MUNI
FACULTY
OF INFORMATICS

BACHELOR THESIS
READER'S REVIEW

Title: *Approximation Techniques for Binary Decision Diagrams*

Author: *Tomáš Kocián*

Supervisor: *prof. RNDr. Jan Strejček, Ph.D.*

Reader: *RNDr. Samuel Pastva, Ph.D.*

The submitted thesis focuses on the study and evaluation of approximation techniques for Binary Decision Diagrams (BDDs). BDDs are widely used in formal verification and SMT solving, where memory exhaustion is a common bottleneck. The author explores the problem of underapproximation—reducing the size of a BDD while sacrificing as few minterms as possible. The work includes an implementation of an optimal brute-force baseline, an analysis of several established heuristic algorithms (Heavy branch subsetting, Short path subsetting, UnderApprox, and Remap) available in the CUDD library, and the proposal of a novel under-approximation technique called SAB (Subsetting Across Branches). Finally, the methods are experimentally evaluated on a dataset derived from SMT-LIB benchmarks.

Overall, the text is generally understandable and well-structured, successfully introducing the reader to the core concepts of BDDs and subsetting. However, the text occasionally lacks mathematical and analytical rigor, particularly concerning the computational complexity discussions, which are somewhat superficial. Some sections appear to have been written in haste, resulting in certain concepts being defined only informally. Furthermore, while the algorithms are thoroughly presented, the pseudocode provided in the text is quite technical. For instance, the algorithms explicitly include the low-level handling of complemented (negation) edges. While I completely understand that complementation is a critical implementation detail in a framework like CUDD, its inclusion in the pseudocode adds a significant layer of complexity that distracts from the fundamental algorithmic concepts. The presentation would have benefited greatly from a slightly higher level of abstraction.

The proposed SAB algorithm introduces an interesting generalization of branch replacement. Although the experimental results demonstrate that SAB does not achieve a meaningful improvement compared to state-of-the-art methods (often creating "bigger steps" and struggling to estimate size reduction correctly), the algorithm itself is sound and complex. It is commendable that the student was able to design and implement this custom technique directly within the architecture of an existing, heavily optimized CUDD library (even though I understand major changes to CUDD were not needed).

The thesis also introduces an optimal brute-force algorithm to establish a theoretical baseline. Unsurprisingly, due to exponential explosion, this optimal approach was only computationally feasible on very small examples (BDDs with fewer than 30 minterms). While it is still a valuable theoretical addition to the text, in my experience, small BDDs are rarely representative of larger structural trends. Specifically, in a small BDD, a single node often covers a relatively large fraction of the total minterms, whereas, for larger real-world BDDs, this fraction can be very small for nodes deep within the diagram, changing how heuristics perform. I do, however, fully realize that overcoming the complexity of the optimal approach is an inherent limitation of the subsetting problem and not something that can be easily solved within the scope of this work.

Regarding the experimental evaluation, the results section provides a robust relative comparison among the methods. The metrics used, such as total AUC scores and rank points, are good for ranking the heuristics against one another. However, I would very much appreciate more general guidance in the text regarding how effective the approximations actually were overall. It is somewhat difficult to discern from these relative metrics whether the tested approximation methods yielded reductions that would be genuinely useful in a practical verification setting. To better contextualize the findings, it would be highly beneficial to include more descriptive statistics about the practical dataset of the tested BDDs—such as their initial node sizes, number of variables, and node density prior to the approximation process.

Despite my reservations regarding the theoretical rigor and evaluation context, the goals of the thesis have been satisfied and I recommend awarding the thesis with the grade **passed**.

Suggested questions for the thesis defense:

- Were the evaluated approximation methods practically viable overall across your testing set? In other words, was it typically possible to find a significant size reduction while still retaining a meaningful portion of the original function for most of the tested BDDs?
- Although it was not the primary focus of the evaluation metrics, how do the evaluated methods compare in terms of computational performance and execution time?
- During the benchmark data preparation phase, a 40-second timeout was utilized to filter files. Is this timeout appropriate, or does it inherently bias the evaluation by restricting the testing predominantly to relatively small or structurally simple BDDs?
- Did you validate the correctness of the approximation techniques, particularly your newly proposed SAB algorithm? Specifically, did you verify that the resulting subset BDD is indeed a strict, correct logical underapproximation of the original BDD?

Brno 14.06.2026 Samuel Pastva