

Advanced Datapath Synthesis using Graph Isomorphism

Cunxi Yu¹, Mihir Choudhury², Andrew Sullivan², Maciej Ciesielski¹
ECE Department, University of Massachusetts, Amherst¹
*IBM T.J Watson Research Center²
ycunxi@umass.edu, choudhury@us.ibm.com

Abstract - This paper presents an advanced DAG-based algorithm for datapath synthesis that targets area minimization using logic-level resource sharing. The problem of identifying common specification logic is formulated using unweighted graph isomorphism problem, in contrast to a weighted graph isomorphism using AIGs. In the context of gate-level datapath circuits, our algorithm solves the unweighted graph isomorphism problem in linear time. The experiments are conducted within an industrial synthesis flow that includes the complete *high-level synthesis*, *logic synthesis* and *placement and route* procedures. Experimental results show a significant runtime improvements compared to the existing datapath synthesis algorithms.

Keywords

Logic synthesis, datapath synthesis, resource sharing, graph isomorphism

1. INTRODUCTION

Due to a large demand for computing, the complexity of hardware systems have been significantly increasing, raising the challenges in design, verification and synthesis to a new level. In last ten years, there has been a push to make changes in optimization algorithms of EDA tools to improve performance in terms of timing, area and power. Datapath modules in modern microprocessors and embedded systems play an important role in computations, which puts new demands on logic synthesis. Traditional datapath synthesis flow includes extraction of arithmetic operations from RTL code, high-level synthesis (HLS), logic synthesis, and technology mapping [1][2]. Datapath synthesis techniques have been mainly discussed in the context of traditional high-level synthesis research, such as *resource sharing*, *scheduling* and *binding*, relied on Data Flow Graph (DFG) representation [3][4][5]. The arithmetic operations such as addition, multiplication, shifting and comparison, and control logic are extracted and are modeled as block modules. At the same time, methods such as carry prefix, and recoded par-

tial product based techniques are applied for delay optimization [6]. The remain part of the design flow produces the technology mapped netlist for a given standard-cell library.

Despite most of the datapath synthesis effort spent in high-level synthesis stage, there are many opportunities in *bit-level* optimization that could overcome limitations of high-level synthesis. Recently, applying high-level optimization technique, such as *resource sharing*, has been applied in logic synthesis. These techniques overcome some limitation of resource sharing encountered traditional datapath synthesis for standard-cell designs. Specifically, a Directed-acyclic-graph (DAG) based logic synthesis technique that targets area minimization of datapath designs was proposed in [7]. It is a structural optimization technique implemented using *And-Inv-Graphs* (AIGs) [8], which offers bit-level resource sharing. Its method includes three steps: 1) identifying sub-circuit candidates by searching *multiplexer-equivalent* AIG structure; 2) identifying common specification logic using graph isomorphism; and 3) finalizing the optimization by relocating the multiplexers. The most critical part of the technique that is step 2, which solves the problems of identifying common logic and performing Boolean matching. In fact, finding isomorphism in AIG is a *weighted* graph isomorphism problem [7]. This is because, to represent an arbitrary Boolean network using AND nodes, the edges are required to represent inversion or a wire, which classifies that AIG is a *weighted* graphs. Note that solving graph isomorphism in *weighted* graphs is much more complex than in the *unweighted* graphs [9].

Although the technique of [7] offers new direction in datapath synthesis and promises area reduction, it has some limitations. First, the complexity of general graph isomorphism problem belongs to *NP*, but is not known if it is *P* or *NP-complete*. Despite the reduction in complexity offered by DAGs, solving a weighted DAG isomorphism could still cause memory and runtime explosion. Second, since that technique is implemented based on AIG, it always requests transformations between gate-level network and AIG representations, to produce the technology mapped netlist. These transformations could affect the optimization solutions done among the previous synthesis procedures.

In this work, we develop new algorithms to overcome these limitations. Specifically, we make the following contributions:

- 1) We propose a novel algorithm for identifying common specification logic that directly supports arbitrary standard-cell netlist, without using AIG, which maintains the optimizations performed by other synthesis techniques.

2) Instead of solving *weighted* graph isomorphism problem, the proposed algorithm formulate the problem as *unweighted* graph isomorphism, which significantly reduces the complexity of solving the problem.

3) The runtime complexity comparison between the AIG-based algorithm [7] and those proposed, is provided using illustrative examples (Section 3.1), and demonstrated using large datapath designs (Figure 6).

4) The proposed algorithm allows *approximate isomorphism classes* to be optimized (Section 3.2).

5) The proposed approach has been evaluated in two complete IBM synthesis flows, including the complete flow of *high-level synthesis*, *logic synthesis* and *placement and route* (P&R), which allows it to make meaningful comparison with other techniques. The experiments used 14nm technology library.

2. BACKGROUND

2.1 Boolean network

A Boolean network can be represented using directed acyclic graph (DAG) with nodes representing logic gates and directed edges representing wires connecting the gates. If the network is sequential, the memory elements are assumed to be D flip-flops with known initial states. In this work, we only consider combinational logic optimization, which means the flip-flops are considered as primary inputs (PI) and primary outputs (PO) for the sub-circuits.

In the AIGs [8], each node has either 0 or two incoming edges. A node with no incoming edges is a primary input. Primary outputs are represented using special output nodes. Each node represents a Boolean AND function. The combinational logic of an arbitrary Boolean network can be transformed into an AIG [10], while the edges can optionally provide inversions. Hence, AIG is considered as a *weighted* DAG.

Alternatively, the Boolean network can be directly represented using the gate-level netlist. The primary inputs, primary outputs, and flip-flops are constructed based on standard-cell netlist. Each logic gate is a vertex in the DAG. The logic gates with the same corresponding logic function are considered as the same vertex type. This DAG has only one type of edge, i.e. *unweighted* DAG, and provides more uniqueness for checking isomorphism. The comparison between AIG and our representation is shown in Figure 1. The actual gate-level netlist, including one AOI21 and two NAND2 gates, is shown in Figure 1(a), and its AIG representation is shown in Figure 1(b). AIG requires four AIG nodes with four inversion edges and five non-inversion edges to represent this netlist. In contrast, the proposed representation in Figure 1(c), has three nodes in two types, and all edges are identical. There are two advantages that we use the representation shown in Figure 1(c): 1) avoid the transformations between different Boolean network to maintain the original structural, which maintains the optimizations done in previous stages; 2) convert the weighted graph isomorphism problem into unweighted graph isomorphism problem to improve the runtime for identifying *common specification logic*.

2.2 Common Specification Logic

Two combinational circuits are considered as *common specification logic* if they have the same specification [11]. In

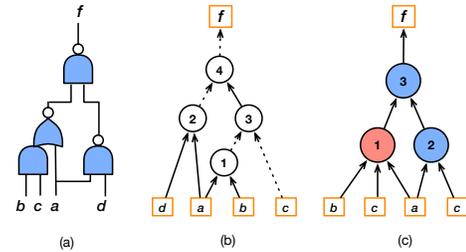


Figure 1: (a) Gate-level netlist $f = \overline{bc} + a \cdot \overline{ad}$; (b) AIG representation, $f = \overline{f_2 f_3}$, $f_2 = \overline{f_1 a}$, $f_3 = ad$, and $f_1 = bc$; (c) the proposed representation, *node1* representing AOI21, and *node2* and *node3* representing the two NAND2s.

this work, common specification logic has to be identified in the following context: given the output boundaries of two logic cones, find the input boundaries that give maximum common logic, and match the signals of the input boundaries. Most techniques for checking if two designs conform to common specification logic is *combinational equivalence checking* (CEC). This problem has been addressed by BDDs [12], SAT[13][14], AIG[10], etc. However, those methods cannot be applied. This is because, in this work, 1) the input boundaries of the designs are unknown; and 2) if the input boundaries are detected, the correspondences (Boolean matching) of those inputs are unknown. Besides, it is well known that the functional methods such as BDDs and SAT, are not scalable for gate-level arithmetic designs, such as multipliers.

2.3 Graph Isomorphism

In graph theory, an isomorphism of graphs G and H is a bijection between the vertex sets $V(G)$ and $V(H)$, $f: V(G) \rightarrow V(H)$, such that any two vertices u and v of G are adjacent in G iff $f(u)$ and $f(v)$ are adjacent in H . Besides the mathematical research on graph isomorphism, the algorithmic approach to graph isomorphism has been widely used in computer engineering, e.g. Boolean matching [15] and program similarity checking [16]. In general, graph isomorphism is understood to be undirected unlabeled unweighted graphs. It is known to be a *NP* problem, but neither a *NP-complete* nor a *P* problem using a deterministic algorithm. However, in the context of Boolean network, this problem could be solved efficiently by heuristic algorithms. In this work, we propose a novel algorithm that reduces the number of reordering operations by employing fanin-fanout information of each node (i.e. standard cell) for checking the existence of an isomorphism between two directed acyclic graphs.

3. APPROACH

The overall methodology of our approach is in three steps. *Vector multiplexer* is a set of 2-to-1 multiplexers with the identical control signals. First, they are collected by first structurally reverse engineering all the 2-to-1 multiplexers from gate-level netlist [7], and then being classified based on their control signals. Note that the multiplexers are eliminated from the collection if any of their data inputs has a fanout. In case of large multiplexers, such as 64-to-1, they are decomposed into 2-to-1 multiplexers [17]. Second, a set

of sub-circuits is created based on these vector multiplexers. Each sub-circuit is a combinational logic cone whose primary outputs are the outputs of all multiplexers in the vector multiplexer. These two procedures are *pre-processing* step. Third, a multiplexer relocation function is applied to each output of the sub-circuit iteratively. The order of applying multiplexer relocation is sorted by the number of logic gates per multiplexer in the sub-circuit. The original design will be updated if the area of the sub-circuit is improved by relocating the multiplexers, i.e., moving the multiplexers backward without changing the functionality of the design. The resulting updated standard-cell netlist, and will be subjected to the remaining logic synthesis steps and eventually to physical design.

3.1 Exact Isomorphism Determination

Even though the multiplexer relocation is applied to a sub-circuit that includes vector multiplexers at the primary outputs, the actual relocation is done individually for each multiplexer. The goal of multiplexer relocation is to share the common specification logic that are the input cones of the multiplexers, by moving the multiplexers backward. The main challenge is to identify the common specification logic in the sub-circuits created by pre-processing step. Specifically, this requires identifying common structure identification and Boolean matching. According to the definition of *graph isomorphism*, the algorithm proposed in [7] determines the isomorphism boundary between two graphs using breath-first-search. To obtain the maximum common logic, a look-ahead heuristic is applied in case of there are multiple identical choices of constructing isomorphism. This could potentially cause an exponential runtime and memory explosion problem, especially in the design with many reconvergent fanouts. In this section, we introduce a novel algorithm to improve the runtime and scalability for identifying common specification logic.

3.1.1 Standard-cell based DAG advantages

Instead of using AIG representation, the standard-cell based representation gives two advantages: 1) some optimization efforts in other stages of the synthesis flow, that may disappear during the transformations between AIG and standard-cell netlist are maintained; 2) standard-cell representation significantly reduces the possible choices for checking the existence of isomorphism. For this advantage, there are three reasons: **(a)** in each topological level, the total possible pairing choices is reduced; **(b)** edge type is no longer necessary to be considered, which makes the isomorphism problem to be *unweighted*; and **(c)** utilizing the number of inputs and outputs of each standard-cell reduces the number of possible choices when checking isomorphism, especially in the representation of logic circuits. We demonstrate these using an example in Figures 2 and 3.

Example 1 (Figure 2) The standard-cell netlist is shown in Figure 2. Signals $data_0$ and $data_1$ are the two inputs to a 2-to-1 multiplexer. Signals a, b, c, d , and e are the primary inputs. In each logic cone, the first two levels logic includes one AOI21 and three NAND2 gates. Each gate is considered as a vertex. The determination process starts with g_0 and g_5 . Then, two vectors of vertices are created using breath-first-search since g_0 and g_5 are the same type vertices. $V_0 = \{g_1, g_4\}$, $V_1 = \{g_7, g_6\}$. To maintain the traversed graphs in the isomorphic class, there exists only one pairing choice, i.e.

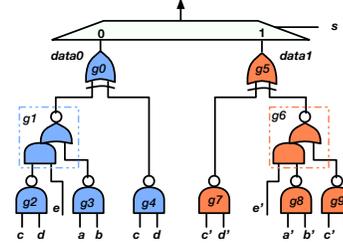


Figure 2: Determine graph isomorphism using standard-cell based DAG.

$(g_1, g_6), (g_4, g_7)$. The two vectors will be updated, $V_0 = \{g_2, e, g_3\}$ and $V_1 = \{g_2, e', g_3\}$. Since x and y are primary inputs, they are paired and eliminated from V_0 and V_1 . Hence, we have two NAND2 vertices in each vector, which has two pairing options, i.e. (g_2, g_8) or (g_2, g_9) . However, in the standard-cell based DAG, only one option remains. This is because AOI21 has two types of inputs, including two inputs for AND and one input for OR/NOR. To maintain the function equivalence, g_2 must pair with g_8 , and so g_3 must pair with g_9 . In summary, the total number of possible attempts for determining isomorphism for the first two level logic is **one**.

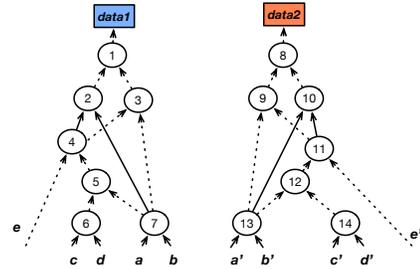


Figure 3: Determining graph isomorphism using AIG.

However, this approach requires much more effort to determine the maximum isomorphism while using AIG representation. The AIG representation of this design is shown in Figure 3. According to the algorithm proposed in [7], the first level logic has two options for pairing, i.e. node 2 with node 9, or node 2 with node 10. The algorithm solves this problem using a look-ahead heuristic, which traverses three levels deeper and picks the pairing that gives more common logic. This situation happens also while checking (node 4 with node 7, and node 11 with node 13), and (node 6 with node 7, and node 13 with node 14). This means that it requires three times look-ahead checking and total of **eight** attempts to identify the same common logic as the one shown in Figure 2.

3.1.2 Including side fanout information

Based on the observation shown in Example 1, we can see that more various types of vertices at each topological level logic can significantly reduce the total number of pairing attempts for isomorphism determination. Thus, we preserve the fanout information of the standard cells in the vertices. This can significantly improve the runtime for a large design that includes many reconvergent fanouts, such as the optimized multipliers.

Example 2 (Figure 4) Assume that each logic cone of

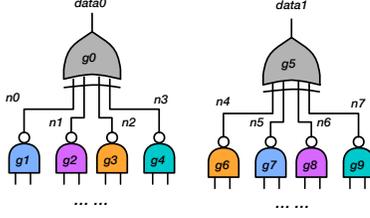


Figure 4: Illustrative example of utilizing fanout information of each vertex.

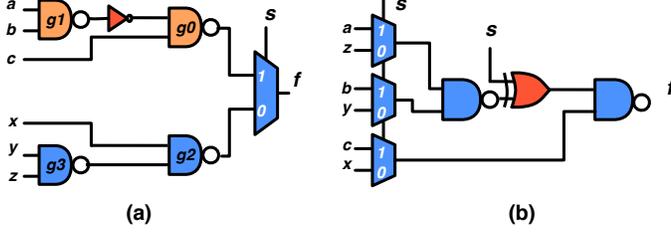


Figure 5: Approximate isomorphism determination by ignoring inverters. a) original design; b) optimized design using extra XOR2 gate.

a 2-to-1 multiplexer includes one XOR4 and four NAND2 gates in the first two levels. Let the number of side fanouts of nets $\{n_0, n_1, n_2, n_3\}$ be $\{3, 2, 1, 0\}$, and the number of side fanouts of nets $\{n_4, n_5, n_6, n_7\}$ be $\{1, 3, 2, 0\}$. Without including the fanout information, the total number of possible pairing is 24 since four vertices in the second level are identical. However, if we consider to pair the vertices according to the number of side fanouts, there will be only one pairing choice, i.e. (g_1, g_7) , (g_2, g_8) , (g_3, g_6) , and (g_4, g_9) . Although, the fanout information can significantly reduce the number of pairing, such case may not always exist. If so, our approach will go through the look-ahead heuristic pairing process.

3.2 Approximate Isomorphism Determination

In addition to considering the exact isomorphism graph as common specification logic, a novel approximate isomorphism determination approach is developed in this work. One observation is that much more common logic exists by ignoring the inversions. For example, in the case of a 2-to-1 multiplexer that selects *less than* operator and *less than or equal to*, there is no common logic that can be identified using both representations while considering inversions. Thus, we propose an approximate isomorphism method to overcome this limitation. Specifically, in the process of identifying common logic, the inverters will be replaced by a 2-input XOR, with an extra input coming from the control signal of the multiplexer, or its complement.

Example 3 (Figure 5) The original netlist is shown in Figure 5(a). Using the approach described in Section 3.1, there will be only one gate in each instance of the common logic, namely g_0 and g_2 . However, we can see that the two logic cones connected to the 2-to-1 multiplexer are identical without considering the inverter. Hence, we continue searching for the common logic by skipping the inverters. In this example, the common logic includes two NAND2 and one in-

verter. To maintain the original function of f , the inverter is replaced by an XOR2, whose extra input is the control signal s . In Figure 5(b), signal s in the XOR2 actually selects the XOR2 to be an inverter or wire, i.e. when $s = 1$, XOR2 is a inverter; and when $s = 0$, XOR2 is a buffer.

3.3 Implementation

The implementation of single multiplexer relocation is shown in Algorithm 2. The multiplexer relocation function of sub-circuit with a vector multiplexer at the primary outputs (line 5 in Algorithm 1), is applying the single relocation function iteratively on each output bit. The input of Algorithm 2 is a sub-circuit with single output bit that is generated by a of 2-to-1 multiplexer. Algorithm 2 operated in three steps:

Algorithm 1 Single Multiplexer Relocation

Input: Pre-processed sub-circuit C
Output: An optimized standard-cell netlist

Single_Mux_Relocate(C)

- 1: $B = \text{RelocationBoundary}(PO)$
- 2: $C \leftarrow$ relocate multiplexer to level B , w/o considering inverters
- 3: $P = \text{inv2xorPosition}(PO, B)$
- 4: $C \leftarrow$ insert XORs to P based on its location
- 5: **return** C

RelocationBoundary(PO)

- 1: $m \leftarrow \text{levels}(PO) - 1$; inverter is considered as 0 level.
- 2: **while** $m \geq 0$ **do**
- 3: $L0_m \leftarrow$ the gates in $(s = 0)$ logic at level m
- 4: $L1_m \leftarrow$ the gates in $(s = 1)$ logic at level m
- 5: **if** $\text{uniqueFanoutPairs}(L0_m, L1_m)$ **then**
- 6: $U0_m, U1_m \leftarrow \text{uniqueFanoutPairs}(L0_m, L1_m)$
- 7: $L0_m \leftarrow L0_m - U0_m$; $L1_m \leftarrow L1_m - U1_m$
- 8: $L0_{m+1}, L1_{m+1} \leftarrow (U0_m, U1_m) + \text{isomorphism}(L0_m, L1_m)$
- 9: **else**
- 10: **if** $\text{isomorphism}(L0_m, L1_m)$ **then**
- 11: $L0_{m+1}, L1_{m+1} \leftarrow \text{isomorphism}(L0_m, L1_m)$
- 12: **else**
- 13: **exit**
- 14: **end if**
- 15: **end if**
- 16: **end while**
- 17: **return** $(\text{level}(PO) - 1 - m), (L0_{m-1}, L1_{m-1})$

inv2xorPosition($PO, \text{boundary}$)

- 1: $P_0 \leftarrow$ the positions of all inverters up to boundary level
 - 2: $P_1 \leftarrow$ the positions of all inverters up to boundary level
 - 3: **return** $P_0 \cap P_1$
-

- The key function of this approach is identifying the maximum common specification logic connected to the multiplexer. The function is described in function **RelocationBoundary** in Algorithm 2. Specifically, our algorithm identifies the boundary logic cut where the isomorphism between two logic cones ends. This function also returns the pairings of the boundary signals that maintains the isomorphism class, which is used for creating the new multiplexers.

We backward traverse the graph from the two inputs of the 2-to-1 multiplexer level by level (lines 1 - 2). The gates at level m are stored in two vectors (lines 3 - 4), depending their selecting signal. As mentioned in Section 3.1.2, our approach benefits significantly from the fanout information. Hence, we first check if there exist unique fanout pairs. If so, we eliminate those

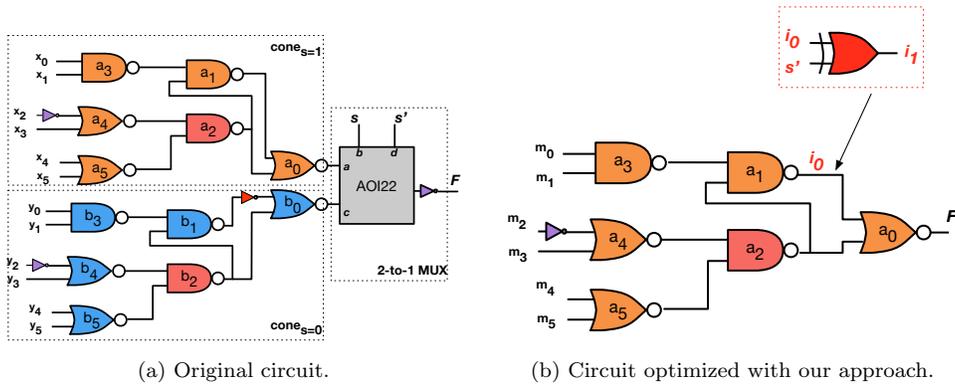


Figure 6: A complete example of multiplexer relocation using the proposed approach.

pairs from the two vectors that store the gates. The rest of the gates in the two vectors will do a regular isomorphism check, with a 3-depth look ahead search [7]. For example, in Figure 6, there are two NAND2 gates in each vector, (a_1, a_2) and (b_1, b_2) . There are two pairing choices at this level, i.e., (a_1, b_1) and (a_2, b_2) , or (a_1, b_2) and (a_2, b_1) . Using the fanout information, there will be only one feasible pairing, i.e., (a_1, b_1) and (a_2, b_2) . This is because a_2 and b_2 have two fanouts, and a_1 and b_1 have only one fanout.

- Relocate the multiplexer across the common specification logic, up to the boundary cut returned by the previous step. The two logic cones between boundary and the multiplexer output have common specification (not functionally equivalent), denoted as $cone_{s=0}$ and $cone_{s=1}$, depending on the select signal of the multiplexer. To relocate the multiplexer, we disconnect all the pins of $cone_{s=1}$ and create a set of multiplexers that select the inputs signals of those two logic cones. For example, in Figure 6, $m_i = x_i s + y_i \bar{s}$, $i = \{1, 2, 3, 4, 5\}$. Then, the inputs of $cone_{s=1}$ will be replaced by the outputs of the new multiplexers. In this case, x_i is replaced by m_i . Finally, the output F will be reconnected to the output of $cone_{s=1}$.
- In the function of **RelocationBoundary**, we do not consider inverter as a gate, or a node in the DAG. This enables the approximate isomorphism determination (Section 3.2). As mentioned earlier, this allows us to identify a larger common logic. For example, if we consider inverter as a node in the graph, the common logic will consist of only two NOR2 gates, a_0 in $cone_{s=1}$ and b_0 in $cone_{s=0}$. To maintain the functionality of the design, we need to insert XOR2 gates with extra input s or \bar{s} depending on which cone the invert belongs to. We first record the locations of all inverters in $cone_{s=0}$ and $cone_{s=1}$, denoted as P_0 and P_1 , up to the boundary cut. The locations that require an XOR2 replacement is included in the result of $P_0 \cap P_1$. This is why the inverters connected to gates a_4 and b_4 do not require XOR2 insertion, since they maintain the two cones in the isomorphism class (Figure 6). The inverter connected to b_0 requires an XOR2 insertion, and it belongs to $cone_{s=0}$. Hence, an XOR2 with extra input \bar{s} is inserted to replace i_0 in Figure 6.

4. EXPERIMENTAL RESULTS

The proposed approach in this Section 3 was implemented in C++ and integrated with the IBM logic synthesis flow [18] and further evaluated with IBM high-level synthesis flow and Place and Route (P&R) flow. Our approach is performed before technology mapping within the logic synthesis flow. The program was tested on a number of datapath designs in SystemC. The datapath designs include large arithmetic operators, such as 64-bit multipliers. All the experimental results are collected at the end of the complete production design flow. This demonstrates that our approach successfully overcomes the limitations of the existing logic synthesis and high-level synthesis techniques reviewed in Section 1. All of our experimental results are obtained using high-performance 14nm technology library. To demonstrate the runtime improvement compared to the work of [7], we examine the runtime using a set of designs, including a multiplier circuit up to 64 bits. Our experiments were conducted on a machine with Intel(R) Xeon CPU 7560 v6 2.20 GHz x32 with 4 TB memory.

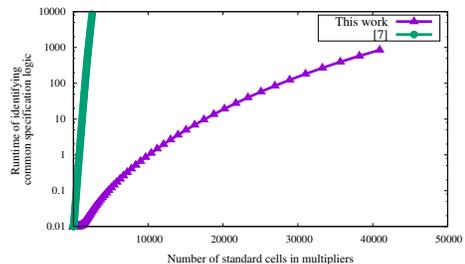


Figure 7: Evaluation of CPU runtime using designs with multipliers compared to [7].

We first evaluate our approach using a set of arithmetic designs in which there are two arithmetic operators selected by control signals. The results are shown in Table 1. The first column indicates the bit-width of the arithmetic operators and the type of the two operators. These designs are implemented in SystemC using "if then else" statement. The second and third columns show the area and logic level results produced by the original IBM synthesis flow. The fourth and fifth columns show the results produced by the original flow with combinational AIG optimization [10]. The last two columns show the results produced by original flow

(n-bit) Operators	IBM flow		IBM flow with AIG Opt		IBM flow with our approach	
	Area	Lev	Area	Lev	Area	Lev
(64), A<B:A<C	2280	11	2124	13	1855	15
(64), A+B, A+C	10162	17	9333	15	5787	20
(64), A+B:A-C	8697	19	8104	25	8062	21
(64) A<B:A<=B	2464	12	2126	13	2198	12
*(64) A×B:A×C	182917	83	482811	211	91245	89
A×B/C[7:0]:A×B/C[15:8]	3626	26	5606	26	1760	27
(32) A×B+C:B×C+A	52943	58	108402	120	26709	58
(6) dec(A):dec(B)	1319	5	667	5	549	7
	1	+0 lev	1.106	+1.16 lev	0.658	+2 lev

Table 1: Results of arithmetic test cases using the original IBM synthesis Flow, IBM synthesis flow with AIG optimization, and original IBM synthesis flow with the proposed approach. (*This design is not used for comparison.)

Benchmarks	Flow1		Flow1 with our approach		Flow2		Flow2 with our approach	
	Area	Delay	Area	Delay	Area	Delay	Area	Delay
ibm1	3622	216.45	2223	255.81	4587	295.16	2235	255.81
ibm2	5454	314.84	3361	354.19	6879	432.90	3366	383.71
ibm3	9115	501.77	5526	501.77	11463	688.71	5610	541.13
ibm4	12782	678.87	7874	649.35	16047	924.84	7854	688.71
ibm5	18323	924.84	11121	787.10	22923	1200.32	12342	875.65
ibm6	27435	1170.81	16843	983.87	34383	1505.32	16803	1023.23
ibm7	31069	1288.87	19083	1082.26	38967	1603.71	19074	1082.26
	1	1	0.613	0.970	1	1	0.487	0.767

Table 2: Evaluation of our approach in the complete production Flow using industry designs in 14nm technology. Flow1 is the IBM synthesis flow with AIG optimization; Flow2 is the original IBM synthesis flow.

with our approach. The last row shows the average improvement gain or loss. Specifically, the increase or decrease area is measured in percentage of the original flow, and the change of logic level is measured in the number of levels. Based on Table 1, we can see that:

- Our approach gives on average 34% area reduction compared to the other two flows. Note that the flows include complete high-level and logic-level optimizations techniques.
- Our approach can handle large complex arithmetic operators, such as datapath with large multipliers. With approximate isomorphism determination, we can optimize the design with various combinations of two different operators.

We then evaluate our approach using seven industrial designs implemented in SystemC. Two synthesis flows are used for experiments: *Flow1* is the IBM synthesis flow with AIG optimization; *Flow2* is the original IBM synthesis flow. The results are shown in Table 2. The second and third columns show the results produced by Flow1, and fourth and fifth columns are produced by Flow1 with our approach. The sixth to seventh columns show the results produced by Flow2. We compare the average improvement of the area and the delay at the last row. We can see that both area and delay have been improved in these experiments. Specifically, using Flow1 the area on average reduces by 39%, and the delay on average reduces 3%, and Flow2 offers 51% area reduction with 23% delay improvement on average. Note that the delay improvements are not provided directly by our approach. The delays are improved because our approach enables other optimization techniques. Specifically, for those benchmarks, an Adder optimization technique [6] implemented in the IBM synthesis flow is enabled and significantly improves the delay after relocating the multiplexers.

Additionally, we evaluate our approach using four designs, *ibm1*, *ibm2*, *ibm4*, *ibm6*, with placement and route (P&R).

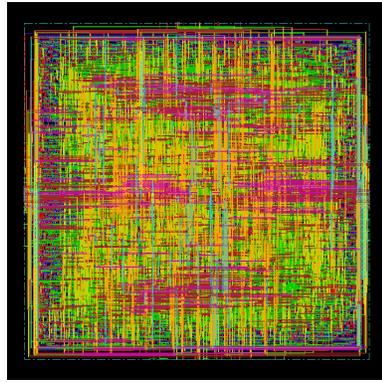
Table 3: Comparing the PnR results with multiplexer relocation with the original flow.

Benchmarks	Route Length	Power	Worst-case delay
ibm1	0.73	0.45	0.95
ibm2	0.79	0.61	0.97
ibm4	0.92	0.71	1.06
ibm6	1.23	0.78	1.10

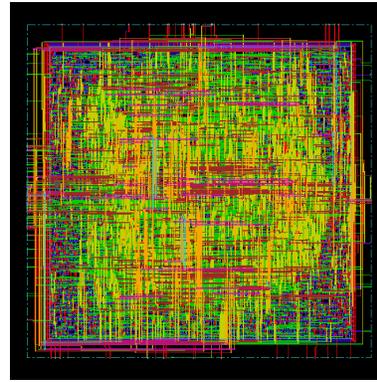
The inputs of P&R process are the designs produced by Flow1 with AIG optimization (4th and 5th columns in Table 2). The routing length, power and worst-case delay are included in Table 3. The improvements of the area of placing the standard cells remain the same as shown in Table 2 with the same density. The P&R results of *ibm2* are shown in Figure 8. We can see that except *ibm6*, the designs are improved successfully using our approach without delay overhead. Particularly, we observe that the power has been significantly improved compared to the original designs. Moreover, we can see that the improvements of *ibm4* and *ibm6* gained after P&R are less than in the other two designs. The possible reasons for that are:

- There are large (≥ 32) fanout signals generated by multiplexer relocation in those two designs.
- A large number of the extra multiplexers have been placed tightly, which decreases routability.

The reason why we didn't compare our approach to the work of [7] in the experiments shown in Table 1 and Table 2 is the following: 1) that algorithm can't be successfully applied on all of the design within eight hours; 2) For the designs that on which the algorithm runs successfully, the results are worse, e.g., 3rd and 4th designs in Table 1. To demonstrate that our approach significantly improves the CPU runtime compared to the existing algorithm in the cases of datapaths with multipliers, the experimental results



(a) P&R result of design *ibm2* without multiplexer relocation.



(b) P&R result of design *ibm2* with multiplexer relocation.

Figure 8: Comparing the P&R results using design *ibm2* with and without our approach.

are provided in Figure 7. The designs used for the experimental results shown in Figure 7 vary from 4-bit to 64-bit. In Figure 7, the x-axis represents the number of standard cells in the design, and the y-axis represents the CPU runtime of the multiplexer relocation algorithm in logarithmic scale. It is clear that our algorithm performs much faster than the AIG-based algorithm [7].

5. CONCLUSION

This paper presents an advanced DAG-based algorithm that targets area minimization using logic-level resource sharing. The common specification logic identification is formulated as *unweighted* graph isomorphism problem, in contrast to a *weighted* graph isomorphism problem using AIG. In addition to exact isomorphism classes, an *approximate* isomorphism relocation approach is proposed in this paper. The proposed approach demonstrates that it can significantly reduce area, and potentially reduce delay on industrial designs, within a complete design flow. The runtime has been reduced from exponential to linear comparing to the existing algorithms. Future work will focus on improving function of identifying common specification logic.

6. REFERENCES

- [1] L. Stok, “Data path synthesis,” *Integration, the VLSI journal*, vol. 18, no. 1, pp. 1–71, 1994.
- [2] G. D. Micheli, *Synthesis and Optimization of Digital Circuits*. McGraw-Hill Higher Education, 1994.
- [3] M. Potkonjak and J. Rabaey, “Optimizing Resource Utilization using Transformations,” *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 13, no. 3, pp. 277–292, 1994.
- [4] M. B. Srivastava and M. Potkonjak, “Optimum and Heuristic Transformation Techniques for Simultaneous Optimization of Latency and Throughput,” *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 3, no. 1, pp. 2–19, 1995.
- [5] J. Cong and J. Xu, “Simultaneous FU and Register Binding-based on Network Flow Method,” in *Design, Automation and Test in Europe, 2008. DATE’08*. IEEE, 2008, pp. 1057–1062.
- [6] S. Roy, M. Choudhury, R. Puri, and D. Z. Pan, “Towards optimal performance-area trade-off in adders by synthesis of parallel prefix structures,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 33, no. 10, pp. 1517–1530, 2014.
- [7] C. Yu, M. J. Ciesielski, M. Choudhury, and A. Sullivan, “Dag-aware logic synthesis of datapaths,” in *Proceedings of the 53rd Annual Design Automation Conference, DAC 2016, Austin, TX, USA, June 5-9, 2016*, 2016, pp. 135:1–135:6.
- [8] A. Mishchenko, S. Chatterjee, and R. Brayton, “DAG-aware AIG Rewriting A Fresh Look at Combinational Logic Synthesis,” in *43rd DAC*. ACM, 2006, pp. 532–535.
- [9] S. Umeyama, “An eigendecomposition approach to weighted graph matching problems,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 10, no. 5, pp. 695–703, 1988.
- [10] A. Mishchenko *et al.*, “ABC: A System for Sequential Synthesis and Verification (2007),” URL <http://www.eecs.berkeley.edu/alanmi/abc>, 2010.
- [11] E. Goldberg, “Equivalence Checking of Dissimilar Circuits II,” Technical report, Tech. Rep., 2004.
- [12] R. E. Bryant, “Graph-based Algorithms for Boolean Function Manipulation,” *Computers, IEEE Transactions on*, vol. 100, no. 8, pp. 677–691, 1986.
- [13] A. Kuehlmann and F. Krohm, “Equivalence checking using cuts and heaps,” in *Proceedings of the 34th annual Design Automation Conference*. ACM, 1997, pp. 263–268.
- [14] E. Goldberg, M. Prasad, and R. Brayton, “Using sat for combinational equivalence checking,” in *Proceedings of the conference on Design, automation and test in Europe*. IEEE Press, 2001, pp. 114–121.
- [15] M. Soeken, B. Sterin, R. Drechsler, and R. Brayton, “Simulation graphs for reverse engineering,” in *Proceedings of the 15th Conference on Formal Methods in Computer-Aided Design*. FMCAD Inc, 2015, pp. 152–159.
- [16] W. Li, H. Saidi, H. Sanchez, M. Schäfer, and P. Schweitzer, “Detecting similar programs via the weisfeiler-leman graph kernel,” in *International Conference on Software Reuse*. Springer, 2016, pp. 315–330.
- [17] S. Mitra, L. J. Avra, and E. J. McCluskey, “Efficient multiplexer synthesis techniques,” *IEEE Design & Test of Computers*, vol. 17, no. 4, pp. 90–97, 2000.
- [18] L. Stok, D. Kung, and *et al.*, “BooleDozer: Logic Synthesis for ASICs,” *IBM Journal of Research and Development*, vol. 40, no. 4, pp. 407–430, 1996.