

**Computer Science Department Technical Report
University of California
Los Angeles, CA 90024-1596**

**DAG-MAP: GRAPH BASED FPGA TECHNOLOGY MAPPING
FOR DELAY OPTIMIZATION**

**K. C. Chen
J. Cong
Y. Ding
A. Kahng
P. Trajmar**

**March 1992
CSD-920003**

DAG-Map: Graph Based FPGA Technology Mapping For Delay Optimization

Jason Cong, Yuzheng Ding, Andrew Kahng, Peter Trajmar¹

*Department of Computer Science
University of California, Los Angeles, CA 90024*

Kuang-Chien Chen

*Fujitsu America, Inc.
3055 Orchard Drive, San Jose, CA 95134*

Abstract

In this paper, we present a graph based technology mapping algorithm, called DAG-Map, for delay optimization in lookup-table based FPGA designs. Our algorithm carries out technology mapping and delay optimization on the entire Boolean network, instead of decomposing the network into fanout-free trees and mapping each tree separately as in most previous FPGA technology mapping algorithms. Moreover, as a preprocessing step of DAG-Map, we introduce a general algorithm for transforming an arbitrary n -input network into a two-input network with only $O(1)$ factor increase in the network depth; previous transformation procedures may result in an $O(\log n)$ factor increase in the network depth. Finally, we present a graph matching based technique which performs area optimization without increasing the network delay; this is used as a postprocessing step for DAG-Map. We implemented the DAG-Map algorithm and tested it on the MCNC logic synthesis benchmarks. Compared with previous FPGA technology mapping algorithms for delay optimization (Chortle-d and MIS-pga), DAG-Map reduces both the network depth and the number of lookup-tables.

¹ Current address: Zycad Inc., 1380 Willow Road, Menlo Park, CA 94025.

1. Introduction

The Field Programmable Gate Array (FPGA) is a relatively new technology which allows the circuit designer to produce ASIC chips without going through the fabrication process. The fast turn-around time and low manufacturing cost have led to increasing interest in FPGA technology for system prototyping and low- or medium-volume production. An FPGA chip usually consists of three components: programmable logic blocks, programmable interconnections, and programmable I/O blocks. Current technology implements programmable logic blocks using either K-input RAM/ROM lookup-tables (K-LUTs) [Xi89] or programmable multiplexers [Ca86, Hs87, El89]. Programmable interconnections consist of one-dimensional segmented channels [GrKE90] or two-dimensional routing grids with switch-matrices [Xi89, BrRV90]. Programmable I/O blocks provide a user configurable interface between internal logic blocks and I/O pads. The design process for FPGAs is similar to that for conventional gate arrays or standard cells. Given a high-level design specification, the design process includes logic synthesis, technology mapping, placement, and routing. However, the end result is not a set of masks for fabrication, but rather a configuration matrix which sets the values of all the programmable elements in a FPGA chip. In this paper, we study the technology mapping problem for delay optimization of lookup-table based FPGAs.

The technology mapping problem is to implement a synthesized Boolean network using logic cells from a given cell family. Much work has been done on the technology mapping problem for conventional gate array or standard cell designs [Gr86, Ka86, Ke87, De87, LiBK88]. In particular, it was shown that a Boolean network can be decomposed into a set of fanout-free trees and that the technology mapping problem can be solved optimally for each tree independently using a dynamic programming approach [Ke87]. However, these methods do not apply immediately to the technology mapping problem for FPGAs since a K-LUT can implement any one of 2^{2^K} K-input logic gates, and consequently the equivalent cell family is too large to be manipulated efficiently.

Recently, a number of technology mapping algorithms have been proposed for area optimization in lookup-table based FPGA designs. The MIS-pga program developed by Murgai et al. [Mu90] first decomposes a given Boolean network into a feasible network using Roth-Karp decomposition and kernel extraction so that the number of inputs at each node is bounded. MIS-pga then enumerates all the possible realizations of each network node and solves the binate covering problem to get a mapping solution using the least number of lookup-tables. In the improved MIS-pga (new) [MuSB91b] more decomposition techniques are incorporated, including bin-packing, cofactoring, and AND-OR decomposition. The covering problem is solved more efficiently via certain preprocessing operations. The Chortle program and its successor Chortle-crf, developed by Francis, Rose and Vranesic [FrRC90, FrRV91a], decompose a given Boolean network into a set of fanout-free trees and then carry out technology mapping on each tree using the dynamic programming approach. Bin-packing heuristics are used in Chortle-crf for gate-level decomposition, yielding significant improvement over its predecessor in the quality of solutions and the running time. The Xmap program developed by Karplus [Ka91a]

transforms a given Boolean network into an if-then-else DAG representation and then goes through a simple marking process to determine the final mapping. Another FPGA technology mapping algorithm was proposed by Woo [Wo91], who introduces the notion of invisible edges to denote the edges which do not appear in the resulting network after mapping. A given network is first partitioned into subgraphs of reasonable size, and then an exhaustive procedure is used to determine the invisible edges in each subgraph. In the meantime, several technology mapping algorithms have also been proposed for multiplexor based FPGAs [Mu90, ErDe91, Ka91b]; some of them use techniques similar to those used for lookup-table based FPGAs. The main objective in all these FPGA technology mapping algorithms is to minimize the number of programmable logic blocks in the mapping solution.

The emphasis of this paper is on delay optimization in FPGA designs, since performance is the main consideration in many applications that use FPGA technology. We set area optimization as the secondary objective, i.e., we minimize the number of lookup-tables after we have obtained a mapping solution with minimum delay.

Previous work on FPGA mapping for delay optimization consists of Chortle-d, developed by Francis, Rose and Vranesic [FrRV91b, FrRV91c], and an extension of MIS-pga, developed by Murgai et al. [MuSB91a]. The basic approach used in Chortle-d is similar to that in Chortle-crf, i.e., decompose the network into fanout-free trees and then use dynamic programming and bin-packing heuristics to map each tree independently. However, the objective at each step in Chortle-d is to minimize the depth of the node being processed. Their method indeed reduced the depths of mapping solutions considerably (29% less compared to MIS-pga and 35% less compared to Chortle-crf). However, the method has two drawbacks. First, it decomposes the network into a set of fanout-free trees. Although it guarantees the optimal depth for each tree (when the input limit of each lookup-table is no more than six), this *prior* decomposition usually results in sub-optimal depth for the overall network. Second, Chortle-d uses many more lookup-tables than are used by area optimization algorithms (MIS-pga and Chortle-crf).

The MIS-pga extension of [MuSB91a] contains two phases, mapping and placement/routing. The mapping phase first computes a delay-optimized 2-input network, then traverses the network from the primary inputs, collapsing the nodes in the longest paths into their fanouts to reduce the network depth. During this procedure various decomposition techniques are used to dynamically resynthesize the network, so this method uses a reduced number of look-up tables. The advantage of this approach is that it takes layout information into consideration at the technology mapping stage. However, on average it yields larger network depth than Chortle-d, especially for large networks, and requires much more computation time.

In this paper, we present a graph based technology mapping algorithm, called DAG-Map, for delay optimization in lookup-table based FPGA design. Our algorithm carries out technology mapping and delay optimization on the entire Boolean network, instead of decomposing it into fanout-free trees as in Chortle-d. Our algorithm is optimal for trees for any K-LUTs while Chortle-d is optimal for trees only when K is no more than six [FrRV91c]. Moreover, in a preprocessing phase of DAG-Map, we introduce a general algorithm for transforming an arbitrary

n -node network into a two-input network with only $O(1)$ factor increase in the network depth, while the previous transformation procedure may result in $O(\log n)$ factor increase in the network depth. Finally, we present a matching based technique for area optimization without increasing the network delay, which are used as a postprocessing step for DAG-Map. We tested DAG-Map on the set of MCNC logic synthesis benchmarks and compared it with the previous FPGA mapping algorithms. Our experimental results show that on average, DAG-Map reduces both the network delay and the number of lookup-tables when compared with either Chortle-d or the mapping phase of MIS-pga delay optimization.

The remainder of this paper is organized as follows. Section 2 gives the precise problem formulation. Section 3 describes our DAG-Map algorithm in detail. Experimental and comparative results are presented in Section 4.

2. Problem Formulation

A Boolean network can be represented as a directed acyclic graph (DAG) where each node represents a logic gate and there is a directed edge (i, j) if the output of gate i is an input of gate j . A primary input (PI) node has no incoming edge and a primary output (PO) node has no outgoing edge. We use $input(v)$ to denote the set of nodes which supply inputs to gate v . Given a subgraph H of the Boolean network, $input(H)$ denotes the set of *distinct* nodes which supply inputs to the gates in H . For a node v in the network, a K -feasible cone at v , denoted C_v , is a subgraph consisting of v and predecessors² of v such that any path connecting a node in C_v and v lies entirely in C_v and $|input(C_v)| \leq K$. The *level* of a node v is the length of the longest path from any PI node to v . The level of a PI node is zero. The *depth* of a network is the largest node level in the network.

We assume that each programmable logic block in an FPGA is a K -input lookup-table (K-LUT) that can implement any K -input Boolean function (this is true for the FPGA chips produced by Xilinx and AT&T [Xi91, Hi91, Wo91]). Thus, each K-LUT can implement any K -feasible cone of a Boolean network. The technology mapping problem for K-LUT based FPGAs is then to cover a given Boolean network with K -feasible cones³. A technology mapping solution S is a DAG where each node is a K -feasible cone (equivalently, a K-LUT) and the edge (C_u, C_v) exists if u is in $input(C_v)$. Our goal is to compute a mapping solution that results in small circuit delay and, secondarily, uses small chip area. The delay of a FPGA circuit is determined by two parts: delay in K-LUTs and delay in the interconnection paths. Since layout information is not available at this stage, we assume that each edge in the mapping solution S contributes a constant delay. Hence, the circuit delay is determined by the depth of S since each K-LUT contributes a constant delay (the access time) independent of the function it implements. Therefore, the main objective of our algorithm is to determine a mapping solution S with depth as small as possible.

² u is a predecessor of v if there is a directed path from u to v .

³ Note that we do not require the covering to be disjoint since we allow nodes in the network to be replicated, if necessary, as long as the resulting network is logically equivalent to the original one. In fact, our algorithm is capable of replicating nodes automatically, when necessary, in order to achieve delay optimization. See Section 3.2 for more details.

Since routing resource requirements are not known during technology mapping, our secondary objective is to minimize the number of K-LUTs in the technology mapping solution.

3. The DAG-Map Algorithm

Our DAG-Map algorithm consists of three major steps. The first step transforms an arbitrary Boolean network into a two-input network. The second step maps the two-input network into a K-LUT FPGA network with minimum delay. The third step performs area optimization of the FPGA network without increasing the network delay. This section describes these three steps in detail.

3.1. Transforming Arbitrary Networks into Two-Input Networks

As in [FrRV91a, FrRV91b], we assume that each node in the given Boolean network is a simple gate (i.e. AND, OR, NAND, or NOR gate).⁴ The first step of our algorithm is to transform the given Boolean network of simple gates into a two-input network (i.e. each gate in the network has at most two inputs). There are two reasons for carrying out such a transformation. First, we want to limit the number of inputs of each gate to be no more than K so that we do not have to decompose gates during technology mapping. Second, if we think of FPGA technology mapping as a process of packing gates in a given network into K-LUTs, then, intuitively, smaller gates will be more easily packed, with less wasted space in each K-LUT.

A straightforward way to transform an n -node arbitrary network into a two-input network is to replace each m -input gate ($m \geq 3$) by a balanced binary tree⁵. Fig. 1(a) shows a 4-input gate v (where the numbers beside the nodes indicate their levels) and Fig. 1(b) shows the result of replacing it by a balanced binary tree. We see that the level of v increases from 7 to 8. In general, such a straightforward transformation may increase the network depth by as much as an




Fig. 1. Transforming a multi-input network into a two-input network.

⁴ If the network has complex gates, we can represent each complex gate in the sum-of-products form and then replace it with by two levels of simple gates. In particular, we use the technology decomposition command `tech_decomp -o 1000 -a 1000` in MIS [BrRS87], which realizes such a transformation.

⁵ The gate type of each node in the binary tree is the same as the gate type of the original multiple-input node. Such a transformation maintains logical equivalence as long as the gate function is associative.

$O(\log n)$ factor as we will show later. However, if we replace v by the binary tree shown in Fig. 1(c), the level of v remains 7. Our goal is to replace each multiple-input node by a binary tree so that the height of the resulting network is as small as possible.

Given an arbitrary Boolean network G , we transform G into a two-input network G' as follows. We process the nodes in G in topological order starting from the PI nodes. At each multiple-input node v , we construct a binary tree $T(v)$ rooted at v using an algorithm similar to Huffman's algorithm for constructing a prefix code of minimum average length [Hu52]. Writing $input(v) = \{u_1, u_2, \dots, u_m\}$, note that nodes u_1, u_2, \dots, u_m have already been processed by the time we process v ; their levels $level'(u_i)$ ($1 \leq i \leq m$) in the new network G' have been determined. Intuitively, we want to combine nodes with smaller levels first when we construct the binary tree $T(v)$. Our algorithm is as follows.

Algorithm: decompose-multi-input-gate (DMIG)
 let $V = input(v) = \{u_1, u_2, \dots, u_m\}$;
 while $|V| > 2$ do
 let u_i and u_j be the two nodes of V with smallest levels;
 introduce a new node x ;
 $input(x) = \{u_i, u_j\}$;
 $level'(x) = \max(level'(u_i), level'(u_j)) + 1$;
 $V = (V - \{u_i, u_j\}) \cup \{x\}$;
 end-while
 Connect the only two nodes left in V to v as its inputs;
 Return the binary tree $T(v)$ rooted at v ;
end-algorithm.

If we apply the DMIG algorithm to the example in Fig. 1(a), we indeed obtain the binary tree shown in Fig. 1(c). An algorithm similar to DMIG was proposed by Wang [Wa89] for timing-driven decomposition in the synthesis of multi-level Boolean network. We have proven the following theoretical results showing that the DMIG algorithm increases network depth by at most a constant factor. First, we have the following lemma.

Lemma Let $V = \{u_1, u_2, \dots, u_m\}$ be the set of inputs of a multi-input node v in the initial network G . Then, after applying the DMIG algorithm we have

$$2^{level'(v)} \leq \sum_{i=1}^m 2^{level'(u_i)+1}$$

where $level'(x)$ is the level of node x in the two-input network G' .

Proof It is easy to see that the DMIG algorithm will introduce $m - 2$ new nodes in processing v . (For any binary tree, the number of leaves equals the number of internal nodes (including the root) plus one.) Let $X_i = \{x_{i,1}, x_{i,2}, \dots, x_{i,m-i}\}$ be the set of nodes left in V after the DMIG algorithm introduces the i -th new node. Clearly, $X_0 = \{u_1, u_2, \dots, u_m\}$. For convenience, we define $X_{m-1} = \{v\}$. Without loss of generality, assume that $level'(x_{i,1}) \leq level'(x_{i,2}) \leq \dots \leq level'(x_{i,m-i})$. Then, the $(i+1)$ -th new node has $x_{i,1}$ and $x_{i,2}$ as its inputs, and its level is given by $level'(x_{i,2}) + 1$. Therefore, we have

$$\begin{aligned} \sum_{x \in X_{i+1}} 2^{\text{level}'(x)} &= 2^{\text{level}'(x_{i,2})+1} - 2^{\text{level}'(x_{i,1})} - 2^{\text{level}'(x_{i,2})} + \sum_{x \in X_i} 2^{\text{level}'(x)} \\ &= 2^{\text{level}'(x_{i,2})} - 2^{\text{level}'(x_{i,1})} + \sum_{x \in X_i} 2^{\text{level}'(x)} \end{aligned}$$

Taking the sum of both sides of the last equation from $i = 0$ to $m - 2$, we have

$$\sum_{i=0}^{m-2} \sum_{x \in X_{i+1}} 2^{\text{level}'(x)} = \sum_{i=0}^{m-2} (2^{\text{level}'(x_{i,2})} - 2^{\text{level}'(x_{i,1})}) + \sum_{i=0}^{m-2} \sum_{x \in X_i} 2^{\text{level}'(x)}$$

Subtracting $\sum_{i=1}^{m-2} \sum_{x \in X_i} 2^{\text{level}'(x)}$ from both sides, we get

$$2^{\text{level}'(v)} = \sum_{i=0}^{m-2} (2^{\text{level}'(x_{i,2})} - 2^{\text{level}'(x_{i,1})}) + \sum_{x \in X_0} 2^{\text{level}'(x)}$$

Note that

$$\sum_{i=0}^{m-2} (2^{\text{level}'(x_{i,2})} - 2^{\text{level}'(x_{i,1})}) = \sum_{i=0}^{m-2} (2^{\text{level}'(x_{i,2})} - 2^{\text{level}'(x_{i+1,1})}) + 2^{\text{level}'(x_{m-2,2})} - 2^{\text{level}'(x_{0,1})}$$

Moreover, $2^{\text{level}'(x_{i,2})} - 2^{\text{level}'(x_{i+1,1})} \leq 0$ for any $0 \leq i \leq m-2$ (since $x_{i+1,1}$ is either the $(i+1)$ -th new node or a node in $X_i - \{x_{i,1}, x_{i,2}\}$; in either case we have $\text{level}'(x_{i+1,1}) \geq \text{level}'(x_{i,2})$) and $2^{\text{level}'(x_{m-2,2})} = \frac{1}{2} \cdot 2^{\text{level}'(v)}$. It follows that

$$2^{\text{level}'(v)} \leq \frac{1}{2} \cdot 2^{\text{level}'(v)} - 2^{\text{level}'(x_{0,1})} + \sum_{x \in X_0} 2^{\text{level}'(x)}$$

Therefore,

$$2^{\text{level}'(v)} \leq 2 \cdot \left(\sum_{x \in X_0} 2^{\text{level}'(x)} - 2^{\text{level}'(x_{0,1})} \right) < \sum_{x \in X_0} 2^{\text{level}'(x)+1} = \sum_{i=1}^m 2^{\text{level}'(u_i)+1} \quad \square$$

From this lemma, we can show that the DMIG algorithm results in at most a constant factor increase in network depth.

Theorem 1 For an arbitrary Boolean network G of simple gates⁶, let G' be the network obtained by applying the DMIG algorithm to each multi-input gate in topological order starting from the PI nodes. Then

$$\text{depth}(G') \leq \log 2d \cdot \text{depth}(G) + \log l$$

where d is the maximum degree of fanout in G and l is the number of PI nodes in G .

Proof Let H denote $\text{depth}(G)$. Let L_i denote the set of nodes $\{x \mid x \in G, \text{level}(x) = i\}$. (Note that $\text{level}(x)$ is the level of node x in the initial network G .) Let A_i denote the set of nodes

⁶ Any complex gate in the network can be decomposed into a two-level AND-OR subnetwork so that the increase of network depth is bounded by a factor of two.

x in G such that $level(x) \leq i$ and x has at least a fanout y with $level(y) > i$. We shall prove by induction that

$$\sum_{v \in A_i} 2^{level'(v)} \leq (2d)^i \cdot I. \quad (*)$$

Since A_0 is the set of PI nodes in G , the inequality (*) holds for $i = 0$. Suppose that the inequality holds for $i - 1$, we want to show that it also holds for i . According to the definition of A_i , it is not difficult to see that

$$A_i \subseteq (A_i \cap A_{i-1}) \cup L_i$$

Moreover, each node v in $A_i \cap A_{i-1}$ has at most $d - 1$ fanouts in L_i . According to the lemma, we have

$$\begin{aligned} \sum_{v \in A_i} 2^{level'(v)} &\leq \sum_{v \in A_{i-1} \cap A_i} 2^{level'(v)} + \sum_{v \in L_i} 2^{level'(v)} \\ &\leq \sum_{v \in A_{i-1} \cap A_i} 2^{level'(v)} + (d - 1) \cdot \sum_{v \in A_{i-1} \cap A_i} 2^{level'(v)+1} + d \cdot \sum_{v \in A_{i-1} - A_i} 2^{level'(v)+1} \\ &\leq 2d \cdot \sum_{v \in A_{i-1}} 2^{level'(v)} \end{aligned}$$

By induction hypothesis, we have

$$\sum_{v \in A_i} 2^{level'(v)} \leq 2d \cdot [(2d)^{i-1} \cdot I] = (2d)^i \cdot I$$

It concludes that the inequality (*) holds for any $0 \leq i \leq H$. Let w be node in G that achieves the maximum level in G' , then all the inputs of w are in A_{H-1} . According to the lemma and the inequality (*), we have

$$2^{level'(w)} < \sum_{v \in input(w)} 2^{level'(v)+1} \leq \sum_{v \in A_{H-1}} 2^{level'(v)+1} \leq 2 \cdot (2d)^{H-1} \cdot I \leq (2d)^H \cdot d.$$

Therefore,

$$depth(G') = level'(w) \leq \log[(2d)^H \cdot I] \leq \log 2d \cdot H + \log I. \quad \square$$

Similar analysis was carried out by Hoover, Klawe, and Pippenger when bounding the maximum degree of fanout in a Boolean network [HoKP84]. Since in practice d is bounded by a constant (fanout limit of any output), the depth of the two-input network G' is increased by just a small constant factor $\log 2d$ away from $depth(G)$.⁷ For example, if $d \leq 4$, the depth increase is bounded by a factor of 3. Balanced binary tree based transformation may increase the depth of an n -node network by as much as an $O(\log n)$ factor, even when d is bounded by a constant, therefore resulting in depth much larger than that obtained by our algorithm, especially when

⁷ Here we assume that $depth(G) = \Omega(\log I)$, which is true for most networks in practice. This excludes the unrealistic case where $\log I$ is the dominating term in the right-hand side of the inequality in Theorem 1.

$depth(G)$ is large. Fig. 2 shows a pathological example for the balanced binary tree based transformation. The initial network of size n is shown in Fig. 2(a), which is a fanout free circuit of depth $\sqrt{n-1}$ (assuming that the primary inputs are of level 0). The two-input network after the balanced binary tree based transformation is shown in Fig. 2(b), which has depth $d_{BBT} = \frac{1}{2} \log(n-1) \cdot \sqrt{n-1}$, even with $d=1$ in this case. Fig. 2(c) shows the DMIG transformation result, which has depth $d_{DMIG} = \frac{1}{2} \log(n-1) + \sqrt{n-1} - 1$. Clearly d_{DMIG} is much smaller than d_{BBT} when n is large.

The experimental results in Section 4 show that the 2-input networks obtained using our transformation procedure lead to smaller network depths and better mapping solutions than those obtained using the transformation procedure in MIS [BrRS87].

3.2. Technology Mapping for Delay Minimization

After we obtain a two-input Boolean network, we carry out technology mapping directly on the entire network. We use a method similar to that of Lawler, Levitt, and Turner for module clustering to minimize delay in digital networks [LaLT69]. Our algorithm consists of two phases. We first label the network to determine the level of each node in the final mapping solution. We

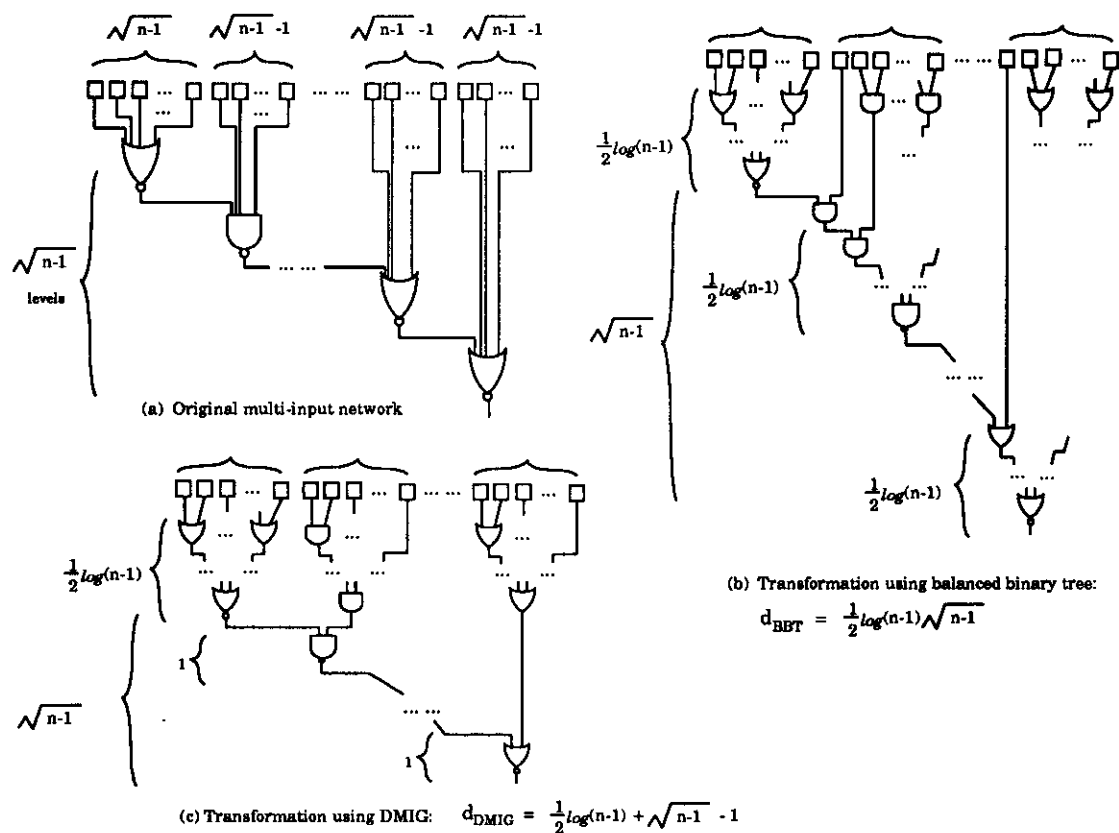


Fig. 2. A pathological example for the balanced binary tree based transformation (assume $n = 2^{2m} + 1$ for some m).

then generate the logically equivalent network of K-LUTs.

The first step assigns a label $h(v)$ to each node v of the two-input network, with $h(v)$ equal to the level of the K-LUT containing v in the final mapping solution. Clearly, we want $h(v)$ to be as small as possible in order to achieve delay minimization. We label the nodes in a topological order starting from the PI nodes. The label of each PI node is zero. If node v is not a PI node, let p be the maximum label of the nodes in $input(v)$ (note that because of the topological ordering all nodes in $input(v)$ have already been labeled). We use $N_p(v)$ to denote the set of predecessors of v with label p . Then, if $input(N_p(v) \cup \{v\}) \leq K$, we assign $h(v) = p$; otherwise, we assign $h(v) = p + 1$. With this labeling, it is evident that $N_{h(v)}(v)$ forms a K-feasible cone at v for each node v in the network⁸.

The second phase of our algorithm is to generate K-LUTs in the mapping solution. Let L represent the set of outputs which are to be implemented using K-LUTs. Initially, L contains all the PO nodes. We process the nodes in L one by one. For each node v in L , we remove v from L and generate a K-LUT v' to implement the function of gate v such that $input(v') = input(N_{h(v)}(v))$. (Recall that $N_{h(v)}(v)$ forms a K-feasible cone at v .) Then, we update the set L to be $L \cup input(v')$. The second phase stops when L consists of only PI nodes in the original network. It is clear that at the end of execution we get a network of K-LUTs which is logically equivalent to the original network.

The algorithm can be summarized as follows.

```

algorithm: DAG-Map
  /* phase 1: labeling the network */
  for each PI node  $v$  do
     $h(v) = 0$ ;
   $T =$  list of non-PI nodes in topological order;
  while  $T$  is not empty do
    remove the first node  $v$  from  $T$ ;
    let  $p = \max\{h(u) \mid u \in input(v)\}$ ;
    if  $|input(N_p(v) \cup \{v\})| \leq K$  then  $h(v) = p$ 
    else  $h(v) = p + 1$ ;
  end-while
  /* phase 2 : generate K-LUTs */
   $L =$  list of PO nodes;
  while  $L$  contains non-PI nodes do
    remove a non-PI node  $v$  from  $L$ , i.e.  $L = L - \{v\}$ ;
    introduce a K-LUT  $v'$  to implement the function of  $v$  such that
       $input(v') = input(N_{h(v)}(v))$ ;
     $L = L \cup input(v')$ ;
  end-while
end-algorithm

```

⁸ Note that $v \in N_{h(v)}(v)$ since v is a predecessor of itself.

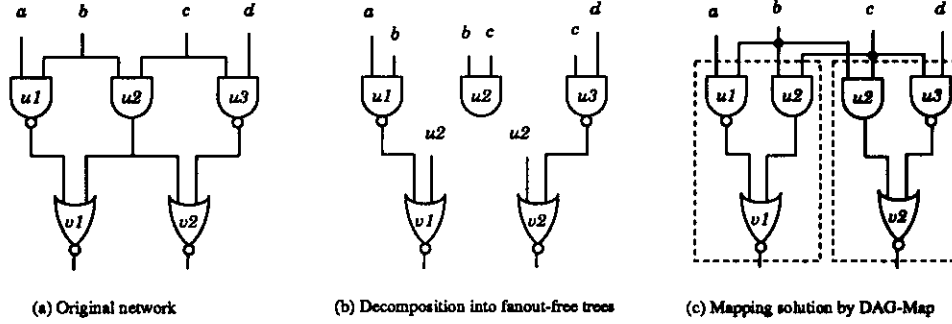


Fig. 3. A mapping example for the case $K = 3$.

The DAG-Map algorithm has several advantages:

- (1) It works on the entire network without decomposing it into fanout-free trees; this usually leads to better mapping solutions. For example, decomposing the two-input network shown in Fig. 3(a) into fanout-free trees (as shown in Fig. 3(b)) yields a two-level mapping solution with three look-up-tables. However, the DAG-Map algorithm gives an one-level mapping solution with two look-up-tables (as shown in Fig. 3(c)).
- (2) The DAG-Map algorithm can replicate nodes, if necessary, to minimize the network delay in the mapping solution. For the solution shown in Fig. 3(c), node u_2 is replicated to get an one-level mapping solution. Note that if node u_2 is not replicated, the depth of the mapping solution is at least two.
- (3) The DAG-Map algorithm is optimal when the initial network is a tree.

Theorem 2 For any integer K , if the given Boolean network is a tree with fanin no more than K at each node, the DAG-Map algorithm produces a minimum depth mapping solution for K -LUT based FPGAs.

Proof It is easy to see that given a tree T , if the fanin limit of each node is K , the DAG-Map algorithm can successfully label all the nodes T . Moreover, for any node v in T , the label $h(v)$ is the level of LUT_v in the mapping solution produced by DAG-Map, where LUT_v is the K -LUT containing v . We shall show that for any mapping solution M , the level of any node v satisfies $level_M(LUT_v) \geq h(v)$, where $level_M(LUT_v)$ is the level of the K -LUT LUT_v in M .

Assume toward a contradiction that M is a mapping solution such that $level_M(LUT_v) < h(v)$ for some node v . Furthermore, let v be the node with the lowest level in T such that $level_M(LUT_v) < h(v)$. Then, for any predecessor w of v , we have $level_M(LUT_w) \geq h(w)$. Let u be the predecessor of v with the maximum label $h(u) = p$. Since $level_M(LUT_v) \geq level_M(LUT_u) \geq h(u) = p$, and $h(v) \leq p + 1$ according to the labeling procedure of DAG-Map, we conclude that $level_M(LUT_v) = p$ and $h(v) = p + 1$. Note that $level_M(LUT_v) = p$ implies that $C_v \supseteq N_p(v) \cup \{v\}$, and $h(v) = p + 1$ implies that $|input(N_p(v) \cup \{v\})| > K$

according to the labeling procedure in the algorithm, where C_v is the K -feasible cone at v which is contained in LUT_v in M . However, since T is a tree, we have

$$|input(C_v)| \geq |input(N_p(v) \cup \{v\})| > K,$$

which contradicts the fact that C_v is K -feasible. \square

Recall that a similar result was shown for Chortle-d in [FrRV91b], but the Chortle-d result holds only for $K \leq 6$ since the bin-packing heuristics are no longer optimal for $K > 6$.⁹

Although the DAG-Map algorithm is optimal for trees, it may not be optimal for general networks. Fig. 4 shows an example where DAG-Map produces a three-level mapping solution while there exists an optimal two-level mapping solution. It is interesting to point out, however, that DAG-Map would be optimal if the mapping constraint for each programmable logic block is monotone. As defined in [LaLT69], a constraint X is *monotone* if a network H satisfying X implies that any subgraph of H also satisfies X . For example, if we assume that the constraint for each programmable logic block is the number of gates it may cover in the original network, it becomes a monotone constraint and the DAG-Map algorithm would produce an optimal mapping solution. Unfortunately, limiting the number of distinct inputs of each programmable logic block is not a monotone constraint. For example, in Fig. 5 the whole network has three distinct inputs, but the subnetwork consisting of t , v and w has four distinct inputs. Although the optimality of DAG-Map is not guaranteed for general networks, the experimental results in Section 4 show that

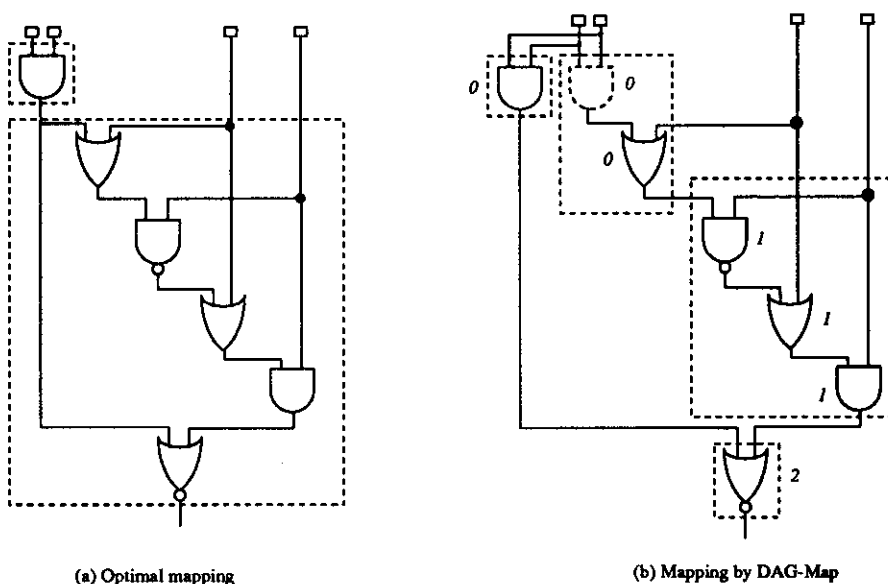


Fig. 4. A pathological example for the DAG-Map algorithm.
 (Assume $K = 3$. Numbers are labels corresponding to node levels in mapping solution.)

⁹ Note that Chortle-d does not require the fanin limit of each node in the tree to be no more than K , since it carries out node decomposition during the bin-packing process.




Fig. 5. Constraint on number of inputs of LUT is not monotone (Assume $K=3$).

it produces very satisfactory mapping solutions with respect to delay optimization for all benchmark circuits.

3.3. Area Optimization Without Increasing Delay

Since the main objective of the DAG-Map algorithm is optimization of the depth of the mapping solution, minimizing the number of K -LUTs is not a consideration. For this reason, we have developed two operations for area optimization which are used in post-processing steps after we obtain a mapping solution of small depth. These operations reduce the number of K -LUTs in the mapping solution without increasing the network depth. Note that in this sub-section, each node in the network is a K -LUT instead of a simple gate as in the preceding sub-sections.

The first operation is called *gate decomposition*, which is inspired by the gate decomposition concept used in Chortle-crf [FrRV91a]. The basic idea is as follows. If node v is a simple gate of multiple inputs in the mapping solution, for any two of its inputs u_i and u_j , if u_i and u_j are single fanout nodes, we can decompose v into two nodes v_{ij} and v' such that v' is of the same type as v and v_{ij} is of the same type as v in *non-negated form*, and $input(v_{ij}) = \{u_i, u_j\}$ and $input(v') = input(v) \cup \{v_{ij}\} - \{u_i, u_j\}$ (intuitively, u_i and u_j are fed into v_{ij} first and then v_{ij}




Fig. 6: Gate decomposition for area optimization (assume $K = 5$).

replaces u_i and u_j as an input to v'). Such a decomposition produces a logically equivalent network because of the associativity of the simple functions. In this case, if $|\text{input}(u_i) \cup \text{input}(u_j)| \leq K$, then we can implement u_i , u_j and v_{ij} using one K-LUT. The result is that the number of K-LUTs is reduced by one and the decomposed node v has one fewer inputs (which is beneficial to subsequent gate decomposition and predecessor packing). Figure 6 illustrates the gate decomposition operation, where the number of nodes, as well as the number of fanins of node v (v' after the operation), is reduced by one.

This method can be generalized to the case where the decomposed node v implements a complex function. In this case we apply the Roth-Karp decomposition [RoKa62] to determine if the node can be feasibly decomposed to v_{ij} and v' as in the preceding paragraph. Given a Boolean function $F(X, Y)$, where X and Y are Boolean vectors, the Roth-Karp decomposition determines if there is a pair of Boolean functions G and H such that $F(X, Y) = G(H(X), Y)$, and generates such G and H if they exist.¹⁰ In our case, F is the function implemented by v , $X = (u_i, u_j)$, and Y consists of the remaining inputs of v . The details of the Roth-Karp decomposition can be found in [RoKa62]. Although the Roth-Karp decomposition may run in exponential time in general, it takes only constant time in our algorithm, since the number of fanins of a K-LUT is bounded by a small constant K . If the Roth-Karp decomposition succeeds on a pair of inputs u_i and u_j of node v , and $|\text{input}(u_i) \cup \text{input}(u_j)| \leq K$, then the gate-decomposition operation is applicable. In this case, we say u_i and u_j are *mergeable* and we call v the *base* of the merge.

Another post-processing operation for area optimization is called *predecessor packing*. The concept behind this method is simple. For each node, examine all of its input nodes. If $|\text{input}(v) \cup \text{input}(u_i)| \leq K$ for some input node u_i , and u_i has only a single fanout, then v and u_i are merged into a single K-LUT. In this case we also say that node u_i and v are *mergeable*, and call v the *base* of the merge. This operation reduces the number of K-LUTs by one. Unlike the gate decomposition method where the number of inputs to the current node v is reduced by one, with this method the number of inputs is actually increased by $|\text{input}(u_i) - \text{input}(v)|$. Although it is less conducive to subsequent gate-decomposition or predecessor packing, the number of instances to which this operation applies is large. Fig. 7 shows an example of the predecessor packing operation. In this example predecessor packing leads to a solution with the same depth as the original network but one fewer K-LUT.

There are usually many pairs of mergeable nodes in a network, but not all of these merge operations can be performed at the same time. We thus use a graph matching approach to avoid merging nodes in arbitrary order; this achieves a globally good result. We construct an undirected graph $G = (V, E)$, where the vertex set V represents the nodes of the K-LUT network, and an edge (v_i, v_j) is in the edge set E if and only if v_i and v_j are mergeable. Clearly a maximum cardinality matching in G corresponds to a maximum set of merge operations that can be applied simultaneously. Therefore we find a maximum matching in G and apply the merge operations corresponding to the matched edges. We then re-construct the graph G for the reduced network

¹⁰ This actually is a special case of Roth-Karp decomposition.




Fig. 7. Predecessor packing for area optimization (assume that $K = 5$).

and repeat the above procedure until we are unable to construct a non-empty E . The experimental results show that this matching based merge algorithm usually converges after only one or two iterations. Since the maximum graph matching problem can be solved in $O(n^3)$ time [Ga76], our area optimization procedure can be implemented efficiently.¹¹

Note that in the above discussion concerning these two operations we assume that each node in a mergeable pair has only a single fanout, unless it is also the base of the merge for predecessor packing. This is because the resulting K -LUT must have only one output. If a node u in a mergeable pair is not the base of the merge and has multiple fanouts, the application of the merge operation requires u to be replicated so that the copy involved in the merge operation is fanout free. However unless every fanout node of u is a base of a merge operation that involves u , we cannot reduce the number of nodes in the network, since there will always be a remaining copy of u which is not merged to any of its fanout nodes.

We say a node u is *removable* if and only if for *each* of its fanout nodes v_i , either u and v_i are mergeable via predecessor packing, or there is another fanin node of v_i , say u_j , such that u and u_j are mergeable via gate decomposition. Fig. 8 shows three different cases where node u can be shown as a removable node. For a removable node u , each of its fanout nodes is a base of a merge operation involving u , and u is removed if all these merge operations are applied simultaneously. Therefore for a removable node u , we define a *mergeable set* of u , denoted as R_u , to be a set of nodes involved in removing u . More precisely, R_u contains u itself and *exactly one* nodes for each fanout node v_i of u , which is selected in the following ways: (1) if v_i is the base of a predecessor packing operation involving u , then we can select v_i as u_i ; or (2) if v_i is the base of a gate decomposition operation involving u , then we can select u_i to be the node other than u involved in this gate decomposition. Note that a removable node may have more than one mergeable set. For example, in Fig. 8 node u has mergeable sets $\{u, u_1, u_2\}, \{u, v_1, v_2\}, \{u, u_1, v_2\}$, and $\{u, v_1, u_2\}$ (the last one is not shown in the figure). If u is fanout free, then a mergeable set of u is a mergeable pair defined previously. Therefore, a mergeable pair is a special case of a

¹¹ We used a standard procedure for maximum cardinality matching in undirected graphs, written by Ed Rothberg, that implements Gabow's algorithm [Ga76].

mergeable set.

In order to reduce the number of K-LUTs as much as possible, we want to determine a maximum collection of mergeable sets for which merge operations can be performed independently. The graph matching based approach is also suitable in this case, but the graph has to be a *hypergraph* [Be89] which allows edges of size larger than two. We construct a hypergraph $H=(V,S)$ for the K-LUT network, where the vertices in V represent the nodes of the network, and the hyperedges in S represent the mergeable sets. Note that H contains the simple graph $G=(V,E)$, which we constructed for merging fanout free nodes, as a subgraph. We will call the edges in E the *simple* edges, and call the edges in $S-E$, each of which contains more than two vertices, the *non-simple* edges. A *matching* in H is defined to be a set of disjoint edges in S . It is easily seen that a maximum matching in H yields the maximum number of K-LUTs in the network that can be removed. However, the maximum matching problem in a hypergraph is NP-complete [GaJo79]. Instead of solving this problem optimally, we compute an approximate solution using the following algorithm. First, we construct the hypergraph $H=(V,S)$ for the K-LUT network as described above. Then, we identify the subgraph $G=(V,E)$ of H , where $E \subseteq S$ consists of all the simple edges in H . Next, we find a maximum cardinality matching $M_2 \subseteq E$ on G using Gabow's algorithm [Ga76]. This matching will be included in the approximate solution for hypergraph matching. Let S_m be the set of non-simple edges in H that are disjoint from the edges in M_2 . We use an exhaustive search procedure to find a maximum matching $M_m \subseteq S_m$ and return $M \cup M'$ as the approximate maximum matching solution in H .

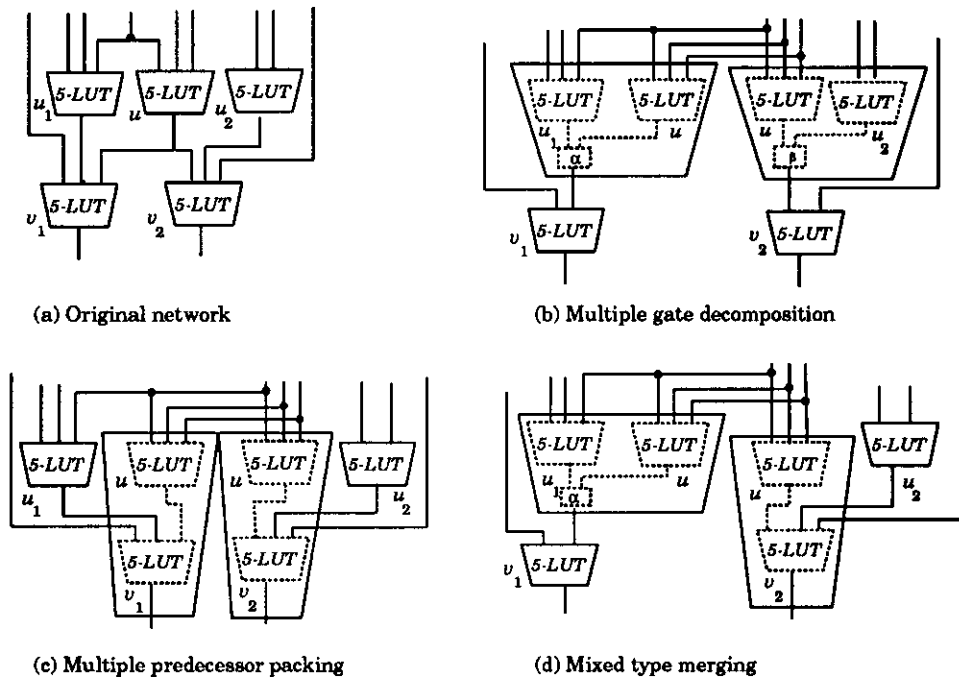


Fig. 8. Merge operations on multi-fanout node (assume $K = 5$).

In practice, $|S_m|$ is quite small. For example, for all the benchmark circuits we used in our experiments, $|S_m|$ never exceeds 10. Therefore, M_m can be computed efficiently.

It is obvious that this algorithm finds a maximal hypergraph matching.¹² Although it may not be a maximum matching, we have the following bound.

Theorem 3 Given a hypergraph H , let M^* be a maximum matching of H , and M be the matching computed using the above algorithm, then $|M^*| \leq 2|M|$.

Proof Let $M = M_2 \cup M_m$, where M_2 consists of all the simple edges in M and M_m consists of all the non-simple edges of M . Let $M^+ = M_2 - M^*$, which is the set of simple edges that are in M but not in M^* , and $M^- = M^* - M_2$. If M^+ is not empty, $M^* \cup M^+$ is not a matching since M^* is a maximum matching. But because M_2 is also a matching, and $M^+ \subseteq M_2$, we can always find a set $S \subseteq M^-$ such that $M' = (M^* \cup M^+) - S$ is a maximal matching. Note that $M_2 \subseteq M'$. The size of the matching has decreased by $|M^*| - |M'| = |S| - |M^+|$. On the other hand, by adding a simple edge to any matching, at most two edges need to be removed to retain the matching property. Therefore, $|S| \leq 2|M^+|$, so $|M^*| - |M'| \leq |M^+|$.

Since M_2 forms a maximum matching among all simple edges in H , M' cannot contain simple edges other than those in M_2 . So $M' - M_2$ consists of a maximal matching of non-simple edges which are disjoint from edges in M_2 . According to the construction of M_m , we have $|M' - M_2| \leq |M_m|$, which leads to $|M'| \leq |M|$.

Therefore, we have $|M^*| - |M| \leq |M^+|$. Note that $|M^+| \leq |M|$, hence we can conclude that $|M^*| \leq 2|M|$. \square

For general hypergraphs, this bound is *tight*. However for hypergraphs that contain only a few non-simple edges, we can obtain a better bound. Notice that in the proof we have $|M^*| - |M| \leq |S| - |M^+|$. Because M_2 is a maximum matching of simple edges, S may contain no more than $|M^+|$ simple edges. Therefore, $|S| - |M^+|$ is no more than the number of non-simple edges in S . Hence, we have

Corollary 1 If the number of non-simple edges in a hypergraph H is $h(H)$, and the matching M obtained by the above algorithm contains $k(H)$ non-simple edges, then $|M^*| - |M| \leq h(H) - k(H)$.

Proof It is clear that the number of non-simple edges in S is no more than $h(H) - k(H)$. According to the above discussion in the last paragraph, $|S| - |M^+| \leq h(H) - k(H)$. From the proof of Theorem 3 we have $|M^*| - |M'| = |S| - |M^+|$, so $|M^*| - |M| \leq h(H) - k(H)$. \square

For all the benchmark circuits we used in our experiments, we applied Corollary 1 and find out that the error bound $|M^*| - |M|$ is no more than 4.

¹² A matching is *maximal* if there is no matching that contains it as a proper subset. A maximal matching is not necessarily a maximum matching.

4. Experimental Results

We implemented the DAG-Map algorithm using the C language on Sun SPARC workstations. We integrated our program as an extension of the MIS system so that we could exploit input/output routines and other functions provided by MIS. DAG-Map was tested on a large number of MCNC benchmark examples and results were compared with both those produced by Chortle-d [FrRV91c], and those produced by the mapping phase of the MIS-pga delay optimization algorithm [MuSB91a]). The specific experimental procedure was as follows.

We chose the size of the K-LUT was to be $K = 5$, reflecting, e.g. the XC 3000 FPGA family produced by Xilinx [Xi89]. For each benchmark example, we first minimized the initial network using a standard MIS script provided by Francis [Fr91]. Next, we applied the DMIG algorithm to transform the network into a two-input network. We then used DAG-Map to map into a 5-LUT network. Finally, the matching based post-processing step was performed.

Table 1 shows the comparison of the results of our algorithm with those of the Chortle-d algorithm (quoted from [FrRV91c]) and those of the mapping phase of MIS-pga delay optimization algorithm (quoted from [MuSB91a]). We only include those benchmarks for which

Technology Mapping for 5-LUT FPGAs									
	chortle-d			MIS-pga (d)			DAG-Map		
	LUTs	depth	time	LUTs	depth	time	LUTs	depth	time
<i>5xp1</i>	26	3	0.1	21	2	3.5	22	3	1.1
<i>9sym</i>	63	5	0.2	7	3	15.2	60	5	2.3
<i>9symml</i>	59	5	0.1	7	3	9.9	55	5	2.5
<i>C499</i>	382	6	1.8	199	8	58.8	68	4	12.2
<i>C880</i>	329	8	0.9	259	9	39.0	128	8	6.3
<i>alu2</i>	227	9	0.7	122	6	42.6	156	9	7.8
<i>alu4</i>	500	10	0.3	155	11	15.4	272	10	16.5
<i>apex6</i>	308	4	0.8	274	5	60.0	246	5	10.9
<i>apex7</i>	108	4	0.2	95	4	8.4	81	4	3.0
<i>count</i>	91	4	0.1	81	4	5.1	31	5	1.4
<i>des</i>	2086	6	9.2	1397	11	937.8	1423	5	91.2
<i>duke2</i>	241	4	0.4	164	6	16.4	177	4	4.9
<i>misex1</i>	19	2	0.1	17	2	1.7	16	2	0.7
<i>rd84</i>	61	4	0.2	13	3	9.8	46	4	2.5
<i>rot</i>	326	6	1.0	322	7	50.0	246	7	11.1
<i>vg2</i>	55	4	0.1	39	4	1.7	29	3	0.9
<i>z4ml</i>	25	3	0.1	10	2	2.1	5	2	0.3
total	4906	87	16.3	3182	90	1277.4	3062	85	175.6
comparison	+60%	+2%	-	+4%	+6%	-	1	1	-

Table 1. Comparison with Chortle-d and MIS-pga (delay optimization) programs.

Comparison of 2-Input Network Transformation Methods								
	Before Mappings				After 5-LUT Mappings			
	mis tech_decomp		DMIG algo		mis tech_decomp		DMIG algo	
	gates	depth	gates	depth	gates	depth	gates	depth
<i>5xp1</i>	88	9	88	9	22	3	22	3
<i>9sym</i>	201	16	201	13	65	5	60	5
<i>9symml</i>	199	17	199	13	61	5	55	5
<i>C499</i>	392	25	392	25	66	4	68	4
<i>C880</i>	347	37	347	35	131	8	128	8
<i>alu2</i>	371	36	371	31	159	10	156	9
<i>alu4</i>	664	40	664	34	263	11	272	10
<i>apex6</i>	651	16	651	15	250	6	246	5
<i>apex7</i>	201	14	201	13	80	4	82	4
<i>count</i>	112	20	112	19	31	5	31	5
<i>des</i>	3049	19	3049	16	1461	6	1423	5
<i>duke2</i>	325	16	325	11	177	5	177	4
<i>misex1</i>	49	6	49	6	19	2	16	2
<i>rd84</i>	153	14	153	11	44	4	46	4
<i>rot</i>	539	27	539	21	256	7	246	7
<i>vg2</i>	72	15	72	10	29	4	29	3
<i>z4ml</i>	27	10	27	10	5	2	5	2
total	7440	337	7440	292	3119	91	3062	85
comparison	+0%	+15%	1	1	+2%	+7%	1	1

Table 3. Comparison of two-input network transformation algorithms.

the test data from all the programs are available. The running time (sec.) of our algorithm (including transformation, mapping, and post-processing) was recorded on a Sun SPARC IPC (15.8 MIPS), while the running time of the other two algorithms is quoted from [MuSB91a], whose authors used a DEC5500 machine (28 MIPS). Overall, the solutions of Chortle-d used 60% more lookup-tables and had 2% larger network depth, and the solutions of MIS-pga with delay optimization used 4% more lookup-tables and had 6% larger network depth.¹³ In all cases, the running time of our algorithm is no more than 100 seconds.

In order to judge the effectiveness of our DMIG algorithm for transforming the initial network into a two-input network, we compared it with the transformation procedure in MIS. Both our DMIG algorithm and the MIS decomposition command `tech_decomp -a 2 -o 2` were applied to the same initial networks¹⁴. We also ran the DAG-Map algorithm on each set of the resulting two-input networks. In Table 3, the first four columns compare the number of

¹³ If we compute the improvement for each circuit and take the average over all circuits, we found that the solutions of Chortle-d used 99% more lookup-tables and had 5% larger network depth, and the solutions of MIS-pga with delay optimization used 19% more lookup-tables and had 8% larger network depth.

Effectiveness of Post-Processing Step for Depth Minimized 5-Input Lookup Table Mappings				
	original		after post-processing	
	LUTs	depth	LUTs	depth
<i>5xp1</i>	25	3	22	3
<i>9sym</i>	76	5	60	5
<i>9symml</i>	68	5	55	5
<i>C499</i>	80	4	68	4
<i>C880</i>	137	8	128	8
<i>alu2</i>	169	9	156	9
<i>alu4</i>	301	10	272	10
<i>apex6</i>	313	5	246	5
<i>apex7</i>	101	4	82	4
<i>count</i>	43	5	31	5
<i>des</i>	1674	5	1423	5
<i>duke2</i>	196	4	177	4
<i>misex1</i>	20	2	16	2
<i>rd84</i>	51	4	46	4
<i>rot</i>	275	7	246	7
<i>vg2</i>	32	3	29	3
<i>z4ml</i>	5	2	5	2
total	3566	85	3062	85
comparison	+16%	+0%	1	1

Table 2. Effect of the post-processing steps for area minimization.

gates and the depth of the two-input networks produced by the two algorithms, while the last four columns compare the number of 5-LUTs and the depth of the 5-LUT network after DAG-Map is applied to the different two-input networks produced by the two algorithms. In all cases, the DMIG procedure resulted in smaller or the same depths in both the two-input networks after decomposition and the 5-LUT networks after mapping, and on average it used fewer lookup-tables. (Since both algorithms decompose a network into a binary tree, the number of gates in the resulting two-input networks is always the same.)

Finally, we also tested the effectiveness of the post-processing procedure for area optimization used in our algorithm, and the results are shown in Table 2. The first two columns show the statistics for the mapping solutions produced by DAG-Map without any post-preprocessing for area optimization. The last two columns describe the same solutions after the post-processing. The total number of lookup-tables is reduced by 16%.

¹⁴ Again, the initial networks were optimized using the MIS minimization script as in the preceding experiment, and in addition to this we also used `tech_decomp -a 1000 -o 1000` to transform it into simple gate network.

5. Conclusions

In this paper, we have presented a graph based technology mapping algorithm for delay optimization in lookup-table based FPGA design. It carries out technology mapping and delay optimization on the entire Boolean network. Our algorithm consists of three main steps: transformation of an arbitrary network into a two-input network, technology mapping on the entire two-input network for delay minimization, and area optimization in the mapping solution. We have also presented several theoretical results which show the effectiveness of our algorithm. The algorithm has been tested on a large set of benchmark examples and it gives satisfactory results.

6. Acknowledgments

We thank Dr. Masakatsu Sugimoto for his support of this research. We thank Professor Jonathan Rose and Robert Francis for providing the Chortle programs and necessary assistance for our comparative study.

References

- [Be89] Berge, C., *Hypergraphs*, Elsevier Science Publishers, Amsterdam, The Netherlands (1989).
- [BrRS87] Brayton, R. K., R. Rudell, and A. L. Sangiovanni-Vincentelli, "MIS: A Multiple-Level Logic Optimization," *IEEE Transactions on CAD*, pp. 1062-1081, November 1987.
- [BrRV90] Brown, S., J. Rose, and Z. Vranesic, "A Detailed Router for Field-Programmable Gate Arrays," *Proc. IEEE Int'l Conference on Computer-Aided Design*, pp. 382-385, 1990.
- [Ca86] Carter, W. et al, "A User Programmable Reconfigurable Gate Array," *Proc. Custom Integrated Circuits Conference*, pp. 233-235, 1986.
- [De87] Detjens, E., et al, "Technology Mapping in MIS," *Proc. IEEE Int'l Conference on Computer-Aided Design*, pp. 116-119, 1987.
- [El89] El Gamel, A. et al, "An Architecture for Electrically Configurable Gate Arrays," *IEEE J. Solid-State Circuits*, Vol. 24, pp. 394-398, April 1989.
- [ErDe91] Ercolani, S. and G. De Micheli, "Technology Mapping for Electrically Programmable Gate Arrays," *Proc. 28th ACM/IEEE Design Automation Conference*, pp. 234-239, 1991.
- [Fr91] Francis, R. J., *personal communication*. May 1991.
- [FrRC90] Francis, R. J., J. Rose, and K. Chung, "Chortle: A Technology Mapping Program for Lookup Table-Based Field Programmable Gate Arrays," *Proceedings 27th ACM/IEEE Design Automation Conference*, pp. 613-619, 1990.
- [FrRV91a] Francis, R. J., J. Rose, and Z. Vranesic, "Chortle-crf: Fast Technology Mapping for Lookup Table-Based FPGAs," *Proceedings 28th ACM/IEEE Design Automation Conference*, pp. 613-619, 1991.

- [FrRV91b] Francis, R. J., J. Rose, and Z. Vranesic, "Technology Mapping for Delay Optimization of Lookup Table-Based FPGAs," *MCNC Logic Synthesis Workshop*, 1991.
- [FrRV91c] Francis, R. J., J. Rose, and Z. Vranesic, "Technology Mapping of Lookup Table-Based FPGAs for Performance," *Proc. Int'l Conf. Computer-Aided Design*, pp. 568-571, Nov., 1991.
- [Ga76] Gabow, H., "An Efficient Implementation of Edmonds' Algorithm for Maximum Matching on Graphs," *Journal of the ACM*, Vol. 23, pp. 221-234, Apr. 1976.
- [GaJo79] Garey, M. and D. Johnson, *Computer and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco (1979).
- [Gr86] Gregory, D., et al, "Socrates: a system for automatically synthesizing and optimizing combinational logic," *Proc. 23th ACM/IEEE Design Automation Conference*, pp. 79-85, 1986.
- [GrKE90] Greene, J., V. Roychowdury, S. Kaptanoglu, and A. El Gamal, "Segmented Channel Routing," *Proc. 27th ACM/IEEE Design Automation Conference*, pp. 567-572, 1990.
- [Hi91] Hill, D., "A CAD System for the Design of Field Programmable Gate Arrays," *Proc. ACM/IEEE Design Automation Conference*, pp. 187-192, 1991.
- [HoKP84] Hoover, H. J., M. M. Klawe, and N. J. Pippenger, "Bounding Fan-out in Logic Networks," *Journal of Association for Computing Machinery*, Vol. 31, pp. 13-18, Jan. 1984.
- [Hs87] Hsieh, H. et al, "A Second Generation User Programmable Gate Array," *Proc. Custom Integrated Circuits Conference*, pp. 515-521, 1987.
- [Hu52] Huffman, D. A., "A method for the construction of minimum redundancy codes," *Proc. IRE 40*, pp. 1098-1101, 1952.
- [Ka86] Kahrs, M., "Matching a parts library in a silicon compiler," *Proc. IEEE Int'l Conference on Computer-Aided Design*, pp. 169-172, 1986.
- [Ka91a] Karplus, K., "Xmap: A Technology Mapper for Table-lookup Field-Programmable Gate Arrays," *Proc. 28th ACM/IEEE Design Automation Conference*, pp. 240-243, 1991.
- [Ka91b] Karplus, K., "Amap: A Technology Mapper for Selector-based Field-Programmable Gate Arrays," *Proc. 28th ACM/IEEE Design Automation Conference*, pp. 244-247, 1991.
- [Ke87] Keutzer, K., "DAGON: Technology Binding and Local Optimization by DAG Matching," *Proc. 24th ACM/IEEE Design Automation Conference*, pp. 341-347, 1987.
- [LaLT69] Lawler, E. L., K. N. Levitt, and J. Turner, "Module Clustering to Minimize Delay in Digital Networks," *IEEE Transactions on Computers*, Vol. C-18(1) pp. 47-57, January 1969.
- [LiBK88] Lisanke, R., F. Brglez, and G. Kedem, "McMAP: A Fast Technology Mapping Procedure for Multi-Level Logic Synthesis," *Proc. IEEE Int'l Conference on Computer Design*, pp. 252-256, 1988.
- [Mu90] Murgai, R., et al, "Logic Synthesis Algorithms for Programmable Gate Arrays," *Proc. 27th ACM/IEEE Design Automation Conf.*, pp. 620-625, 1990.

- [MuSB91a] Murgai, R., N. Shenoy, R. K. Brayton, and A. Sangiovanni-Vincentelli, "Performance Directed Synthesis for Table Look Up Programmable Gate Arrays," *Proc. Int'l Conf. Computer-Aided Design*, pp. 572-575, Nov., 1991.
- [MuSB91b] Murgai, R., N. Shenoy, R. K. Brayton, and A. Sangiovanni-Vincentelli, "Improved Logic Synthesis Algorithms for Table Look Up Architectures," *Proc. Int'l Conf. Computer-Aided Design*, pp. 564-567, Nov., 1991.
- [RoKa62] Roth, J. P. and R. M. Karp, "Minimization Over Boolean Graphs," *IBM Journal of Research and Development*, pp. 227-238, April 1962.
- [Wa89] Wang, A., "Algorithms for Multi-level Logic Optimization," *U.C. Berkeley Memorandum No. UCB/ERL M89/50*, April 1989.
- [Wo91] Woo, N.-S., "A Heuristic Method for FPGA Technology Mapping Based on the Edge Visibility," *Proc. 28th ACM/IEEE Design Automation Conference*, pp. 248-251, 1991.
- [Xi89] Xilinx, *The Programmable Gate Array Data Book*, Xilinx, San Jose (1989).
- [Xi91] Xilinx, *User Guide and Tutorials*, Xilinx, San Jose (1991).