot \omega(hi(u)) & \text{decision} \end{cases}$$

where $n_0 = |Vars(u)| - |Vars(lo(u))| - 1$ and $n_1 = |Vars(u)| - |Vars(hi(u))| - 1$.

Proof. The time complexity is immediate by using dynamic programming. We prove the equation can compute the weight correctly by induction on the number of variables of the $\text{ADD}[\wedge]$ rooted at u . It is obvious that the weight of a terminal node is the real value labeled. For the case of the \wedge node, since the variables of the child nodes are all disjoint, it can be easily seen from Definition 3. Next, we will prove the case of the decision node. Assume that when $|Vars(u)| \leq n$, this proposition holds. For the case where $|Vars(u)| = n + 1$, we use Y_0 and Y_1 to denote $Vars(lo(u))$ and $Vars(hi(u))$, and we have $|Y_0| \leq n$ and $|Y_1| \leq n$. Thus, $\omega(lo(u))$ and $\omega(hi(u))$ can be computed correctly. According to Definition 3, $w(u) = \sum_{\sigma \in 2^{Vars(u)}} \omega(\sigma, u)$. The assignments over $Vars(u)$ can be divided into two categories:

- The assignment $\sigma \models \neg var(u)$: It is obvious that $\omega(\sigma, u) = \omega(\sigma \downarrow_{Y_0}, lo(u))$. Each assignment over Y_0 can be extended to exactly 2^{n_0} different assignments over $Vars(u)$ in this category. Thus, we have the following equation:

$$\sum_{\sigma \in 2^{Vars(u)} \wedge \sigma \models \neg var(u)} \omega(\sigma, u) = 2^{n_0} \cdot \omega(lo(u)).$$

- The assignment $\sigma \models var(u)$: This case is similar to the above case.

To sum up, we can obtain that $\omega(u) = 2^{n_0} \cdot \omega(lo(u)) + 2^{n_1} \cdot \omega(hi(u))$. ◀

We now present how $\text{ADD}[\wedge]$ computes Shannon entropy in polynomial time.

► **Proposition 5.** *Given an $\text{ADD}[\wedge]$ rooted at u , if $\omega(u) = 0$, we define its entropy $H(u)$ as 0, and otherwise its entropy can be recursively computed in polynomial time as follows:*

$$H(u) = \begin{cases} 0 & \text{terminal} \\ \sum_{v \in Ch(u)} H(v) & \text{decomposition} \\ p_0 \cdot (H(lo(u)) + n_0 - \log p_0) + p_1 \cdot (H(hi(u)) + n_1 - \log p_1) & \text{decision} \end{cases}$$

where $n_0 = |Vars(u)| - |Vars(lo(u))| - 1$, $n_1 = |Vars(u)| - |Vars(hi(u))| - 1$, $p_0 = \frac{2^{n_0} \cdot \omega(lo(u))}{\omega(u)}$, and $p_1 = \frac{2^{n_1} \cdot \omega(hi(u))}{\omega(u)}$.

Proof. According to Proposition 4, $\omega(u)$ can be computed in polynomial time, and therefore the time complexity in this proposition is obvious. Next we prove the correctness of the computation method. The case of terminal nodes is obviously correct. The case of decomposition follows directly from the additivity property of entropy. Next, we show the correctness of the case of decision.

Let $H_0(u)$ be $-\sum_{\sigma \models \neg sym(u)} p(u, \sigma) \log p(u, \sigma)$ and $H_1(u)$ be $-\sum_{\sigma \models sym(u)} p(u, \sigma) \log p(u, \sigma)$. Similar to proposition 4, we can obtain $H(u) = 2^{n_0} \cdot H_0(u) + 2^{n_1} \cdot H_1(u)$. The assignments over $Vars(u)$ can be divided into two categories:

- The assignment $\sigma \models \neg var(u)$: According to definition 3, the probability $p(\sigma, \omega)$ satisfies $p(\sigma, \omega) = \frac{\omega(\sigma, lo(u))}{\omega(u)}$. Given $p_0 = \frac{2^{n_0} \cdot \omega(lo(u))}{\omega(u)}$, it follows that $\frac{\omega(lo(u))}{\omega(u)} = \frac{p_0}{2^{n_0}}$. Substituting this into the expression for $p(\sigma, u)$, we derive $p(\sigma, u) = \frac{\omega(\sigma, lo(u))}{\omega(lo(u))} \cdot p_0 \cdot 2^{-n_0} = p(\sigma, lo(u)) \cdot p_0 \cdot 2^{-n_0}$. $H_0(u)$ then expands as $H_0(u) = -\sum_{\sigma \models \neg sym(u)} p(\sigma, u) \log p(\sigma, u) =$

- $-\sum_{\sigma \models \neg \text{sym}(u)} p(\sigma, lo(u)) \cdot p_0 \cdot 2^{-n_0} \cdot (\log p(\sigma, lo(u)) + \log p_0 - n_0) = p_0 \cdot 2^{-n_0} \cdot$
 $[-\log p_0 \cdot \sum_{\sigma \models \neg \text{sym}(u)} p(\sigma, lo(u)) + n_0 \cdot \sum_{\sigma \models \neg \text{sym}(u)} p(\sigma, lo(u)) - \sum_{\sigma \models \neg \text{sym}(u)} p(\sigma, lo(u)) \cdot$
 $\log p(\sigma, lo(u))].$ Noting that $\sum_{\sigma \models \neg \text{sym}(u)} p(\sigma, lo(u)) = 1$, $-\sum_{\sigma \models \neg \text{sym}(u)} p(\sigma, lo(u)) \cdot$
 $\log p(\sigma, lo(u)) = H(lo(u))$, we simplify $H_0(u) = p_0 \cdot 2^{-n_0} \cdot (-\log p_0 + n_0 + H(lo(u)))$.
- The assignment $\sigma \models \text{var}(u)$: This case is similar to the above case. It is easy to obtain $H_1(u) = p_1 \cdot 2^{-n_1} \cdot (-\log p_1 + n_1 + H(hi(u)))$.
- To sum up, we can obtain that $H(u) = p_0 \cdot (H(lo(u)) + n_0 - \log p_0) + p_1 \cdot (H(hi(u)) + n_1 - \log p_1)$ \blacktriangleleft

We conclude this section by explaining why ordering is used in the design of $\text{ADD}[\wedge]$. In fact, Propositions 4–5 remain valid even when we use only the more general read-once property, where each variable appears at most once along any path from the root of an $\text{ADD}[\wedge]$ to a terminal node. First, our experimental results indicate that the linear ordering determined by the minfill algorithm in our tool PSE outperforms the dynamic orderings employed in the state-of-the-art model counters, where the former imposes the orderedness and the latter imposes the read-once property. Second, $\text{ADD}[\wedge]$ can provide tractable equivalence checking between probability distributions beyond this study.

4 PSE: Scalable Precise Entropy Computation

In this section, we introduce our tool PSE, designed to compute the Shannon entropy of a given circuit CNF formula with respect to its output variables. PSE, as presented in Algorithm 1, takes as input a CNF formula φ , an input set X , and an output set Y , and returns the Shannon entropy $H(\varphi)$ of the formula. Like other tools for computing Shannon entropy, PSE follows a two-stage process: the Y -stage (corresponding to outputs) and the X -stage (corresponding to inputs). In the X -stage (lines 3–4), we perform multiple optimized model counting operations on sub-formulas over variables in X , where the leaves of $\text{ADD}[\wedge]$ are implicitly generated. The optimization technique is discussed in Section 4.1. In the Y -stage (the remaining lines), we conduct a search within the $\text{ADD}[\wedge]$ framework to precisely compute the Shannon entropy, where the internal nodes of $\text{ADD}[\wedge]$ are implicitly generated. The following observation states the input of each recursive call is still a circuit formula and the two input formulas of a call corresponding to a decision node in $\text{ADD}[\wedge]$ have the same output variables.

► **Observation 6.** *Given a circuit formula $\varphi(X, Y)$ and a partial assignment σ without any input variables, we have the following properties:*

- $\varphi[\sigma](X, Y \setminus \text{Vars}(\sigma))$ is a circuit formula;
- Each $\psi_i(X, Y \cap \text{Vars}(\psi_i))$ is a circuit formula if $\varphi = \bigwedge_{i=1}^m \psi_i$ and for $1 \leq i \neq j \leq m$, $\text{Vars}(\psi_i) \cap \text{Vars}(\psi_j) = \emptyset$;
- If $\varphi[\sigma] \equiv \text{true}$, σ contains each output variable.

Proof. The first two properties obviously hold when φ is unsatisfiable. Thereby, we assume φ is satisfiable. For the first property, let φ' be $\varphi[\sigma](X, Y \setminus \text{Vars}(\sigma))$. For each $\sigma', \sigma'' \in \text{Sol}(\varphi)$, $\sigma'_{\downarrow X} = \sigma''_{\downarrow X}$ implies $\sigma' = \sigma''$. $\text{Sol}(\varphi')$ can be seen as a subset of $\text{Sol}(\varphi)$. Consequently, for each $\sigma', \sigma'' \in \text{Sol}(\varphi')$, $\sigma'_{\downarrow X} = \sigma''_{\downarrow X}$ still implies $\sigma' = \sigma''$, concluding that φ' is also a circuit formula.

For the second property, each solution of ψ_i can be obtain from a solution of φ . Let σ', σ'' be two solutions of φ . We only need to prove that $\sigma'_{\downarrow X \cap \text{Vars}(\psi_i)} = \sigma''_{\downarrow X \cap \text{Vars}(\psi_i)}$ implies $\sigma'_{\downarrow (X \cup Y) \cap \text{Vars}(\psi_i)} = \sigma''_{\downarrow (X \cup Y) \cap \text{Vars}(\psi_i)}$. We construct another solution of φ , $\sigma''' = \sigma'_{\downarrow (X \cup Y) \cap \text{Vars}(\psi_i)} \cup \sigma''_{\downarrow (X \cup Y) \setminus \text{Vars}(\psi_i)}$. Then $\sigma' = \sigma'''$, which implies $\sigma'_{\downarrow (X \cup Y) \cap \text{Vars}(\psi_i)} = \sigma'''_{\downarrow (X \cup Y) \cap \text{Vars}(\psi_i)} = \sigma''_{\downarrow (X \cup Y) \cap \text{Vars}(\psi_i)}$.

We prove the last property by contradiction. Suppose that $\varphi[\sigma] \equiv \text{true}$, and σ is a partial assignment with only one free variable $y \in Y$. Then the value of y can take either *false* or *true*. That is, $\sigma \cup \{y = \text{false}\}$ and $\sigma \cup \{y = \text{true}\}$ are solutions of φ , which contradicts the definition of circuit formula. \blacktriangleleft

■ **Algorithm 1** $\text{PSE}(\varphi, X, Y)$.

Input: A circuit CNF formula φ with input variables X and output variables Y
Output: the entropy of φ

```

1 if  $\text{Cache}_H(\varphi) \neq \text{nil}$  then return  $\text{Cache}_H(\varphi)$ 
2 if  $Y = \emptyset$  then
3    $\text{Cache}_\#(\varphi) \leftarrow \text{CountModels}(\varphi)$ 
4   return  $\text{Cache}_H(\varphi) \leftarrow 0$ 
5  $\Psi = \text{Decompose}(\varphi)$ 
6 if  $|\Psi| > 1$  then
7    $\text{Cache}_H(\varphi) \leftarrow \sum_{\psi \in \Psi} \text{PSE}(\psi, X, Y \cap \text{Vars}(\psi_i))$ 
8    $\text{Cache}_\#(\varphi) \leftarrow \prod_{\psi \in \Psi} \text{Cache}_\#(\psi)$ 
9   return  $\text{Cache}_H(\varphi)$ 
10  $y \leftarrow \text{PickGoodVar}(Y)$ 
11  $\varphi_0 \leftarrow \varphi[y \mapsto \text{false}]; \varphi_1 \leftarrow \varphi[y \mapsto \text{true}]$ 
12  $H_0 \leftarrow \text{PSE}(\varphi_0, X, Y \setminus \{y\})$ 
13  $H_1 \leftarrow \text{PSE}(\varphi_1, X, Y \setminus \{y\})$ 
14  $\text{Cache}_\#(\varphi) \leftarrow \text{Cache}_\#(\varphi_0) + \text{Cache}_\#(\varphi_1)$ 
15  $p_0 = \frac{\text{Cache}_\#(\varphi_0)}{\text{Cache}_\#(\varphi)}; p_1 = \frac{\text{Cache}_\#(\varphi_1)}{\text{Cache}_\#(\varphi)}$ 
16  $H \leftarrow p_0 \cdot (H_0 - \log p_0) + p_1 \cdot (H_1 - \log p_1)$ 
17 return  $\text{Cache}_H(\varphi) \leftarrow H$ 

```


In line 1, if the formula φ is cached, its corresponding entropy is returned. If the current set Y is empty (in line 2), this indicates that a satisfiable assignment has been found under the restriction of the output set Y . We do not explicitly handle the case where φ evaluates to *true*, as this naturally implies that Y is empty, as indicated by Observation 6. Consequently, the scenario in which the set Y is empty inherently encompasses the case where φ evaluates to *true*. Lines 3–4 perform model counting on the residual formula and compute its entropy H , corresponding to the terminal case of Proposition 5. We invoke the **Decompose** function in line 5 to determine whether the formula φ can be decomposed into multiple components. In lines 6–9, if φ can be decomposed into multiple sub-components, we compute the model count and entropy of each component ψ , and subsequently derive the entropy of the formula φ . In this case, computing the model count and computing the entropy correspond respectively to the \wedge cases in Propositions 4 and 5. When there is only one component, we select a variable from Y in line 10. The **PickGoodVar** function operates as a heuristic algorithm designed to select a variable from set Y , with the selection criteria determined by the specific heuristic employed. Moving forward, line 11 generates the residual formulas φ_0 and φ_1 , corresponding to assigning the variable y to *false* and *true*, respectively. Subsequently, lines 12 and 13 recursively compute the entropy H for each derived formula. Since φ is a circuit formula, all residual formulas generated in the recursive process after making decisions on variables in Y remain circuit formulas. It follows from Observation 6 that when computing the Shannon entropy of the circuit formula, $n_0 = n_1 = 0$. The model count of φ is cached in line 14,

corresponding to the decision node case in Proposition 4. Finally, in lines 15–16, we compute the entropy of φ (corresponding to the third case in Proposition 5), store it in the cache, and return it as the result in line 17.

► **Example 7.** Consider the following circuit CNF formula with input variables $X = \{x_1, x_2, x_3, x_4, x_5\}$ and output variables $Y = \{y_1, y_2, y_3, y_4, y_5\}$:

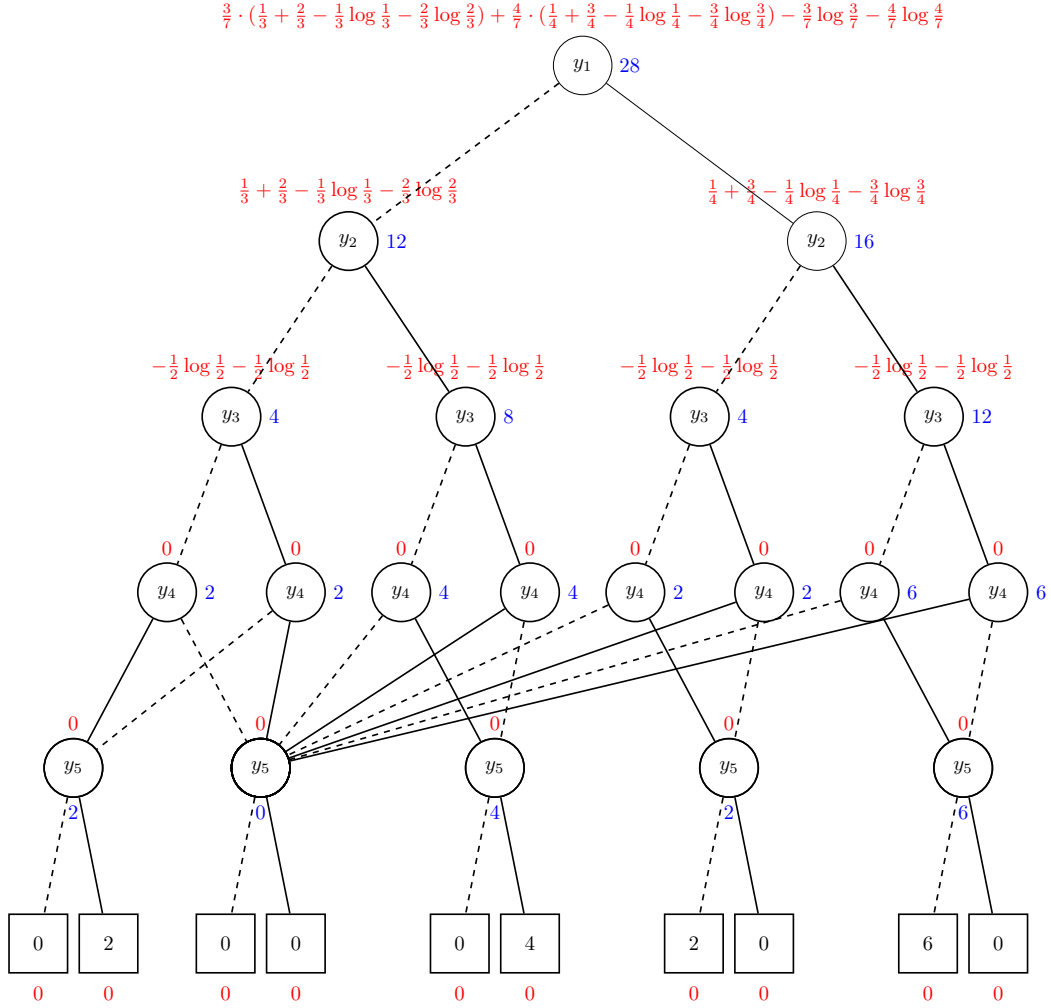
$$\begin{aligned} \varphi(X, Y) = & (x_2 \vee x_3 \vee y_3) \wedge (\neg y_3 \vee \neg y_4) \wedge (x_2 \vee y_3) \wedge (\neg x_2 \vee y_4) \wedge (\neg x_1 \vee \neg y_1) \wedge (x_1 \vee y_1) \wedge \\ & (\neg x_4 \vee x_5 \vee y_2) \wedge (x_4 \vee \neg x_5 \vee y_2) \wedge (\neg x_4 \vee \neg x_5 \vee \neg y_2) \wedge (x_4 \vee x_5 \vee y_2) \wedge \\ & (\neg y_1 \vee \neg y_5) \wedge (y_1 \vee y_5) \wedge (y_1 \vee x_4 \vee x_5) \wedge (y_1 \vee y_3 \vee y_4) \end{aligned}$$

Figure 2 illustrates the execution trace of PSE taking in φ with the variable ordering $y_1 \prec y_2 \prec y_3 \prec y_4 \prec y_5$, which is an implicit $\text{ADD}[\wedge]$. If we do not perform decomposition in line 5, the search trace is depicted in Figure 3, an ADD structure. It is evident that $\text{ADD}[\wedge]$ and ADD yield consistent results, both in terms of Shannon entropy computation and model counting. After merging identical terminal nodes, the $\text{ADD}[\wedge]$ contains 14 nodes, which is fewer than the 24 nodes in the ADD . A comparison between Figure 2 and Figure 3 demonstrates the succinctness of the $\text{ADD}[\wedge]$ structure.



■ **Figure 2** The execution example of PSE on Example 7 follows the variable order of $y_1 \prec y_2 \prec y_3 \prec y_4 \prec y_5$, where the corresponding computational trajectory is represented as an $\text{ADD}[\wedge]$. The entropy computation process performed by PSE is explicitly annotated in red font. The calculation process of weight (number of models) is presented in blue font.

From the aforementioned example, we observe that the search space of PSE corresponds to an $\text{ADD}[\wedge]$ that represents the weights of assignments for the output variables. Thus, Propositions 4–5 and Observation 6 ensure that the entropy of the original formula is obtained from the root call of PSE.



■ **Figure 3** An ADD structure, constructed in Example 7, follows the variable order of $y_1 \prec y_2 \prec y_3 \prec y_4 \prec y_5$. According to Proposition 4, weight is marked in blue font, and according to Proposition 5, entropy is marked in red font.

4.1 Implementation

We now discuss the implementation details that are crucial for the runtime efficiency of PSE. Specifically, leveraging the tight interplay between entropy computation and model counting, our methodology integrates a variety of state-of-the-art techniques in model counting.

In the X-stage of algorithm 1, we have the option to employ various methodologies for the model counting query denoted by `CountModels` in line 3. The first method involves individually employing state-of-the-art model counters, such as SharpSAT-TD [20], Ganak [32], and ExactMC [25]. The second method, known as `ConditionedCounting`, requires the preliminary construction of a representation for the original formula φ to support linear model counting. The knowledge compilation languages that can be used for this method include d-DNNF [8], OBDD[\wedge] [24], and SDD [7]. Upon reaching line 3, the algorithm executes conditioned model counting, utilizing the compiled representation of the formula and incorporating the partial assignment derived from the ancestor calls. The last method, `SharedCounting`, also relies on exact model counters but, unlike the first method, it shares the component cache across all

model counting queries using a strategy called **XCACHE**. To distinguish it from the caching approach used in the *X*-stage, the caching method in the *Y*-stage is referred to as **YCache**. Our experimental observations indicate that the **SharedCounting** method is the most effective within the PSE framework.

Conjunctive Decomposition. We employed dynamic component decomposition (well-known in model counting and knowledge compilation) to divide a formula into components, thereby enabling the dynamic programming calculation of their corresponding entropy, as stated in Proposition 5.

Variable Decision Heuristic. We implemented the current state-of-the-art model counting heuristics for picking variables from *Y* in the computation of Shannon entropy, including VSADS [31], minfill [9], the SharpSAT-TD heuristic [20], and DLCP [25]. Our experiments consistently demonstrate that the minfill heuristic exhibits the best performance. Therefore, we adopt the minfill heuristic as the default option for our subsequent experiments.

Pre-processing. We have enhanced our entropy tool, PSE, by incorporating an advanced pre-processing technique that capitalizes on literal equivalence in model counting. This idea is inspired by the work of Lai et al. [25] on capturing literal equivalence in model counting. Initially, we extract equivalent literals to simplify the formula. Subsequently, we restore the literals associated with the variables in set *Y* to prevent the entropy of the formula from becoming non-equivalent after substitution. This targeted restoration is sufficient to ensure the equivalence of entropy calculations. The new pre-processing method is called **Pre** in the following. This pre-processing approach is motivated by two primary considerations. Firstly, preprocessing based on literal equivalence can simplify the formula and enhance the efficiency of subsequent model counting. Secondly, and more crucially, it can reduce the treewidth of tree decomposition, which is highly beneficial for the variable heuristic method based on tree decomposition and contributes to improving the solving efficiency.

5 Experiments

We implemented a prototype of PSE in C++ and performed evaluations in order to understand its performance. We experimented on benchmarks from the same domains as the state-of-the-art Shannon entropy tool EntropyEstimation [15], that is, QIF benchmarks, plan recognition, bit-blasted versions of SMTLIB benchmarks, QBFEval competitions, program synthesis, and combinatorial circuits [27]². EntropyEstimation reported results only for 96 successfully solved benchmarks (denoted Suite₁), which we found insufficient for scalability testing. To ensure a rigorous evaluation, we extended Suite₁ as follows:

- Suite₂ (399 benchmarks): Suite₁ is from the benchmarks that were used to test a well-known model counter called Ganak³; thereby, we added each circuit formula in the aforementioned domain but not in Suite₁ from the Ganak benchmarks.
- Suite₃ (459 benchmarks): Incorporated 60 additional combinatorial circuits⁴ from [27] on the basis of Suite₂.

² The paper of EntropyEstimation [15] does not mention the domains of program synthesis and combinatorial circuits but actually presents benchmarks in these two domains.

³ The benchmarks are available at <https://github.com/meelgroup/ganak>

⁴ The additional benchmarks are available at <https://github.com/nianzelee/PhD-Dissertation>

All experiments were run on a computer with Intel(R) Core(TM) i9-10920X CPU @ 3.50GHz and 32GB RAM. Each instance was run on a single core with a timeout of 3000 seconds and 4GB memory, the same setup adopted in the evaluation of EntropyEstimation.

Through our experiments, we sought to answer the following research questions:

RQ1: How does the runtime performance of PSE compare to the state-of-the-art Shannon entropy tools with (probabilistic) accuracy guarantee?

RQ2: How do the utilized methods impact the runtime performance of PSE?

5.1 RQ1: Performance of PSE

Golia et al. [15] have already demonstrated that their probably approximately correct tool EntropyEstimation is significantly more efficient than the state-of-the-art precise Shannon entropy tools. The comparative experiments between PSE and the state-of-the-art precise tools are presented in the appendix. We remark that PSE significantly outperforms the precise baseline (the baseline was able to solve only 18 benchmarks, whereas PSE solved 332 benchmarks). This marked improvement is attributed to the linear entropy computation capability of ADD[\wedge] and the effectiveness of various strategies employed in PSE.

Table 1 presents the performance comparison between PSE and EntropyEstimation across the three benchmark suites. For Suite₁, EntropyEstimation solved two more instances than PSE, indicating a slight advantage. However, among the 94 instances that both solved, PSE demonstrated higher efficiency. Moreover, PSE achieved a lower PAR-2⁵ score than EntropyEstimation, suggesting that PSE holds an overall performance advantage. We remark that in the computation of the PAR-2 scores, we *did not perform additional penalization* for each successful run of EntropyEstimation as its output was very close to the true entropy. For Suite₂, PSE solved 44 more instances than EntropyEstimation and achieved a significantly lower PAR-2 score, further demonstrating its superior performance. For Suite₃, PSE solved 56 more instances than EntropyEstimation. Additionally, in terms of overall performance, PSE achieved a significantly lower PAR-2 score than EntropyEstimation, reinforcing its advantage.


■ **Table 1** Detailed performance comparison of PSE and EntropyEstimation. Unique represents the number of instances that can only be solved by a specific tool. Fastest represents the number of instances that a tool solves with the shortest time.

Suite	Tool	Solved Instances			PAR-2 score
		Unique	Fastest	Total	
Suite ₁	EntropyEstimation	2	0	96	149.63
	PSE (ours)	0	94	94	126.04
Suite ₂	EntropyEstimation	2	0	268	2132.52
	PSE (ours)	46	266	312	1314.59
Suite ₃	EntropyEstimation	2	0	276	2536.53
	PSE (ours)	58	274	332	1666.11

⁵ The PAR-2 scoring scheme gives a penalized average runtime, assigning a runtime of two times the time limit for each benchmark that the tool fails to solve

Figure 4 demonstrates the detailed performance comparison between PSE and EntropyEstimation on Suite3. More intuitively, among all the benchmarks that both PSE and EntropyEstimation are capable of solving, in 98% of those benchmarks, the efficiency of PSE surpasses that of EntropyEstimation by a margin of at least ten times. For all the benchmarks where PSE and EntropyEstimation did not timeout and took more than 0.1 seconds, the mean speedup is 506.62, which indicates an improvement of more than two orders of magnitude.

The aforementioned results clearly indicate that PSE outperforms EntropyEstimation in the majority of instances. This validates a positive answer to **RQ1**: PSE outperforms the state-of-the-art Shannon entropy tools with (probabilistic) accuracy guarantee. We remark that EntropyEstimation is an estimation tool for Shannon entropy with probabilistic approximately correct results [15]. PSE consistently performs better than a state-of-the-art entropy estimator across most instances, highlighting that our methods significantly enhance the scalability of precise Shannon entropy computation.

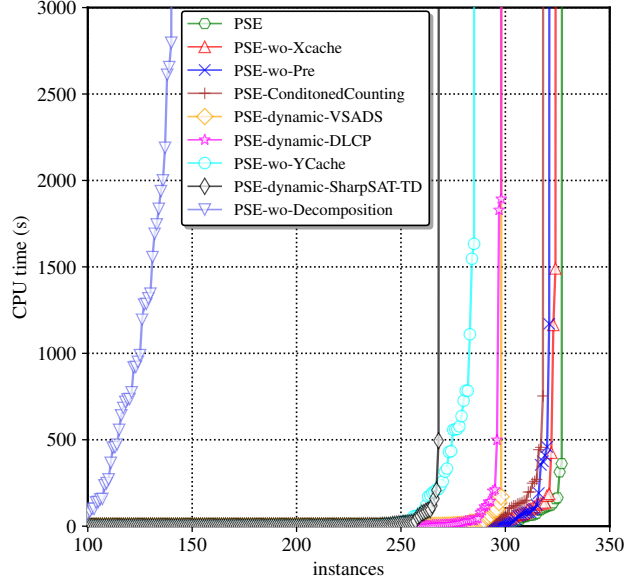


■ **Figure 4** Scatter Plot of the running time Comparison between PSE and EntropyEstimation.

5.2 RQ2: Impact of algorithmic configurations

To better verify the effectiveness of the PSE methods and answer **RQ2**, we conducted a comparative study on all the utilized methods, including methods for the *Y*-stage: **Conjunctive Decomposition**, **YCache**, **Pre**, variable decision heuristics (**minfill**, **DLCP**, **SharpSAT-TD heuristic**, **VSADS**), and methods for the *X*-stage: **XCache** and **ConditionedCounting**. In accordance with the principle of control variables, we conducted ablation experiments to evaluate the effectiveness of each method, ensuring that each experiment differed from the PSE tool by only one method. The cactus plot for the different methods is shown in Figure 5, where PSE represents our tool. PSE-wo-Decomposition indicates that the **ConjunctiveDecomposition** method is disabled in PSE, which means that its corresponding trace is **ADD**. PSE-wo-Pre means that **Pre** is turned off in PSE. PSE-ConditionedCounting indicates that PSE employed the **ConditionedCounting** method rather than **SharedCounting** in the *X*-stage. PSE-wo-XCache

indicates that the caching method is turned off in PSE in the X -stage. PSE-wo-YCache indicates that the caching method is turned off in PSE in the Y -stage. PSE-dynamic-SharpSAT-TD means that PSE replaces the minfill static variable order with the dynamic variable order: the variable decision-making heuristic method of SharpSAT-TD (all other configurations remain identical to PSE, with only the variable heuristic differing). Similarly, PSE-dynamic-DLCP and PSE-dynamic-VSADS respectively indicate the selection of dynamic heuristic DLCP and VSADS.



■ **Figure 5** Cactus plot comparing different methods.

The experimental results highlight the significant effects of conjunctive decomposition. Caching also demonstrates significant benefits, consistent with findings from previous studies on knowledge compilation. It can also be clearly observed that **Pre** improves the efficiency of PSE. Among the heuristic strategies, it is evident that **minfill** performs the best. In the technique of the X -stage, the **ConditionedCounting** method performs better than **SharedCounting** without **XCache**, but not as well as the **SharedCounting** method. This comparative experiment indicates that the shared component caching is quite effective. The **ConditionedCounting** method's major advantage is its linear time complexity [24]. However, a notable drawback is the requirement to construct an **OBDD**[\wedge] (or other knowledge compilation languages such as **d-DNNF**, **SDD**, etc.) based on a static variable order, which can introduce considerable time overhead for more complex problems. Although the **ConditionedCounting** method is not the most effective, we believe it is still a promising and scalable method. In cases where an **ADD**[\wedge] can be efficiently constructed based on a static variable order, the **ConditionedCounting** method may be more effective than the **SharedCounting** method, especially when modeling counting in the X -stage is particularly challenging. Finally, PSE utilizes the **SharedCounting** strategy in the X -stage, and incorporates **ConjunctiveDecomposition**, **YCache**, **Pre**, and the **minfill** heuristic method in the Y -stage.

Finally, we analyze the effectiveness of algorithmic configurations across benchmarks. In terms of the number of solved instances, PSE either solves the most instances or ties with other configurations across all domains. Regarding the **PAR-2** score, on **QBF** benchmarks, PSE-dynamic-VSADS has the lowest score, while in other domains, PSE has the lowest

scores. Among all the instances, there are a total of two instances⁶ which PSE failed to solve within the specified time limit, but were solved by PSE-wo-Pre. In PSE, we use the minfill heuristic to construct a tree decomposition for a given circuit formula. We also observed that the resulting treewidth strongly correlates with compilation size – smaller treewidth in a benchmark typically leads to more efficient PSE execution.

6 Related work

Our work is based on the close relationship between QIF, model counting, and knowledge compilation. We introduce relevant work from three perspectives: (1) quantitative information flow analysis, (2) model counting, and (3) knowledge compilation.

Quantified information flow analysis. At present, the QIF method based on model counting encounters two significant challenges. The first challenge involves constructing the logical postcondition Π_{proc} for a program *proc* [34]. Although symbolic execution can achieve this, existing symbolic execution tools have limitations and are often challenging to extend to more complex programs, such as those involving symbolic pointers. The second challenge concerns model counting, a key focus of our research. For programs modeled by Boolean clause constraints, Shannon entropy can be computed via model counting queries, enabling the quantification of information leakage. Golia et al. [15] have made notable contributions to this field. They proposed the first efficient Shannon entropy estimation method with PAC guarantees, utilizing sampling and model counting. Their approach focuses on reducing the number of model counting queries by employing sampling techniques. Nevertheless, this method yields only an approximate estimation of entropy. Our research is motivated by the work of Golia et al., but diverges in its approach and optimization strategy. We enhance the existing model counting framework for precise Shannon entropy by reducing the number of model counting queries and concurrently improving the efficiency of model counting solutions. Inspired by Golia et al.’s work, our research differs in approach and optimization strategy. We improve the existing model counting framework for precise Shannon entropy by reducing the number of model counting queries and enhancing solution efficiency.

Model counting. Since the computation of entropy relies on model counting, we reviewed advanced techniques in this domain. The most effective methods for exact model counting include component decomposition, caching, variable decision heuristics, pre-processing, and so on. In our research, these methods can all be optimized and improved for application in Shannon entropy computation. The fundamental principle of disjoint component analysis involves partitioning the constraint graph into separate components that do not share variables. The core of ADD[\wedge] lies in leveraging component decomposition to enhance the efficiency of construction. We also utilized caching techniques in the process of computing entropy, and our experiments once again demonstrated the power of caching techniques. Extensive research has been conducted on variable decision heuristics for model counting, which are generally classified into static and dynamic heuristics. In static heuristics, the minfill [9] heuristic is notably effective, while in dynamic heuristics, VSADS [31], DLCP [25], and SharpSAT-TD heuristic [20] have emerged as the most significant in recent years. Lagniez et al. [22] offer a comprehensive review of preprocessing techniques in model counting.

⁶ The two instances are in bit-blasted versions of SMTLIB benchmarks with names `blasted_case_0_ptb_1` and `blasted_TR_b12_1_linear`.

Knowledge compilation. The motivation for knowledge compilation lies in transforming the original representation into a target language to enable efficient solving of inference tasks. Darwiche et al. first proposed a compiler called c2d [8] to convert the given CNF formula into Decision-DNNF. Lai et al. proposed two extended forms of OBDD: Ordered Binary Decision Diagram with Implied Literals (OBDD-L [23]), which is developed by extracting implied literals recursively; OBDD[\wedge] [24], which is proposed by integrating conjunctive decomposition. Both forms aim to reduce the size of OBDD. Exploiting literal equivalence, Lai et al. [25] proposed a generalization of Decision-DNNF, called CCDD, to capture literal equivalence. They demonstrate that CCDD supports model counting in linear time and design a model counter called ExactMC based on CCDD. In order to compute the Shannon entropy, the focus of this paper is to design a compiled language that supports the representation of probability distributions. Numerous target representations have been used to concisely model probability distributions. For example, d-DNNF can be used to compile relational Bayesian networks for exact inference [6]; Probabilistic Decision Graph (PDG) is a representation language for probability distributions based on BDD [18]. Macii and Poncino [28] utilized knowledge compilation to calculate entropy, demonstrating that ADD enables efficient and precise computation of entropy. However, the size of ADD often grows exponentially for large scale circuit formulas. To simplify ADD size, we propose an extended form, ADD[\wedge]. It uses conjunctive decomposition to streamline the graph structure and facilitate cache hits during construction.

7 Conclusion

In this paper, we propose a new compilation language, ADD[\wedge], which combines ADD and conjunctive decomposition to optimize the search process in the first stage of precise Shannon entropy computation. In the second stage of precise Shannon entropy computation, we optimize model counting queries by utilizing the shared component cache. We integrated preprocessing, heuristics, and other methods into the precise Shannon computation tool PSE, with its trace corresponding to ADD[\wedge]. Experimental results demonstrate that PSE significantly enhances the scalability of precise Shannon entropy computation, even outperforming the state-of-the-art entropy estimator EntropyEstimation in overall performance. We believe that PSE has opened up new research directions for entropy computing in Boolean formula modeling.

References

- 1 Michael Backes, Matthias Berg, and Boris Köpf. Non-uniform distributions in quantitative information-flow. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, pages 367–375, 2011. doi:10.1145/1966913.1966960.
- 2 Michael Backes, Boris Köpf, and Andrey Rybalchenko. Automatic discovery and quantification of information leaks. In *2009 30th IEEE Symposium on Security and Privacy*, pages 141–153. IEEE, 2009. doi:10.1109/SP.2009.18.
- 3 R Iris Bahar, Erica A Frohm, Charles M Gaona, Gary D Hachtel, Enrico Macii, Abelardo Pardo, and Fabio Somenzi. Algebraic decision diagrams and their applications. *Formal methods in system design*, 10:171–206, 1997. doi:10.1023/A:1008699807402.
- 4 Randal E Bryant. Graph-based algorithms for boolean function manipulation. *Computers, IEEE Transactions on*, 100(8):677–691, 1986. doi:10.1109/TC.1986.1676819.
- 5 Pavol Cerný, Krishnendu Chatterjee, and Thomas A Henzinger. The complexity of quantitative information flow problems. In *2011 IEEE 24th Computer Security Foundations Symposium*, pages 205–217. IEEE, 2011.

- 6 Mark Chavira, Adnan Darwiche, and Manfred Jaeger. Compiling relational Bayesian networks for exact inference. *International Journal of Approximate Reasoning*, 42(1-2):4–20, 2006. doi:10.1016/J.IJAR.2005.10.001.
- 7 Arthur Choi, Doga Kisa, and Adnan Darwiche. Compiling probabilistic graphical models using sentential decision diagrams. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty: 12th European Conference, ECSQARU 2013, Utrecht, The Netherlands, July 8-10, 2013. Proceedings 12*, pages 121–132. Springer, 2013. doi:10.1007/978-3-642-39091-3_11.
- 8 Adnan Darwiche. New advances in compiling CNF to decomposable negation normal form. In *Proc. of ECAI*, pages 328–332. Citeseer, 2004.
- 9 Adnan Darwiche. *Modeling and reasoning with Bayesian networks*. Cambridge university press, 2009.
- 10 Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17:229–264, 2002. doi:10.1613/JAIR.989.
- 11 Dorothy Elizabeth Robling Denning. *Cryptography and data security*, volume 112. Addison-Wesley Reading, 1982.
- 12 Jeffrey Dudek, Vu Phan, and Moshe Vardi. ADDMC: weighted model counting with algebraic decision diagrams. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1468–1476, 2020. doi:10.1609/AAAI.V34I02.5505.
- 13 Hélène Fargier, Pierre Marquis, Alexandre Niveau, and Nicolas Schmidt. A knowledge compilation map for ordered real-valued decision diagrams. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, 2014.
- 14 Daniel Fremont, Markus Rabe, and Sanjit Seshia. Maximum model counting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- 15 Priyanka Golia, Brendan Juba, and Kuldeep S Meel. A scalable Shannon entropy estimator. In *International Conference on Computer Aided Verification*, pages 363–384. Springer, 2022. doi:10.1007/978-3-031-13185-1_18.
- 16 James W Gray III. Toward a mathematical foundation for information flow security. *Journal of Computer Security*, 1(3-4):255–294, 1992.
- 17 Jesse Hoey, Robert St-Aubin, Alan Hu, and Craig Boutilier. SPUDD: Stochastic planning using decision diagrams. *arXiv preprint*, 2013. arXiv:1301.6704.
- 18 Manfred Jaeger. Probabilistic decision graphs – Combining verification and AI techniques for probabilistic inference. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 12(supp01):19–42, 2004. doi:10.1142/S0218488504002564.
- 19 Vladimir Klebanov, Norbert Manthey, and Christian MuiSe. SAT-based analysis and quantification of information flow in programs. In *International Conference on Quantitative Evaluation of Systems*, pages 177–192. Springer, 2013. doi:10.1007/978-3-642-40196-1_16.
- 20 Tuukka Korhonen and Matti Järvisalo. Integrating tree decompositions into decision heuristics of propositional model counters. In *27th International Conference on Principles and Practice of Constraint Programming (CP 2021)*, pages 8:1–8:11. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CP.2021.8.
- 21 Marta Kwiatkowska, Gethin Norman, and David Parker. Stochastic model checking. *Formal Methods for Performance Evaluation: 7th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2007, Bertinoro, Italy, May 28-June 2, 2007, Advanced Lectures 7*, pages 220–270, 2007. doi:10.1007/978-3-540-72522-0_6.
- 22 Jean-Marie Lagniez and Pierre Marquis. On preprocessing techniques and their impact on propositional model counting. *Journal of Automated Reasoning*, 58:413–481, 2017. doi:10.1007/S10817-016-9370-8.
- 23 Yong Lai, Dayou Liu, and Shengsheng Wang. Reduced ordered binary decision diagram with implied literals: A new knowledge compilation approach. *Knowledge and Information Systems*, 35:665–712, 2013. doi:10.1007/S10115-012-0525-6.

- 24 Yong Lai, Dayou Liu, and Minghao Yin. New canonical representations by augmenting obdds with conjunctive decomposition. *Journal of Artificial Intelligence Research*, 58:453–521, 2017. doi:10.1613/JAIR.5271.
- 25 Yong Lai, Kuldeep S Meel, and Roland HC Yap. The power of literal equivalence in model counting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3851–3859, 2021.
- 26 Yong Lai, Zhenghang Xu, and Minghao Yin. Pbcounter: weighted model counting on pseudo-boolean formulas. *Frontiers of Computer Science*, 19(3):193402, 2025. doi:10.1007/S11704-024-3631-1.
- 27 Nian-Ze Lee, Yen-Shi Wang, and Jie-Hong R Jiang. Solving exist-random quantified stochastic boolean satisfiability via clause selection.c. In *IJCAI*, pages 1339–1345, 2018. doi:10.24963/IJCAI.2018/186.
- 28 Enrico Macii and Massimo Poncino. Exact computation of the entropy of a logic circuit. In *Proceedings of the Sixth Great Lakes Symposium on VLSI*, pages 162–167. IEEE, 1996. doi:10.1109/GLSV.1996.497613.
- 29 Ziyuan Meng and Geoffrey Smith. Calculating bounds on information leakage using two-bit patterns. In *Proceedings of the ACM SIGPLAN 6th Workshop on Programming Languages and Analysis for Security*, pages 1–12, 2011.
- 30 Quoc-Sang Phan, Pasquale Malacaria, Oksana Tkachuk, and Corina S Păsăreanu. Symbolic quantitative information flow. *ACM SIGSOFT Software Engineering Notes*, 37(6):1–5, 2012. doi:10.1145/2382756.2382791.
- 31 Tian Sang, Paul Beame, and Henry Kautz. Heuristics for fast exact model counting. In *Theory and Applications of Satisfiability Testing: 8th International Conference, SAT 2005, St Andrews, UK, June 19-23, 2005. Proceedings 8*, pages 226–240. Springer, 2005. doi:10.1007/11499107_17.
- 32 Shubham Sharma, Subhajit Roy, Mate Soos, and Kuldeep S Meel. GANAK: A Scalable Probabilistic Exact Model Counter. In *IJCAI*, volume 19, pages 1169–1176, 2019. doi:10.24963/IJCAI.2019/163.
- 33 Geoffrey Smith. On the foundations of quantitative information flow. In *International Conference on Foundations of Software Science and Computational Structures*, pages 288–302. Springer, 2009. doi:10.1007/978-3-642-00596-1_21.
- 34 Ziqiao Zhou, Zhiyun Qian, Michael K Reiter, and Yinqian Zhang. Static evaluation of noninterference using approximate model counting. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 514–528. IEEE, 2018. doi:10.1109/SP.2018.00052.

A Comparison with precise Shannon entropy computing methods

In the appendix, we compare PSE with the state-of-the-art precise methods of computing Shannon entropy. The existing precise Shannon entropy tools do not use the techniques in the state-of-the-art model counters. Just like [15], we implemented the precise Shannon entropy baseline with state-of-the-art model counting techniques. In the baseline, we enumerate each assignment $\sigma \in \text{Sol}(\varphi)_{\downarrow Y}$ and compute $p_\sigma = \frac{|\text{Sol}(\varphi(Y \mapsto \sigma))|}{|\text{Sol}(\varphi)_{\downarrow X}|}$, where $\text{Sol}(\varphi(Y \mapsto \sigma))$ denotes the set of solutions of $\varphi(Y \mapsto \sigma)$ and $\text{Sol}(\varphi)_{\downarrow X}$ denotes the set of solutions of φ projected to X . As can be seen from the previous proposition, $|\text{Sol}(\varphi)_{\downarrow X}|$ can be replaced by $|\text{Sol}(\varphi)|$. Finally, entropy is computed as $H(\varphi) = \sum_{\sigma \in 2^Y} -p_\sigma \log p_\sigma$. For a formula with an output set size of m , 2^m model counting queries are required. For model counting queries, we have adopted two different methods. One is to directly invoke the currently state-of-the-art model counters, and our experiment, SharpSAT-TD, Ganak, and ExactMC are employed. The other method involves utilizing knowledge compilation. Firstly, we construct an offline knowledge compilation language that supports linear model counting, and then perform online conditioning based on each assignment over the Y variables. The knowledge compilation

language in our experiment is (OBDD[\wedge] via KCBox), and this method corresponds to the baseline-Panini in Table 2. Panini is an efficient compilation tool that supports the compilation of CNF formulas into the form of OBDD[\wedge] to enable efficient model counting.

■ **Table 2** Entropy computation performance of baselines and PSE. “—” represents that the entropy cannot be computed within the specified time limit.

instance	$ X $	$ Y $	baseline-SharpSAT-TD		baseline-ExactMC		baseline-Ganak		baseline-Panini		PSE	
			Entropy	Time(s)	Entropy	Time(s)	Entropy	Time(s)	Entropy	Time(s)	Entropy	Time(s)
blasted_case102.cnf	11	23	8	81.29	8	0.39	8	35.42	8	1.82	8	0.16
s27_15_7.cnf	7	25	3.3	142.77	3.3	0.15	3.3	0.25	3.3	0.31	3.3	0.14
small-bug1-fixpoint-5.cnf	66	21	12.81	79.51	12.81	2.99	12.81	104.17	12.81	8.79	12.81	0.18
small-bug1-fixpoint-6.cnf	79	25	15.31	1406.47	15.31	51.9	15.31	1846.36	15.31	127.70	15.31	0.21
blasted_case144.cnf	138	627	—	—	—	—	—	—	—	—	77.62	35.42
s1423a_15_7.cnf	91	773	—	—	—	—	—	—	—	—	88.17	232.65
s382_15_7.cnf	24	326	—	—	—	—	—	—	—	—	23.58	0.22
CVE-2007-2875.cnf	752	32	—	—	—	—	—	—	—	—	32	0.72
10.sk_1_46.cnf	47	1447	—	—	—	—	—	—	—	—	13.58	0.18

Our experimental results indicate that all four representative state-of-the-art exact Shannon entropy baselines can only solve 18 benchmarks within the time limit of 3000 seconds, whereas PSE can solve 332 benchmarks. Table 2 shows the comparison between baselines and PSE on some instances. Notably, although some instances have similar sizes of X and Y sets, their computation times vary significantly (e.g., blasted_case144.cnf vs. s1423a_15_7.cnf). To clarify, computation times depend on multiple parameters, such as an exponential relationship with treewidth in addition to problem size. We employ the minfill heuristic to compute tree decompositions, guiding the entropy calculation. Our experimental results show that blasted_case144.cnf has a minfill treewidth of 22, whereas s1423a_15_7.cnf has a minfill treewidth of 27. The results show a significant improvement in the efficiency of PSE for computing the precise Shannon entropy. We remark that the poorer performance of these baselines is due to the exponential size of 2^Y .