

Multipartite Tables in JBits for the Evaluation of Functions on FPGA

Jérémie Detrey, Florent de Dinechin

► **To cite this version:**

Jérémie Detrey, Florent de Dinechin. Multipartite Tables in JBits for the Evaluation of Functions on FPGA. [Research Report] RR-4305, LIP RR-2001-44, INRIA, LIP. 2001. inria-00072282

HAL Id: inria-00072282

<https://hal.inria.fr/inria-00072282>

Submitted on 23 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Multipartite Tables in JBits
for the Evaluation of Functions on FPGA***

Jérémie Detrey, ENS-Lyon

Florent de Dinechin, LIP/ENS-Lyon

No 4305

Novembre 2001

_____ THÈME 2 _____

 ***R**apport
de recherche*

Multipartite Tables in JBits for the Evaluation of Functions on FPGA

Jérémie Detrey, ENS-Lyon
Florent de Dinechin, LIP/ENS-Lyon

Thème 2 — Génie logiciel
et calcul symbolique
Projet Arénaire

Rapport de recherche n° 4305 — Novembre 2001 — 11 pages

Abstract: This paper presents a core generator for arbitrary numeric functions on Xilinx Virtex FPGAs. The cores use the state-of-the-art multipartite table method, which allows input and output precisions in the range of 8 to 24 bits on current Virtex chips. The implementation uses the JBits API to embed elaborate optimisation techniques in the description of the hardware.

Key-words: Function generator, FPGA core, multipartite method, JBits

(Résumé : tsvp)

This text is also available as a research report of the Laboratoire de l'Informatique du Parallélisme
<http://www.ens-lyon.fr/LIP>.

Des tables multipartites en JBits pour l'évaluation de fonctions sur FPGA

Résumé : Cet article présente un générateur de cœurs de circuits pour le calcul de fonctions numériques arbitraires sur FPGA Xilinx Virtex. Les cœurs obtenus utilisent la méthode multipartite, qui permet des précisions d'entrée et de sortie de 8 à 24 bits sur les circuits Virtex actuels. L'implémentation profite de l'API JBits pour intégrer, dans le code décrivant les cœurs, des techniques d'optimisation sophistiquées.

Mots-clé : Générateur de fonction, cœur FPGA, méthode multipartite, JBits

1 Introduction

Many FPGA applications require the evaluation of some unary functions such as trigonometric (sine/cosine, tangent...), algebraic (square or square root, cube or cubic root, ...) or transcendental (exponential or logarithm). The fastest implementation consists in tabulating all the values, which is impractical for precisions higher than a few bits, because the size of the table is exponential in the input size. There exist compact – but slow – implementations for some functions, such as CORDIC-like algorithms for trigonometric and exponential/logarithm. More general methods for function evaluation, based on polynomial approximation, rely on multipliers and are therefore less suited for FPGA implementations.

1.1 Table-based methods

Pioneered by DasSarma and Matula [6], these methods reduce the size of a straightforward table implementation by exploiting the property of continuity of the function. The latest developments in this area is the multipartite method by Dinechin and Tisserand [1], for precisions of 8 to 24 bits. For precisions higher than 24 bits, Piñeiro *et al* recently demonstrated a method based on a second degree polynomial approximation using a squarer unit and a multiplier, which is more area-efficient and probably faster [5].

This paper therefore focuses on the FPGA implementation of the multipartite method, for arbitrary functions, and for precisions of 8 to 24 bits. Our first contribution is to show that this method is well suited to FPGAs of the Xilinx family: In these circuits, the tables can be built efficiently out of the FPGA LUTs, and the built-in fast carry logic optimises the additions.

1.2 JBits

The second contribution of this paper is the use of the JBits framework [11] for implementing function evaluator cores on FPGAs. JBits is an Application Programming Interface (API) developed by Xilinx for programming FPGAs of the 4000 and Virtex series. It comes as a set of Java classes. One of its main strong points is to allow a detailed, structural description of a circuit down to the CLB level. Its main drawback is that it imposes this structural description, even requiring the user to specify all the placement. However, the modern object-oriented features of Java, and the availability of a router [2] lighten the task to a very acceptable level.

Another strong point of the JBits framework is a tight and natural integration in the same language of the hardware and software parts of an application. Originally, this integration is mainly aimed at run-time reconfigurability: it allows to embed hardware objects in a software which reconfigures the hardware on the fly [7, 4, 8]. We did exploit this for testing our cores. However we also exploited it the opposite way, to embed elaborate optimisation techniques in the code for our hardware. We are thus able to produce cores where both the abstract architecture, and its FPGA implementation are optimised for the required function and precision.

To our knowledge, the function generator that we have developed is the most complex core ever developed in JBits, along with the DES implementation presented by Patterson [4]. We hope that our experience may be of benefit for prospective users of this technology.

1.3 Outline of the paper

Section 2 describes the multipartite method used to generate architectures for arbitrary numeric functions. Section 3 discusses the JBits implementation details, focusing on the opportunities and constraints of this framework. Section 4 describes and comments the results of this methodology. Section 5 draws conclusions and suggests future works.

2 Function evaluation using the multipartite method

2.1 The bipartite method

First presented by Das Sarma and Matula [6] in the specific case of the reciprocal function, this method consists in approximating the function by affine segments, as illustrated on Figure 1. The 2^α segments are selected by the α most significant bits of the input word. Instead of tabulating all the 2^{w_I} values of the function (where w_I is the width of the input word), it is possible, for each segment, to tabulate one initial value, and to construct the other values by adding, to this initial value, an offset computed by an linear approximation using the $w_I - \alpha$ least significant bits of the input word. The idea behind the bipartite method is to group the segments into 2^γ (with $\gamma < \alpha$) larger intervals (4 on the figure) such that the slope of the segments is considered constant on each larger interval. Now there are only 2^γ tables of offsets, each containing 2^β offsets. Altogether, we thus need to store $2^\alpha + 2^{\gamma+\beta}$ values instead of $2^{\alpha+\beta}$. The corresponding architecture is depicted on Figure 1.




Figure 1: An example of bipartite approximation and architecture, for 6-bit input: $w_I = 6$, $\alpha = 4$, $\gamma = 2$, $\beta = 2$.

2.2 The multipartite method

It is possible to exploit the bipartite idea more than one time, splitting the input word into more subwords. Schulte and Stine [10] and Muller [3] independently found two different ways to do it, and Dinechin and Tisserand [1] unified both approaches in an algorithm that explores the whole implementation space, leading to maximal table reduction. The typical multipartite architecture is presented on Fig. 2.

The algorithm presented in [1] (too complex to be presented here) ensures that the cumulated approximation and rounding errors sum up to less than one LSB. Table 1 presents some table reductions achieved by this algorithm. To measure the significance of this method, one should note that the sine cores offered by Xilinx, which use a naive table approach, have therefore a precision limit of 10 bits due to their size. In comparison, Figure 8 shows that a 16 bit sine core using our method occupies a small fraction of a Virtex 1000.

Precision	8 bits	12 bits	16 bits	20 bits
Size uncompressed	2,048	49,152	1,048,576	20,971,520
Number of tables	3	4	4	4
Total size compressed	224	1,552	8,960	50,176

Table 1: Table compression for the sine function on $[0, \pi/4]$ (sizes in bits)




Figure 2: A multipartite architecture. The output of the tables are summed by a multi-operand adder. The XORs are a trick due to Schulte and Stine [10] that allows to halve the size of each table: Exploiting the symmetry of each segment to the function (see Figure 1), one of the input bit can be considered as a sign bit.

2.3 Multipartite tables and FPGAs

The initial implementation of the multipartite method was able to output VHDL which was then synthesised with Leonardo. Among the interesting results, it was remarked that:

- Architectures like that shown on Figure 2 lend itself to efficient implementation on Virtex devices, thanks to their LUT-based structure and the availability of fast adder circuitry;
- The size compression is so drastic that usually, it also entails a speed improvement over the naive table implementation in spite of the adders;
- The VHDL synthesiser was able to compress the tables even further using logic optimisation techniques. This was measured as a *bit per LUT* factor [1] which could be as high as 18 (a LUT holds 16 bits, and some of them are used as multiplexers or adders).

These observations suggested that a specific core implementation should be developed to improve on these results. The JBits framework was chosen for this purpose. This is the subject of the rest of this paper.

3 JBits implementation

This section presents our JBits implementation of the previous multipartite method, first describing the general structure of the core, then focusing on a table compression heuristic and concluding with the floor-planning problem.

3.1 Overview of the architecture

Looking back at Figure 2, it can be seen that most area of the core will be dedicated to lookup tables, with the multi-operand adder and the rows of XORs occupying little area.

In Virtex devices, the best option for the multi-operand adder is also the simplest, a row of simple adders using the built-in fast carry logic. This way the multi-operand adder is a convenient rectangular area. The XORs are also naturally built as columns of LUTs.

The rest of this section focuses on implementing a look-up table. Two JBits classes have been written for this purpose. The first one implements uncompressed, and therefore regularly placed and theoretically faster [9] tables. The other one compresses the table using Virtex-specific binary optimisation techniques. Its drawback is that the placement is no longer regular.

3.2 Uncompressed tables

To build a table addressed by w_I bits, we first consider independently each output bit, and fill the leaves of a multiplexer tree with the values of the bit according to the address. The multiplexer tree is a simple binary tree, each of its node being a multiplexer controlled by a bit of the address, and its leaves being LUTs used as small 4-bit addressed memories. The Virtex architecture provides multiplexers (called F5 and F6) specifically designed for arranging 2 or 4 LUTs in a CLB as a bigger look-up table with 5 or 6 input bits. Thus the level of the tree closest to the leaves uses F5 multiplexers, and the next level uses F6 multiplexers. Subsequent levels use LUTs as multiplexers, as shown on Figure 3.

The placement of the CLB tree is then performed recursively, by placing each CLB between its sons. This way, routing is simple and tables assemble in big rectangles.




Figure 3: A multiplexer tree in the Virtex architecture




Figure 4: Placement of the multiplexer tree. This row corresponds to one output bit.

The tables built by this method are convenient to place and route. However Dinechin and Tisserand have shown that a VHDL synthesiser is able to compress them significantly [1]. The next step was therefore to study table compression. We developed a heuristic that gives excellent results with all kinds of tables. This heuristic is described in 3.3. The main drawback of such optimisation is that it produces a tree of CLBs that we cannot organise as regularly as previously. In VHDL we would simply leave the placement of this tree to the back-end tools. In JBits, however, we had to write an ad-hoc placer, described in 3.4.

3.3 Table compression heuristic

Classically, each output bit of the table can be expressed as a boolean function of w_I variables. Well-known polynomial reduction algorithms may thus be applied to this function. Besides, in our implementation, these algorithms have been targeted to the specific architecture of the Virtex.

3.3.1 From tables to polynomials

The first step is to get a polynomial of the function from the (truth) table, and then, thanks to Karnaugh maps, a minimal expression of it: we build the boolean hypercube, then label each vertex by the value of the function at this point, and look for the biggest sub-hypercubes whose all vertices are labelled by 1. See

Figure 5 for an example. This is obviously an exponential algorithm, but given the small size of the tables, this complexity is not noticeable. For 20-bit tables, it only lasts a couple of minutes.




Figure 5: Reduction example

3.3.2 From polynomials to multiplexer trees

Classical reduction techniques consist in trying to minimise the number of monomials in the polynomial. To target the particular architecture of our FPGAs, we will adapt these known algorithms so that they produce the multiplexer trees with 4-input LUTs at the leaves that we already used previously.

To this effect we will focus on splitting a polynomial by a given variable x_K , which we call a *key variable*. Splitting can be achieved by Shannon's theorem:

$$p(x_0, \dots, x_K, \dots, x_{w_I-1}) = x_K \cdot p_1(x_0, \dots, x_{K-1}, x_{K+1}, \dots, x_{w_I-1}) + \overline{x_K} \cdot p_0(x_0, \dots, x_{K-1}, x_{K+1}, \dots, x_{w_I-1})$$

The architectural interpretation of this equation is a multiplexer, controlled by x_K , with two sub-trees implementing the polynomials p_0 and p_1 .

After splitting the polynomial, we simplify the resulting p_0 and p_1 using the Quine-McCluskey algorithm, and split them using another key variable.

Eventually this algorithm produces functions of 4 variables which are leaves of the multiplexer tree, and will be mapped to LUTs in the FPGA.

The clever part of the heuristic is now to chose the key variables and their order in order to get the minimal multiplexer tree.

3.3.3 Key variable selection

At each step, our heuristic is to chose the key variable that ensures the largest amount of simplification in the Quine-McCluskey simplification of p_0 and p_1 . To achieve that we want to chose the key variable which plays a part in the largest number of smallest monomials.

We therefore count the occurrences of each variable in the smallest monomials of the polynomial, and choose one of the most used ones. An example of this heuristic is given in Figure 6.

3.4 Floorplanning

Compressed tables cannot be placed as regularly as uncompressed ones, due to unbalances in the trees. Our solution was to write a *CLB provider* class that works with a bitmap of the FPGA's area and centralises CLB allocation requests. It tries to avoid any empty space, by allocating the first free suitable space for a given order.

At the moment, the order established between the CLBs is simply a linear order. For instance, you can see in Figure 7 an example of this *CLB provider* placing a CLB tree.

This algorithm does not even try to group CLBs from the same tree. In spite of this lack of regularity, however, the JBits router is always able to obtain a good, congestion free, routing. For this reason we did not try to improve on this rough placement. According to a study by Singh [9], however, more careful placement could lead to an increase in performance by up to 30%, and this will be the subject of further investigation.




Figure 6: Reduction example




Figure 7: Placement example

3.4.1 Graphical interface

In order to synthesise rectangular cores, we have jointly developed a graphical interface to allow the user to specify the dimensions of the bounding rectangle, as it can be seen in the screenshot Figure 8. This interface actually comes on top of a previous interface to the multipartite method.

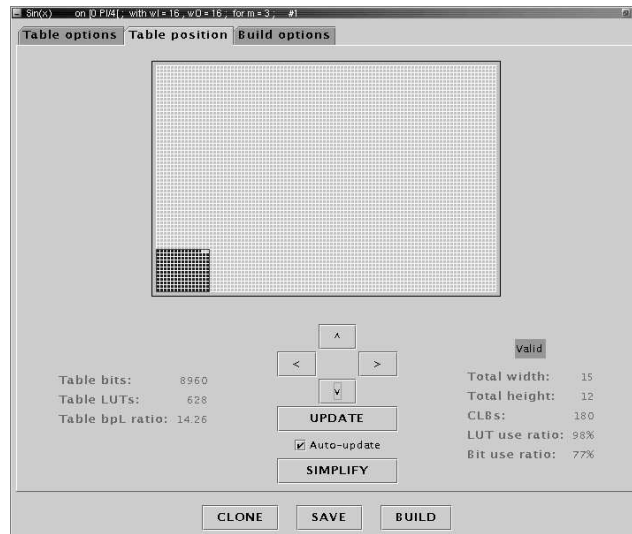


Figure 8: Interface example: the implantation of a 16-bit sine operator on a Virtex XCV1000

3.4.2 Class hierarchy

To sum up this section, Figure 9 describes the class hierarchy of our core generator. The classes have been written in such a way to allow easy reutilisation. Altogether, this represents 3 month of work for a postgraduate student.

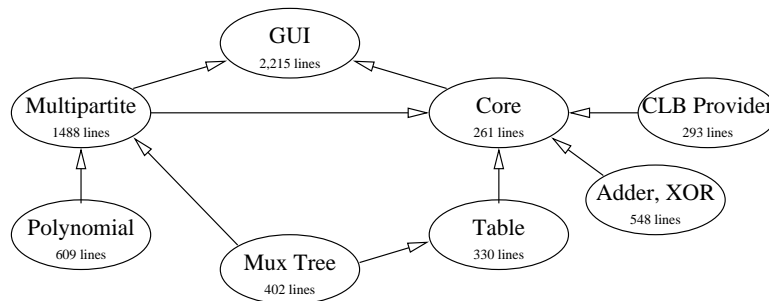


Figure 9: The class hierarchy with the size of each class

4 Results

Some results of our multipartite core generator are given by Table 2. Synthesis was performed on a Pentium 400 with 512 MB of memory. For technical reasons, although we were able to test our cores on a Celoxica RCP1000 board, we were unable to time them on this board so far. The frequencies given in this table are therefore only approximations obtained using the Xilinx FPGA Timing Analyser tool.

function	12 bits sine		16 bits sine	
	uncompressed	compressed	uncompressed	compressed
memory	1,552 bits		8,960 bits	
size of tables	97 LUT	84 LUT	628 LUT	473 LUT
bits per LUT ratio	16.00	18.48	14.27	18.94
total size	160 LUT	147 LUT	710 LUT	555 LUT
frequency	52 MHz	39 MHz	36 MHz	35 MHz
reduction time	—	1''	—	15''
synthesis time	4''	4''	35''	20''
function	16 bits exp		20 bits sine	
	uncompressed	compressed	uncompressed	compressed
memory	11,520 bits		50,176 bits	
size of tables	810 LUT	643 LUT	3,573 LUT	2,546 LUT
bits per LUT ratio	14.22	17.91	14.04	19.70
total size	901 LUT	734 LUT	3,682 LUT	2,655 LUT
frequency	35 MHz	30 MHz	21 MHz	27 MHz
reduction time	—	20''	—	5'40''
synthesis time	40''	30''	9'	4'

Table 2: Timings, area and synthesis time of multipartite cores

Three points need to be underlined:

- The compression algorithm provides an improvement of at least 20% in area, which is better than the results obtained by Leonardo. The reason for that is probably that we performed binary optimisation with the target architecture in mind, instead of separating optimisation and technology mapping.
- The total implantation time is still acceptable, although slightly longer for reduced tables, it . Most of this time is spent in routing using JRoute. Unoptimised tables could be routed much faster by hand-writing the routing algorithm, but this would add a lot to the development time. It is doubtful, however, that cores using the multipartite method will qualify as run-time reconfigurable cores for precisions larger than 12 bits.
- The frequencies shows that the “smaller is faster” effect overcomes the negative effect of bad placement for all cores except the largest ones. This is consistent with Singh’s findings [9].

5 Conclusion and future work

There are three main conclusions to this work.

- Arbitrary numeric functions can be implemented efficiently in LUT-based FPGAs using the multipartite method. The core generator we developed can build a function evaluator for any function with up to 16 bits precision in seconds, and the resulting core will need only a fraction of the FPGA resources. This is in itself a great improvement over currently available cores.
- The JBits API is a great tool for developing such cores, because it allows one to integrate any kind of optimisation within the hardware description of the core. Obviously there is nothing in what we presented that could not have been done using a combination of Java, VHDL, a VHDL synthesiser, backend tools and makefiles, but the JBits approach is much more elegant for this specific problem.
- We are very glad we did not have to route the cores ourselves, and we wish we did not have to place them by ourselves. Manual placement is desirable to improve performance, it is easy when the architecture is regular, but it should be avoidable in the other cases. We hope the next generations of JBits will offer the facility of our CLB provider class, but with a configurable amount of optimisation.

Two natural directions of future work concern the aspect of performance.

- The speed of the bigger cores can very probably benefit from better placement.
- Pipelining should also be explored.

We hope to distribute these cores under the GNU Public Licence soon on www.ens-lyon.fr/LIP/Arenaire/.

Acknowledgements

This work was partially supported by an ACI Jeunes Chercheurs of the French government, and Xilinx University Programme.

References

- [1] Florent de Dinechin and Arnaud Tisserand. Some improvements on multipartite table methods. In Neil Burgess and Luigi Ciminiera, editors, *15th IEEE Symposium on Computer Arithmetic*, pages 128–135, Vail, Colorado, June 2001. Also available as LIP research report 2000-38.
- [2] Eric Keller. JRoute: A run-time routing API for FPGA hardware. In *7th Reconfigurable Architectures Workshop*, Cancun, Mexico, May 2000. LNCS 1800.
- [3] J.M. Muller. A few results on table-based methods. *Reliable Computing*, 5(3):279–288, 1999.
- [4] Cameron Patterson. High performance DES encryption in Virtex(tm) FPGAs using JBits(tm). In *IEEE Symposium on FPGAs for Custom Computing Machines*, Napa Valley, USA, 2000.
- [5] J.A. Piñeiro, J.D. Bruguera, and J.-M. Muller. Faithful powering computation using table look-up and a fused accumulation tree. In Neil Burgess and Luigi Ciminiera, editors, *15th IEEE Symposium on Computer Arithmetic*, pages 40–47, Vail, Colorado, June 2001.
- [6] D. Das Sarma and D.W. Matula. Faithful bipartite ROM reciprocal tables. In S. Knowles and W.H. McAllister, editors, *12th IEEE Symposium on Computer Arithmetic*, pages 17–28, Bath, UK, 1995. IEEE Computer Society Press.
- [7] J. Scalera and M. Jones. A run-time reconfigurable plug-in for the winamp MP3 player. In *IEEE Symposium on FPGAs for Custom Computing Machines*, Napa Valley, USA, 2000.
- [8] J. Scalera and M. Jones. Cores and anti-cores: Using JBits as part of a mainstream design flow. In *8th Reconfigurable Architectures Workshop*, San Francisco, USA, April 2001.
- [9] Satnam Singh. Death of the RLOC? In *IEEE Symposium on FPGAs for Custom Computing Machines*, Napa Valley, USA, 2000.
- [10] J.E. Stine and M.J. Schulte. The symmetric table addition method for accurate function approximation. *Journal of VLSI Signal Processing*, 21(2):167–177, 1999.
- [11] Xilinx Corporation. *The JBits 2.7 SDK for Virtex*, 2001.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399