Can We Trust PRA ?

Qui a vist Paris et noun Cassis, ren a vist.
If one has seen Paris, but not Cassis, one has seen nothing.
--- an old Provencal expression

A. Rauzy
ARBoost Technologies
24, Allée Chabrier
13008 Marseille, France
antoine.rauzy@arboost.com

1 Introduction

Katatsuburi
soro-soro nobore
fuji no yama
oh snail
climb Mount Fuji,
but slowly, slowly
--- *Issa*

The Fault Trees/Event Trees method is widely used in industry. Probabilistic Risk Assessment in the nuclear industry relies worldwide almost exclusively on this technique. Several tools are available to assess event tree models. Almost all of them implement what we call the “classical” approach: first, event tree sequences are transformed into Boolean formulae. Then, after possibly applying some rewriting rules, minimal cutsets of these formulae are determined. Finally, various probabilistic measures are assessed from the cutsets (including probabilities and/or frequencies of sequences, importance factors, sensitivity analyzes, …). This approach is broadly accepted. However, it comes with several approximations:

- In order to assess probabilistic quantities from the cutsets, the rare event approximation is applied. Under certain conditions the min-cut upper bound approximation can be used, but only when the boolean equation does not have negation and all basic event probabilities are quite low, at least smaller than 10^{-2}.

- In order to minimize cutsets, and therefore avoiding combinatorial explosion, probability truncation (hereafter referred to as simply truncation) is applied.
- Finally, in order to handle success branches, various recipes more or less mathematically justified are applied.

Since, up to now, all of the assessment tools rely on the same technology (with some variations indeed), it was not possible to verify whether the above approximations are accurate for large real-life models, especially since to compute error bounds, the exact solution is necessary

In the beginning of the nineties, a new technology was introduced to handle Boolean models: Bryant's Binary Decision Diagrams (BDD for short) [Bry86,Bry92]. One of the major advantages of the BDD technology is that it provides exact values for probabilistic measures [Rau93,DR00]. It does not need any kind of truncation or approximations. BDDs are however highly memory consuming. Very large models, such as event trees of the nuclear industry, were beyond their reach. Nevertheless, the methodology can be improved by means of suitable variable heuristics and formula rewritings.

Recently, we were given a rather large event tree model (coming from the nuclear industry). We designed a strategy, i.e. a sequence of rewritings, that made it possible to handle all of the 181 sequences of the model within reasonable running times and memory consumptions. For one of the first times, it was possible to compare results of the classical approach with those of the BDD approach, i.e. with exact results.

As the epigram to this section intimates, we should not draw definitive conclusions from a single test case. But a single example suffices to ring the alarm bell: the classical approach gives wrong results in a significant proportion of cases.

The remainder of this article is organized as follows. Sections 2 and 3 present respectively the classical and the BDD approaches. Section 4 gives some insights on the test case we used for this study. Section 5 reports comparative results for the computation of sequence frequencies. Section 6 considers briefly the complexity, runtime, and space considerations when trying to solve large problems. Finally, section 7 presents our preliminary conclusions.

2 The classical approach to assess event trees

Da Vinci was so steeped in his own tradition that with each step he took, walked a bit beyond it.

--- Scott Buchanan, EMBERS of the WORLD

2.1 Principle

By construction, sequences of event trees are mutually exclusive. Therefore, they can be treated separately, at least for what concerns the computation of their probabilities.

The classical approach to assess event trees works as follows.

- First, sequences are compiled as explained above.
- Second, some rewriting is performed on the formula associated with each sequence (e.g. modularization) in order to facilitate their treatment.
- Third, minimal cutsets of each sequence (or group of sequences) are determined. Classical algorithms to compute the minimal cutsets work either top-down (e.g. [FV72, Rau03]) or bottom-up (e.g. [JK98,JHH04]).
- Fourth, probabilities/frequencies of sequences are assessed from the cutsets. More generally, cutsets are used to get various measures of interest such as importance factors of components, sensitivity to variations in basic event probabilities, …

In this process, three kinds of approximations are used:

- Sequences, including success branches, are quantified by means of minimal cutsets (which, by definition, do not embed negations).
- Truncation is applied to limit the process, and therefore reduce the possibility of combinatorial explosion.
- Probabilities are evaluated using one of two first order approximations: the rare event approximation or min-cut upper bound.

2.2 Truncation in minimal cutsets determination

In general, sequences of large event trees admit huge numbers of minimal cutsets. Therefore, only a subset of the latter's can be considered (the most important ones, in terms of probability, one expects). Algorithms to compute minimal cutsets apply truncation to keep only few thousands cutsets (beyond computations are intractable).

The choice of the right truncation value is a result of trade-offs between accuracy of the computation and resource (time and memory) consumption. Expert knowledge about the expected probability of the sequence plays also an important role in that choice.

It remains that, by applying truncation, one gets an optimistic approximation. Moreover, there is no way to ensure that this approximation is accurate. For instance, if we keep a thousand cutsets of probability 10^{-9} and by the way we ignore a million cutsets of order 10^{-11}, then we underestimate the risk by a factor 10. This problem is largely ignored by most of the practitioners.

2.3 Quantification of success branches

But the main problem in the classical approach stands in the way success branches are (badly or even not at all) taken into account. None of the classical algorithms are actually able to deal with negations, for two main reasons. First, by definition, minimal cutsets do not contain negative literals. Therefore, the functions they encode are coherent. The notion of minimal solutions of general (coherent or non-coherent) functions exists (this is the notion of prime implicants), but that's another (very different) story. Second, truncation and minimality tests and reduction rules used by classical algorithms are not compatible with negations. The interested reader should see [Rau01] for a detailed discussion on that topics, including theoretical computational complexity arguments.

Some authors propose process success branches as follows. First, negations are pushed down toward variables, using de Morgan's Laws. Second, new variables are introduced to represent negative literals. Third, minimal cutsets of the rewritten formula are computed. Finally, those that contain both a variable and its (encoded) negation are eliminated. This attempt is interesting. However, it cannot work correctly because of truncation.

3 The BDD approach to assess event trees

If a man does not keep pace with his companions, perhaps it is because he hears a different drummer. Let him step to the music which he hears, however measured or far away.

Henry David Thoreau, WALDEN

Bryant's Binary Decision Diagrams [Bry86], BDD for short, are now a well-known and widely used technique [Bry92]. In this section, we recall briefly the basics of this technique and we discuss its use to assess event trees (J. Andrews initiated this work in [AD00]).

3.1 Binary Decision Diagrams

The Binary Decision Diagram of a formula is a compact encoding of the truth table of this formula. From a BDD, it is possible to perform efficiently all of the probabilistic quantifications (top event probability, importance factors,…). The BDD representation is based on the Shannon decomposition.

By choosing a total order over the variables and applying recursively the Shannon decomposition, the truth table of any formula can be graphically represented as a binary tree. The nodes are labelled with variables and have two outedges (a *then*-outedge, pointing to the node that encodes *F[v←1]*, and an *else*-outedge, pointing to the node that encodes *F[v←0]*). The leaves are labelled with either 0 or 1. The value of the formula for a given variable assignment is obtained by descending along the corresponding branch of the tree. The Shannon tree for the formula $F = ab + \overline{a}c$ and the lexicographic order is pictured Fig. 2 (dashed lines represent *else*-outedges).

3.2 Application to Fault Trees/Event Trees assessment

Thanks to the Shannon decomposition, the probability of a formula *F* can be computed efficiently from the BDD that encodes *F* (and the probabilities of basic events).

It is easy to derive a recursive algorithm from equality (8) [Rau93]. This algorithm is linear in the size of the BDD and gives exact results. It needs no truncations and makes no approximation. Importance factors can also be computed efficiently and exactly from BDD [DR00].

By slightly modifying the semantics of nodes, BDD can also be used to compute and to encode minimal cutsets (see [Rau93, Rau01]). BDDs that encode minimal cutsets are called ZBDD, from the name given by in its Minato's seminal article [Min93]. Truncation can be applied to keep only the most relevant cutsets.

Hence, the BDD approach to assess event trees works as follows.

- First, sequences are compiled as explained above.
- Second, some rewriting is performed on the formula associated with each sequence in order to facilitate their treatment and to select a good variable ordering. We shall discuss this very important issue in the next section.
- Third, the BDD that encode the sequence is computed.
- Fourth, the exact value of the probability (or the frequency) of the sequence is computed from the BDD. More generally, importance factors of components, as well as sensitivity to variations in basic event probabilities are assessed from the BDDs in a exact way.

As a fourth or fifth step and for the sake of the verification of the model, minimal cutsets can be extracted. However, this is not necessary. Moreover, since minimal cutsets are used only for verification purposes, one need only consider very few of them. In fact, for an analyst to consider more than a few hundred cutsets may be cognitively infeasible.

4 A Case Study

Mais cher Woody, on doit goûter le vin délicatement, si on l'apprécier à sa juste valeur.

But my dear Woody, one must place wine gently in one's mouth, if one wishes to make an informed judgement.

--- CoCo to Woody, at Maison d'H

4.1 The Model

The basis of this study is an actual event tree coming from the nuclear industry. This event tree is made of 181 sequences, with a total of 2259 basic events. Broken down by sequence, the smallest is made of 78 gates and 158 basic events. The largest one is made of 1128 gates and 1745 basic events. The mean numbers of gates and basic events are respectively 857 and 1455 per sequence. Among the 181 sequences, 171 lead to core damage.

Ten fault trees, representing the top events of the event trees, are used to build the sequences (plus 8 individual events). The smallest of these fault trees is made 74 gates and 155 basic events. The biggest one is made of 561 gates and 946 basic events. The mean numbers of gates and basic events are respectively 277 and 490.

4.2 Efficiency of the BDD approach

For each sequence, we computed (with the strategy discussed in section 4.3) the following data structures and quantities.

- The formula rewritten by the strategy.
- The BDD that encodes this formula.
- The probability of the sequence computed from the BDD.
- The minimal cutsets of the sequence. However, it is not possible to compute all of the minimal cutsets (for most of the sequences there are more than 10^9 of them). The BDD approach can efficiently count the minimal cutsets, and the prime implicants, even though none need be listed.

To limit the number of minimal cutsets generated, we used a probability truncation limit defined thusly:

*(Cutset Value) >= (BDD Value of the Sequence) * 10^{-4}*

So if the probability of a sequence is 10^{-9}, as calculated by BDD, we limited the cutsets to those whose probabilities are greater than 10^{-13}.

It is worth noting that we used BDDs to calculate the exact results as well as to create ZBDDs, a data structure from which we can extract the cutsets [Min93]. The cutsets we obtained are the same as those that would have been obtained with a classical bottom-up or a top-algorithm. However, to extract them from the ZBDD is much faster.

5 Comparison between the classical and the BDD approaches

[T]hese discoveries clearly confute the Ptolemaic system, and they agree admirably with this other position and confirm it.

--- Galileo, in a letter to the Grand Duchess Christina of Lorraine

5.1 Experimental protocol

In this section, we compare the probabilities of sequences we obtained with the BDD approach with those that would have been obtained with a classical approach. The questions we aim to answer are the following.

Question 1: Is the classical approach, the rare event and mincut upper bound approximations, good enough?

Question 2: Is the approximation Q_1 accurate (Q_1 consists in ignoring success branches)?

Question 3: Is the approximation Q_2 accurate (Q_2 consists in correcting Q_1 by multiplying it with the probability of success branches)?

Question 4: Is the approximation Q_3 accurate (Q_3 consists in computing the probability of the sequence from its minimal cutsets and employing “delete term”).

In order to answer questions 1 to 4, we computed, for each sequence, the following quantities:

- The exact probability of sequence, computed from the BDD.
- The first term of the Sylvester-Poincaré development (Q_3) computed from the minimal cutsets of the sequence (recall that we keep only cutsets whose contribution is at least 10^{-4} times the probability of the sequence as calculated by BDD).

- The difference between the first and the second terms of the Sylvester-Poincaré development, still computed from the cutsets.
- The exact probability of the conjunction of the failure branches of the sequence computed from the BDD that encodes this conjunction. It is worth noticing that this approximation should be better than Q_1 as we defined it section 2. However, for the sake of the simplicity, we shall denote it Q_1.
- The exact probability of the conjunction of the failure branches times one minus the exact probability of the disjunction of the success branches (both obtained by the BDDs that encode them). For the same reason as previously, this quantity should be a better approximation than Q_2 but we shall denote it Q_2.

Except the first sequence, that contains only success branches and whose probability is 0.999817, probabilities of sequences range rather log-uniformly from $4.99\ 10^{-5}$ to $2.87\ 10^{-18}$. It would be an error to concentrate on most probable sequences for each sequence corresponds to a different situation. The less frequent sequences may also be those with the most severe consequences for the environment. Table 5 gives a distribution of the sequences according to their probabilities (this distribution is a bit arbitrary for the reasons we just gave).

5.2 Analysis of the results

Question 1: the question 1 is easy to answer. the range given by the two first terms of the Sylvester-Poincaré development is narrow for all of the sequences. The relative difference second-term / first-term never exceeds 6%. This means that the rare event approximation is accurate, at least for what concerns the quantity it assesses.

Question 2: To answer this question, we compute the relative difference $[Q_1(S)-p(S)]/p(S)$ which represents the relative error one makes by considering Q_1 (note that Q_1 is always bigger than the exact probability).

For only one sixth of the sequences Q_1 is pessimistic by a factor less than 2! For half of the sequences, Q_1 is pessimistic by at least two orders of magnitude! Sequence number 13 has the “gold medal” with a relative error of $6.53\ 10^{7}$ for a probability $1.72\ 10^{-12}$.

Question 3: Q_2 corrects a bit Q_1. However, it gives very pessimistic results for two thirds of the sequences and is still pessimistic by two orders of magnitude for half of the sequences. Sequence number 13 has again the “gold medal” with a relative error of 3.98 10^6.

Question 4: The answer to this question is a bit more complex for Q_3 gives sometimes optimistic results. However, the greater the number of minimal cutsets considered, the less the expectation to be optimistic. So, on the one hand, one may argue that we didn’t take into account enough cutsets. On the other hand, the truncation has to be put somewhere in order to avoid prohibitively long run times. By setting it to 10^{-4} the probability of the sequence, we adopted a quite conservative attitude. If we had used the min-cut upper bound approximation, the results would have been even more optimistic.

Q_3 is thus optimistic in more than a quarter of the sequences. It is optimistic by a factor 2 in at least one sequence whose probability is around 10^{-9} and by a factor 4 for at least one sequence whose probability is around 10^{-12}.

Q_3 is thus pessimistic by a factor *2* or more for 104 sequences among 181 and by a factor 10 or more for one third of the sequences. For instance, it is pessimistic by a factor 14 for the sequence number whose exact probability 7.42 10^{-6}.

In order to confirm these results, we calculated the cutsets whose fractional contributions to the probability of the sequence is greater than 10^{-6} (rather than 10^{-4}). Indeed, running times and numbers of cutsets increase quite a lot (running times are up to 20 minutes and for some of the sequences up to 200,000 cutsets show up). With such a low truncation, Q_3 is pessimistic on all but 7 sequences. On the latter sequences, it is optimistic by at most 7%, which is quite acceptable. Table 10 gives the distribution of sequences according to the relative error made by Q_3, when Q_3 is pessimistic. It is worth noticing that Q_3 is now pessimistic by a factor greater than 10 for more than a third of the sequences and by a factor greater than 80 for 18 of them.

Last, but not least: the sum of the probabilities of core damage sequences computed with the BDD approach is 2.27 10^{-5}. With the classical approach we obtain, for both cutoffs (10^{-4} times the probabilty of the sequences and 10^{-6} times the probability of the sequences), 1.29 10^{-4}. It has been pointed out by M. Barrett that just summing the probabilities obtained for each sequence may be incorrect because some cutsets may be duplicated (or even subsumed). With the BDD approach this problem does not exist since the sequences are exclusive to one another by construction. In order to confirm our observations, we performed the following experiment. For different absolute cut-offs, we computed the cutsets for each core-damage sequence, we collected all of these cutsets together, we removed those duplicated and subsumed, and then we computed the probability of a core damage from the resulting set. Table 11 gives the results obtained for absolute cut-offs of 10^{-10}, 10^{-11}, 10^{-12} and 10^{-13} (i.e. we kept only those cutsets whose probabilities are greater than the above cut-offs). For all of these cut-offs, the probability of a core damage is 1.21 10^{-4}. Therefore, no matter which way it is applied, the classical approach overestimates by almost a factor 5 the likehood of core damage. Recall that even in the small example at the end of Section 3.2, all approximations overstated results by a factor of 2.

6 Runtime, Space, and Complexity

Faire de la bonne cuisine demande un certain temps. Si on vous fait attendre, c'est pour mieux vous servir, et vous plaire.

Good cooking takes the time it takes. If we are making you wait, it is to better serve you, and to please you.

--- Menu of Restaurant Antoine, New Orleans

Our total run times are obtained by accumulating the running times to rewrite the formula, to build the BDD, to compute the probability from the BDD and to build the ZBDD that encodes the minimal cutsets (with a 10^{-4} relative cut-off). They were observed on a laptop computer running Windows 2000 with a processor speed of 1.8 Ghz and with a 1 gigabyte of RAM. It takes at most 156.1 seconds to handle a sequence. All but 2 sequences are treated in less than 75 seconds. On average it takes 16.03 seconds to fully quantify a sequence. These running times can surely be improved by using larger hashtables, a faster computer, or by improving the strategy. More importantly, the resulting BDD can be cached for subsequent quantification at a later time with different variable probabilities. It takes at most 4.27 seconds to compute the probability of a sequence from its BDD (and 0.81 seconds on average). The running time to assess a probability from a BDD doesn't depend on the probabilities of basic events.

It takes at most 4,86 millions nodes to build the BDD and the ZBDD that encodes the cutsets. On average these computations require 1,28 millions nodes. It other words, given the size of the hashtables, the most difficult sequence is handled within around 200 megabytes.

The past 20 years have seen the movement of crucial engineering programs from mainframe and mini-computers to personal computers, and dramatic increases in computation speed and memory limits. We all remember the PC-AT in 1986 which increased the clock speed of the original PC from 4.77 Mhz to 6 Mhz, while now off-the-shelf hardware can run at 2.6 Ghz.

While all of this is well and good, it has also given us unrealistic notions of what runtimes and memory requirements should be when solving NP-hard problems, such as the one described in this article. Certainly computing from BDDs will take longer and need more space than simple bottom-up algorithms with truncation. But when solving calculations to help us understand the risk to and safety of populations and environments, the runtime difference between 5 seconds and 50 seconds can rationally be ignored when the extra 45 seconds will produce exactly correct answers.

Another interesting question is how one measures the complexity of an event tree made up of fault trees? Does one count the number of gates and basic events? Does one count the number of levels in the structure? How does one score the combination of operators so as to distinguish the difference in complexity between $F = -a+b+(c*d)$ and $G = -(a*(b+c) \text{ xor } d)$? Does one count the number of independent sub-trees, the number of branches for each gate, the number of negations? This is not simply an academic question. By understanding the complexity of the problem space, a set of heuristics can be chosen quickly, instead of randomly trying one after another.

In computer science, the analogous problem exists in measuring the “size” of a program. There are many putative measures being used: SLOC, McCabe’s Cyclomatic Complexity, NPATH, Halstead Software Science are examples of some standard metrics. The authors are investigating such metrics, which will be detailed in a forthcoming technical report.

7 Conclusion

> *There is no single development, in either technology or management technique, which by itself promises even one-order of magnitude improvement within a decade in productivity, in reliability, in simplicity.*
>
> *--- Fredrick P. Brooks, Jr., NO SILVER BULLET*

In this article, we have reported the results of a comparative study of two technologies to assess risk models: the classical approach, widely used and trusted, based on minimal cutsets and the BDD approach, improved by means of heuristics, that in making no approximations, gives exact results. The study is based on an actual linked-fault tree model representing an event tree coming from the nuclear industry. We used the Aralia computation engine which implements both approaches as well as many heuristics and formula rewriting strategies.

Indeed, definitive conclusions cannot be drawn from a single example. However, our test case is sufficiently large and representative and the results are sufficiently clear to make the following observations.

- The approximation that consists in taking into account failure branches only should be avoided. Our experiments show that, even corrected by a factor obtained from success branches, this approximation overestimates, very often by orders of magnitude, the probability of the sequence.
- The assessment of the probabilities of the sequences through the minimal cutsets should be considered with care. In a significant number of cases, this approximation gives optimistic results, because of truncations. Such an underestimation of the risk is not acceptable. Moreover, the same truncation that gives an optimistic result in one sub-system may give a very pessimistic result in another sub-system, within the same top event! With a truncation set to 10^{-4}

times the probability of the sequence, we observed, among the 181 sequences of our test case, results that are optimistic by a factor 4 together with results that are pessimistic by a factor 96.

- Such variations make the ranking of sequences according to their contributions to the overall risk delicate, if not dubious.
- The classical approach overestimates the likehood of a core damage by almost a factor 5.
- Because of imprecision on the values of probabilities (when computed from the cutsets), the rankings of basic events induced by importance factors should be considered with care. This remark is especially important for the so-called risk achievement worth that can miss important basic events.

The above observations do not mean that existing PRA studies based on event trees must be discarded. Nevertheless, they are a stone in the garden of the classical approach based on minimal cutsets. On the other hand, we don't claim that Binary Decision Diagrams are the universal panacea. This technique still suffers from the exponential explosion of memory requirements. We have shown that heuristics can be designed which improve dramatically its efficiency. They are however hard to tune. More experiments, more efforts are necessary to make our approach able to deal with all the existing models.

One final note: as Fred Brooks' quotation states at the beginning of this section, we should expect no silver bullet to slay the werewolf of complex computations. What is important is that any solution is (1) productive, that it allows us to work orderly and rationally, (2) reliable, that it gives us the correct answers with explicit knowledge of the error bounds, and (3) simple, that it allows us to confirm the problem we are solving and its solution. This study is a step towards this goal.

8 Bibliography

[AD00] J.D. Andrews and S.J. Dunnett, Event Tree Analysis using Binary Decision Diagrams, *IEEE Transactions on Reliability*, Vol 49, No 2, June 2000, pp230-239.

[ARb00] Arboost Technologies. *ARALIA 4.2e Users' Manual*, February, 2002.

[BW96] B. Bollig and I. Wegener. Improving the Variable Ordering of OBDDs is NP-Complete. *IEEE Trans. on Software Engineering*, 45(9):993–1001, Sep. 1996.

[BA01] E. Borgonovo and G.E. Apostolakis, A new importance measure for risk-informed decision making *Reliability Engineering and System Safety*, Volume 72, Issue 2 , pp 193-212, May 2001.

[Bry86] R. Bryant. Graph Based Algorithms for Boolean Function Manipulation. *IEEE Transactions on Computers*, 35(8):677–691, August 1986.

[BRB90] K. Brace, R. Rudell, and R. Bryant. Efficient Implementation of a BDD Package. In *Proceedings of the 27th ACM/IEEE Design Automation Conference*, pages 40–45. IEEE 0738, 1990.

[Bry92] R. Bryant. Symbolic Boolean Manipulation with Ordered Binary Decision Diagrams. *ACM Computing Surveys*, 24:293–318, September 1992.

[BRKM91] K.M. Butler, D.E. Ross, R. Kapur, and M.R. Mercer. Heuristics to Compute Variable Orderings for Efficient Manipulation of Ordered BDDs. In *Proceedings of the 28th Design Automation Conference, DAC'91*, June 1991.

[DR00] Y. Dutuit and A. Rauzy. Efficient Algorithms to Assess Components and Gates Importances in Fault Tree Analysis. *Reliability Engineering and System Safety*, 72(2):213–222, 2000.

[ER04] S. Epstein and A. Rauzy. LukeTreeWalker: Specifications. Technical Report TR2004-01. ARBoost Technologies.

[FFK88] M. Fujita, H. Fujisawa, and N. Kawato. Evaluation and Improvements of Boolean Comparison Method Based on Binary Decision Diagrams. In *Proceedings of IEEE International Conference on Computer Aided Design, ICCAD'88*, pages 2–5, 1988.

[FV72] J.B. Fussel and W.E. Vesely. A New Methodology for Obtaining Cut Sets for Fault Trees. *Trans. Am. Nucl. Soc.*, 15:262–263, June 1972.

[KH96] H. Kumamoto and E. J. Henley. *Probabilistic Risk Assessment and Management for Engineers and Scientists*. IEEE Press. 1996. ISBN 0-7803-6017-6.

[JK98] W.S. Jung and D.K. Kim. FORTE: a fast new algorithm for risk monitors and PSA. *Proceedings of the Fourth International Conference on Probabilistic Safety Assessment and Management*. New York, USA, p. 1221, 1998.

[JHH04] W. S. Jung, S. H. Han and J. Ha. A fast BDD algorithm for large coherent fault trees analysis. *Reliability Engineering and System Safety*. Vol. 83, pp 369-374, 2004

[MIY90] S. Minato, N. Ishiura, and S. Yajima. Shared Binary Decision Diagrams with Attributed Edges for Efficient Boolean Function Manipulation. In L.J.M Claesen, editor, *Proceedings of the 27th ACM/IEEE Design Automation Conference, DAC'90*, pages 52–57, June 1990.

[Min93] S. Minato. Zero-Suppressed BDDs for Set Manipulation in Combinatorial Problems. In *Proceedings of the 30th ACM/IEEE Design Automation Conference, DAC'93*, pages 272–277, 1993.

[MJF99] R. Murgai, J. Jain and M.Fujita, Efficient Scheduling Techniques for ROBDD Construction, *Proceedings of the International Conference on VLSI Design*, pp 394-401, 1999.

[Pap98] I.A. Papazoglou. Mathematical foundations of event trees. *Reliability Engineering and System Safety*, 61:169–183, 1998.

[Rau93] A. Rauzy. New Algorithms for Fault Trees Analysis. *Reliability Engineering & System Safety*, 05(59):203–211, 1993.

[Rau01] A. Rauzy. Mathematical Foundation of Minimal Cutsets. *IEEE Transactions on Reliability*, volume 50, number 4, pages 389-396*,* 2001.

[Rau03] A. Rauzy. Towards an Efficient Implementation of Mocus. *IEEE Transactions on Reliability*, vol. 52:2, pp 175-180, 2003.

[Rud93] R. Rudell. Dynamic Variable Ordering for Ordered Binary Decision Diagrams. In *Proceedings of IEEE International Conference on Computer Aided Design, ICCAD'93*, pages 42–47, November 1993.

[WW96] I.B. Wall and D.H. Worledge. Some perspectives on risk importance measures. In *Proceedings of the international conference on Probabilistic Safety Assessment, PSA'96*, pages 203–207, 1996.

Acknowledgements:

The authors would like to thank K. Kawai, J. Sakaguchi, and K. Nishi from MHI for their help, D. Johnson and D. Wakefield from ABS Consulting, R. Cooke from TU Delft, K. Fleming from Technical Insights Inc., and M. Barrett from Duke Energy for careful reviews of this article and their support, and A. Soma from ABS Consulting Japan for her help.

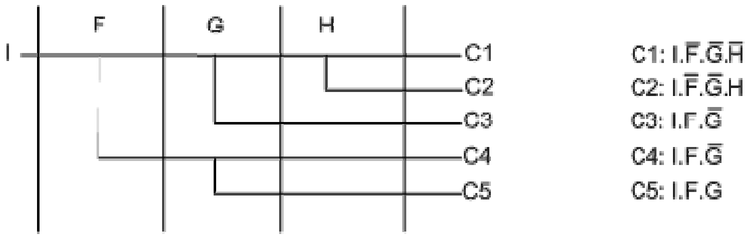
F
G
H
I
C1
C2
C3
C4
C5
C1: $I.\overline{F}.\overline{G}.\overline{H}$
C2: $I.\overline{F}.\overline{G}.H$
C3: $I.F.\overline{G}$
C4: $I.F.\overline{G}$
C5: $I.F.G$

Fig. 1. An event tree and Boolean formulae associated with its sequences.

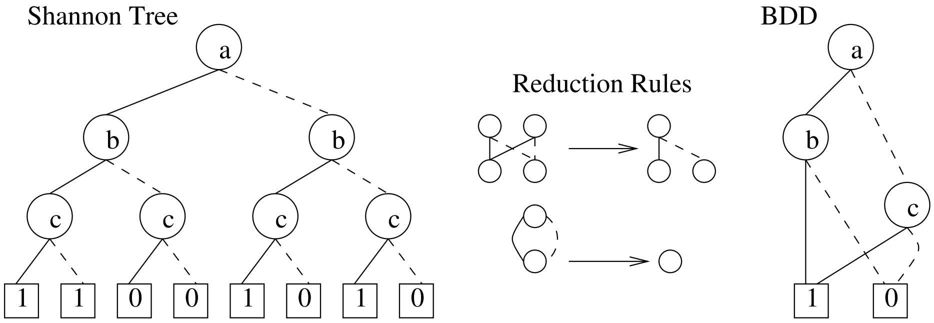

Figure 2. From the Shannon Tree to the BDD.