

Spectral-Based Multi-Way FPGA Partitioning

Pak K. Chan*, Martine D.F. Schlag,[†] and Jason Y. Zien

Computer Engineering
University of California, Santa Cruz
Santa Cruz, California 95064 USA

November 21, 1994

Abstract— Recent research on FPGA partitioning has focussed on finding minimum cuts between partitions without regard to the routability of the partitioned subcircuits. In this paper we develop a spectral approach to multi-way partitioning in which the primary goal is to produce routable subcircuits while maximizing FPGA device utilization. To assist the partitioner in assessing the routability of the partitioned subcircuits, we have developed a theory to predict the routability of the partitioned subcircuits prior to partitioning. Advancement over the current work is evidenced by results of experiments on the standard MCNC benchmarks.

I FPGA partitioning

The design flow using commercial FPGA tools involves technology mapping, placement and routing. Since FPGA devices have relatively low density, the use of multiple FPGAs is often required to implement a large circuit. A large circuit has to be decomposed or partitioned into subcircuits for a multiple-FPGA realization, as shown in Fig. 1.

Modifications of standard iterative mincut-based partitioning algorithms have been applied to FPGA partitioning. In [15], Kuznar *et al* considered the problem of partitioning a circuit for heterogeneous FPGA systems. Their cost function was the total-dollars to implement the circuit. This approach is perhaps suitable from the standpoint of developing a new multiple-FPGA board to realize the given circuit. However, it doesn't consider the (labor) cost to layout the FPGA systems, where regularity is one of the major design issues. In addition, most existing multiple FPGA systems are homogeneous systems; for example, the Quickturn emulators and reconfigurable FPGA-based computing engines [19, 14]. So there is a need for partitioning algorithms for homogeneous FPGA systems.

Given a large circuit, a partitioner generates many subcircuits. Only when *all* the subcircuits have been successfully placed and routed, will the multiple FPGAs realize the large circuit. An important issue that has not been considered in the literature is the routability of the partitioned subcircuits. In this paper, we shall see that the manner in which a circuit is partitioned can determine whether *all* the subcircuits can be automatically routed. We present a theory to predict the routability of partitioned subcircuits prior to partitioning: pre-partitioning routability prediction. The routability predictor can assist a user in determining the number of partitions, and the selection of FPGA devices (when possible) to realize the large circuit.

Knowing the parameters that would affect the routability of the subcircuits, we devise a spectral-based partitioning algorithm to decompose a large circuit into subcircuits. Routability of the partitioned subcircuits is the primary concern. The inputs to our partitioner include partition size and cut constraints. Although spectral-based partitioning algorithms have not been recognized for their ability to realize hard constraints, we shall demonstrate that spectral-based partitioning algorithms are excellent candidates for routability reasons.

In this paper, we shall present two main results on spectral-based partitioning targeted for homogeneous FPGA systems:

1. a theory to predict the routability of the partitioned subcircuits to assist the partitioner in assessing the routability of the partitioned circuits prior to partitioning.
2. a k -way, spectral partitioning method which handles both partition-size and cut-size constraints. The partitioner's primary goal is to generate routable subcircuits while maximizing logic utilization.

*Supported in part by NSF Grant MIP-9196276 & MIP-9223740 & MICRO Program of University of California

[†]Supported in part by NSF Grant MIP-9223740 & MICRO Program of University of California.

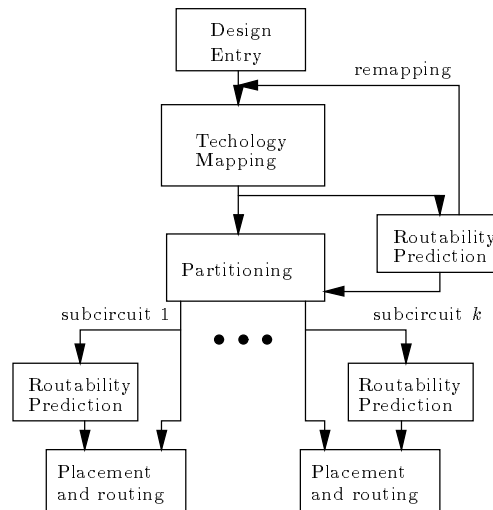


Figure 1: Multiple FPGAs design flow.

We compare our partitioning results to those in the literature using the standard MCNC Partitioning 1993 benchmarks. Our partitioner yields partitions which are predictable in the routability sense, while still producing fewer partitions in comparison to results reported in the literature.

II Routability of partitioned subcircuits

A routability barometer has been suggested by the authors in [5] to predict the routability of a single FPGA circuit. We attempt to extend the technique to predict the routability of the partitioned subcircuits. There are two ways to approach this. The straightforward approach is to partition the circuit, and apply the single-FPGA routability predictor to the partitioned subcircuits. We call this post-partitioning routability prediction. A more ambitious approach is to predict the routability of the subcircuits prior to partitioning, and we call this pre-partitioning routability prediction. This approach has the benefit that it can assist a user in determining the number of partitions, and the suitable FPGA devices required to generate routable subcircuits. Given a single FPGA circuit, a single-FPGA routability predictor benefits from the fact that the structure of the circuit can be extracted, hence the average wire length can be estimated with accuracy. Pre-partitioning routability prediction for multiple FPGAs is harder because the structures of the partitioned circuits can only be estimated before partitioning, and the predictions might not be as accurate as the single-FPGA case.

Figure 1 depicts the roles of both routability predictions in the multiple-FPGA design flow. Given a large circuit to be partitioned, the design flow also suggests two possible methods to control the routability of the partitioned subcircuits:

1. remap the given circuit to achieve 100 percent routable subcircuits, and
2. partition the given circuit to facilitate the routability of the subcircuits.

The first method has been the subject of several papers [22, 21, 4]. Here, we shall focus on the second method.

Simply put, the routability prediction for a single FPGA circuit involves estimating the channel width requirement of the circuit for routing completion [10]. The parameters of the circuit involved are:

1. γ : the average number of pins emanating from each logic block, we shall refer to γ as the pins-per-CLB ratio,
2. β : the average number of pins on a net (we refer to β as the pins-per-net ratio), and
3. L : the average wire length of the circuit after placement.

The *average* channel width requirement for routing completion is [10, 5]:

$$W = \frac{1}{2} \left[\gamma \cdot \left(1 + \frac{\beta - 2}{\beta} \right) L \right] . \quad (1)$$

Routability is determined by comparing W with the FPGA device's channel width. We present a theory to predict the *average* channel width requirement of subcircuits before partitioning by predicting the average pins-per-CLB and pins-per-net ratios of the partitioned subcircuits (before the circuit is partitioned). The inputs to the routability predictor are:

1. the total number of CLB pins of the (original) circuit,
2. the total number of nets of the circuit,
3. the average wire lengths of the subcircuits,
4. the number of partitions k , and
5. the average number of Input/Outputs pins per partition.

The first two parameters can be calculated from the (unpartitioned) circuit. The average wire lengths of the subcircuits can only be estimated (see Equation (9) later). The last two parameters depend on the devices of the homogeneous FPGA system. Assuming that the pins-per-CLB ratios do not vary widely among the partitioned subcircuits, on the average the pins-per-CLB ratio of a partitioned subcircuit is:

$$\bar{\gamma}_p \approx \gamma = \text{total CLB pins} / \text{total CLBs} . \quad (2)$$

Similarly, the average pins-per-net ratio of the partitioned subcircuits is

$$\bar{\beta}_p = \text{CLB and IO pins in a partition} / \text{nets in a partition} . \quad (3)$$

The average pins-per-net ratio $\bar{\beta}_p$ can be estimated with good accuracy, as we shall see. Let

1. $pins$ denote total CLB pins before partitioning,
2. $pins_p$ denote CLB pins in partition p ,
3. $nets$ denote total nets before partitioning,
4. $nets_p$ denote (uncut) nets (with one or more pins) contained in partition p ,
5. IO_b denote total IOs before partitioning, this is the sum of all IO pins in the original circuit,
6. IO_a denote total IOs after partitioning, this is the total number of IO pins in all the partitioned subcircuits, and
7. io_p denote the number of IO pins in partition p .

The number of nets before partitioning and after partitioning can be related as:

$$nets = \sum_{p=1}^k nets_p - (IO_a - IO_b) + \text{nets_cut_by_partitioner} \quad (4)$$

which essentially states that "nets" are conserved. Equation (4) follows from the observation that when a net is partitioned into t pieces exactly t new I/Os will be introduced if the net had no I/O pin to begin with, while only $t-1$ new I/Os will be introduced if it had an I/O pin (if it was already cut). The "nets_cut_by_partitioner" term only counts the nets cut by the partitioner which were not connected to the original IOs. Equation (4) is exact, and on the average:

$$\overline{nets}_p = (nets + (IO_a - IO_b) - \text{nets_cut_by_partitioner}) / k \quad (5)$$

so the average pins-per-net ratio is

$$\begin{aligned} \bar{\beta}_p &\approx \frac{\text{average_number_of_pins in a partition}}{\text{average_number_of_nets_in_a_partition}} \\ &= (\overline{pins}_p + \overline{io}_p) / \overline{nets}_p \\ &= k \times (\overline{pins}_p + \overline{io}_p) / (nets + (IO_a - IO_b) - \text{nets_cut_by_partitioner}) \\ &= (pins + k \times \overline{io}_p) / (nets + (IO_a - IO_b) - \text{nets_cut_by_partitioner}). \end{aligned} \quad (6)$$

We'll need to make some approximation of Equation (4) before Equation (6) can be useful. We approximate:

$$\text{nets_cut_by_partitioner} \approx (IO_a - IO_b) / \min(k, \beta) \quad (7)$$




Figure 2: Tradeoffs between routability and number of FPGAs used to implement the partitioned subcircuits.

since these are the nets that are not connected to the IOs. On the average these nets fanout to k or β chips, which ever is smaller. Also, $IO_a = k \times \overline{io_p}$. So from Equations (6) and (7), we deduce

$$\overline{\beta}_p \approx \frac{(pins + k \times \overline{io_p})}{nets + (k \times \overline{io_p} - IO_b)(1 - 1/\min(k, \beta))} \quad (8)$$

Given a circuit, we can extract the number of pins, the number of nets, and the number IO pins, so we know β . If we know the number of partitions and the average number of IO pins used after partitioning, we should be able to predict the average pins-per-net ratio $\overline{\beta}_p$ prior to the actual partitioning. A good guess of $\overline{io_p}$ is the number of I/O pins available on the FPGA device. From Equations (1), (2), and (8), we can predict the routability of the partitioned circuits even before partitioning if the average wire length can be estimated. Our approach is to calculate L based on the logic utilization of the FPGA devices and the FPGA device used:

$$L = \sqrt{\text{logic_utilization}} \times L_{device_type} \quad (9)$$

where L_{device_type} depends solely on the FPGA device type used. We shall validate our routability theory and Equation (8) through experiments presented in Section IV.

Equation (8) identifies an important aspect of partitioning for routability; the number of IOs ($\overline{io_p}$) has some bearing on the pins-per-net ratio (hence the routability) of the partitioned subcircuits. Equation (8) indicates that $\overline{\beta}_p$ decreases with increasing $\overline{io_p}$, this suggests that a routability-driven partitioner should maximize the IO utilization, whenever it is possible. Intuitively, a partitioner decomposes the nets of a circuit into inter-partition net and intra-partition nets. Equation (4) expresses the law of ‘‘conservation of nets.’’ Large cut sizes imply smaller fanout inside the FPGA devices, making the subcircuits more routable.

Figure 2 illustrates the difficult task of a routability-driven partitioner. There are tradeoffs between routability and the cost of devices to implement the partitioned subcircuits. Routability generally increases with reduced utilization, whereas cost decreases with utilization. The art of routability-driven partitioning is to seek the fine balance between routability and cost. It is not obvious that mincut-based algorithms can handle this seemingly paradoxical set of constraints. Typically, mincut-based partitioning algorithms minimize inter-partition cuts [23]. We have devised a spectral-based partitioner that maximizes the IO utilization while minimizing the number of FPGA devices. The partitioner will be discussed in the next section.

III Spectral k -way partitioning

We provide some background to understand spectral-based partitioning. We shall also discuss the subtlety involved in choosing the proper graph matrix (Laplacian or the adjacency matrix) in order to handle both partition-size and cut-size constraints.

III-A Formulation

An instance of the graph partitioning problem consists of a graph, $G = (\mathcal{V}, \mathcal{E})$ with vertices, $\mathcal{V} = \{\nu_1, \nu_2, \dots, \nu_n\}$, and weighted edges where the weight of edge $e = (\nu_i, \nu_j)$, represents the cost of putting ν_i and ν_j in separate partitions. The problem is to find a partition of the set of vertices $\mathcal{P} = \{P_1, P_2, P_3, \dots, P_k\}$ for a given k , which

optimizes some cost criterion based on the weights of the edges cut and/or the sizes of the partitions. The partitioning results can be represented by an $n \times k$ *ratioed assignment matrix* $R = [r_{ih}]$ where

$$r_{ih} = \begin{cases} \frac{1}{\sqrt{|P_h|}} & \text{if } \nu_i \in P_h \\ 0 & \text{if } \nu_i \notin P_h \end{cases} \quad (10)$$

The rows do not necessarily sum to 1, and column h sums to $\sqrt{|P_h|}$. The related $n \times n$ *ratioed partition matrix* $P_R = [rp_{ij}]$ is defined by

$$rp_{ij} = \begin{cases} \frac{1}{|P_g|} & \text{if } \nu_i \text{ and } \nu_j \text{ both belong to } P_g \\ 0 & \text{otherwise} \end{cases}$$

For a given partition $\mathcal{P} = \{P_1, P_2, \dots, P_k\}$ we have that $P_R = RR^T$.

The *adjacency matrix* of G is the $n \times n$ matrix $A(G) = [a_{ij}]$ where a_{ij} is the weight of the edge between nodes ν_i and ν_j . The *degree matrix* of G is the $n \times n$ matrix $D(G) = [d_{ij}]$ defined by,

$$d_{ij} = \begin{cases} \sum_{k=1}^n a_{ik} & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

The *Laplacian* of G is the $n \times n$ symmetric matrix $Q(G) = D(G) - A(G)$.

Spectral-based partitioning methods extract global information about the structure of a graph from the eigenvalues/eigenvectors of graph matrices. The relation between the properties of a graph and its *spectrum* (the eigenvalues/eigenvectors of its associated matrices) has been an area of active research for several years [13, 8, 16, 20]. The spectra of the adjacency matrix A or the Laplacian Q of a graph are the basis for both partitioning [3, 9, 12] and placement techniques [11].

III-B Spectral ratio-cut partitioning using the Laplacian Q

Spectral partitioning forms clusters of vertices based on the embedding implied by the eigenvectors V of a graph matrix, which can be the Laplacian Q , or the adjacency matrix A of the graph.

To minimize the ratio-cut cost metric, researchers used the Laplacian Q to obtain an embedding [18, 2] of the eigenvectors. This is the correct approach since spectral embedding of the Laplacian appears to be closely related to the ratio-cut cost metric.

Let Π be the set of all k -way partitions of a graph G , and E_h be the total weight of the edges in G having exactly one endpoint in partition P_h . In [6] the authors show that,

$$\min_{X^T X = I} \text{trace}(X^T Q X) \leq \min_{\mathcal{P} \in \Pi} \sum_{h=1}^k \frac{E_h}{|P_h|} \quad (11)$$

which is a lower bound on k -way ratio-cut cost metric. The first k eigenvectors V_k of the Laplacian Q satisfy this inequality and the eigenvectors can be used to form the projector, VV^T , as an approximation of the ratioed partition matrix P_R .

Note that minimization of the ratio-cut cost metric implies that partition size and partition cut constraints *are not* imposed simultaneously on the partitions. This is a consequence of the definition of the ratio-cut cost metric.

III-C Spectral partitioning using the adjacency matrix A

The adjacency matrix $A(G)$ of a graph G is a less popular choice of the matrix to obtain an embedding. Based on Donath and Hoffman's result [9], Rendl and Wolkowicz [20] derived an upper bound on the weight of the edges uncut (E_{uncut}) by a partition satisfying pre-determined partition sizes. If $m_1 \geq m_2 \geq \dots \geq m_k$ are the given partition sizes, $M = \text{diag}(m_1, \dots, m_k)$, and $\lambda_{n-k+1} \leq \lambda_{n-k+2} \leq \dots \leq \lambda_n$ are the largest k eigenvalues of the adjacency matrix, then

$$\frac{1}{2} \sum_{i=n-k+1}^n m_i \lambda_i(A) = \max_{X^T X = I} \left\{ \frac{1}{2} \text{trace}(M X^T A X) \right\} \geq E_{uncut} \quad (12)$$




Figure 3: (a) Partitions induced by the Laplacian Q ; (b) Partitions induced by the adjacency matrix A .

where m_i is the partition size of partition i . If all the partitions have the same size, $m_1 = m_2 = \dots = m_k = m$, then Equation (12) can be simplified to:

$$\frac{1}{2} \sum_{i=n-k+1}^n \lambda_i(A) = \max_{X^T X=I} \left\{ \frac{1}{2} \text{trace}(X^T A X) \right\} \geq \frac{E_{uncut}}{m} \quad (13)$$

which is an upper bound on the number of edges uncut. The last k eigenvectors V_k of the adjacency matrix A satisfy this inequality and the eigenvectors can be used to form the projector $V V^T$. It is an approximation of the ratioed partition matrix P_R . Consequently the partitions will be equal-sized, roughly speaking.

We shall use a small example to illustrate the differences in partitioning using A and Q . Figure 3 shows a circuit with 3 natural “partitions,” which are unbalanced. We want to partition this circuit into three subcircuits. The partitions induced by using the first 3 eigenvectors of the Laplacian Q are shown in Fig. 3(a). It clearly shows 3 unbalanced partitions which reflect the nature of Q to minimize the ratio-cut cost metric. On the other hand, the partitions induced by using the last three eigenvectors of the adjacency matrix A is shown in Fig. 3(b), reflecting the nature of A to find “balanced” partitions. Hence, the adjacency matrix A is a better choice than the Laplacian Q to partition a circuit for a homogeneous FPGA system.

Equation (8) states that the pins-per-net ratio decreases with increasing number of I/O pins used. So, it is meaningful for a partitioner to find suboptimal cuts for routability purposes. First, a spectral-based partitioner has the inherent property that it can generate a wide selection of cuts. It is easy to see why. For example, from Equation (11), if we exclude the first eigenvector v_1 from the minimization, then the lower bound will be increased:

$$\min_{\substack{X^T X=I \\ X^T v_1=0}} \text{trace}(X^T Q X) = \sum_{h=2}^{k+1} \lambda_h(Q) \geq \sum_{h=1}^k \lambda_h(Q). \quad (14)$$

This has the implication that more cuts (consequently more I/O pins) would be obtained by using intermediate eigenvectors to induce the partitions (Similar argument holds for the adjacency matrix). Second, some net models used in transforming hypergraphs into graphs are known to produce more cuts than the others, we may use those net models to increase the cuts. As the last resort, we can use a different graph matrix. For example, use the Laplacian in place of the adjacency matrix.

III-D From Spectral Embeddings to Partitions

Given the logic capacity and I/O capacity constraints, we present a procedure to use the eigenvector embedding of a graph matrix to construct partitions that satisfy the constraints. As in [6], our approach to k -way

```

KPF_Partition(I/O constraint, logic constraint) {
  Remove high-fanout nets in the hypergraph and transform it to a graph
  Find  $k$  eigenvectors of this graph,  $V$ 

  Associate each row of  $V$  with its vertex in the original hypergraph
  Mark all vertices unallocated
   $h=0$ ;
  Select vertex with largest-magnitude vector as first prototype

  while (there are unallocated vertices) {
     $h++$ ;
    if  $h \neq 1$  select vertex most orthogonal to previous prototype
      and use as it as the prototype for partition  $h$ 
    Build heap based on the ranking function, Eq. (15)
    while partition  $h$  has not reached logic capacity {
      remove largest cost vertex,  $\nu$ , from heap
      assign  $\nu$  to partition  $h$ 
      update cut tally for each net to which  $\nu$  belongs
      update heap cost for all vertices adjacent to  $\nu$ 
      put  $\nu$  on Rollback list if I/O capacity is exceeded
      reset Rollback list if I/O capacity is not exceeded
    }
    Remove vertices on Rollback list from  $h$  and mark them unallocated
  }
  Print partition results
}

```

Figure 4: Pseudo-code for spectral **KPF** partitioning method. Both I/O and logic capacity constraints are provided by the user.




Figure 5: Ranking heuristic based on Cuts.

partitioning is to “reverse engineer” the partitions from the embeddings implied by the last eigenvectors of A , $V = [v_{n-k+1}, \dots, v_n]$ (or the intermediate eigenvectors $V = [v_{n-k+1-q}, \dots, v_{n-q+1}]$, for some $q \leq n$ if so desired).¹

Given the eigenvectors, we measure the cosine of the angle between the two row vectors i and j of V (or the column vectors of V^T). These *directional cosines* provide a measure of the proximity of the vertices relative to each other. This strategy, in effect, identifies all the ‘1’s in the partitioning matrix P implied by the projector VV^T .

The first step is to determine a vertex of largest magnitude to serve as the *prototype* (seed) for the first partition. The first prototype is selected by magnitude, and the rest are determined by their relative (anti)orthogonality with respect to the existing prototypes.

A vertex is said to be *allocated* if it has been assigned and committed to a partition. After the determination of the prototype, the rest of the unallocated vertices are ranked based on their anti-orthogonality with respect to the prototype. The ranked vertices are sorted in a heap data structure, which is updated as vertices are extracted from the heap.

An additional factor to be considered in the ranking of the vertices is the number of *cuts* incurred by including a vertex in a partition. The ranking function for the vertices is:

$$\text{rank} = \text{anti_orthogonality_wrt_prototype} - f \times \text{cuts_incurred} \quad (15)$$

¹ We use the first k eigenvectors of the Laplacian graph matrix, and the last k eigenvectors of the adjacency matrix.

Circuit	Kuznar <i>et al's</i> results	Kuznar % CLB Util	KPF results # of chips	CLB Util %	IOB Util %	elapsed time in sec.
c3540	{0,0,3,0,0}	66	{0,0,3,0,0}	66	99	38
c5315	{2,1,2,0,0}	73	{0,0,4,0,0}	65	94	64
c6288	{0,0,4,2,0}	81	{0,0,0,4,0}	93	77	84
c7552	{0,0,4,0,0}	85	{0,0,4,0,0}	85	99	83
s5378	{0,0,1,0,1}	82	{0,0,0,0,2}	60	72	38
s9234	{0,0,0,1,1}	77	{0,0,0,0,2}	71	69	69
s13207	{3,5,4,0,0}	72	{0,0,11,0,0}	58	98	202
s15850	{0,0,2,2,1}	83	{0,0,0,0,4}	66	91	171
s38584	{0,5,15,4,1}	75	{0,0,0,16,0}	81	95	1710

Circuit	Chou (SC) <i>et al's</i> results	KPF results
s13207	{0,0,0,0,6}	{0,0,0,0,6}(routable)
s15850	{0,0,0,0,3}	{0,0,0,0,3} (unroutable); {0,0,0,0,4} (marginal)
s38417	{0,0,0,0,10}	{0,0,0,0,8} (unroutable); {0,0,0,0,9} (marginal)
s38584	{0,0,0,0,14}	{0,0,0,0,11}(unroutable); {0,0,0,0,12}(marginal)

Table 1: Summary of **KPF** (k -way partitioning for homogeneous FPGA systems). Partitioning results using a SUN IPC workstation with 32 Mbyte of memory. The number of chips generated by the partitioners is presented as {# of XC3020, # of XC3030, # of XC3042, # of XC3064, # of XC3090} Xilinx XC3000 FPGAs. Kuznar *et al's* result [15] is used as a reference, their partitioner is targeted for heterogeneous FPGA systems. Chou *et al's* partitioner (SC) [7] is targeted for homogeneous FPGA hardware emulators.

where f is a weight parameter associated with the heuristics. We use $0.5 < f < 2.5$; a typical value is 1. We have two different strategies. We call them the *pcut* and *pcluster* heuristics, respectively. Figure 5 illustrates a vertex (in solid) which, with the *pcut* heuristics, is considered to have saved $4 - 3 = 1$ cut if incorporated into the current partition; whereas with the *pcluster* heuristics, it is considered to have *saved* 4 cuts (or incur -4 cuts).

The logic capacity and I/O constraints are satisfied by using the procedure **KPF** as outlined in Fig. 4. In essence, the partitions are built one at a time. A prototype is the first entry of a partition. One at a time, vertices are extracted from the heap and tentatively assigned to the current partition. If the I/O capacity is exceeded, then the vertex's ID is entered into a rollback list. The rollback list is reset if the number of I/Os drops within the bound. Upon each extraction of a vertex from the heap, the heap is (incrementally) updated according to the cost function in Equation (15). The extraction is repeated until the logic capacity constraint is exceeded. Rollback commences if there are entries in the rollback list, it rolls back to the last point where the I/O constraint is met. Then all vertices in the feasible partition are marked "allocated."

IV Experimental results

IV-A KPF K -way partitioning for FPGAs

We implement our k -way partitioning algorithm with both clustering heuristics *pcuts* and *pcluster*, and the best results of the two heuristics are presented. We refer to our k -way spectral-based partitioning algorithm as **KPF**.

We ran the graphs derived from the MCNC FPGA partitioning 1993 benchmarks for the Xilinx XC3000 series. The hypergraphs of the benchmarks were transformed into graphs by using Frankle's clique expansion net model [11], or the star graph net model. Unlike the clique net model, the star graph net model produces a very sparse graph matrix, which accelerates the eigensolver. On the other hand, this hypergraph model generates auxiliary vertices that have to be filtered out from the eigenvector embedding. We have observed that the star graph net model tends to produce more partitions than the clique net model. Unless otherwise stated, the partitioning results reporting are produced by Frankle's clique net model.

Also, the hypergraphs of the benchmarks were pre-processed to remove high-fanout nets of degree greater than 99. Nets whose degree were greater than 99 were removed in order to reduce storage space and processing time. This step is essential in facilitating the eigensolver (Scott/Parlett implementation of the Lanczos algorithm [17]) to run on a low-end SPARC station with only 32 Mbyte of memory. The high-fanout nets are

Circuit	pins/club predicted by eqn (2)	pins/net predicted by eqn (8)	pins/club actual average	pins/net actual average	CLB % util	IOB % util	# of partitions	# of unrouted nets in subcircuits		pre-partition prediction
								Average	Minimum	
c3540	5.56	2.94	5.58	3.05	66	99	3	0,0,0	0,0,0	✓
c5315	5.59	3.29	5.68	3.15	65	94	4	0,0,0,0	0,0,0,0	✓
c6288	4.53	2.46	4.53	2.50	93	77	4	0,0,0,0	0,0,0,0	✓
c7552	5.36	3.10	5.35	3.05	85	99	4	0,0,0,0	0,0,0,0	✓
s5378	5.24	3.13	5.23	3.19	60	72	2	0,4.55	0,0	M
s9234	5.11	3.11	5.05	3.15	71	69	2	0,0.5	0,0	M
s13207	4.70	2.68	4.83	2.79	58	98	11	0,0,...,0	0,0,...,0	✓
s15850	4.87	2.93	4.90	3.07	66	91	4	2.8,0,0,0	0,0,0,0	M
s38584	5.08	3.27	5.15	3.38	81	95	16	0,0,...,0	0,0,...,0	✓
alu4	6.45	3.20	6.45	3.18	58	80	2	3,15	1,11	M
misex3	6.49	3.36	6.49	3.30	90	89	2	0,10.55	0,6	M

Table 2: Routability results of the subcircuits produced by the spectral k -way partitioner **KFP** for homogeneous FPGA systems using the clique hypergraph model. Reporting the predicted and measured (average) pins/net ratios of the benchmark circuits; also reporting are the prediction of the pre-partitioning routability predictor (✓=routable, M=Marginal, U=Unroutable). Subcircuits with zero average unrouted nets indicate successful *first-time* placement and routing completion. Nonzero average unrouted nets are the averages of 10 placement and routing runs.

only excluded during the calculation of the eigenvectors. Once the eigenvectors are computed, all nets are considered in determining the partitions.

A user supplies the desired FPGA package type, number of eigenvectors (h) desired, a balancing constraint, the circuit, and the net model desired to the partitioner. The balancing constraint determines the ratio of the partition sizes of the last two partitions. The partitioner transforms the circuit to a graph (adjacency) matrix according to the net model. The Lanczos algorithm is used to find the last h eigenvectors; the partitioner then generates the partitions. The number of partitions might be different from the number of eigenvectors specified.

The partitioned subcircuits are translated to the Xilinx **map** format. The routability predictor **kop** was applied to all the subcircuits. All the subcircuits are subsequently placed and routed using the vendor's tool **apr** (version XACT 5.0) with the default options. Routability results are presented in the next section.

Overall, the results in Table 1 demonstrate that our spectral-based graph partitioner **KPF** is capable of handling hard constraints and producing good partitions in reasonable time. We used the results from [15, 7] as references. The partitioner of Kuznar *et al* [15] is targeted for heterogeneous systems, while our partitioner **KPF** is targeted for homogeneous systems. We only list our homogeneous partitioning results in Table 1. However, our partitioner is also capable of generating very competitive heterogeneous results. For example, in [15] the authors use a monetary cost function, their partitioner generates a device distribution of

$$\{0, 0, 1, 0, 1\}$$

meaning one XC3042 and one XC3090 FPGA to implement the benchmark circuit **s5378** (routability results are not reported in [15]). But without balancing constraints, our partitioner generates a better device distribution

$$\{1, 0, 0, 0, 1\}$$

meaning one XC3020 and one XC3090 FPGAs to implement the circuit, and the average CLB utilization is 99.6%. Also, our result is less expensive than [15] with a monetary cost function. But the vendor's placement and routing tool **apr** [1] cannot complete the routing of this XC3090 subcircuit, the tool reports 50 unrouted nets! This demonstrates that partitioning without consideration for routability could be an exercise of futility.

Even though several parallel processing researchers reported that in terms of cuts, spectral-based partitioning results can be further improved by applying mincut-based partitioning algorithm such as Kernighan and Lin; we do not resort to this post-processing step for routability considerations.

The elapsed times (not CPU time) reported in Table 1 do not include the time taken by an **awk** script to translate the partitioning results generated by **KPF** into the Xilinx **map** format.

IV-B Routability and routability prediction results

The routability of all the subcircuits are predicted before they are placed and routed, and the accuracy of the pre-partitioning routability predictor is reported in the last column of Table 2.

Circuit	net model	number of partitions	eigenvectors	# of CLBs in subcircuits	# of IOBs in subcircuits	avg. unrouted nets in subcircuits
s15850	clique	—	last 2	—	—	—
s15850	clique	3	last 3	300,299,243	144,143,142	10.3,2.5,0
s15850	clique	4	last 4	303,214,150,175	144,144,96,122	7.7,0,0,0
s15850	clique	4	last 2nd to 4th	294,191,230,127	144,144,144,92	1.4,0,0,0
s15850	clique	4	last 3rd to 5th	249,165,197,231	144,144,137,122	0,0,0,0
s15850	star	5	last 3	203,193,127,116,203	143,144,143,143,141	0,0,0,0,0

Table 3: Partitioning and routability results of subcircuits of benchmark circuit **s15850** generated by using different net models and eigenvectors of the graph adjacency matrix.

Circuit	package	net model	subcircuit1		subcircuit2		avg. unrouted nets (subcircuit1,subcircuit2)
			IOBs	CLBs	IOBs	CLBs	
alu4	3042PG132	clique	73	72	81	95	(3.0, 15.0)
alu4	3042PG132	star	90	77	96	90	(0.6, 11.9)
misex3	3020PC84	clique	56	51	57	64	(0.0, 10.6)
misex3	3020PC84	star	64	54	64	61	(0.0, 6.2)

Table 4: Routability results of subcircuits of **alu4** and **misex3** after 10 placement and routing runs, showing the effect of increased cuts on routability. The FPGA package XC3042PG132 has a maximum of 144 logic blocks (CLBs) and 96 I/O blocks (IOBs), and the XC3020PC84 package has a maximum of 64 CLBs and 64 IOBs.

As we see from the second column in Table 2, the pins-to-cell ratios of the MCNC benchmark circuits are relatively low. All subcircuits predicted to be “routable” automatically completed the routing the first time by **apr**. On the other hand, subcircuits predicted to be “marginal” didn’t complete the routing the first time by **apr**. Ten placement and routing runs were applied to those subcircuits and the average number of unrouted nets are recorded.

The majority of the subcircuits are predicted to be routable and this was later verified, as given in Table 2. Of particular interest is that circuit **s15850** with 3 partitions was predicted to be *unroutable*, so we increased the number of partitions to 4, and the predictor predicts marginal routability. The prediction was verified to be correct, as depicted in the first five entries, last column of Table 3. This table also illustrates that spectral-based partitioners can generate a wide variety of cuts, and hence partitioning results. For instance, one of the 4 subcircuits generated by using the last 4 eigenvectors (of the adjacency matrix) has 7.7 unrouted (average) nets after 10 **apr** runs, and this subcircuit has higher cell utilization than the others. Instead, by using the 2nd to 4th eigenvectors our partitioner generates slightly more balanced and routable subcircuits, as indicated in the 4th entry of Table 3. Last, the star net model generates 5 very routable subcircuits, but the utilizations of the subcircuits are low.

We also include two hard-to-route circuits, **alu4** and **misex3** from the MCNC combinational circuit benchmark suite to test our partitioner and routability predictor. Initially, the circuit **alu4** uses two XC3042PG132’s and circuit **misex3** uses two XC3020PC84’s to implement the subcircuits, respectively. The routability predictor predicts the subcircuits to be marginally routable, and this was verified, as shown in Table 2.

We repartition both circuits using the star net model, the results are shown in Table 4. The number of cuts (IOBs) produced by using the star net model is higher than the clique net model. The subcircuits are more “routable” with increased IOB usage, as indicated by the drop in the average number of unroutable nets in Table 4.

We also applied the pre-partitioning routability predictor to check to see if the circuit **misex3** would be routed using a different package type, say XC3030PQ100, which has 100 CLBs and 80 IOBs. The predictor predicts positive results and this was verified to be the case.

V Conclusion

In this paper, we have presented partitioning results of a spectral-based partitioner targeted for homogeneous FPGA systems. The partitioner handles both partition-size and cut-size constraints. Routability of the partitioned subcircuits is the primary concern. To assist the partitioner in assessing the routability of the partitioned subcircuits, we have developed a theory to predict the routability of the partitioned subcircuits prior to partitioning. We feel that routability prediction is very valuable and should play an important role in the partitioning for multiple-FPGA systems.

References

- [1] XILINX: *The Programmable Gate Array Data Book*. 2100 Logic Drive, San Jose, CA 95124, 1993.
- [2] C. J. Alpert and A. B. Kahng. Multi-way partitioning via spacefilling curves and dynamic programming. In *ACM IEEE 31st Design Automation Conference Proceedings*, pages 652–657, San Diego, California, June 1994.
- [3] E. R. Barnes. An algorithm for partitioning the nodes of a graph. *SIAM Journal on Algorithm and Discrete Method*, 3(4):541–550, Dec. 1982.
- [4] N. Bhat and D. Hill. Routable Technology Mapping for LUT FPGAs. In *1992 IEEE International Conference on Computer Design: VLSI in Computers and Processors*, pages 95–98, October 1992.
- [5] P. K. Chan, M. Schlag, and J. Zien. On routability prediction for field-programmable gate arrays. In *ACM IEEE 30th Design Automation Conference Proceedings*, pages 326–330, Dallas, Texas, June 1993.
- [6] P. K. Chan, M. Schlag, and J. Zien. Spectral k -way ratio-cut partitioning and clustering. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 13(9):1088–1096, Sept. 1994.
- [7] N.-C. Chou, L.-T. Liu, C.-K. Cheung, W.-J. Dai, and R. Lindelof. Circuit partitioning for huge logic emulation systems. In *ACM IEEE 31st Design Automation Conference Proceedings*, pages 244–249, San Diego, California, June 1994.
- [8] D. M. Cvetkovic, M. Doob, and H. Sachs. *Spectra of Graphs: Theory and Application*. Academic Press, Inc., New York, New York, 1979.
- [9] W. Donath and A. Hoffman. Lower bounds for the partitioning of graphs. *IBM Journal of Research and Development*, pages 420–425, 1973.
- [10] A. El Gamal. Two-Dimensional Stochastic Model for Interconnections in Master Slice Integrated Circuits. *IEEE Transactions on Circuits and Systems*, 28(2):127–138, Feb. 1981.
- [11] J. Frankle and R. M. Karp. Circuit placements and cost bounds by eigenvector decomposition. In *IEEE International Conference on Computer-Aided Design ICCAD-86*, pages 414–417, Santa Clara, California, November 1986.
- [12] S. W. Hadley, B. L. Mark, and A. Vanelli. An Efficient Eigenvector Approach for Finding Netlist Partitions. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, CAD-11(7):885–892, July 1992.
- [13] K. M. Hall. An r -dimensional quadratic placement algorithm. *Management Science*, 17(3):219–229, Nov. 1970.
- [14] IDA Supercomputing Research Center, 17100 Science Drive, Bowie, MD 20715. *Splash2: Programmer's Manual Version 0.7*, 1992.
- [15] R. Kuznar, F. Brglez, and K. Kozminski. Cost minimization of partition into multiple devices. In *ACM IEEE 30th Design Automation Conference Proceedings*, pages 315–320, Dallas, Texas, June 1993.
- [16] B. Mohar and S. Poljak. *Eigenvalues in Combinatorial Optimization*. University of Ljubljana, Jadranska 19, 61 111 Ljubljana, Slovenia, 1993.
- [17] B. N. Parlett and D. S. Scott. The Lanczos algorithm with selective orthogonalization. *Mathematics and Computations*, 33(11):217–238, 1979.
- [18] A. Pothen, H. D. Simon, and K.-P. Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM Journal of Matrix Analysis and Applications*, 11(3):430–452, 1990.
- [19] Quickturn Design Systems, Inc. *Quickturn Emulation System*. 325 East Middlefield Road, Mountain View, CA 94043, 1993.
- [20] F. Rendl and H. Wolkowicz. A projection technique for partitioning the nodes of a graph. Technical report, University of Waterloo, Apr. 1991.
- [21] M. Schlag, J. Kong, and P. K. Chan. Routability-driven technology mapping for lookup table-based FPGAs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, (1):13–26, Jan. 1994.
- [22] S. Trimberger and M.-R. Chene. Placement-Based Partitioning for Lookup-Table-Based FPGAs. In *1992 IEEE International Conference on Computer Design: VLSI in Computers and Processors*, pages 91–94, October 1992.
- [23] N.-S. Woo and J. Kim. An efficient method of partitioning circuits for multiple-FPGA implementation. In *ACM IEEE 30th Design Automation Conference Proceedings*, pages 202–207, Dallas, Texas, June 1993.