

# Design and Implementation of Move-Based Heuristics for VLSI Hypergraph Partitioning

Andrew E. Caldwell, Andrew B. Kahng and Igor L. Markov  
{caldwell,abk,imarkov}@cs.ucla.edu

---

Supported by a grant from Cadence Design Systems, Inc.

---

We summarize the techniques of implementing move-based hypergraph partitioning heuristics and evaluating their performance in the context of VLSI design applications. Our first contribution is a detailed software architecture, consisting of seven reusable components, that allows flexible, efficient and accurate assessment of the practical implications of new move-based algorithms and partitioning formulations. Our second contribution is an assessment of the modern context for hypergraph partitioning research for VLSI design applications. In particular, we discuss the current level of sophistication in implementation know-how and experimental evaluation, and we note how requirements for real-world partitioners – if used as motivation for research – should affect the evaluation of prospective contributions. Two “implicit decisions” in the implementation of the Fiduccia-Mattheyses heuristic are used to illustrate the difficulty of achieving meaningful experimental evaluation of new algorithmic ideas.

Categories and Subject Descriptors: D.2.11 [**Software Engineering**]: Software Engineering—*Software architecture*; J.6 [**Computer Applications**]: Computer-Aided Engineering

General Terms: Algorithms, Experimentation, Measurement, Performance

Additional Key Words and Phrases: balanced min-cut hypergraph partitioning, VLSI CAD

---

## 1. INTRODUCTION: HYPERGRAPH PARTITIONING IN VLSI DESIGN

Given a hyperedge- and vertex-weighted hypergraph  $H = (V, E)$ , a *k-way partitioning* of  $V$  assigns the vertices to  $k$  disjoint nonempty partitions. The *k-way partitioning problem* seeks to minimize a given cost function  $c(P^k)$  whose argu-

---

Affiliation: UCLA Computer Science Dept.  
Address: Boelter Hall, Los Angeles, CA 90095-1596

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works, requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept, ACM Inc., 1515 Broadway, New York, NY 10036 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

ments are partitionings. A standard cost function is *net cut*,<sup>1</sup> which is the sum of weights of hyperedges that are cut by the partitioning (a hyperedge is *cut* exactly when not all of its vertices are in one partition).

Constraints are typically imposed on the partitioning solution, and make the problem difficult. For example, certain vertices can be fixed in particular partitions (*fixed constraints*). Or, the total vertex weight in each partition may be limited (*balance constraints*), which results in an NP-hard formulation [21]. Thus, the cost function  $c(P^k)$  is minimized over the set of *feasible solutions*  $S_f$ , which is a subset of the set of all possible  $k$ -way partitionings. Effective move-based heuristics for  $k$ -way hypergraph partitioning have been pioneered by Kernighan and Lin [37], Fiduccia and Mattheyses [20] and Bui et al. [9], with refinements given by Krishnamurthy [39], Sanchis [43], Hagen and Kahng [27], Liu et al. [41], Dutt and Deng [18], Alpert, Hagen and Kahng [4], Cong et al. [12], Hauck and Borriello [26], Karypis et al. [35], Dutt and Theny [19] as well as many others. A comprehensive survey of partitioning formulations and algorithms, centered on VLSI applications and covering move-based, spectral, flow-based, mathematical programming-based, etc. approaches, is given in a comprehensive survey [5] by Alpert and Kahng. A recent update on balanced partitioning in VLSI physical design is provided by a short survey [32] by Kahng.

### 1.1 The VLSI Design Context

VLSI design has long provided driving applications and ideas for hypergraph partitioning heuristics. For example, the methods of Kernighan-Lin [37] and Fiduccia-Mattheyses [20] form the basis of today’s move-based approaches. The method of Goldberg-Burstein [25] presaged the multilevel approaches recently popularized in the parallel simulation [23; 28; 33] and VLSI [3; 35; 4] communities. As Alpert and Kahng notice in their survey [5], applications in VLSI design include test, simulation and emulation, design of systems with multiple field-programmable devices, technology migration and repackaging, and top-down floorplanning and placement.

Depending on the specific VLSI design application, a partitioning instance may have directed or undirected hyperedges, weighted or unweighted vertices, etc. However, in all contexts the instance represents – at the transistor-level, gate-level, cell-level, block-level, chip-level, or behavioral description module level – a human-designed system. Such instances are highly non-random. Many efforts (e.g., [31], [14], [24], [8]) have used statistical attributes of real-world circuit hypergraphs (based on Rent’s parameter [40], shape, depth, fanout distribution, etc.) to generate random hypergraphs believed relevant to evaluation of heuristics. These efforts have not yet met with wide acceptance in the VLSI community, mostly because generated instances do not guarantee “realism”. Hence, the current practice remains to evaluate new algorithmic ideas against suites of benchmark instances.

In the VLSI partitioning community, performance of algorithms is typically evaluated on the ACM/SIGDA benchmarks now maintained by the Collaborative Benchmarking Laboratory at North Carolina State University <http://www.cb1.ncsu.edu/benchmarks>.<sup>2</sup>

<sup>1</sup>Or simply *cut*, as in *minimum cut partitioning*. Note that in the VLSI context, a circuit hypergraph is called a *netlist*; a hyperedge corresponds to a *signal net*, or *net*; and a vertex corresponds to a *module*.

<sup>2</sup>These benchmarks are typically released by industry or academic designers at various workshops and conferences (e.g., LayoutSynth90, LayoutSynth92, Partitioning93, PDWorkshop93, ...).

Alpert [2] noted that many of these circuits no longer reflect the complexity of modern partitioning instances, particularly in VLSI physical design; this motivated the release of eighteen larger benchmarks produced from internal designs at IBM [1].<sup>3</sup>

Salient features of benchmark (real-world) circuit hypergraphs include

- size: number of vertices can be up to one million or more (instances of all sizes are equally important).
- sparsity: average vertex degrees are typically between 3 and 5 for cell-, gate- and device-level instances; higher average vertex degrees occur in block-level design.
- number of hyperedges (nets) typically between 0.8x and 1.5x of the number of vertices (each module typically has only one or two outputs, each of which represents the source of a new signal net).
- average net sizes are typically between 3 to 5.
- a small number of very large nets (e.g., clock, reset, test) connect hundreds or thousands of vertices.

Statistics for circuit benchmarks used in our work are given in Tables 4, 5 and 6.

Partitioning heuristics must be highly efficient in order to be useful in VLSI design.<sup>4</sup> As a result – and also because of their flexibility in addressing variant objective functions – fast and high-quality iterative move-based partitioners based on the approach of Fiduccia-Mattheyses [20] have dominated recent practice.

## 1.2 The Fiduccia-Mattheyses Approach

The Fiduccia-Mattheyses (FM) heuristic for bipartitioning circuit hypergraphs [20] is an iterative improvement algorithm. Its neighborhood structure is induced by single-vertex, partition-to-partition moves.<sup>5</sup> FM starts with a possibly random solution and changes the solution by a sequence of moves which are organized as *passes*. At the beginning of a pass, all vertices are free to move (*unlocked*), and each possible move is labeled with the immediate change in total cost it would cause; this is called the *gain* of the move — positive gains reduce solution cost, while negative gains increase it (see example in Figure 1). Iteratively, a move with highest (but not necessarily positive) gain is selected and executed, which may result in a lower, higher or unchanged solution cost. The moved vertex is *locked*, i.e., is not allowed to move again during that pass. Since moving a vertex can change gains of adjacent vertices, all affected gains are updated. Selection and execution of a best-gain move, followed by gain update, are repeated until every

<sup>3</sup>While those benchmarks are now used in most partitioning papers, we would like to stress that they present considerably harder partitioning problems than earlier available benchmarks available from <http://www.cb1.ncsu.edu/benchmarks>, primarily due to more varied distributions of node degrees and weights. See, e.g., Tables 4, 5 and 6.

<sup>4</sup>For example, a modern top-down standard-cell placement tool might perform timing- and routing congestion-driven recursive min-cut bisection of a cell-level netlist to obtain a “coarse placement”, which is then refined into a “detailed placement” by stochastic hill-climbing search. The *entire* placement process in currently released tools from companies like Avant!, Cadence, CLK CAD, Gambit, etc. takes approximately 1 CPU minute per 6000 cells on a 300MHz Sun Ultra-2 uniprocessor workstation with adequate RAM. The implied partitioning runtimes are on the order of 1 CPU second for netlists of size 25,000 cells, and 30 CPU seconds for netlists of size 750,000 cells [16]. Of course, we do not advocate performance tuning to match industrial-strength runtimes. However, absent other justifications, “experimental validation” of heuristics in the wrong runtime regimes (say, hundreds of CPU seconds for a 5000-cell benchmark) has no practical relevance.

<sup>5</sup>By contrast, the stronger Kernighan-Lin (KL) heuristic [37] uses a pair-swap neighborhood structure.

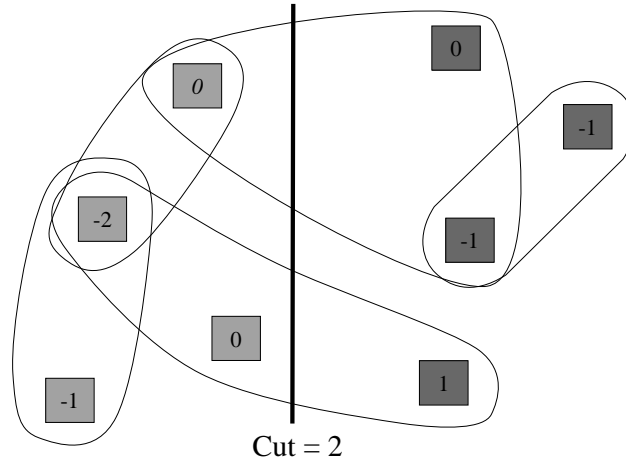


Fig. 1. Gains in hyperedge cut associated with each potential vertex move.

vertex is locked. After every vertex has moved exactly once, we get a partitioning solution symmetric to the one at the beginning of the pass, thus having the same cost (see Figure 2). Then, the best solution seen during the pass is adopted as the starting solution of the next pass. The algorithm terminates when a pass fails to improve solution quality.



Fig. 2. Characteristic profile of solution cost change during one FM pass. The lowest-cost solution seen during the pass (best prefix of moves) is highlighted; this is adopted as the initial solution for the next pass.

The FM algorithm can be easily seen to have three main operations: (1) the computation of initial gain values at the beginning of a pass; (2) the retrieval of the best-gain (feasible) move; and (3) the update of all affected gain values after a move is made. The contribution of Fiduccia and Mattheyses lies in observing

that circuit hypergraphs are sparse, so that any move gain is bounded between two and negative two times the maximal vertex degree in the hypergraph (times the maximal edge weight, if edge weights are used). This allows hashing of moves by their gains: all affected gains can be updated in linear time, yielding overall linear complexity per pass. In [20], all moves with the same gain are stored in a linked list representing a “gain bucket”.

### 1.3 Contributions of This Paper

In this paper, we discuss the implementation and evaluation of move-based hypergraph partitioning heuristics, notably the FM heuristic, in the context of VLSI design applications. Our first contribution is a detailed software architecture, consisting of seven reusable components, that allows flexible, efficient and accurate assessment of the practical implications of new move-based algorithms and partitioning formulations. Our second contribution is an assessment of the modern context for hypergraph partitioning research for VLSI design applications. In particular, we discuss the current level of sophistication in implementation know-how and experimental evaluation, and we note how requirements for real-world partitioners – if used as motivation for research – should affect the evaluation of prospective contributions. Two “implicit decisions” in the implementation of the FM heuristic are demonstrated to illustrate the difficulty of achieving meaningful experimental evaluation of new algorithmic ideas. Since our architecture allows one to capture and study such previously unexplored degrees of freedom, we believe that it is conducive to algorithm innovation and leading-edge quality of results.

## 2. ARCHITECTURE OF A MOVE-BASED PARTITIONING TESTBENCH

In this section, we describe a seven-component software architecture for implementation of move-based partitioning heuristics, particularly those based on the FM approach. By way of example, we reword the Fiduccia-Mattheyses algorithm in terms of these seven software components. By carefully dividing responsibilities among components we attempt to provide the implementation flexibility and run-time efficiency that is needed to evaluate the practical impact of new algorithmic ideas and partitioning formulations.

### 2.1 Main Components

**Partitioning Interface.** Formally describes the input and output to partitioners without mentioning internal structure and implementation details. All partitioner implementations then conform to this input/output specification.

**Initial Solution Generator.** Generates partitionings that satisfy given constraints, typically using randomization in the construction.

**Incremental Cost Evaluator.** Evaluates the cost function for a given partitioning and dynamically maintains cost values when the partitioning is changed by applying moves. Updates typically should be performed in constant time.

**Legality Checker.** Verifies whether a partitioning satisfies a given constraint. The Legality Checker is used to determine the legality of a move. Multiple constraints may be handled with multiple legality checkers.

**Gain Container.** A general container for moves, optimized for efficient allocation, retrieval and queueing of available moves by their gains. Moves can be retrieved by, e.g., the index of the vertex being moved, and/or the source or destination partition. A Gain Container supports quick updates of the gain for a move, and fast retrieval of a move with the highest gain. The Gain Container is also independent of the incremental cost evaluator and legality checker; it is populated and otherwise managed by the Move Manager.

**Move Manager.** Responsible for choosing and applying one move at a time. It may rely on a Gain Container to choose the best move, or randomly generate moves. It can undo moves on request. If used in pass-based partitioners, it incrementally computes the change in gains due to a move, and updates the Gain Container.

The Move Manager maintains “status information”, such as the current cost and how each partition is filled. It may be controlled by the caller via parameter updates before every move selection (e.g. a temperature parameter in simulated annealing).

**Pass-Based Partitioner (proper).** Solves “partitioning problems” by applying incrementally improving passes to initial solutions. A pass consists of legal moves, chosen and applied by the move manager. Within a pass, a partitioner can request that the Move Manager *undo* some of the moves, i.e., perform inverse moves. The Pass-Based Partitioner is an implementation of the *Partitioning Interface*.

This modularity allows for separate benchmarking and optimization of most components. It also provides flexibility to use multiple alternative implementations relevant to special cases.<sup>6</sup> A fundamental facility enabling such modularity is a common efficient hypergraph implementation.<sup>7</sup>

## 2.2 Component Descriptions

We now give somewhat more detailed component descriptions, omitting three components for which implementation choices are less critical.

**2.2.1 Incremental Cost Evaluator.** Initialized with a hypergraph, the Incremental Cost Evaluator is responsible for evaluating the cost function for a given partitioning, and incrementally maintaining this value when the partitioning changes (i.e., a

<sup>6</sup>For example, many optimizations for 2-way partitioning from the general  $k$ -way case can be encapsulated in the evaluator. On the other hand, in our experience optimizing a Gain Container for 2-way is barely worth maintaining separate pieces of code.

<sup>7</sup>A generic hypergraph implementation must support I/O, statistics, various traversals and optimization algorithms. However, no such implementation will be optimal for all conceivable uses. We decided to implement our own reusable components based on the Standard Template Library and optimize them for our use models.

Features directly supported by the hypergraph component include memory management options, conversions, I/O, various construction options such as ignoring hyperedges of size less than 2 or bigger than a certain threshold, lazily executed calls for sorting nodes or edges in various orders etc. Many trade-off had to be made, e.g. the hypergraph objects used in critical pieces of code have to be unchangeable after their initial construction so as to allow for very efficient internal data structures.

None of the many existing generic implementations we reviewed was sufficiently malleable to meet our requirements without overwhelming their source code by numerous compiler `#defines` for adapting the code to a given use model. Having complete control over the source code and internal interfaces also allows for maximal code reuse in implementing related functionalities.

vertex is moved). When the cost function is computed as sum of hyperedge costs, those costs should also be maintained and available.

Efficient implementations typically maintain an internal state, e.g. relevant statistics, for each hyperedge. This facilitates efficient constant-time cost updates when single moves are performed. An Evaluator whose values are guaranteed to be from a small (esp. finite) range should be able to exploit this range to enable faster Gain Container implementations (e.g. using buckets *versus* priority queues).

**Interface:**

- Initialize (internal structures) with a hypergraph and a partitioning solution.
- Report current cost (total or of one net) without changing internal state.
- Complete change of internal state (re-initialization) for all vertices and nets.
- Incremental change of internal state (for all nets whose cost is affected or for a given net) due to one elementary move without updating the costs.<sup>8</sup>

**2.2.2 Gain Container.** The Gain Container stores all moves currently available to the partitioner (these may not all be legal at a given time) and prioritizes them by their gains (i.e., the immediate effect each would have on the total cost). A container of move/gain pairs is defined by the interface below, which allows quick updates to each and any move/gain pair after a single move.

A Gain Container should be able to find the move with highest gain quickly, possibly subject to various constraints such as a given source or destination partition, and may provide support for various *tie-breaking schemes* in order to choose the best move among the moves with highest gain. A Gain Container does not initiate gain updates by itself and is not aware of Cost Evaluators, the Move Manager, or how the gains are interpreted. Gain Containers do not need to determine the legality of moves. This makes them reusable for a range of constrained partitioning problems. Faster implementations (e.g. with buckets) may require that the maximal possible gain be known.

**Interface:**

- Add a move to the Container, given the gain.
- Get the gain for a move.
- Set the gain for a move (e.g. to update).
- Remove a move from the Container.
- Find a move of highest gain.
- Invalidate current highest gain move, in order to request the next highest gain move.<sup>9</sup> Typically applied if the current highest gain move appears illegal.
- Invalidate current highest gain bucket to access the next highest gain bucket.

The primary constituents of a Gain Container are a *repository* and *prioritizers*.

**Repository for gain/move pairs** handles allocation and deallocation of move/gain pairs, and supports fast gain lookups given a move.

<sup>8</sup>Changing the state of one net will, in general, make the overall state of the evaluator inconsistent. This can be useful, however, for “what-if” cost lookups when a chain of incremental changes can return to the original state.

<sup>9</sup>Note that this does not remove the move from the Gain Container.

**Prioritizer** finds a move with highest gain. In addition, may be able to choose best-gain moves among moves with certain properties, such as a particular destination or source partition. Updates gains and maintains the queue containing the moves, in particular, is responsible for tie-breaking schemes.

We say that some moves stored in the repository are *prioritized* when they participate in the prioritizer’s data structures. Not prioritizing the moves affecting a given vertex corresponds to “locking” the vertex, as it will never be chosen as the highest-gain move. The standard FM heuristic locks a given cell as soon as it is moved in a pass; however, variant approaches to locking have been proposed [15].

**2.2.3 Move Manager.** A Move Manager handles the problem’s move structure by choosing and applying the best move (typically, the best *legal* move), and incrementally updates the Gain Container that is used to choose the best move. The Move Manager reports relevant “status information” after each move, e.g. current cost and partition balance, which allows the caller to determine the best solution seen during the pass. In order to return to such best solution, the move manager must perform undo operations on request.

**Interface:**

- Choose one move (e.g., the best feasible) and apply it. Ensure all necessary updates (gain container, incremental evaluator).
- Return new “status info”, e.g., total cost, partition balances, etc.
- Undo a given number of moves (each move applied must be logged to support this).

**2.2.4 Pass-Based Partitioner (Proper).** Recall that a Pass-Based Partitioner applies incrementally improving passes to initial solutions and returns best solutions seen during such passes.<sup>10</sup> A pass consists of moves, chosen and applied by move manager. After a pass, a Partitioner can request that the Move Manager perform undo operations to return to the best solution seen in that pass. A Partitioner decides when to stop a pass, and what intermediate solution within the pass to return to, on the basis of its control parameters and status information returned by the Move Manager after each move. The Partitioner can have access to multiple combinations of Incremental Cost Evaluators, Move Managers and Gain Containers, and can use them flexibly at different passes to solve a given partitioning problem. Note that a Partitioner is not necessarily aware of the move structure used: this is a responsibility of Move Managers.

**Interface:**

- Takes a “partitioning problem” and operational parameters on input.
- Returns all solutions produced, with the best solution marked.

## 2.3 A Generic Component-based FM Algorithm

A “partitioning problem” consists of

- hypergraph

---

<sup>10</sup>While every pass as a whole must not worsen current solution, individual moves within a pass may do so.



- arrays of solution placeholders (“buffers”), possibly populated with initial solutions
- information representing relevant constraints, e.g., fixed assignments of vertices to partitions and maximum total vertex area in each partition
- additional information required to evaluate the cost function, e.g., geometry of partitions for wirelength-driven partitioning in the top-down placement context.

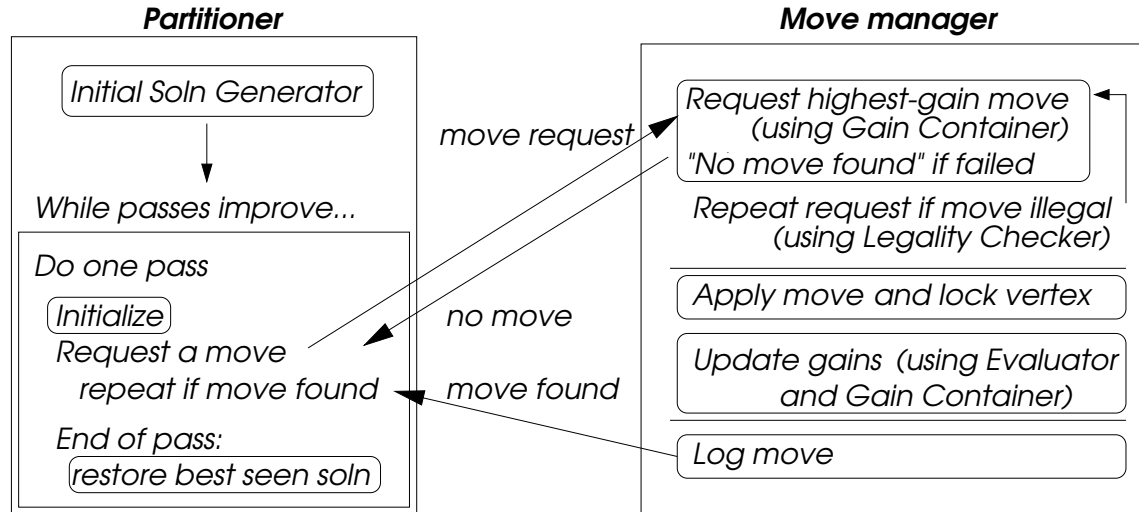


Fig. 3. Component-based FM algorithm improves a given solution or a randomly generated one. Improvement is performed in passes. A pass consists of legal moves with the best possible gain — positive, zero or negative. Every vertex can be only moved once which is enforced by “locking” every moved vertex, i.e. removing its moves from the gain container. The choice of the best move is facilitated by updates to the gain container after every move. At the end of the pass, the best seen solution is restored using a move log (every move is logged).

The Partitioner goes over relevant places in solution buffers and eventually writes good partitioning solutions into them. An existing solution may thus be improved, but if a place is empty, an initial solution generator will be called. A relevant Move Manager must be instantiated and initialized; this includes instantiation of the constituent Evaluator and Gain Container. The Partitioner then performs successive passes as long as the solution can be improved.

Our generic component-based FM algorithm is described in Figure 3. At each pass, the Partitioner repeatedly requests the Move Manager to pick one [best] move and apply it, and processes information about the new solutions thus obtained. Since no vertex can be moved twice in a pass, no moves will be available beyond a certain point (*end of a pass*). Some best-gain moves may increase the solution cost, and typically the solution at the end of the pass is not as good as the best solutions seen during the pass. The Partitioner then requests that the Move Manager undo a given number of moves to yield a solution with best cost.

While the reinitialization of the Move Manager at the beginning of each pass seems almost straightforward, picking and applying one move is subtle. For example, note that the Move Manager requests the best move from the gain container and can keep on requesting more moves until a move passes legality check(s). As the Move Manager applies the chosen move and locks the vertex, gains of adjacent vertices may need to be updated.

In performing “generic” gain update, the Move Manager walks all nets incident to the moving vertex and for each net computes gain updates (*delta gains*) for each of its vertices due to this net (these are combinations of the given net’s cost under four distinct partition assignments for the moving and affected vertices; see Section 3.4). These partial gain updates are immediately applied through Gain Container calls, and moves of affected vertices may have their priority within the Gain Container changed. Even if the delta gain for a given move is zero, removing and inserting it into the gain container will typically change tie-breaking among moves with the same gain.

In most implementations the gain update is the main bottleneck, followed by Gain Container construction. Numerous optimizations of generic algorithms exist for specific cost functions, hyperedge cut being particularly amenable to such optimizations.

### 3. EVALUATING PROSPECTIVE ADVANCES IN PARTITIONING

VLSI design presents many different flavors of hypergraph partitioning. Objective functions such as ratio-cut [48], scaled cost [11], absorption cut [47] sum of degrees, number of vertices on the cut line [29], etc. have been applied for purposes ranging from routability-driven clustering to multilevel annealing placement. In top-down coarse placement, partitioning involves fixed or “propagated” terminals [17; 46], tight partition balance constraints (and non-uniform vertex weights), and an estimated-wirelength objective (e.g., sum of half-perimeters of net bounding boxes). By contrast, for logic emulation the partitioning might have all terminals unfixed, loose balance constraints (with uniform vertex weights), and a pure min-cut objective. The partitioning can also be multi-way instead of 2-way [46; 43; 30], “multi-dimensional” (e.g., simultaneous balancing of power dissipation and module area among the partitions), timing-driven, etc. With this in mind, partitioners are best viewed as “engines” that plug into many different phases of VLSI design. Any prospective advance in partitioning technology should be evaluated in a range of contexts.

In recent VLSI CAD partitioning literature, comparisons to previous work are made using as wide a selection of benchmark instances as practically possible; using uniform vs. non-uniform vertex weights; and using tight vs. loose partition balance constraints (typically 49-51% and 45-55% constraints for bipartitioning).<sup>11</sup>

---

<sup>11</sup> VLSI CAD researchers also routinely document whether large nets were thresholded, the details of hypergraph-to-graph conversions (e.g., when applying spectral methods), and other details necessary for others to reproduce the experiments. The reader is referred to [5; 2] for discussions of reporting methodology.

### 3.1 Need For “Canonical” Testbench

The components described in Section 2 yield a testbench, or “framework”, that can be recombined and reused in many ways to enable experiments with

- multiple objective functions, e.g., ratio cut [48], absorption [47], the number of boundary vertices [29], the “ $k - 1$  objective” [13] etc.
- multiple constraint types ([34])
- variant formulations, e.g., multi-way [43; 33; 15], replication-based [38] etc.
- new partitioning algorithms and variations

The component-based framework allows seamless replacement of old algorithms by improved ones in containing applications. Even more important, a solid testbench is *absolutely essential* to identify algorithmic improvements “at the leading edge” of heuristic technology. I.e., it is critical to evaluate proposed algorithm improvements not only against the best available implementations, but also *using a competent implementation*. This is the main point we wish to make.

In our experience, new “improvements” often look good if applied to weak algorithms, but may actually worsen strong algorithms. Only after an improvement has been thoroughly analyzed, implemented and confirmed empirically, can it be turned on by default and be applied to all evaluations of all subsequent proposed improvements. On the other hand, one often encounters pairs of conflicting improvements of which one, if applied by itself, dominates the other while the combination of the two is the worst. Therefore, interacting improvements must be implemented as options, and tested in all possible combinations, and a flexible testbench is necessary to advance the understanding of algorithm performance by conclusive and revealing experiments.

### 3.2 Algorithm and Implementation Improvements

In this subsection we note the existence of several types of implementation decisions for optimization metaheuristics, and illustrate such decisions for the Fiduccia-Mattheyses heuristic [20] and its improvements.

Of particular interest are *implicit decisions* – underspecified features and ambiguities in the original algorithm description that need to be resolved in any particular implementation. Examples for the Fiduccia-Mattheyses heuristic include:

- tie-breaking* in choosing highest gain bucket (see Subsection 3.4)
- tie-breaking* on where to attach new element in gain bucket, i.e., LIFO versus FIFO versus random [27]<sup>12</sup>
- whether to update, or skip the updating, when the delta gain of a move is zero (see Subsection 3.4)
- breaking ties when selecting the best solution during the pass — choose the first or last one encountered, or the one that is furthest from violating constraints.

---

<sup>12</sup>In other words, gain buckets can be implemented as stacks or as queues, or even as random priority queues (where the chances of all elements to be selected are equal at all times). [27] demonstrated that stack-based gain containers (i.e., LIFO) are superior.

Below, we show that the effects of such implicit decisions can far outweigh claimed improvements in solution quality due to algorithm innovation. Other types of implementation decisions, details of which are beyond our present scope, include:

- *Modifications of the algorithm*: important changes to steps or the flow of the original algorithm as well as new steps and features. Among the more prominent examples are “lookahead” tie-breaking [39] among same-gain moves; the *multiple unlocking* heuristic of [15] which allows vertices to move more than once during a pass; and the CLIP heuristic of [18] which chooses moves according to “updated” gains (i.e., the actual gain minus the gain at the beginning of the pass) instead of actual gains.
- *Tuning that can change the result*: minor algorithm or implementation changes, typically to avoid particularly bad special cases or pursue only “promising computations”. Examples include thresholding large nets from the input to reduce run time; “loose net” removal [12] where gain updates are performed only for [loose] nets that are likely to be uncut; and allowing of illegal solutions during a pass (to improve hill-climbing ability of the algorithm) [19].
- *Tuning that can not change the result*: minor algorithm or implementation changes to simplify computations in critical or statistically significant special cases. Examples include skipping nets which cannot have non-zero delta gains (updates); code optimizations that are specific to the netcut objective; and code optimizations that are specific to 2-way partitioning.

#### 4. EMPIRICAL ILLUSTRATION

We now illustrate how “implicit implementation decisions” can severely distort the experimental assessment of new algorithmic ideas. Uncertainties in the description of the Fiduccia-Mattheyses algorithm have been previously analyzed, notably in [27], where the authors show that inserting moves into gain buckets in LIFO order is much preferable to doing so in FIFO order (also a constant-time insertion) or at random. Since the work of [27], all FM implementations that we are aware of use LIFO insertion.<sup>13</sup> In our experiments, we consider the following two implicit implementation decisions:

- **Zero delta gain update.** Recall that when a vertex  $x$  is moved, the gains for all vertices  $y$  on nets incident to  $x$  must potentially be updated. In all FM implementations, this is done by going through the incident nets one at a time, and computing the changes in gain for vertices  $y$  on these nets. A straightforward implementation computes the change in gain (“delta gain”) for  $y$  by adding and subtracting four cut values for the net under consideration,<sup>14</sup> and immediately updating  $y$ ’s position in the gain container.

Notice that sometimes the delta gain can be zero. An implicit implementation decision is whether to reinsert a vertex  $y$  when it experiences a zero delta gain

<sup>13</sup>A series of works in the mid-1990s retrospectively show that the LIFO order allows vertices in “natural clusters” to move together across the cutline. The CLIP variant of [18] is a more direct way of moving clusters.

<sup>14</sup>These four cut values correspond to: (a)  $x$ ,  $y$  in their original partitions; (b)  $x$  in original partition,  $y$  moved; (c)  $x$  moved,  $y$  in original partition; and (d)  $x$  and  $y$  both moved. (a) – (b) is the original gain for  $y$  due to the net under consideration; (c) – (d) is the new gain for  $y$  due to the same net. The difference ((a) – (b)) – ((c) – (d)) is the delta gain. See [30] for a discussion.

move (“**All $\Delta_{gain}$** ”), or whether to skip the gain update (“**Nonzero**”). The former will shift the position of  $y$  within the same gain bucket; the latter will leave  $y$ 's position unchanged. The effect of zero delta gain updating is not immediately obvious.<sup>15</sup>

—**Tie-breaking between two highest-gain buckets in move selection.**

When the gain container is implemented such that available moves are segregated, typically by source or destination partition, there can be more than one nonempty highest-gain bucket. Notice that when the balance constraint is anything other than “exact bisection”, it is possible for all the moves at the heads of the highest-gain buckets to be legal. The FM implementer must choose a method for dealing with this situation. In our experiments, we contrast three approaches:<sup>16</sup> (i) choose the move that is not from the same partition as the last vertex moved (“**away**”); (ii) choose the move in partition 0 (“**part0**”); and (iii) choose the move from the same partition as the last vertex moved (“**toward**”).

Our experimental testbench allows us to test an FM variant in the context of flat LIFO (as described in [27]), flat CLIP (as described in [18]), and multilevel LIFO and multilevel CLIP (as described in [4]). Our implementations are in C++ with heavy use of STL3.0; we currently run in the Sun Solaris 2.6 and 2.7 and Sun CC4.2 environment (as well as g++ 2.95.2 on RedHat Linux 6.0 and 6.1). We use standard VLSI benchmark instances available on the Web at [1] and several older benchmarks from <http://www.cbl.ncsu.edu/benchmarks>. Node and hyperedge statistics for the benchmarks are presented in Tables 4, 5 and 6. Our tests are for bipartitioning only. We evaluate all partitioning variants using actual vertex areas and unit vertex areas, incorporating the standard protocols for treating pad areas described in [2]. We also evaluate all partitioning variants using both a 10% balance constraint (i.e., each partition must have between 45% and 55% of the total vertex area) as well as a 2% balance constraint (results are qualitatively similar; we therefore report only results for 2% balance constraints). All experiments were run on Sun Ultra workstations, with runtimes normalized to Sun Ultra-1 (140MHz) CPU seconds. Each result represents a set of 100 independent runs with random initial starting solutions; Tables 4 and 4 report triples of form “**average cut** (average CPU sec)” (see also Figures 4 and 5). From the data, we make the following observations.

- The average cutsizes for a flat partitioner can increase by rather stunning percentages if the worst combination of choices is used instead of the best combination. Such effects stronger for actual areas and weaker for unit areas, but in all cases far outweigh the typical solution quality improvements reported for new algorithm ideas in the partitioning literature.
- Moreover, we see that one wrong implementation decision can lead to misleading conclusions with respect to other implementation decisions. For example, when zero delta gain updates are made (a wrong decision), the “part0” biasing choice

<sup>15</sup>The gain update method presented in [20] has the *side effect* of skipping all zero delta gain updates. However, this method is both netcut- and two-way specific; it is not certain that a first-time experimenter with FM will find analogous solutions for  $k$ -way partitioning with a general objective.

<sup>16</sup>These approaches are described for the case of bipartitioning. Other approaches can be devised for  $k$ -way partitioning.

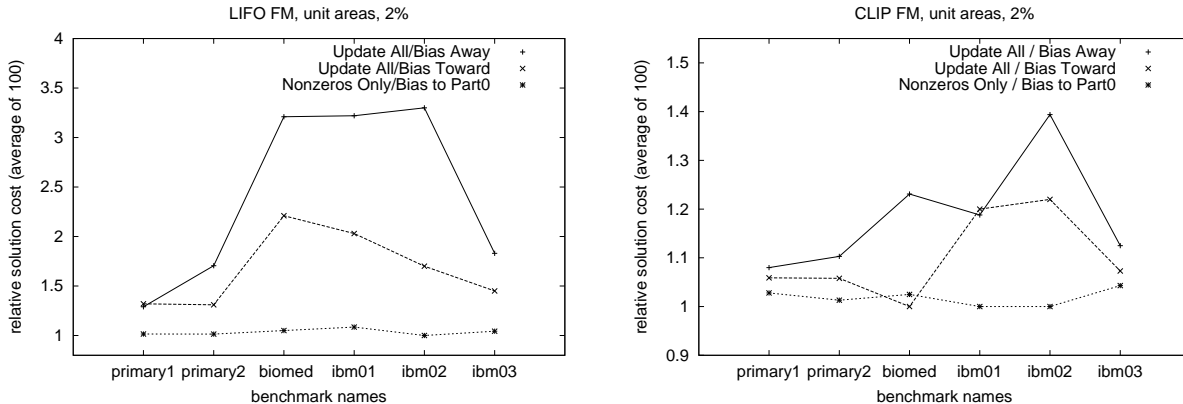


Fig. 4. Performance of three implementation variants of LIFO FM and CLIP FM on six benchmarks, with unit areas and partitioning tolerance 2%. We plot performance relative to the best solution cost seen over all runs of all variants. The worst results seen are up to 3.5 times worse than the best seen. Additional implementation options are analyzed in Table 4, which gives more detailed information about the same experiment. The plots clearly show that a poor implementation choice (Update All) is masked by an overall stronger heuristic (CLIP FM).

appears significantly worse than the “toward” choice. However, when zero delta gain updates are skipped, “part0” is as good as or even slightly better than “toward”.<sup>17</sup>

- Stronger optimization engines (order of strength: ML CLIP > ML LIFO > flat CLIP > flat LIFO) can tend to decrease the “dynamic range” for the effects of implementation choices. This is actually a danger: e.g., developing a multilevel FM package may hide the fact that the underlying flat engines are badly implemented. At the same time, the effects of a bad implementation choice are still apparent even when that choice is wrapped within a strong optimization technique (e.g., ML CLIP).

## 5. CONCLUSIONS

The results reported in the previous section are for a “detuned” or “generalized” version of our testbench, where we deliberately re-enabled the zero delta gain update and gain bucket choice as options. In our current testbench, these are not options, i.e., our FM-based engines always skip zero delta gain updates and always choose the “toward” gain bucket in case of ties. Our current testbench is also able to invoke several speedups that exploit the nature of the netcut objective and the two-way partitioning context. Its performance is summarized in Table 5 and can be used for benchmarking other implementations. The flexibility of our testbench is particularly valuable; it allows to address new objectives, neighborhood structures, constraint types and certainly justifies the small (10 – 15%) runtime overhead due to object-oriented structure, conditional tests and the like.

<sup>17</sup>We have observed other similar reversals, e.g., in our experience multiple unlocking is less valuable than reported in [15].

ALGORITHM		TESTCASES with <b>unit</b> areas and <b>2%</b> balance					
Updates	Bias	<b>primary1</b>	<b>primary2</b>	<b>biomed</b>	<b>ibm01</b>	<b>ibm02</b>	<b>ibm03</b>
<b>Flat LIFO FM</b>							
All $\Delta_{gain}$	Away	<b>102</b> (0.247)	<b>486</b> (1.6)	<b>459</b> (29.3)	<b>1778</b> (16.4)	<b>1810</b> (36.5)	<b>4175</b> (34.1)
All $\Delta_{gain}$	Part0	<b>102</b> (0.27)	<b>465</b> (2.06)	<b>422</b> (37.8)	<b>1673</b> (19.8)	<b>1570</b> (43.3)	<b>4064</b> (35)
All $\Delta_{gain}$	Toward	<b>102</b> (0.264)	<b>374</b> (1.94)	<b>316</b> (36.5)	<b>1030</b> (15.4)	<b>931</b> (30)	<b>3323</b> (39.6)
Nonzero	Away	<b>80.5</b> (0.222)	<b>285</b> (1.31)	<b>166</b> (22.8)	<b>543</b> (10.2)	<b>549</b> (18)	<b>2304</b> (29.2)
Nonzero	Part0	<b>80</b> (0.236)	<b>289</b> (1.47)	<b>150</b> (26.3)	<b>551</b> (11.8)	<b>549</b> (18.6)	<b>2383</b> (31.1)
Nonzero	Toward	<b>78.8</b> (0.219)	<b>291</b> (1.41)	<b>143</b> (24.8)	<b>508</b> (10.9)	<b>551</b> (19.3)	<b>2285</b> (33.1)
<b>Flat CLIP FM</b>							
All $\Delta_{gain}$	Away	<b>69.8</b> (0.351)	<b>246</b> (2.22)	<b>149</b> (53.9)	<b>555</b> (21.1)	<b>707</b> (46.2)	<b>1747</b> (45.5)
All $\Delta_{gain}$	Part0	<b>68</b> (0.319)	<b>242</b> (2.3)	<b>138</b> (52)	<b>562</b> (21.3)	<b>636</b> (45.5)	<b>1686</b> (51.7)
All $\Delta_{gain}$	Toward	<b>68.4</b> (0.35)	<b>236</b> (2.14)	<b>121</b> (44)	<b>561</b> (19.3)	<b>619</b> (44)	<b>1664</b> (43.2)
Nonzero	Away	<b>66.3</b> (0.284)	<b>231</b> (2.02)	<b>130</b> (29.1)	<b>488</b> (17.5)	<b>534</b> (46.8)	<b>1650</b> (41.9)
Nonzero	Part0	<b>66.4</b> (0.305)	<b>226</b> (2.2)	<b>124</b> (33.7)	<b>467</b> (17.1)	<b>507</b> (37.4)	<b>1618</b> (41.2)
Nonzero	Toward	<b>64.6</b> (0.304)	<b>223</b> (2.13)	<b>125</b> (31.4)	<b>474</b> (15.8)	<b>559</b> (36.9)	<b>1551</b> (41.9)
<b>ML LIFO FM</b>							
All $\Delta_{gain}$	Away	<b>64.2</b> (0.818)	<b>183</b> (4.95)	<b>147</b> (36.8)	<b>280</b> (29.7)	<b>399</b> (70.6)	<b>1043</b> (138)
All $\Delta_{gain}$	Part0	<b>63.8</b> (0.871)	<b>182</b> (5.21)	<b>145</b> (33.5)	<b>282</b> (31.6)	<b>406</b> (66.8)	<b>999</b> (113)
All $\Delta_{gain}$	Toward	<b>63.2</b> (0.836)	<b>180</b> (4.74)	<b>145</b> (33.6)	<b>272</b> (30.8)	<b>421</b> (66.1)	<b>1035</b> (118)
Nonzero	Away	<b>62.9</b> (0.807)	<b>182</b> (4.69)	<b>142</b> (28.1)	<b>274</b> (29.7)	<b>396</b> (60.4)	<b>1037</b> (123)
Nonzero	Part0	<b>61.8</b> (0.849)	<b>176</b> (5.15)	<b>144</b> (26.8)	<b>270</b> (28.8)	<b>403</b> (60.2)	<b>1048</b> (121)
Nonzero	Toward	<b>61.4</b> (0.879)	<b>181</b> (4.91)	<b>143</b> (24.8)	<b>271</b> (29.1)	<b>415</b> (56)	<b>1015</b> (119)
<b>ML CLIP FM</b>							
All $\Delta_{gain}$	Away	<b>63.4</b> (0.912)	<b>180</b> (5.27)	<b>146</b> (32.6)	<b>276</b> (30.2)	<b>400</b> (68.7)	<b>1042</b> (138)
All $\Delta_{gain}$	Part0	<b>62.8</b> (0.849)	<b>180</b> (5.18)	<b>145</b> (39.1)	<b>278</b> (29.1)	<b>408</b> (67.8)	<b>1022</b> (129)
All $\Delta_{gain}$	Toward	<b>63.3</b> (0.871)	<b>178</b> (5.36)	<b>141</b> (36.7)	<b>270</b> (29.4)	<b>425</b> (64.4)	<b>1032</b> (101)
Nonzero	Away	<b>61.2</b> (0.887)	<b>178</b> (5.12)	<b>142</b> (26.8)	<b>268</b> (28.4)	<b>392</b> (57)	<b>1037</b> (116)
Nonzero	Part0	<b>62.8</b> (0.878)	<b>176</b> (5.14)	<b>142</b> (29.3)	<b>275</b> (28.6)	<b>409</b> (55.6)	<b>1026</b> (111)
Nonzero	Toward	<b>63</b> (0.924)	<b>179</b> (4.86)	<b>142</b> (28.5)	<b>271</b> (28.2)	<b>409</b> (52.9)	<b>1034</b> (115)

Table 1. Average cuts in partitioning with **unit** areas and **2%** balance tolerance, over 100 independent runs. Average CPU time per run in Sun Ultra-1 (140MHz) seconds is given in parentheses.

The replicability of reported results is a well-known problem in the VLSI hypergraph partitioning community [5], as well as in experimental algorithms research as a whole. Certainly, such issues as experimental protocols, statistical significance tests, data reporting methodology and comparison with previous results are all critical aspects of scientific communication that have been addressed by previous literature [6; 8; 22]. In this work, we have focused on the need to describe proposed algorithms and algorithm improvements in sufficient detail, notably with respect to the documentation of all significant implementation decisions. We have presented evidence suggesting the appropriate level of detail for reporting algorithm improvements; we suggest that even an expert programmer and algorithm researcher may

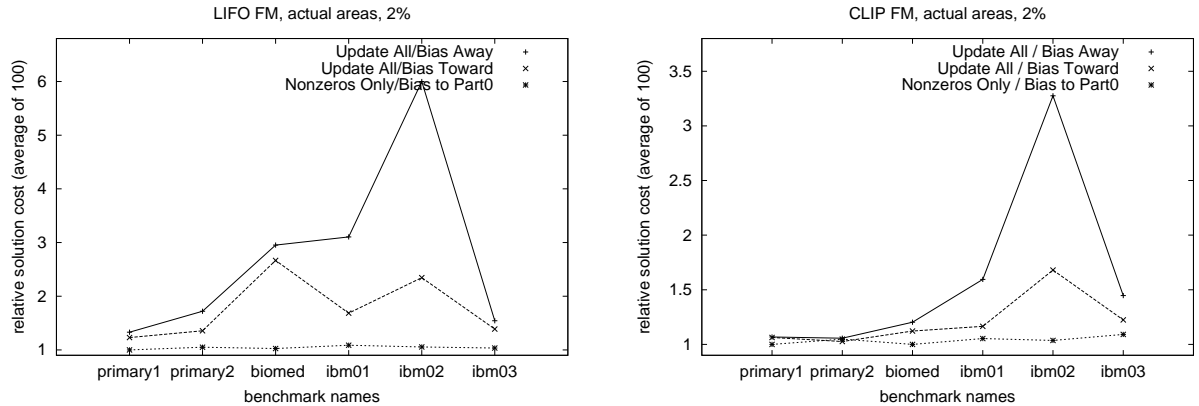


Fig. 5. Performance of three implementation variants of LIFO FM and CLIP FM on six benchmarks, with actual areas and partitioning tolerance 2%. We plot performance relative to the best solution cost seen over all runs of all variants. The worst results seen are up to six times worse than the best seen. Additional implementation options are analyzed in Table 4, which gives more detailed information about the same experiment. The plots clearly show that a poor implementation choices (Updating All gains) is masked by an overall stronger heuristic (CLIP FM).

not otherwise recover a leading edge implementation (i.e., from a poor description). Our main contributions have been the description of a software architecture for partitioning research, a review of the modern context for such research in the VLSI CAD domain, and a sampling of hidden implementation decisions that are crucial to obtaining a useful testbench.

#### Acknowledgments

We thank Max Moroz for design, implementation and support of an efficient STL-based hypergraph package used in the paper.

#### REFERENCES

- [1] C. J. Alpert, “Partitioning Benchmarks for VLSI CAD Community”, Web page, <http://vlsicad.cs.ucla.edu/~cheese/benchmarks.html>
- [2] C. J. Alpert, “The ISPD-98 Circuit Benchmark Suite”, *Proc. ACM/IEEE International Symposium on Physical Design*, April 98, pp. 80-85. See errata at <http://vlsicad.cs.ucla.edu/~cheese/errata.html>
- [3] C. J. Alpert and L. W. Hagen and A. B. Kahng, “A Hybrid Multilevel/Genetic Approach for Circuit Partitioning”, *Proc. IEEE Asia Pacific Conference on Circuits and Systems*, 1996, pp. 298-301.
- [4] C. J. Alpert, J.-H. Huang and A. B. Kahng, “Multilevel Circuit Partitioning”, *ACM/IEEE Design Automation Conference*, pp. 530-533. <http://vlsicad.cs.ucla.edu/papers/conference/c68.ps>
- [5] C. J. Alpert and A. B. Kahng, “Recent Directions in Netlist Partitioning: A Survey”, *Integration*, 19(1995) 1-81.
- [6] R. S. Barr, B. L. Golden, J. P. Kelly, M. G. C. Resende and W. R. Stewart, “Designing and Reporting on Computational Experiments with Heuristic Methods”, *technical report* (extended version of *J. Heuristics* paper), June 27, 1995.



ALGORITHM		TESTCASES with actual areas and 2% balance					
Updates	Bias	primary1	primary2	biomed	ibm01	ibm02	ibm03
<b>Flat LIFO FM</b>							
All $\Delta_{gain}$	Away	<b>105</b> (0.283)	<b>481</b> (2.05)	<b>446</b> (27.8)	<b>1885</b> (11.1)	<b>3256</b> (34.8)	<b>4389</b> (25.4)
All $\Delta_{gain}$	Part0	<b>103</b> (0.311)	<b>467</b> (2.4)	<b>428</b> (34.9)	<b>1909</b> (12.3)	<b>2440</b> (33.5)	<b>4166</b> (27.1)
All $\Delta_{gain}$	Toward	<b>96.8</b> (0.285)	<b>380</b> (2.15)	<b>408</b> (30.7)	<b>1023</b> (11.6)	<b>1274</b> (22.7)	<b>3939</b> (22.3)
Nonzero	Away	<b>80.5</b> (0.258)	<b>291</b> (1.78)	<b>164</b> (27.3)	<b>639</b> (8.86)	<b>551</b> (14.9)	<b>2838</b> (25.2)
Nonzero	Part0	<b>79</b> (0.271)	<b>294</b> (1.74)	<b>157</b> (29.2)	<b>660</b> (7.87)	<b>573</b> (17)	<b>2938</b> (24)
Nonzero	Toward	<b>79.7</b> (0.25)	<b>280</b> (1.79)	<b>153</b> (24.7)	<b>607</b> (7.62)	<b>543</b> (15.8)	<b>2843</b> (25.4)
<b>Flat CLIP FM</b>							
All $\Delta_{gain}$	Away	<b>66.2</b> (0.361)	<b>244</b> (2.71)	<b>148</b> (57)	<b>842</b> (15.3)	<b>1841</b> (27.5)	<b>3623</b> (23.4)
All $\Delta_{gain}$	Part0	<b>66.3</b> (0.395)	<b>242</b> (2.81)	<b>141</b> (54.6)	<b>772</b> (14.9)	<b>1499</b> (32.4)	<b>3543</b> (29.3)
All $\Delta_{gain}$	Toward	<b>65.9</b> (0.393)	<b>237</b> (2.7)	<b>138</b> (58)	<b>615</b> (13)	<b>945</b> (21.5)	<b>3066</b> (25.6)
Nonzero	Away	<b>64.5</b> (0.351)	<b>231</b> (2.66)	<b>130</b> (33.3)	<b>542</b> (12.1)	<b>574</b> (18.5)	<b>2689</b> (22.8)
Nonzero	Part0	<b>62</b> (0.371)	<b>242</b> (2.49)	<b>123</b> (35)	<b>556</b> (12.4)	<b>582</b> (17.8)	<b>2732</b> (23.1)
Nonzero	Toward	<b>63.9</b> (0.377)	<b>233</b> (2.31)	<b>124</b> (34.9)	<b>528</b> (11.8)	<b>562</b> (15.3)	<b>2504</b> (23.2)
<b>ML LIFO FM</b>							
All $\Delta_{gain}$	Away	<b>62.8</b> (0.905)	<b>163</b> (5.27)	<b>145</b> (34.9)	<b>289</b> (27.9)	<b>433</b> (42.9)	<b>958</b> (59.4)
All $\Delta_{gain}$	Part0	<b>62.5</b> (0.954)	<b>166</b> (5.03)	<b>144</b> (38.7)	<b>289</b> (27.2)	<b>429</b> (44.6)	<b>957</b> (56.3)
All $\Delta_{gain}$	Toward	<b>61.1</b> (0.974)	<b>161</b> (5.28)	<b>143</b> (36.3)	<b>289</b> (27.7)	<b>423</b> (47)	<b>971</b> (58.7)
Nonzero	Away	<b>60.4</b> (0.914)	<b>158</b> (4.52)	<b>142</b> (29.9)	<b>287</b> (22.7)	<b>432</b> (39.6)	<b>969</b> (52.2)
Nonzero	Part0	<b>59.9</b> (0.882)	<b>158</b> (4.36)	<b>142</b> (29.3)	<b>282</b> (25.3)	<b>421</b> (44)	<b>952</b> (50.6)
Nonzero	Toward	<b>60.5</b> (0.9)	<b>159</b> (4.5)	<b>142</b> (27.5)	<b>276</b> (25.4)	<b>419</b> (43.2)	<b>959</b> (52.5)
<b>ML CLIP FM</b>							
All $\Delta_{gain}$	Away	<b>63.5</b> (0.88)	<b>163</b> (5)	<b>144</b> (35.9)	<b>283</b> (24.5)	<b>428</b> (41.5)	<b>960</b> (59.3)
All $\Delta_{gain}$	Part0	<b>62.5</b> (0.891)	<b>161</b> (4.24)	<b>143</b> (37.1)	<b>289</b> (25.3)	<b>441</b> (46.5)	<b>969</b> (63.6)
All $\Delta_{gain}$	Toward	<b>61.9</b> (0.927)	<b>162</b> (4.83)	<b>141</b> (37.8)	<b>284</b> (24.9)	<b>425</b> (44.2)	<b>953</b> (62.1)
Nonzero	Away	<b>60.2</b> (0.939)	<b>160</b> (4.66)	<b>144</b> (31.8)	<b>283</b> (23)	<b>414</b> (48.7)	<b>957</b> (50.4)
Nonzero	Part0	<b>61</b> (0.895)	<b>161</b> (4.89)	<b>144</b> (28.1)	<b>285</b> (24.7)	<b>447</b> (41.8)	<b>934</b> (53)
Nonzero	Toward	<b>60.9</b> (0.864)	<b>155</b> (4.7)	<b>144</b> (29.6)	<b>282</b> (22.4)	<b>433</b> (42.6)	<b>959</b> (50.6)

Table 2. Average cuts in partitioning with **actual** areas and **2%** balance tolerance, over 100 independent runs. Average CPU time per run in Sun Ultra-1 (140MHz) seconds is given in parentheses.

- [7] F. Brglez, "ACM/SIGDA Design Automation Benchmarks: Catalyst or Anathema?", *IEEE Design and Test*, 10(3) (1993), pp. 87-91.
- [8] F. Brglez, "Design of Experiments to Evaluate CAD Algorithms: Which Improvements Are Due to Improved Heuristic and Which are Merely Due to Chance?", *technical report CBL-04-Brglez*, NCSU Collaborative Benchmarking Laboratory, April 1998.
- [9] T. Bui, S. Chaudhuri, T. Leighton and M. Sipser, "Graph Bisection Algorithms with Good Average Behavior", *Combinatorica* 7(2), 1987, pp. 171-191.
- [10] A. E. Caldwell, A. B. Kahng and I. L. Markov, *manuscript*, 1998.
- [11] P. K. Chan and M. D. F. Schlag and J. Y. Zien, "Spectral  $K$ -Way Ratio-Cut Partitioning and Clustering", *IEEE Transactions on Computer-Aided Design*, vol. 13 (8), pp. 1088-1096.
- [12] J. Cong, H. P. Li, S. K. Lim, T. Shibuya and D. Xu, "Large Scale Circuit Partitioning with

Configuration	primary1	primary2	biomed	ibm01	ibm02	ibm03
2% unit area	57.1(0.418)	164.7(1.3)	132.8(2.24)	275.0(5.7)	372.1(10.6)	1031.4(11.6)
10% unit area	52.1(0.412)	143.6(1.3)	117.0(2.16)	247.9(5.82)	268.8(10.3)	820.9(11.8)
2% actual area	61.3(0.389)	193.1(1.41)	137.2(2.07)	284.5(6.63)	421.3(14.6)	1079.8(15.8)
10% actual area	54.1(0.421)	177.8(1.07)	125.0(2.23)	244.4(5.36)	275.7(14.7)	1062.1(16.4)

Table 3. Results of applying our optimized multilevel partitioner on 6 test-cases. Solutions are constrained to be within 2% or 10% of bisection. Data expressed as (average cut / average CPU time) over 100 runs, with the latter normalized to CPU seconds on a 140MHz Sun Ultra-1.

Node degree statistics for testcases											
primary1(833)		primary2(3014)		biomed(6514)		ibm01(12752)		ibm02(19601)		ibm03(23136)	
deg:	#	deg:	#	deg:	#	deg:	#	deg:	#	deg:	#
Avg:	3.49	Avg:	3.72	Avg:	3.230	Avg:	3.965	Avg:	4.143	Avg:	4.044
1:	48	1:	43	1:	97	1:	781	1:	1591	1:	363
2:	145	2:	453	2:	792	2:	3722	2:	4448	2:	7093
3:	205	3:	1266	3:	4492	3:	2016	3:	2318	3:	4984
4:	273	4:	519	4:	440	4:	1430	4:	1714	4:	4357
5:	234	5:	402	5:	35	5:	1664	5:	3406	5:	2778
6:	5	6:	23	6:	658	6:	1761	6:	4435	6:	1252
7:	22	7:	260			7:	542	7:	1055	7:	300
9:	1	8:	4			8:	194	8:	303	8:	689
		9:	44			9:	368	9:	319	9:	708
						10:	3	29:	2	10:	50
						13:	220	32:	1	11:	25
						39:	1	51:	4	12:	192
								53:	3	13:	47
								60:	1	14:	7
								69:	1	16:	10
										17:	7
										18-19:	4
										20:	16
										21-23:	8
										24:	136
										25:	101
										84-100:	9

Table 4. Hypergraph node degree statistics. The numbers of nodes in degree ranges are given for each testcase together with the total nodes and average node degree.

Loose/Stable Net Removal and Signal Flow Based Clustering”, *Proc. IEEE International Conference on Computer-Aided Design*, 1997, pp. 441-446.

- [13] J. Cong and S. K. Lim, “Multiway Partitioning with Pairwise Movement”, *Proc. IEEE/ACM International Conference on Computer-Aided Design*, 1998, to appear.
- [14] J. Darnauer and W. Dai, “A Method for Generating Random Circuits and Its Applications to Routability Measurement”, *Proc. ACM/SIGDA International Symposium on FPGAs*, 1996, pp. 66-72.
- [15] A. Dasdan and C. Aykanat, “Two Novel Multiway Circuit Partitioning Algorithms Using Relaxed Locking”, *IEEE Transactions on Computer-Aided Design* 16(2) (1997), pp. 169-178.
- [16] W. Deng, *personal communication*, July 1998.
- [17] A. E. Dunlop and B. W. Kernighan, “A Procedure for Placement of Standard Cell VLSI

Node weight statistics for testcases						
weight range	# nodes					
	primary1	primary2	biomed	ibm01	ibm02	ibm03
1	81	107	97	12749	19589	23126
2				2	3	
3					3	
4	89	367				
5			723			7
6	142	715			1	
7	151	494	2818			1
8	73	398				
9			1323		1	
10						
11	27	278			11	
12			35			
13						
14			3			
15	7	550				
16						
17	1	52				
18			860			
19						
20	262	53	655	1	1	2

Table 5. Hypergraph node weight statistics. The interval between the smallest and largest node weight has been divided into 20 equal ranges for each testcase. For each such range we report the number of nodes with weight in this range.

- Circuits”, *IEEE Transactions on Computer-Aided Design* 4(1) (1985), pp. 92-98
- [18] S. Dutt and W. Deng, “VLSI Circuit Partitioning by Cluster-Removal Using Iterative Improvement Techniques”, *Proc. IEEE International Conference on Computer-Aided Design*, 1996, pp. 194-200
- [19] S. Dutt and H. Theny, “Partitioning Using Second-Order Information and Stochastic Gain Function”, *Proc. IEEE/ACM International Symposium on Physical Design*, 1998, pp. 112-117
- [20] C. M. Fiduccia and R. M. Mattheyses, “A Linear Time Heuristic for Improving Network Partitions”, *Proc. ACM/IEEE Design Automation Conference*, 1982, pp. 175-181.
- [21] M. R. Garey and D. S. Johnson, “Computers and Intractability, a Guide to the Theory of NP-completeness”, W. H. Freeman and Company: New York, 1979, pp. 223
- [22] I. P. Gent, S. A. Grant, E. McIntyre, P. Prosser, P. Shaw, B. M. Smith and T. Walsh, “How Not To Do It”, *research report 97-27*, Univ. of Leeds School of Computer Studies, May 1997.
- [23] M. Ghose, M. Zubair and C. E. Grosch, “Parallel Partitioning of Sparse Matrices”, *Computer Systems Science & Engineering* (1995) 1, pp. 33-40.
- [24] D. Ghosh, “Synthesis of Equivalence Class Circuit Mutants and Applications to Benchmarking”, *summary of presentation at DAC-98 Ph.D. Forum*, June 1998.
- [25] M. K. Goldberg and M. Burstein, “Heuristic Improvement Technique for Bisection of VLSI Networks”, *IEEE Transactions on Computer-Aided Design*, 1983, pp. 122-125.
- [26] S. Hauck and G. Borriello, “An Evaluation of Bipartitioning Techniques”, *IEEE Transactions on Computer-Aided Design* 16(8) (1997), pp. 849-866.
- [27] L. W. Hagen, D. J. Huang and A. B. Kahng, “On Implementation Choices for Iterative Improvement Partitioning Methods”, *Proc. European Design Automation Conference*, 1995, pp. 144-149.

Hyperedge statistics for testcases					
primary1(902)	primary2(3029)	biomed(5742)	ibm01(14111)	ibm02(19568)	ibm03(27401)
deg: #	deg: #	deg: #	deg: #	deg: #	deg: #
Avg: 3.22	Avg: 3.70	Avg: 3.664	Avg: 3.583	Avg: 4.146	Avg: 3.415
2: 494	2: 1835	2: 3998	2: 8341	2: 10692	2: 17619
3: 236	3: 365	3: 870	3: 2082	3: 1934	3: 3084
4: 62	4: 203	4: 427	4: 1044	4: 1951	4: 2155
5: 26	5: 192	5: 184	5: 737	5: 1946	5: 1050
6: 25	6: 120	6: 13	6: 407	6: 376	6: 790
7: 13	7: 52	7: 11	7: 235	7: 332	7: 436
8: 2	8: 14	8: 28	8: 188	8: 256	8: 342
9: 9	9: 83	9: 7	9: 192	9: 424	9: 501
10: 1	10: 14	10: 4	10: 194	10: 431	10: 235
11: 6	11: 35	11: 5	11: 147	11: 498	11: 198
12: 9	12: 5	12: 5	12: 91	12: 46	12: 162
13: 1	13: 3	13: 1	13: 133	13: 52	13: 195
14: 3	14: 10	14: 2	14: 54	14: 52	14: 112
16: 1	15: 3	15: 41	15: 34	15: 85	15: 79
17: 11	16: 1	17: 21	16: 54	16: 94	16: 100
18: 3	17: 72	18: 1	17: 31	17: 143	17: 119
	18: 1	20: 2	18: 17	18: 100	18: 81
	23: 1	21: 65	19: 12	19: 44	19: 41
	26: 1	22: 34	20: 21	20: 15	20: 24
	29: 1	23: 6	21: 18	21: 11	21: 12
	30: 1	24: 6	22: 31	22: 5	22: 16
	31: 1	43: 6	23: 18	24-29: 10	23: 9
	33: 14	656: 4	25: 2	30: 4	24: 3
	34: 1	861: 1	28: 1	31: 11	25: 6
	37: 1		30: 2	32: 4	26: 3
			31: 2	33: 2	27: 2
			32: 5	34: 1	28: 2
			33: 6	35: 5	29: 6
			34: 1	36: 3	30: 1
			35: 7	37: 2	31: 2
			38: 1	38: 2	31: 3
			39: 2	39: 1	32: 3
			42: 1	40-97: 51	33: 5
				107: 1	34: 3
				134: 1	37-55: 5

Table 6. Hyperedge statistics. The numbers of hyperedges in ranges are given for each testcase together with total hyperedges and average hyperedge degree.

- [28] B. Hendrickson and R. Leland, “A Multilevel Algorithm for Partitioning Graphs”, draft, 1995.
- [29] B. Hendrickson and T. G. Kolda, “Partitioning Rectangular and Structurally Nonsymmetric Sparse Matrices for Parallel Processing”, *manuscript*, 1998 (extended version of PARA98 workshop proceedings paper).
- [30] D. J. Huang and A. B. Kahng, “Partitioning-Based Standard Cell Global Placement with an Exact Objective”, *Proc. ACM/IEEE International Symposium on Physical Design*, 1997, pp. 18-25.  
<http://vlsicad.cs.ucla.edu/papers/conference/c66.ps>
- [31] M. Hutton, J. P. Grossman, J. Rose and D. Corneil, “Characterization and Parameterized Random Generation of Digital Circuits”, *In Proc. IEEE/ACM Design Automation Con-*

- ference, 1996, pp. 94-99.
- [32] A. B. Kahng, "Futures for Partitioning in Physical design", *Proc. IEEE/ACM International Symposium on Physical Design*, April 1998, pp. 190-193.  
<http://vlsicad.cs.ucla.edu/papers/conference/c77.ps>
- [33] G. Karypis and V. Kumar, "Multilevel  $k$ -way Partitioning Scheme For Irregular Graphs", Technical Report 95-064, University of Minnesota, Computer Science Department.
- [34] G. Karypis and V. Kumar, "Multilevel Algorithms for Multi-Constraint Graph Partitioning", Technical Report 98-019, University of Minnesota, Department of Computer Science.
- [35] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar, "Multilevel Hypergraph Partitioning: Applications in VLSI Design", *Proc. ACM/IEEE Design Automation Conference*, 1997, pp. 526-529.  
Additional publications and benchmark results for hMetis-1.5 are available at <http://www-users.cs.umn.edu/~karypis/metis/hmetis/main.html>
- [36] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar, "Multilevel Hypergraph Partitioning: Applications in VLSI Domain", *technical report*, University of Minnesota Computer Science Department, March 27, 1998.
- [37] B. W. Kernighan and S. Lin, "An Efficient Heuristic Procedure for Partitioning Graphs", *Bell System Tech. Journal* 49 (1970), pp. 291-307.
- [38] C. Kring and A. R. Newton, "A Cell-Replicating Approach to Mincut-Based Circuit Partitioning", *Proc. IEEE/ACM International Conference on Computer-Aided Design*, 1991, pp. 2-5.
- [39] B. Krishnamurthy, "An Improved Min-cut Algorithm for Partitioning VLSI Networks", *IEEE Transactions on Computers*, vol. C-33, May 1984, pp. 438-446.
- [40] B. Landman and R. Russo, "On a Pin Versus Block Relationship for Partitioning of Logic Graphs", *IEEE Transactions on Computers* C-20(12) (1971), pp. 1469-1479.
- [41] L. T. Liu, M. T. Kuo, S. C. Huang and C. K. Cheng, "A Gradient Method on the Initial Partition of Fiduccia-Mattheyses Algorithm", *Proc. IEEE International Conference on Computer-Aided Design*, 1995, pp. 229-234.
- [42] T. Lengauer, *Combinatorial Algorithms for Integrated Circuit Layout*, Wiley-Teubner, 1990.
- [43] L. Sanchis, "Multiple-way network partitioning with different cost functions", *IEEE Transactions on Computers*, Dec. 1993, vol.42, (no.12):1500-4.
- [44] G. R. Schreiber and O. C. Martin, "Procedure for Ranking Heuristics Applied to Graph Partitioning", *Proc. 2nd International Conference on Metaheuristics*, July 1997, pp. 1-19.
- [45] G. R. Schreiber and O. C. Martin, "Cut Size Statistics of Graph Bisection Heuristics", *manuscript* in submission to *SIAM J. Optimization*, 1997.
- [46] P. R. Suaris and G. Kedem, "Quadrisection: A New Approach to Standard Cell Layout", *Proc. IEEE/ACM International Conference on Computer-Aided Design*, 1987, pp. 474-477.
- [47] W. Sun and C. Sechen, "Efficient and Effective Placements for Very Large Circuits", *Proc. IEEE/ACM International Conference on Computer-Aided Design*, 1993, pp. 170-177.
- [48] Y. C. Wei and C. K. Cheng, "Towards Efficient Design by Ratio-cut Partitioning", *Proc. IEEE International Conference on Computer-Aided Design*, 1989, pp. 298-301.