

Delay Optimization Using SOP Balancing

Alan Mishchenko Robert Brayton

Department of EECS, University of California, Berkeley
{alanmi, brayton}@eecs.berkeley.edu

Stephen Jang

Agate Logic Inc.
sjang@agatelogic.com

Victor Kravets

IBM Corporation
kravets@us.ibm.com

Abstract

Reducing delay of a digital circuit is an important topic in logic synthesis for standard cells and LUT-based FPGAs. This paper presents a simple, fast, and very efficient synthesis algorithm to improve the delay after technology mapping. The algorithm scales to large designs and is implemented in a publicly-available technology mapper. The code is available online. Experimental results on industrial designs show that the method can improve delay after standard cell mapping by 30% with the increase in area 2.4%, or by 41% with the increase in area by 3.9%, on top of a high-effort synthesis and mapping flow. In a separate experiment, the algorithm was used as part of a complete industrial standard cell design flow, leading to improvements in area and delay after place-and-route. In yet another experiment, the algorithm was applied before FPGA mapping into 4-LUTs, resulting in 16% logic level reduction at the cost of 9% area increase on top of a high-effort mapping.

1. Introduction

Delay optimization has been studied extensively since the early days of logic design, as part of both technology independent [22][2][12][21] and technology dependent synthesis [10][15][9][5]. However, existing methods for delay optimization have several known limitations:

- Numerous local changes to the network may be applied, with no guarantee that the delay is globally improved or that additional area has been effectively spent for delay improvements.
- Algorithms of high computational complexity are often used, leading to prohibitive runtime on large designs. Much effort is spent on deciding where to make the changes.
- Structural flexibilities that are available during synthesis and potentially capable of producing a delay improvement may not be exploited by technology mapping.

The method described in this paper overcomes these limitations. Unlike previous methods, it does not perform a sequence of local changes, each one updating the mapped network and then running incremental timing analysis after each change. Instead, the proposed method transforms the subject graph before technology mapping, by minimizing the number of logic levels. A subject graph with structural choices [10][4] can be used as input to the algorithm, resulting in improved quality of results.

The method has been implemented as a straight-forward extension of the publicly available priority-cut-based technology mapper [18]. The extension is described in this paper. The resulting source code is publicly available for unrestricted use and as a benchmark for future comparisons.

The new logic structures for delay optimization are created by transforming logic structure of the cuts in the timing-critical areas. The technology mapper [18] allows for efficient area recovery in the regions where area inevitably grows due to initial logic duplication.

Previous methods in delay-oriented restructuring focused on MUX-based resynthesis, e.g. [2][19], generalized select transform (GST), e.g. [12][21], and various BDD-based techniques, e.g. [5][6]. The proposed method is simpler, scales better, and leads to competitive quality of results. It can also be extended to work for the sequential case, similar to the way delay optimization is done in [23].

The rest of this paper is organized as follows. Section 2 describes the background. Section 3 describes the algorithm. Section 4 reports experimental results. Section 5 concludes the paper and outlines future work.

2. Background

A *Boolean network* is a directed acyclic graph (DAG) with nodes corresponding to logic gates and edges corresponding to wires connecting the gates. The terms Boolean network, netlist, and circuit are used interchangeably in this paper. In this paper, we consider only combinational Boolean networks.

A node n has zero or more *fanins*, i.e. nodes that are driving n , and zero or more *fanouts*, i.e. nodes driven by n . The *primary inputs* (PIs) are nodes without fanins in the current network. The *primary outputs* (POs) are a subset of nodes of the network. A *fanin (fanout) cone* of node n is a subset of all nodes of the network, reachable through the fanin (fanout) edges of the node.

A combinational *And-Inverter Graph* (AIG) is a Boolean network composed of two-input ANDs and inverters. To derive an AIG, the SOPs of the nodes in a logic network are factored, the AND and OR gates of the factored forms are converted into two-input ANDs and inverters using DeMorgan's rule, and these two-input ANDs are added to the AIG manager in a topological order. The *size (area)* of an AIG is the number of its nodes; the *depth (delay)* is the number of nodes on the longest path from the PIs to the POs. The goal of optimization by local transformations of an AIG is to reduce both area and delay.

Structural hashing of AIGs ensures that all constants are propagated and, for each pair of nodes, there is at most one two-input AND with them as fanins (up to a permutation). Structural hashing is performed by hash-table lookups when AND nodes are created and added to an AIG manager. Structural hashing can be applied on-the-fly during AIG construction, which reduces the AIG size.

A *cut* C of a node n is a set of nodes of the network, called *leaves* of the cut, such that each path from a PI to n passes through at least one leaf. Node n is called the *root* of cut C . The cut *size* is the number of its leaves. A *trivial cut* of a node is the cut composed of the node itself. A cut is *K-feasible* if the number of nodes in the cut does not exceed K . A cut is *dominated* if there is another cut of the same node, which is contained, set-theoretically, in the given cut.

Area of a cut is the number of AIG nodes found on the path between the root and the leaves, including the root and excluding the leaves. The concepts of area and the number of AIG nodes are used interchangeably in this paper.

Delay of a cut is the number of AIG nodes on the longest path between the root of the cut and a primary input of the AIG. The concepts of delay, depth, and logic level are used interchangeably in this paper.

A *local function* of an AIG node n , denoted $f_n(x)$, is a Boolean function of the logic cone rooted in n and expressed in terms of the leaves, x , of a cut of n . The *global function* of an AIG node is its function expressed in terms of the PIs of the AIG.

AIGs can efficiently represent both local and global functions. Because of their low memory usage, speed of manipulation and scalability, AIGs have recently emerged as a widely-used data-structure for various applications in logic synthesis and formal verification.

If Boolean functions in some application depend on 16 or fewer inputs, it is often more convenient to use truth tables to represent and manipulate them. For example, a truth table can be efficiently converted into an irredundant Sum-of-Products (ISOP) using a truth-table implementation of the Minato-Morreale algorithm [13][14].

Additional information can be found in the following publications: AIGs [11][3], AIG-based synthesis, [16][17], cut-based technology mapping, delay optimization, and area recovery can be found in [8][7][20][18].

3. Algorithm

This section introduces AND- and SOP-balancing, which are the key ingredients of the proposed algorithm, followed by the overall pseudo-code of the algorithm.

3.1 AND-balancing

AND-balancing of an AIG is a well-known fast transform that reduces the number of AIG levels. AND-balancing is performed in two steps: covering and tree-balancing.

The covering step identifies large multi-input ANDs in the AIG by grouping together two-input ANDs that have no complemented attributes in between and no external fanout, except possibly at the root node of each multi-input AND.

The covering step is illustrated in Figure 3.1.1. The circles stand for two-input ANDs and the small bubbles on the edges stand for the complemented attributes.

The tree-balancing step decomposes each multi-input AND into two-input ANDs while trying to reduce the total number of AIG levels. As the result of this step, a new structure of two-input ANDs is created. This structure is constructed to minimize the delay while taking into account logic levels of the inputs. The tree-balancing step is illustrated in Figure 3.1.2.

It should be noted that the covering step is unique, while the tree-balancing step is not unique and depends on the grouping of the inputs with equal delay, while transforming multi-input ANDs into trees of two-input ANDs.

Because the covering step stops at the multiple-fanout nodes, AND-balancing cannot increase the total number of two-input AND nodes. However, some nodes can be reduced when AND-balancing is applied to a large AIG and logic sharing is created in the process.




Figure 3.1.1: Illustration of the covering step.




Figure 3.1.2: Illustration of the tree-balancing step.




Figure 3.1.3: Illustration of AND-balancing.

Figure 3.1.3 illustrates AND-balancing, which combines covering and tree-balancing. In the above figures, the delays of the PIs are assumed to be 0. The total delay of the AIG in this example is reduced from 5 to 3 levels.

AND-balancing described in this section is implemented in ABC [1] as command *balance*.

3.2 SOP-balancing of a small AIG

In this paper, an AIG is considered *small* if it depends on roughly 10 or less inputs. A small AIG can be converted into an SOP, and then AND-balancing can be applied to each product and the sum. In doing so, the products and the sum are treated as multi-input ANDs and decomposed to minimize the delay of the output node.

Figure 3.2.1 illustrates SOP-balancing for a small AIG, where the delays of the PIs are equal to 0. The total delay of the AIG in this example is reduced from 4 to 3. Note that AND-balancing cannot reduce the delay in this example.




Figure 3.2.1: Illustration of SOP-balancing.

In general, AND-balancing is limited to multi-input ANDs, while SOP-balancing looks at larger functions. As a result, in many cases, SOP-balancing can reduce delay when AND-balancing cannot.

3.3 SOP-balancing of a large AIG

A large AIG, for example, the AIG representing combinational logic of an industrial design, can contain millions of AIG nodes. It is impossible to apply SOP-balancing to such an AIG as a whole, but it is possible to break it down into parts, try SOP-balancing for each part, and if the delay is improved, locally update the large AIG with the structure derived by SOP-balancing.

The latter is, in essence, the SOP-balancing algorithm described in this paper. A self-explanatory pseudo-code is given in Figure 3.3 below.

```

subject_graph performSopBalancing (
  subject_graph S, // S is an And-Inverter Graph
  int K,           // K is the cut size
  int C )         // C is the number of cuts at each node
{
  for each node n in S, in a topological order {
    compute C structural K-input cuts of n;
    for each cut {
      compute truth table;
      compute irredundant SOP;
      perform delay-optimal balancing of the SOP;
      if ( the cut has smaller AIG level than the best cut )
        save the cut as the best cut;
    }
    if ( root node AIG level is reduced using the best cut )
      update AIG structure;
  }
  return S;
}
  
```

Figure 3.3. Pseudo-code of SOP-balancing.

4. Experimental results

4.1 AIG optimization

The proposed algorithm is implemented in ABC [1][3] as command sequence (*if-g -K <num> -C <num>; st*), where

- *if* is the priority-cut-based FPGA mapper [18],
- *-g* enables SOP-balancing for cut evaluation,
- *-K <num>* specifies the cuts size and,
- *-C <num>* is the number of cuts used at a node,
- *st* transforms the mapped network back into an AIG.

The input of the command sequence is an AIG. The output is a delay-optimized AIG, with a reduced number of logic levels on any path from the PIs to the POs.

The following cost functions are used to prioritize the cuts in the priority-cut-based mapper:

- *Delay* of a cut is the root node level, counting from the PIs of the AIG, after SOP-balancing was applied to the Boolean function of the cut.
- *Area* of a cut is the number of two-input ANDs derived after SOP-balancing was applied to the Boolean function of the cut.

Mapping into standard cells was performed by command *map* [4] in ABC. Experiments targeting standard-cell library *mcnc.genlib* from SIS distribution [24] were run on a workstation with Intel Xeon Quad Core CPU and 48Gb RAM. Only one thread and less than 1Gb of RAM were used for the largest design in our experiments. The resulting networks were verified by a SAT-based combinational equivalence checker (command *cec* in ABC).

The experimental results were collected using a suite of industrial designs optimized in three different ways:

- Reference run: (*st; dch; map*)⁴
- Run 1: (*st; if-g -K 6 -C 8*)(*st; dch; map*)⁴
- Run 2: (*st; if-g -K 6 -C 8*)²(*st; dch; map*)⁶

The reference run is a typical synthesis and mapping flow targeting standard-cells. It consists of four iterations. Each iteration derives an AIG (*st*), followed by AIG-based synthesis with choices (*dch*), followed by cut-based technology mapping (*map*). This or a very similar flow is currently used by most of the industrial users of ABC.

Run 1 performs one iteration of delay optimization followed by the reference flow (4 iterations).

Run 2 performs two iterations of delay optimization followed by the 1/2 reference flows (6 iterations). The increased effort of the reference flow was needed to mitigate area increase.

The results for the three experimental runs are reported in Table 4.1.1. Two outlier designs were removed from the table because the delay improvement exceeded 50%, and this would skew the general conclusions. The table shows that, compared to the reference flow, Run 1 reduces delay by 30% with an area increase of 2.4%, Run 2 reduced delay by 41% with an area increase of 3.9%.

Table 4.1.2 shows the detailed break-down of delay improvements for one design in the test suite. The table lists

area and delay after standard-cell mapping, level count in the AIG before mapping and in the resulting mapped network, as well as the runtime, in seconds, for each step of each of the optimization flows (Reference, Run 1, Run 2). This table indicates that the proposed method is very efficient in reducing the total number of AIG levels as well as the number of levels in the mapped network, which leads to delay reduction after technology mapping. The table also shows that area increase can be further reduced by performing more iterations of logic synthesis with choices.

In another experiment, we applied the proposed delay optimization based on SOP-balancing to MCNC benchmarks. The delay improvements were similar to those in Table 4.1.1 for industrial designs, but the area penalty was higher. We speculate that this is because the ratio of the critical path to the total amount of logic is relatively high in these benchmarks. The detailed results for MCNC benchmarks are not reported here because they are not representative of realistic circuits synthesized these days.

Finally, a similar flow was applied to FPGA mapping, but the delay improvements were not as substantial as for standard cells reported in this paper. We speculate that this is due to LUT mapping being less sensitive to the number of levels and more sensitive to the logic density on the critical path.

4.2 Complete synthesis for standard cells

It is important to note that the delay model used by the ABC technology mapping is approximate. Therefore some part of the improvement will probably be lost, when the mapped netlist is post-processed by a typical industrial physical synthesis flow, which performs gate-sizing, buffering, gate-cloning, and other steps, followed by place-and-route. However, given the high margin of improvement, it is likely that some of the delay reduction will persist even after place-and-route. The slack gains of early delay-oriented optimizations may also bolster downstream design closure efforts that target various parameters such as area reduction and wiring congestion.

To evaluate the effect of our optimization, in a separate experiment, we applied the proposed method to several test cases, followed by the full physical synthesis flow, including place-and-route for standard cells. Our delay reduction optimization "if" was applied with $K=6$ just prior to the technology mapping step within Booledozer [25], where basic gates of a netlist (such as NAND2 and INV) are grouped into more complex gate components. The "if" command was interleaved with the ABC area reducing script "resyn2" to minimize its area increase effects due to duplicated logic. The experiment was conducted on a set of high performance macro blocks that have stringent area optimization goals in addition to tight timing constraints. They span logic from various instruction units, and contain data muxing, parity checking, control, etc.

We analyze experimental results that were collected from the runs at both early and late stages of our industrial synthesis flow. Table 4.2.1 describes the effect of AIG

optimizations on technology mapping performed at the early synthesis stage, immediately after AIG transformation. For each of the macro blocks, the table compares the intermediate technology mapped netlists from two parallel runs: baseline and with AIG optimizations. These are compared in terms of gate count (columns "Gates"), gate's average fanin and fanout ("AvFi" and "AvFo"), average number of pins per net ("AvP"), combinational area ("AreaC"), and total area, which includes sequential logic ("Area"). The "Average" row at the bottom suggests that while the AIG optimization increases the gate count, the netlist improves in other parameters. The AIG-optimized netlists tend to exhibit reduced average of gate's fanin and fanout connections, along with pins per net count. These properties make a netlist less entangled, improving its placement and wiring congestion characteristics. The table shows that the early design area of the mapped netlists is also often reduced.

For each of the macro blocks, two parallel synthesis runs were completed. Table 4.2.2 lists data collected at the end of the physical synthesis flow for the two design variants. To show the impact of the proposed optimization on synthesis quality, the data is tabulated in terms of timing and physical properties of a design. For each of the macro blocks, the table lists figure of merit (FOM) for timing, total wire length (TWL), and the active area size. The timing FOM gives cumulative slack across critical paths in a design, TWL uses half-perimeter estimate when adding net-lengths, and area includes sequential elements of a design. In each of the three categories, the AIG optimizations lead to 20% improvement in timing FOM, to 1% reduction in estimated wire reduction, and to 5% smaller area. The pronounced gains in these categories are due to the early delay oriented optimizations and improvements to the underlying AIG netlist topologies, as reflected in Table 4.2.1.

4.3 Mapping into LUTs

This section reports the impact of SOP balancing on the logic level count after FGPA mapping into 4-input LUTs. Mapping is performed by the priority-cut based mapper *if* in ABC [18]. The same benchmarks as in [18] (Table 6.1) are used in this experiment.

The following three mapping flows are compared:

- Baseline: **(st; if -K 4)**
- Choices: **Baseline + (st; dch; if -K 4)⁵**.
- SOP balancing: **Baseline + (st; dch; if -K 4)²; st; if -g -C 8 -K 8; (st; dch; if -K 4)³**.

Using choices in the second flow involved running five rounds of choice computation (**dch**) followed by mapping with choices (**if**). In the last flow, SOP balancing (**if -g**) was applied after two rounds of mapping with choices, followed by the three rounds of mapping with choices.

The results of mapping into 4-LUTs are shown in Table 4.3. The first section of the table lists benchmark statistics (the number of PIs, the number of POs and the number of flip-flops). The last three sections compare the mapping

results after the three experimental flows. In each section, area is expressed in terms of 4-LUTs and delay is expressed in terms of 4-LUT levels were reported.

To ensure that the delay after mapping into 4-LUTs can be reduced, the cut size for SOP balancing was set to 8, which is larger than the size 6 used in the standard-cell experiments. Using the larger cut size resulted in more aggressive optimization, which reduced delay by 16% and increased area by 9%, on top of the high-effort mapping with choices. However, compared to the baseline mapping, there is no area increase. In fact, in this case, both delay and area are reduced by 32% and 3%, respectively.

5. Conclusions

This paper introduces a simple, fast, and efficient algorithm for delay optimization after technology mapping.

The proposed algorithm preprocesses the subject graph represented as an AIG to reduce the number of levels of two-input ANDs. It is implemented as a straight-forward modification of the publicly-available priority-cut-based technology mapper [18] and its runtime is close to that one run of the mapper. The area increase due to logic duplication is relatively small because of the efficient area recovery done as part of the logic synthesis flow.

Future work may include: (a) improving the quality of the algorithm by pre-computing the smallest delay AIG subgraphs, instead of deriving them using SOP balancing. (b) measuring the improvements in delay after place-and-route for FPGAs, (c) extending the algorithm to work for sequential circuits, as suggested in [23].

Acknowledgements

This work has been supported in part by SRC contract 1875.001 and our industrial sponsors: Altera, Atrenta, Cadence, Calypto, IBM, Intel, Jasper, Magma, Mentor Graphics, Microsemi (Actel), Oasys, Real Intent, Synopsys, Tabula, and Verific. Special thanks to Andrew Kennings and Baruch Sterin for suggesting the idea to perform the SOP balancing with the cut size larger than the LUT size in the experiment of Section 4.3.

6. REFERENCES

- [1] Berkeley Logic Synthesis and Verification Group. *ABC: A System for Sequential Synthesis and Verification*. <http://www-cad.eecs.berkeley.edu/~alanmi/abc>
- [2] C. L. Berman, D. J. Hathaway, A. S. LaPaugh, and L. H. Trevillyan, "Efficient techniques for timing correction", *Proc. ISCAS '90*.
- [3] R. Brayton and A. Mishchenko, "ABC: An academic industrial-strength verification tool", *Proc. CAV'10*, LNCS 6174, pp. 24-40.
- [4] S. Chatterjee, A. Mishchenko, R. Brayton, X. Wang, and T. Kam, "Reducing structural bias in technology mapping", *ICCAD '05*.
- [5] L. Cheng, D. Chen, and D.F. Wong, "DDBDD: Delay-driven BDD synthesis for FPGAs", *Proc. DAC '07*, pp. 910-915. <http://www.icims.csl.uiuc.edu/~dchen/ddbdd.pdf>
- [6] M. Choudhury and K. Mohanram, "Bi-decomposition of large Boolean functions using blocking edge graphs", *Proc. ICCAD '10*.
- [7] J. Cong and Y. Ding, "FlowMap: An optimal technology mapping algorithm for delay optimization in lookup-table based FPGA designs", *IEEE Trans. CAD*, vol. 13(1), Jan. 1994, pp. 1-12.
- [8] R. J. Francis, J. Rose, and K. Chung, "Chortle: A technology mapping program for lookup table-based field programmable gate arrays", *Proc. DAC '90*, pp. 613-619.
- [9] V. N. Kravets and P. Kudva, "Implicit enumeration of structural changes in circuit optimization", *Proc. DAC '04*, pp. 438-441.
- [10] E. Lehman, Y. Watanabe, J. Grodstein, and H. Harkness, "Logic decomposition during technology mapping", *IEEE Trans. CAD*, Vol. 16(8), Aug. 1997, pp. 813-833.
- [11] A. Kuehlmann, V. Paruthi, F. Krohm, and M. K. Ganai, "Robust boolean reasoning for equivalence checking and functional property verification", *IEEE TCAD*, Vol. 21(12), Dec. 2002, pp. 1377-1394.
- [12] P. McGeer, R. K. Brayton, A. L. Sangiovanni-Vincentelli, and S. K. Sahn, "Performance enhancement through the generalized bypass transform", *Proc. ICCAD '91*, pp. 184-187.
- [13] S. Minato, "Fast generation of irredundant sum-of-products forms from binary decision diagrams", *Proc. of SASIMI'92 (Synthesis and Simulation Meeting and International Interchange)*, Kobe, Japan, pp. 64-73.
- [14] E. Morreale, "Recursive Operators for Prime Implicant and Irredundant Normal Form Determination", *IEEE Trans. Comp.*, C-19(6), 1970, pp. 504-509.
- [15] A. Mishchenko, X. Wang, and T. Kam, "A new enhanced constructive decomposition and mapping algorithm", *DAC '03*, pp. 143-148.
- [16] A. Mishchenko, S. Chatterjee, and R. Brayton, "DAG-aware AIG rewriting: A fresh look at combinational logic synthesis", *Proc. DAC '06*, pp. 532-536.
- [17] A. Mishchenko and R. K. Brayton, "Scalable logic synthesis using a simple circuit structure", *Proc. IWLS '06*, pp. 15-22.
- [18] A. Mishchenko, S. Cho, S. Chatterjee, R. Brayton, "Combinational and sequential mapping with priority cuts", *Proc. ICCAD '07*, pp. 354-361.
- [19] A. Mishchenko, R. Brayton, and S. Jang, "Global delay optimization using structural choices", *Proc. FPGA'10*, pp. 181-184.
- [20] P. Pan and C.-C. Lin, "A new retiming-based technology mapping algorithm for LUT-based FPGAs", *Proc. FPGA '98*, pp. 35-42.
- [21] A. Saldanha, H. Harkness, P.C. McGeer, R. K. Brayton, and A. L. Sangiovanni-Vincentelli, "Performance optimization using exact sensitization", *Proc. DAC '94*, pp. 425-429.
- [22] K. J. Singh, A. R. Wang, R. K. Brayton, and A. L. Sangiovanni-Vincentelli, "Timing optimization of combinational logic", *Proc. ICCAD '88*, pp. 282-285.
- [23] C. Soviani, O. Tardieu, and S. A. Edwards, "Optimizing sequential cycles through Shannon decomposition and retiming", *IEEE Trans. CAD*, Vol. 26(3), March 2007, pp. 456-467.
- [24] E. M. Sentovich, K. J. Singh, L. Lavagno, C. Moon, R. Murgai, A. Saldanha, H. Savoj, P. R. Stephan, R. K. Brayton, and A. Sangiovanni-vincentelli. "SIS: A system for sequential circuit synthesis." *Technical Report*, UCB/ERI, M92/41, ERL, Dept. of EECS, UC Berkeley, 1992.
- [25] L. Stok et al, "BooleDozer: Logic synthesis for ASICs", *IBM Jour. of Research and Development*, 1996. Vol. 40(4), pp. 407-430.

Table 4.1.1. Experimental evaluation of the proposed algorithm on industrial circuits after standard-cell mapping.

Design	Reference		Run 1		Run 2	
	Area	Delay	Area	Delay	Area	Delay
D01	180978	34.2	180002	30.0	178099	27.7
D02	16296	15.0	16540	12.8	16082	12.3
D03	50431	41.4	56212	38.6	56212	38.6
D04	16296	15.0	16540	12.8	16082	12.3
D05	509984	74.1	554324	31.4	562109	25.7
D06	443913	37.9	443573	23.9	443181	20.2
D07	80939	21.4	82438	19.9	80347	18.6
D08	257609	31.3	263519	20.8	257917	21.4
D09	597980	81.2	620415	48.0	626055	42.8
D10	612608	32.1	621065	22.3	621838	19.8
D11	73191	46.0	74413	19.8	76346	14.5
D12	429761	48.4	443453	32.9	449604	25.2
D13	236783	26.2	239248	17.5	237456	14.4
D14	848678	54.4	885102	40.4	873752	39.3
D15	13066	54.4	13385	34.0	14561	26.0
D16	220757	80.9	216977	56.0	224621	26.3
D17	316893	19.7	314956	18.7	310999	18.6
Geomean	158148	36.9	161990	25.86	180418	21.8
Ratio	1	1	1.024	0.70	1.039	0.59

Table 4.1.2. Detailed breakdown of delay improvement achieved on one design in the test suite.

Experiments performed	Sequence of optimization steps	Final Mapped Area	Final Mapped Delay	Starting AIG Level	Final Mapped Level	Runtime, sec
Reference flow	st; dch; map	224079	92.90	164	89	222
	st; dch; map	221866	82.10	160	75	143
	st; dch; map	220757	80.90	112	71	136
Run 1	st; if -g -K 6 -C 8	n/a	n/a	164	n/a	66
	st; dch; map	230138	45.00	55	40	208
	st; dch; map	221435	44.60	58	39	149
	st; dch; map	220171	44.60	57	29	143
Run 2	st; if -g -K 6 -C 8	n/a	n/a	164	n/a	66
	st; if -g -K 6 -C 8	n/a	n/a	55	n/a	63
	st; dch; map	240809	25.30	30	24	227
	st; dch; map	232301	25.30	36	25	165
	st; dch; map	230393	25.20	40	24	160
	st; dch; map	229464	24.90	39	24	155
	st; dch; map	228302	25.60	38	24	158
	st; dch; map	227636	25.70	39	24	154

Table 4.2.1. Early optimization data for the mapped netlist.

Netlist	Booleadozer Baseline						Relative change with AIG optimization					
	Gates	AvFi	AvFo	AvP	AreaC	Area	Gates	AvFi	AvFo	AvP	AreaC	Area
fx_macro1	5921	2.01	1.99	2.05	21208	53307	1.03	0.95	0.97	0.97	0.99	1.00
id_macro1	4686	2.39	2.28	2.32	18398	28380	1.07	0.93	0.95	0.94	1.02	1.01
if_macro1	13899	2.36	2.41	2.37	52570	84359	1.11	0.93	0.94	0.95	1.06	1.04
if_macro2	7566	2.40	2.33	2.38	27603	68635	0.99	1.01	1.00	1.00	1.03	1.01
if_macro3	10358	2.38	2.28	2.38	35207	87370	0.95	1.02	1.01	1.01	0.98	0.99
if_macro4	5360	2.40	2.29	2.30	22589	38730	1.06	0.95	0.97	0.96	1.01	1.01
if_macro5	8054	2.34	2.35	2.28	31121	60979	0.97	1.00	1.00	1.00	0.97	0.99
if_macro6	4749	2.23	2.16	2.26	17467	37360	1.02	0.96	0.97	0.97	0.96	0.98
is_macro1	3717	2.65	2.53	2.54	18071	36425	1.00	0.97	0.97	0.98	0.98	0.99
is_macro2	7365	2.26	2.21	2.23	26354	45369	0.98	0.97	0.98	0.98	0.97	0.98
is_macro3	5523	2.31	2.28	2.30	21202	42702	1.00	1.00	1.00	1.00	1.00	1.00
is_macro4	5630	2.25	2.23	2.27	20576	37340	0.88	1.07	1.05	1.05	0.97	0.98
is_macro5	5742	2.29	2.23	2.30	22192	44407	1.03	0.92	0.95	0.95	0.97	0.98
is_macro6	6274	2.09	2.08	2.18	19995	62932	1.07	0.93	0.97	0.96	1.02	1.01
ls_macro1	6087	2.22	2.04	2.13	22530	45449	0.97	0.95	0.97	0.97	0.94	0.97
ls_macro2	3726	2.18	2.20	2.22	14571	34870	1.07	0.92	0.96	0.95	1.02	1.01
ls_macro3	5512	2.16	2.02	2.13	20440	48607	1.04	0.95	0.98	0.97	1.01	1.00
ls_macro4	7278	2.06	2.02	2.09	21570	64264	0.97	0.96	0.97	0.98	0.95	0.98
Average							1.012	0.966	0.978	0.978	0.992	0.997

Table 4.2.2. Late physical synthesis data.

Netlist	Placement-driven synthesis baseline			With AIG optimization		
	FOM	TWL	Area	FOM	TWL	Area
fx_macro1	-1443	1679216	77316	-1365	1753354	77137
id_macro1	-5323	1690176	73373	-6103	1663948	69352
if_macro1	-5874	3533386	157582	-4367	3865699	151298
if_macro2	-1720	2503959	108346	-1411	2385904	102720
if_macro3	-8437	3539756	142900	-1993	3285048	128004
if_macro4	-1498	1244580	67572	-1900	1271511	66468
if_macro5	-2184	2510401	112186	-2490	2438322	106536
if_macro6	-8283	1262190	66438	-6097	1280597	64705
is_macro1	-945	1159571	57332	-758	1125653	55099
is_macro2	-3011	2067163	92230	-2434	2134954	85558
is_macro3	-1772	1798489	74781	-1545	1788065	71898
is_macro4	-7031	1731164	71547	-5202	1637196	61609
is_macro5	-1270	1644306	68413	-1822	1597819	63596
is_macro6	-701	2409278	92544	-344	2344632	88279
ls_macro1	-1301	1792453	75101	-1712	1835897	73443
ls_macro2	-3179	1135350	54898	-3684	1115315	54323
ls_macro3	-2755	1949240	76567	-2986	1939440	75609
ls_macro4	-2365	2761289	112878	-1085	2588526	103444
Cumulative				0.80	0.99	0.95

Table 4.3. Delay improvements after mapping into 4-LUTs.

Design	Statistics			Baseline		Choices		SOP balancing	
	PI	PO	FF	LUT	Level	LUT	Level	LUT	Level
alu4	14	8	0	704	10	667	8	689	8
apex2	39	3	0	962	10	728	9	809	8
b14	32	54	245	2243	23	1672	20	1981	14
b15	36	70	449	3294	24	3057	23	3480	19
b17	37	97	1415	10401	35	9114	33	9613	21
b20	32	22	490	4495	24	3453	23	4274	16
b21	32	22	490	4751	24	3473	23	4341	16
b22	32	22	735	6740	25	5157	23	6600	16
clma	383	82	33	3890	18	3792	13	3979	12
des	256	245	0	1244	9	1194	7	1219	7
elliptic	19	2	194	376	12	424	9	475	8
ex5p	8	63	0	507	13	444	7	499	6
frisc	4	116	886	2242	22	2214	20	2575	14
i10	257	224	0	686	20	678	15	711	12
pdc	16	40	0	2238	14	2062	9	2118	9
s38584	13	278	1452	3825	12	3518	9	3510	9
s5378	36	49	161	381	8	361	7	368	6
seq	41	35	0	980	9	863	7	871	7
spla	16	46	0	2167	14	1808	9	1745	9
tseng	52	122	385	788	14	744	14	792	11
Geomean				1731	15.61	1537	12.67	1678	10.62
Ratio1				1	1	0.888	0.812	0.970	0.681
Ratio2						1	1	1.092	0.838