WWW.BDD-PORTAL.ORG

Ch. Meinel, A. Wagner

Abstract

There is a need for specialized World Wide Web portal sites to facilitate access to resources for specific research communities. We present the portal for the Decision Diagram research community we have built, featuring links to important resources, a powerful literature search and a sophisticated system that allows use and evaluation of Decision Diagram heuristics online via the Internet. A discussion of the difficulties involved in such a project is given and our solutions are presented.

1 Introduction

Today the Internet offers the possibility of standardized and global communication without a need for special hardware or expensive infrastructure. While a lot of resources are available on the World Wide Web (WWW), there is no central place to access these resources in a convenient way. There are search engines, but they are indexing only small segments of the web. A recent survey ([LG99]) had the search engine with the best coverage at only 13 percent of all pages. And even where search engines reach, the results of their searches are unstructured and often the user is presented with an all-or-nothing situation, where queries either return far too many hits or none at all. One solution is the recent notion of so called "portal"-sites, that organize WWW contents into categories and in some cases even grade the quality. But what efforts are underway are mostly targeted at the general public and relatively small communities, like research communities, are not commercially interesting to the existing portals.

To address the needs of the Decision Diagram research community, we have created a specialized portal site that provides easy access to researchers, events and literature and other material in the area of Decision Diagrams. Furthermore this portal includes a system for the online "publication" of Decision Diagram functionality by allowing the use of heuristics and tools via a WWW interface. The system is mainly aimed at facilitating the process of determining which heuristic is best suitable for a special application and at allowing easy comparison between heuristics. It is open to all researchers who whish to publicize their results in this way. As applications using Decision Diagrams are usually quite memory and cpu intensive, the system represents a major effort and greatly facilitates the evaluation of research results in Decision Diagram heuristics. Given the importance of Ordered Binary Decision Diagrams (OBDDs [Bry86], for an overview see for example [MT98]) and other types of Decision Diagrams in verification, simulation, synthesis and similar uses, such an effort seems to have been overdue.

In this paper we present our specialized research portal for the ares of Decision Diagrams. While the main focus is on Binary Decision Diagrams, other types are represented as well. Section 2 explains the need for specialized portals by discussing the shortcomings of conventional portals in greater detail. Section 3 presents the specialized Decision Diagram portal. Section 4 discuses the question of online use of Decision Diagram heuristics, including the need for such possibilities. Section 5 presents the part of the portal that allows online use of heuristics and discusses technical aspects of our solution. Section 6 discusses further work and the paper finishes with the conclusion in section 7.

2 Research Portals

The WWW has great potential for becoming something like a "super library", allowing access to all information somebody finds relevant enough to put up on some web pages. Today its size and growth are impressive and the basic technology for accessing the WWW and maintaining a presence on the WWW seem to be adequately developed. But by its very nature as a decentralized, somehow "anarchistic" medium, the *structure* of the WWW is lacking. Despite the existence of search engines it is often still difficult to find specific content. If one solely relies on traditional search engines the situation is likely to get even worse [LG99], and especially so for research content of limited public interest. We will now discuss some of the reasons.

Structure of a Conventional Portal

By a *conventional portal* we mean an usual search site, like for example Yahoo! [YAH], that supports general (raw) searches and in most cases has a set of prearranged categories where handpicked sites are accessible in a more structured way. Figure 1 depicts the basic structure of such a conventional portal. The main information gathering mechanism is shown to the right: A web-robot or "spider", that searches known pages within the WWW for embedded links to other pages and includes the pages linked into the database of known pages. Usually there is also some manual insertion of good starting points for these searches when new popular sites come online. As a third mechanism to insert new pages into the database most conventional portals also offer the possibility that pages might be registered by somebody, usually the site owners.

The central component of a conventional portal is the database with its search mechanism. This component processes the search requests users are submitting. In most cases there is also a database of paid ad-

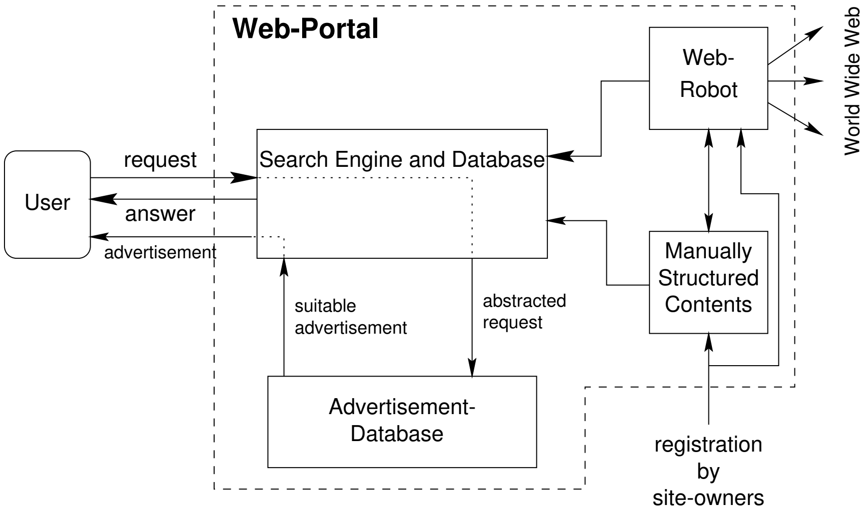

Figure 1: A Conventional Portal

vertisements that are sent back with the pages found for the users requests. Often additional effort is spent to match the advertising to the request the user has made.

Shortcomings of Conventional Portals

Exactly here the inadequacy of a conventional portal manifests itself: The primary reason for operating a conventional portal is to earn money. This is done by being paid for advertisements sent to the user with his answers. The main focus is hence to attract as many people as possible that are interested at least in some of the advertising handed out by the portal. And here the time for manual improvement of the content will be spent. Web pages of interest only to a small group like researchers in a specific field will usually only be added to the database if they are registered manually by their owners. With the number of conventional portals out there, registering pages to most is a significant amount of work.

For these reasons we believe there is need for small, specialized portal sites that support specific research communities. These specialized portals should contain a significant amount of structured and selected contents aimed at the specific needs of the research community addressed.

3 Our Decision Diagram Portal

As our research group is involved in the creation of algorithms and heuristics for Decision Diagrams, we decided to build a specialized portal for this area. This effort is mainly targeted at *Binary* Decision Diagrams, but not exclusively so.

Binary Decision Diagrams (BDDs for short, see for example [MT98]) are a data structure that has had a great impact on the power of tools for simulation, verification and synthesis of digital systems. There are other applications as well, but this is by far the most prominent. Whereas techniques like Model Checking had to enumerate elements of sets more or less explicitely before the discovery of BDDs by R. Bryant ([Bry86]) these techniques can now handle sets in a symbolic way, thus greatly improving the size of problems that can be handled on computers.

The structure of our specialized Decision Diagram Portal can be found in figure 2. The topmost component is a list of links to homepages of active researchers. "Active" means that they have at least one publication in the area of Decision Diagrams. This list aims to be a complete representation of the Decision Diagram research community. Where no homepage could be found, an email address is provided instead.

The second component is a collection of all relevant events in Decision Diagrams, like conferences and workshops with dates, deadlines and links to homepages. This collection is frequently updated.

The third component is a database of literature on Decision Diagrams available on the WWW. This database contains information about technical reports, papers and other texts which can be searched in their full text, even if the documents are only available in Postscript or PDF (portable document format). Search results will include the first lines of text for documents found and links to their original location. The documents for this database are collected by a specialized robot, that is driven of the links to homepages of researchers. The robot is capable of recognizing document formats and creates a searchable index from them. This component represents a significant improvement compared to conventional search engines which can cannot access the contents of non-html documents.

The fourth component is an effort to address some

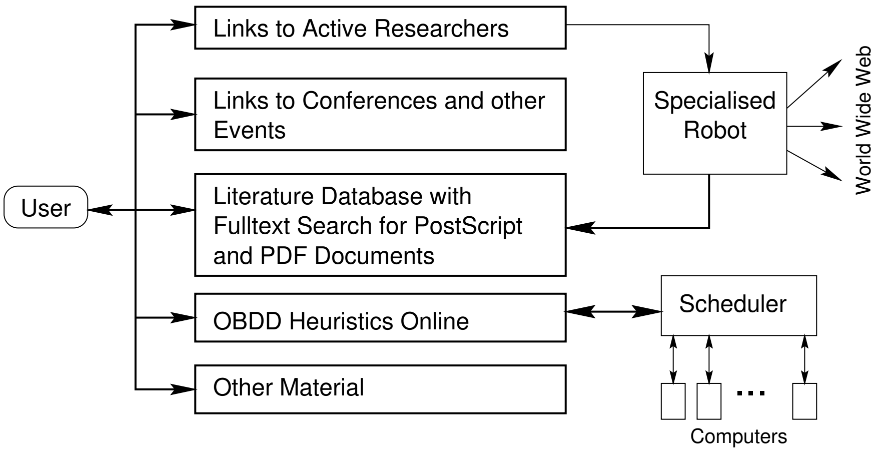

Figure 2: The Decision Diagram Portal

of the specific problems with the publication of results from research in heuristics for Decision Diagrams, namely the specific problems of describing the characteristics of a Decision Diagram heuristic purely on paper. The next section describes this component in detail. Other material, like links to relevant benchmarks, will be added to the site as well.

4 Online Use of Heuristics

When new Decision Diagram heuristics are developed, they are usually added to some existing software package. After optimization, when a sufficient level of performance is reached, the algorithm is presented to the community. Evaluation is usually done using benchmarks from a standard benchmark set first, for example the ISCAS 89 ([BBK89]) set of circuits. Due to the nature of the problem, performance both in speed and memory consumption can vary extremely between different circuits of the benchmark. Predictions about the behavior of a heuristic on other circuits are extremely difficult or impossible. Because of this there often arises the whish to allow others to use the software, to evaluate the suitability of the new heuristic for their own problems.

To Publicize Code is not Optimal

But here the problem manifests itself: Let us consider giving away the source code:

- The code usually only works within the software package it was written for.
- The code has research level quality and may still lack documentation or even worse contain errors.
- Due to some ideas that are not yet published, one may not want to give away all the details of the implementation.

So this is clearly not an appealing option. But giving away executables is no improvement either:

- The software would have to be ported to different platforms.
- There would be need for support for every of these platforms.
- There would be need to develop the software to a higher level of stability.

An other option is reimplementation of the algorithm by interested parties. But this requires a larger effort and special experience. Almost nobody is willing to undertake such an effort for an uncertain outcome.

Further complications arise because the published results are difficult to reproduce as fine-tuning is often done and usually critical for good results. This makes comparisons outside the published results extremely difficult. We think most of these issues can be addressed in a satisfying way by a possibility to access the tools containing heuristics using the Internet.

We believe that some requirements need to be addressed to make an Internet based solution truly usable. Ease of use is achieved by using a WWW interface, but other aspects are important as well. As the computations done are time and memory intensive (i.e. taking hours of cpu-time and up to hundreds of megabytes of memory), just writing some cgi-scripts is clearly inadequate. A significantly larger effort is needed even if the system is only used for the evaluation of heuristics. Usability would be critically lowered if getting computations done would take very long or only very small "toy" examples could be computed.

Flexibility

It should be possible to add tools of various nature to the system with reasonable effort. We do not want to limit the system to a specific tool. If some researcher has a new heuristic and wants to publish it via our system, the amount of customization necessary should be as low as possible, and ideally no more than a recompilation for the computational platform we use should be necessary.

Speed

To achieve an appropriate overall speed it is necessary to distribute the actual computations to several computers. The number of computers should be easily adjustable, and it should even be possible to include computers that are remote and only reachable by an Internet connection. To maintain comparability of the results, it is necessary that the pool of computers can be divided into groups of machines with equal power.

Reliability

As computations can take hours and as there might be a number of pending computations, there should be a mechanism that allows automated recovery from crashes without loss of submitted requests.

5 Solution for Online Use

We have developed a system, called OHO (for OBDD Heuristics Online), that meets the requirements mentioned.

The WWW Interface

The system is accessed by a web interface. After an introductory page containing general information, the user can access a menu that allows the selection of submission forms for individual tool and heuristic types, as well as browsing the specific documentation. When one of the possible types of computation is selected, the user is taken to a submission form where all the relevant information can be entered. This includes an (optional) e-mail address for notification about results, tool options and selection of input data. Options include the choice between the featured heuristic for reordering or a reference heuristic. Input data can consist of a predefined circuit or an user transferred file. File upload is done with the browser, a feature found in most Web browsers.

After submission the request is queued and computations begin as soon as the needed resources become available. If a notification email address has been given, the results will be emailed immediately after completion.

Technical Details

The basic structure of the system is simple. Individual tools have individual user interfaces, that have the task of accepting new requests, giving feedback about the status of a computation submitted previously, and returning the results to the submitter. A central scheduler manages every submitted request during its whole life-cycle. A number of computers perform the actual computations necessary to complete a request. If a computation is terminated because of unavailable resources due to other use of the computer it was started on, the request can be rescheduled on an other computer or restarted later. We use Linux as operating system.

To achieve **flexibility** every contained tool is fitted with a small wrapper script that controls its I/O. As long as the tool communicates via command line, the standard input and output and files, customization of these scripts is quite easy. We believe this approach covers the majority of research tools.

Speed is achieved by distributing the computations, as mentioned earlier. This distribution is managed by the central scheduler, in such a way, that non-dedicated computers can be used. Queuing of requests is performed here as well.

As requests are queued and queues might get longer, **reliability** becomes important. In order to achieve this, the scheduler does frequent dumps of the current system status. In case of a crash the scheduler will automatically recover from these dumps and resume computations without the need for manual intervention.

Available Decision Diagram Tools

At the moment we have integrated current heuristics added to nanotrav (part of the CUDD system, [Som96]) by a number of researchers. These heuristics are not part of the CUDD standard distribution, but rather actual research code by the developers of the specific heuristics. There are as well some heuristics that are part of the CUDD standard distribution for the purpose of comparison. We also have heuristics online that where added recently to the well known model checker SMV ([oCS]). At the moment these are only heuristics implemented by our research group.

We want to incorporate more tools and heuristics in the near future and invite every researcher that wants to publish a specific heuristic or tool within our portal to contact us. As stated above, we do not need access to source code and adding usual research tools is not a problem.

6 Related Work

As far as the authors know there is no other effort to create a specialized portal site for Decision Diagram research and there are only very few sites targeted at specific research communities. The authors also do not know of any effort to make BDD-tool functionality available via the WWW on a comparable scale. There seem to be a few efforts capable of computing only extremely small examples.

There have been efforts to use the WWW as an unified interface to a heterogenous set of EDA tools. In [BBMR96] such an application has been described. The main focus there is the integration of different tools on different platforms with different data formats into a seemingly homogenous environment. The framework is intended to do actual work rather than to allow the comparison of the power of different approaches to a specific problem (like BDD heuristics) as in our case.

In [Ku97] the vision of integrating a great number of different EDA CAD services running in different places into one "global" EDA CAD system using the Internet has been described. The main techniques here are the use of proxies to abstract the actual tool being used. In [SK95] an other approach to the integration of different EDA tools is described. This approach is centered mostly on the notion of "active messages".

7 Conclusion and Future Plans

We have built a WWW portal for the Decision Diagram research community. This portal includes links to researchers and events. It also includes a specialized literature search mechanism far superior to that of current search engines. The collected information greatly facilitates retrieving information in the area of Decision Diagram research from the WWW. No comparable collection of online information is known to the authors at this time. As an important component the portal contributes a powerful and easy to use system that allows the online evaluation of a number of current heuristics for Decision Diagrams in different tools. This system is open to anyone. Inclusion of further tools and heuristics into the system is easy and contributions by researchers in the area of Decision Diagrams are welcome. We believe that this site is an important step in improving the usage of the WWW for Decision Diagram research and may even have some exemplary character for other research areas.

We intend to add links to further WWW resources for Decision Diagrams. We would like to establish this portal as a central announcement and link site for all information in the WWW relevant to Decision Diagrams. We are thinking about a mechanism that allows remote administration of the conference and workshop database by the organizers of the events.

More Decision Diagram heuristics and tools will be added to the section that allows online evaluation. There might also be other kinds of interfaces than the WWW interface that would prove beneficial.

The portal can be accessed at `http://www.bdd-portal.org`

References

[BBK89] F. Brglez, D. Bryan, and K. Kozminski. Combinational profiles of sequential benchmark circuits. In *Proc. IEEE Int. Symposium on Circuits and Systems*, pages 1929–1934, May 1989.

[BBMR96] A. Bogliolo, L. Benini, G. De Micheli, and B. Ricc. Ppp: A gate-level power simulator - a world wide web application. Technical Report Stanford Technical Report No. CSL-TR-96-691, Stanford University, 1996.

[Bry86] R.E. Bryant. Graph Based Algorithms for Boolean Function Manipulation. *IEEE Transaction on Computers*, C-35:677–691, 1986.

[Ku97] D. Ku. Embedded tutorial: Eda and the network. In *Procceedings of ICCAD 97*, 1997.

[LG99] S. Lawrence and C.L. Gilles. Accessibility of information on the web. In *Nature, Volume 400, Number 6740*, 1999.

[MS97] C. Meinel and A. Slobodová. Speeding up variable ordering of obdds. In *Proc. of IEEE ICCD'97, Austin(Texas,USA)*, 1997.

[MS98] C. Meinel and A. Slobodová. Sample method for minimization of obdds. IWLS'98, Granlibakken Resort, Lake Tahoe, USA, 1998.

[MT96] C. Meinel and T. Theobald. Local encoding transformations for optimizing obdd-representations of finite state machines. In *Proc. FMCAD'96, LNCS 1166*, pages 404–418. Springer-Verlag, 1996.

[MT98] C. Meinel and T. Theobald. *Algorithms and Data Structures in VLSI Design*. Springer, 1998.

[MW99] C. Meinel and A. Wagner. Evaluation of obdd-heuristics via the internet. IWLS'99, Granlibakken Resort, Lake Tahoe, USA, 1999.

[oCS] CMU School of Computer Science. Formal methods - model checking. http://www.cs.cmu.edu/˜ modelcheck/.

[SK95] M. J. Silva and R. H. Katz. The case for design using the world wide web. In *Procceedings of the 32th Design Automation Conference*, 1995.

[Som96] F. Somenzi. Colorado university decision diagram package. ftp://vlsi.colorado.edu/pub/, 1996.

[YAH] Yahoo! http://www.yahoo.com.