Multilevel Optimization for Large-Scale Circuit Placement

Tony F. Chan
Mathematics Department
University of California at Los Angeles

Collaborators (UCLA CS & Math Dept): Chin-Chih Chang, ***Jason Cong,*** Tianming Kong, Michalis Romesis, Joseph Shinnerl, Kenton Sze, Min Xie, Xin Yuan, Yan Zhang

Research Supported by Intel, SRC, NSF.

Purpose of Talk

- **Introduce problem and some key math & algorithmic ideas --- minimal details**
- **Overview of our Multilevel framework**
- **Connection to Optimization, Multigrid, Total Variation**
- **Will give general references at the end.**

VLSICAD Design Flow

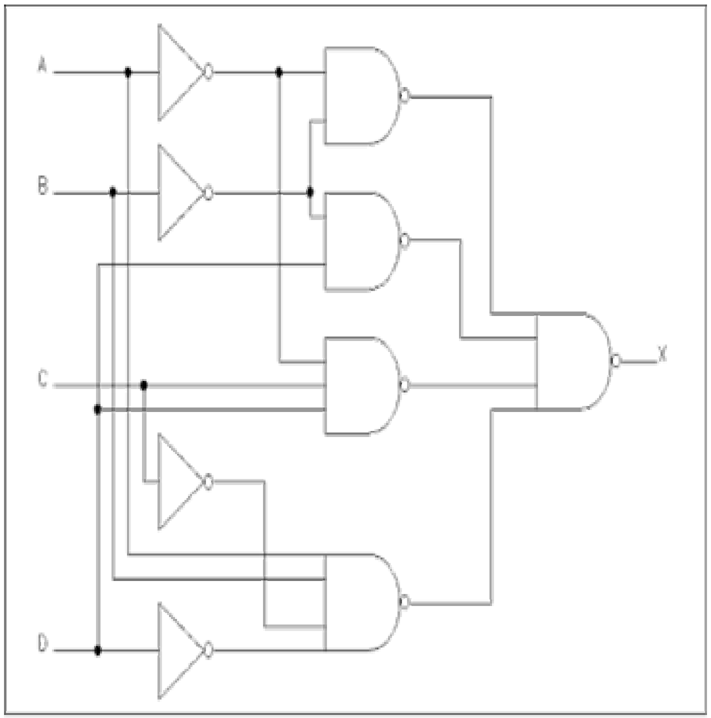

module r1 1 3 30 40
module r2 40 40 10
20 module r3 0 40
10 10 terminal a r1
0 0 terminal b r1 30
20 terminal c r3 0 0
terminal d r2 10 10
net N1 r1 a r3 c
net N2 r1 b r2 d

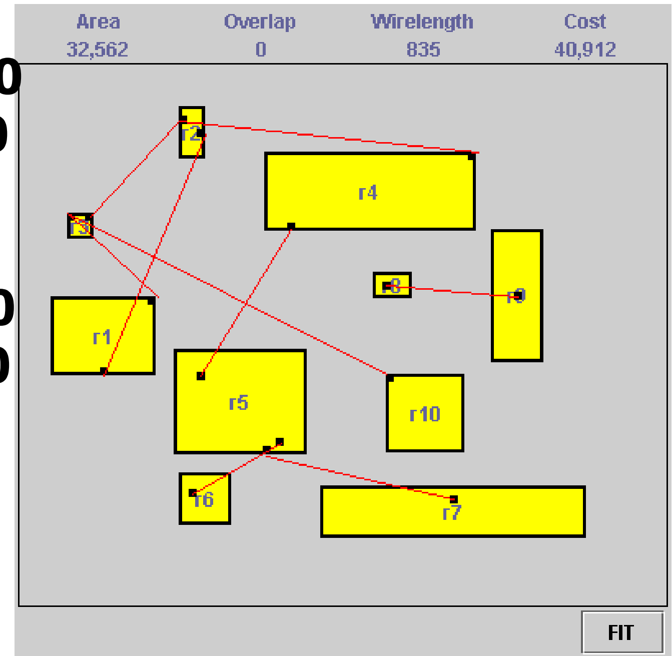

Logical Specification

Geometrical Specification

Hypergraph $H = (V,E)$ model

A NET (hyperedge) is a subset of interconnected CELLS (vertices).

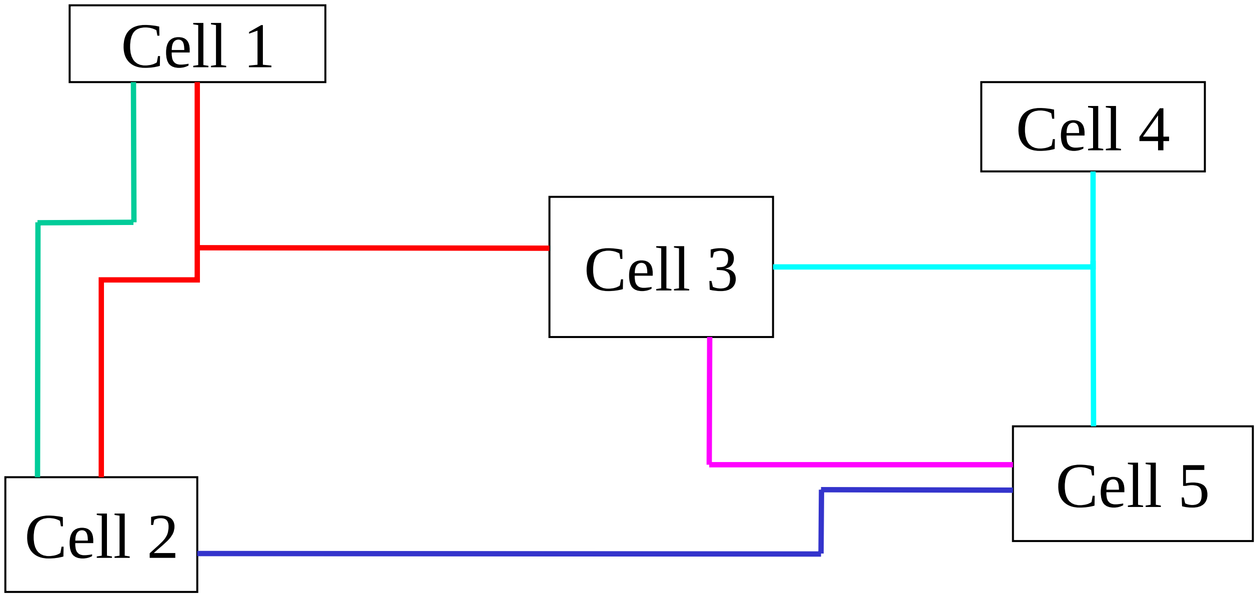

$$\sum_{e \in E} \max_{v,w \in e} (|x(v) - x(w)| + |y(v) - y(w)|)$$

Placement problem: **arrange the cells to minimize total wirelength (= sum of the half perimeter wirelength in each nets).**

Good Placement vs. Bad Placement

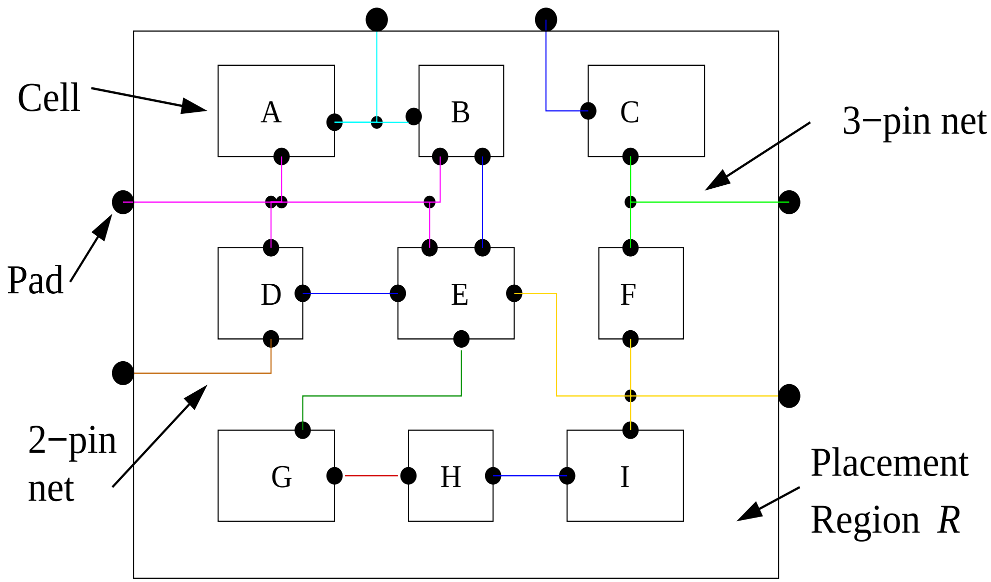

Good

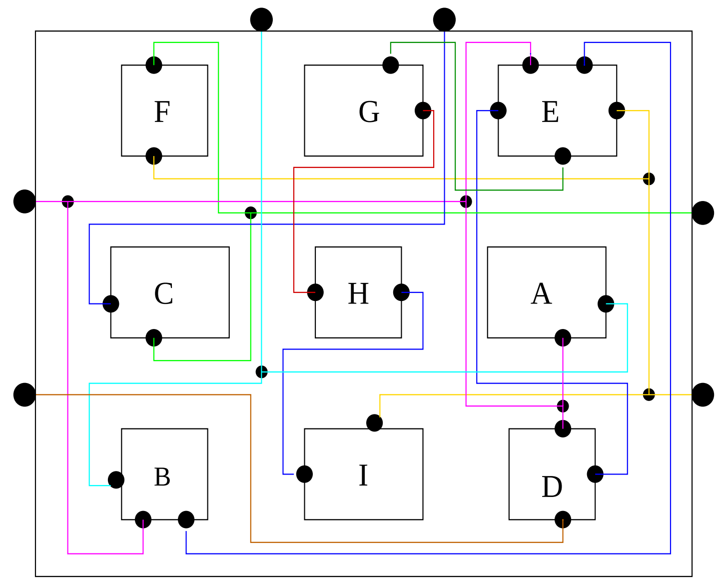

Bad

The Circuit Placement Problem

Given:

- ***N*** **circuits, a.k.a "blocks," "modules," or "cells"**
- **A rectangle ("the chip") in which the circuits must be placed *without overlapping***
- **Connectivity specs (a hypergraph "netlist")**
- **Constraints, e.g., timing, heat dissipation, routability**

Problem: Find an arrangement of the circuits on the chip that minimizes total wirelength subject to all constraints above.

Difficulty:

- **Non-convex and non-differentiable objective function**
- **Modeling all $O(N^2)$ non-overlap constraints when $10^4 \leq N \leq 10^7$**

Smooth Approximation of the Objective Function (Wirelength)

Approx. $$\sum_{e \in E} \max_{v,w \in e}(|x(v)-x(w)|+|y(v)-y(w)|)$$

by $$\sum_{e \in E} \sum_{v,w \in e}[|x(v)-x(w)|^k+|y(v)-y(w)|^k]$$

- **Squared Euclidean distance (k=2) (Quadratic Wirelength)**
  - **Advantage: Involve solving a sparse positive definite linear system.**
  - **But too much penalty on long nets.**
- **Manhattan distance (k=1) (Linear Wirelength)**
  - **Advantage: Good approximation**
  - **But need to solve a sequence of weighted quadratic programming.**

Smoothing the Placement Domain

Two steps in placement:

- Global Placement
  - Relax the placement region. Allow overlapping.
- Detailed Placement
  - Put the cells in standard rows by preserving the global placement as much as possible.

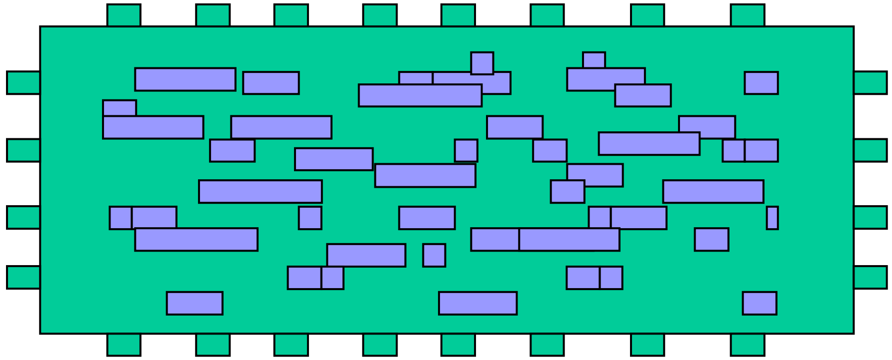

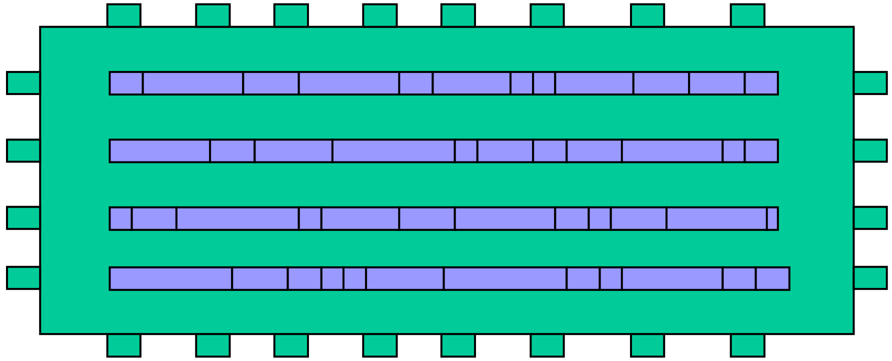

Nonlinear Programming Formulation

$$\begin{aligned} &\textbf{min} \quad f(x) \\ &\textbf{subject to} \quad c(x) \geq 0 \qquad\qquad \text{(NP)} \end{aligned}$$

where $x \in \mathrm{R}^n$.

- f: $\mathrm{R}^n \rightarrow \mathrm{R}$ **objective function** ($n \approx 2N$ **or** $3N$)
- c: $\mathrm{R}^n \rightarrow \mathrm{R}^m$ **constraint functions** ($m \approx N(N\text{-}1)/2 + N$)
- $F \equiv \{x \in \mathrm{R}^n \mid c(x) \geq 0\}$ **feasible region**
- x^* **local solution to NP (KKT conditions)**
- **Assumption:** f **and** c **are "smooth"**
- **Difficulty:**
  - ***Active set*** $A \equiv \{i \mid c_i(x^*) = 0\}$ **is unknown**
  - **A lot of local minimizers**
  - **N is large**

Why Use Multilevel for Placement Problem?

- ***Better Scalability*:**

  need to solve placement problems with millions of cells

- ***Better Global Optimization*:**

  need to find a good local minima in the placement problem

Motivation and Related Work
Multilevel Methods in Scientific Computation

- **Originally developed to solve boundary-value partial differential equation (PDE) problems on continuous field**
- **Discretized elliptic PDE is a structured, positive-definite system of linear equations**
  - **Multi-grid method**
  - **Algebraic Multi-grid (AMG)**

Multilevel Methods in VLSICAD

- **Successfully applied to solve hypergraph partitioning problem:**
  - **Hmetis [G. Karypis 1998]**
  - **MLpart [C. Alpert, J. Huang, A. Kahng 1998]**
- **Our goal:**
  - **Want to apply the Multilevel ideas to solve placement problem directly, not as an equation solver.**

Our Multilevel Placement Framework

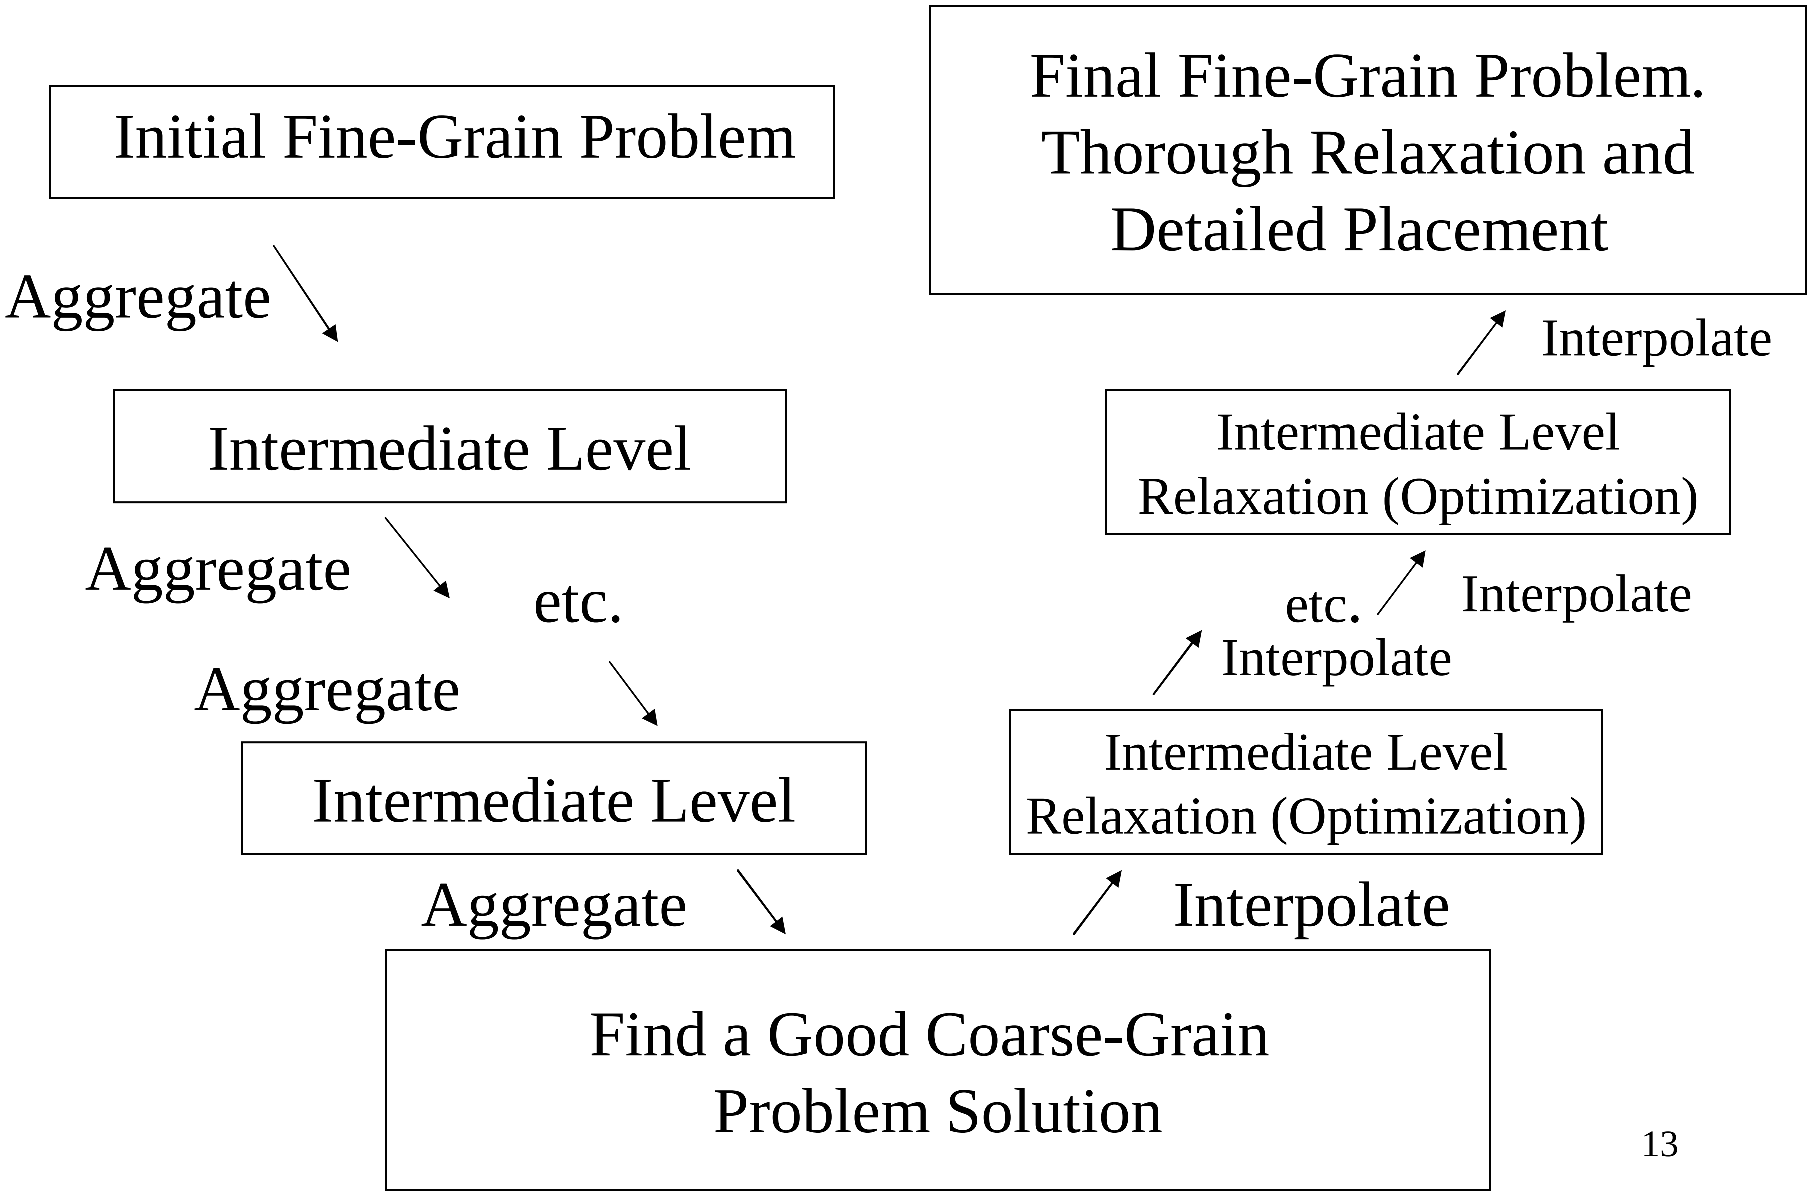

Main Components in Multilevel Framework

1. **Create coarser problems [*aggregation/coarsening/clustering*]**
2. **Optimize coarser problems [*relaxation/smoothing*]**
3. **Transform coarser problem solution to finer level**

   [*Interpolation/declustering*]

Challenge: Blend PDE-based & VLSI-specific

Overview of Our Multilevel Placement

- Coarsening ←
  - Modified First Choice clustering
- Relaxation (Intralevel Optimization)
- Interpolation
- Iterated Multilevel Flow

Coarsening by Recursive Aggregation
First Choice Aggregation [Karypis, 1999]

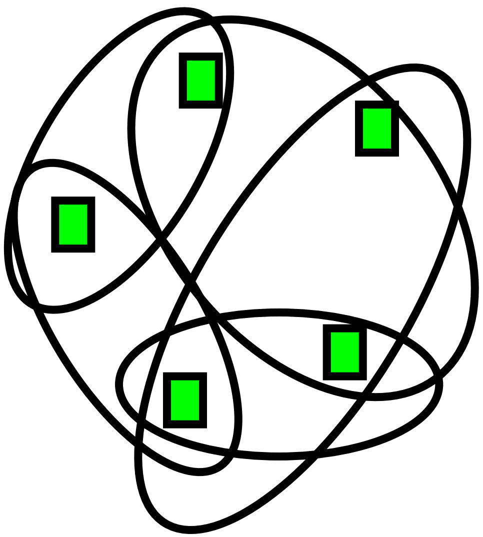

Transform the hypergraph to clique model graph using the weight 1/(|e|-1)

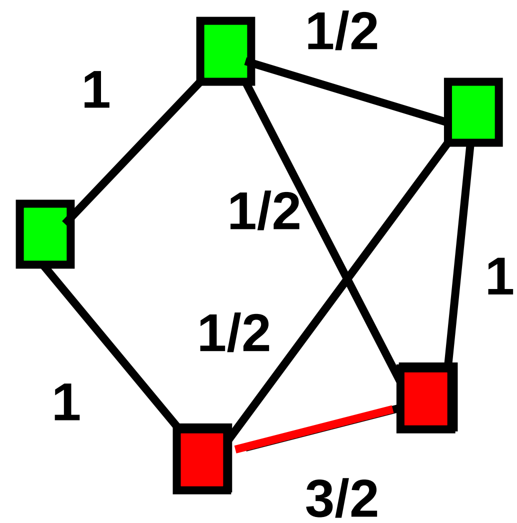

Match each vertex with a neighboring vertex with which it shares the most total hyperedge weight

mPL Coarse-Level Formulation

[Chan, Cong, Kong, Shinnerl; ICCAD 2000]

- **Nonlinear-Programming Formulation**
  - **Direct formulation for the coarse placement problem**
  - **Cells are modeled as circular disks for smoothness**
  - **Quadratic wirelength objective on a clique-model**
  - **Pairwise nonoverlap constraints**
- **Reasonable performance for coarse-level sizes N <= 500 only.**

Overview of Our Multilevel Placement

- **Coarsening**
- **Relaxation (Intralevel Optimization) ←**
  - **Quadratic relaxation on subsets (QRS) + bounded domain ripple-move to relieve area congestion**
- **Interpolation**
- **Iterated Multilevel Flow**

Quadratic Relaxation on Noncontiguous Subsets (QRS)

- **Select a subset *M* of cells to move.**
- **Identify other cells and pads, *F*, connected to *M* by nets in**

$$E_M = \{ e \in E \mid e \cap M \neq \phi \}.$$

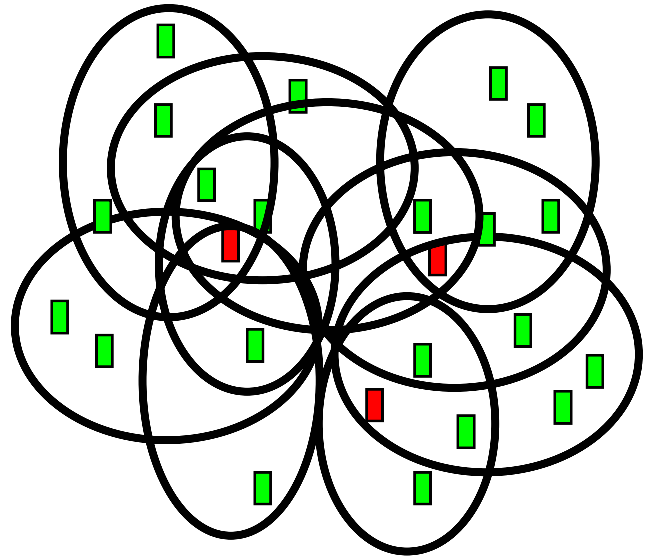

Movable cell

fixed cell

M is obtained as segments of length 3 along a DFS vertex traversal of the netlist, where starting the DFS at a vertex connected with largest wirelength.

Solving the Subproblem

- **Problem formulation (horizontal case):**

$$\min \sum_{e \in E_M} \sum_{v \in e} \frac{(x(v) - x_e)^2}{|x^{(k)}(v) - x_e^{(k)}| + \varepsilon}$$

$$\text{where } \; x_e = \frac{1}{|e|} \sum_{v \in e} x(v) \, , \;\; \varepsilon \; \text{ is small number.}$$

- **Iteratively solve the *weighted quadratic* minimization problem, using the current solution to determine the weight (Gordian-L).**

Ripple-move Legalization [Hur and Lillis, 2000]

Because QRS ignores overlap constraints, post-QRS cell swaps are used to remove the area congestion.

Bound the searching region for scalability

Define a DAG on neighboring bins. Edge cost reflects the best wirelength gain over all cell swaps between two bins.

Calculate a max-gain monotone path on the bin-grid graph

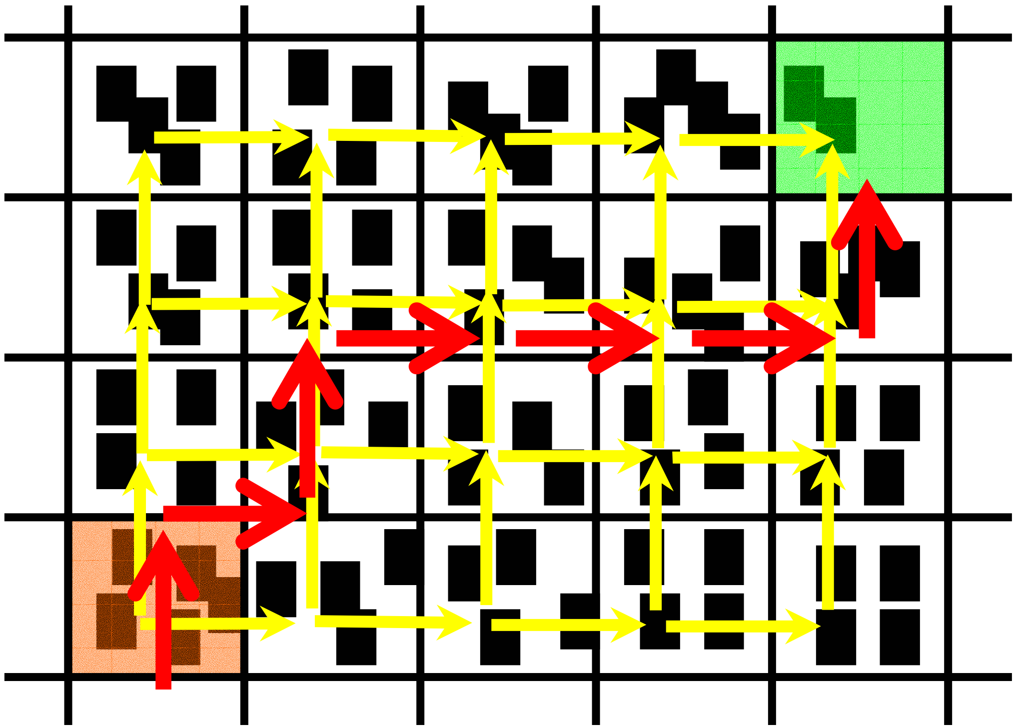

Overview of Our Multilevel Placement

- Coarsening
- Relaxation (Intralevel Optimization)
- Interpolation ←
  - AMG-based weighted disaggregation
- Iterated Multilevel Flow

AMG-based Linear Interpolation [A. Brandt 1986]

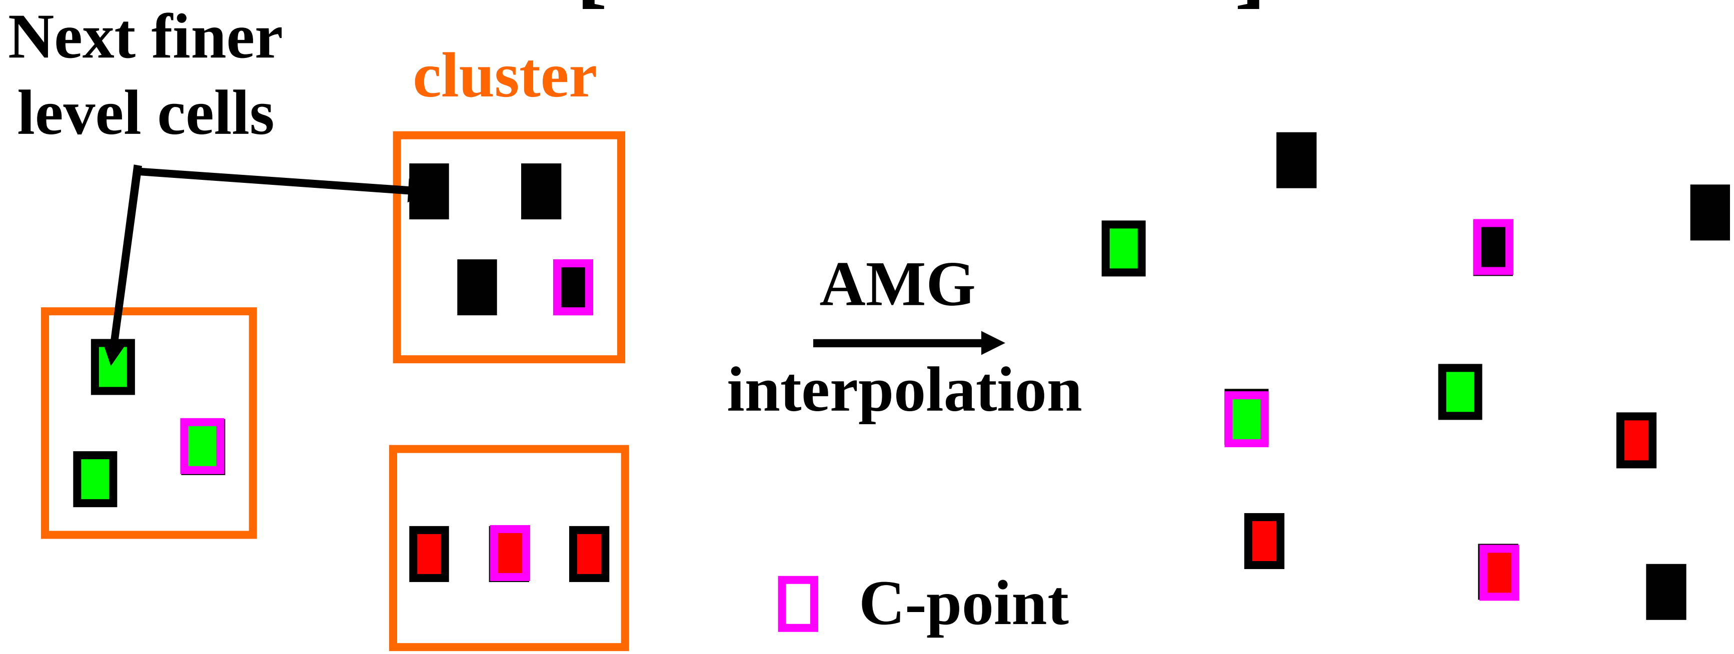

Within each cluster, select the one with *maximum degree* as C-point; others are considered as F-points

$$\boxed{v_i = \sum_{C-\text{points } v_j} a_{ij} v_j + \sum_{F-\text{points } v_j} a_{ij} v_j}$$

AMG-based Interpolation

$$\begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{pmatrix}$$

- **Use the clique-model graph to define connectivity weights *(connectivity matrix)***
- **Within each cluster, select the one with *maximum degree* as a C-point**
- **Each C-point is placed at the cluster's position.**
- **Each F-point is placed at the weighted average of the C-points *and F-points* to which it is *strongly* connected**
- **The F-points' positions can be iteratively improved.**

Overview of Our Multilevel Placement

- **Coarsening**
- **Relaxation (Intralevel Optimization)**
- **Interpolation**
- **Iterated Multilevel Flow ←**

Iterated Multilevel Flow

Make use of placement solution from 1st V-cycle

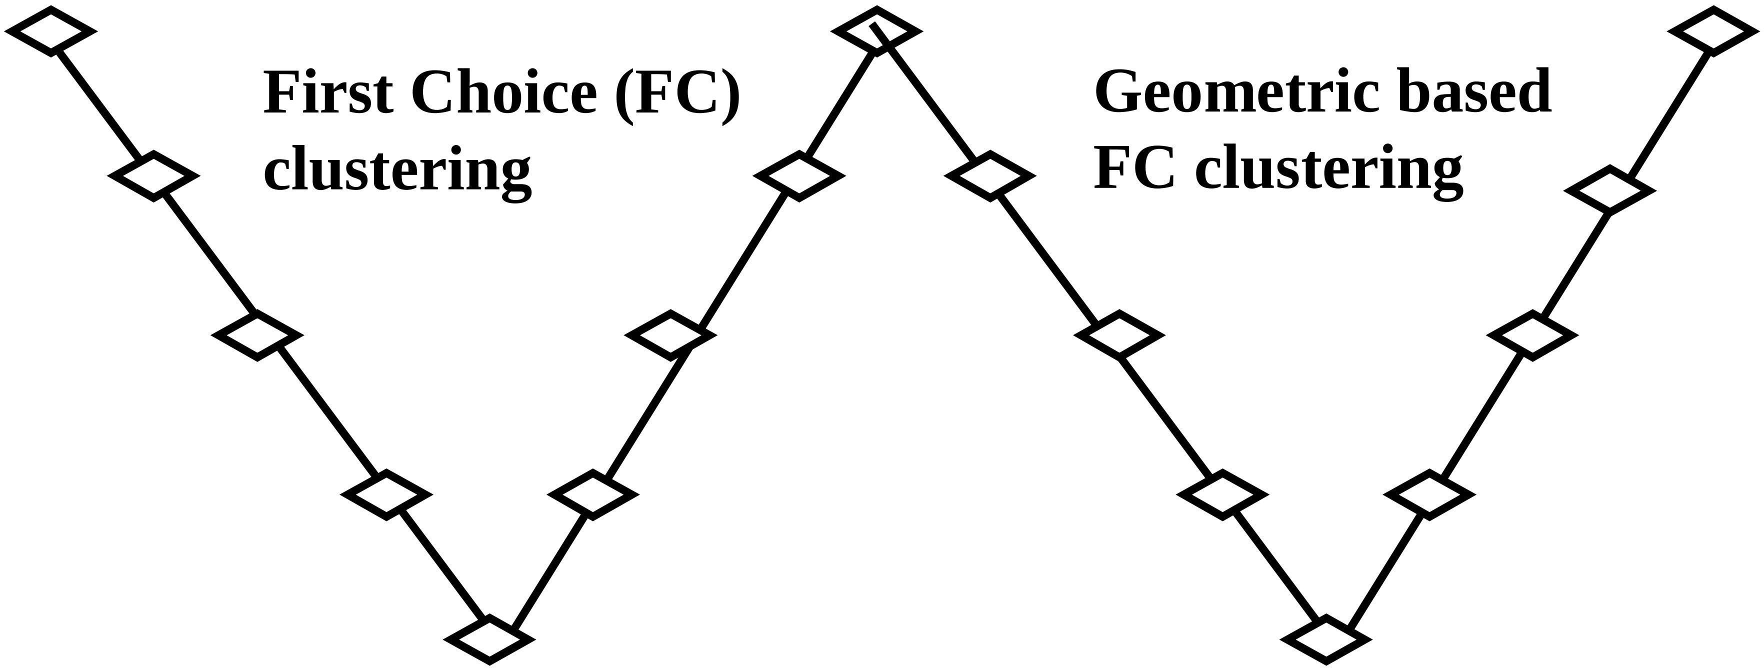

Adjustable Vertex Affinity for Re-aggregation

- **First V-cycle affinity:**

$$r_{ij} = \sum_{\{e \in E | i, j \in e\}} \frac{w(e)}{(|e|-1)\mathrm{area}(e)}$$

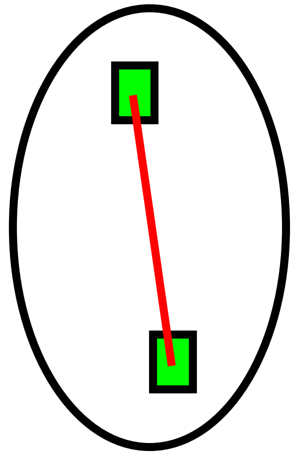

- **Next V-cycle affinity (distance is incorporated):**

$$r_{ij} = \sum_{\{e \in E | i, j \in e\}} \frac{w(e)}{(|e|-1)\mathrm{area}(e) d(v_i, v_j)}$$

Placement: pre-1997 State of the Art

- Simulated Annealing-based methods (SA) can handle complex design constraints, but its runtime does not scale well (Timberwolf/iTools).
- Quadratic Programming-based methods (QP) are very efficient, but they cannot handle complex constraints well (Gordian-L)
  - Force-directed methods (Kraftwerk)
- Recursive Bipartitioning was not competitive with QP and SA.

Multilevel Hypergraph Partitioning

- **1997: hMetis (Karypis et.al.), MLpart (Caldwell et al.)**
- **The first widely successful implementation of multilevel hypergraph partitioning.**
- **10x speedup or more; improved scalability; improved cutsize**
- **Very influential on current algorithms for placement**
  - **Dragon (Sarrafzadeh et al., 2000): top-down 4-way partitioning by hMetis with wirelength improvement at each stage by simulated annealing**
  - **Capo (Kahng et al., 2000) Recursive multilevel hypergraph bipartitioning with carefully chosen cutlines and branch-and-bound on base cases (<= 30 cells).**

mPL3.0 vs. mPL1.0, Capo8.5, Dragon and Gordian-L

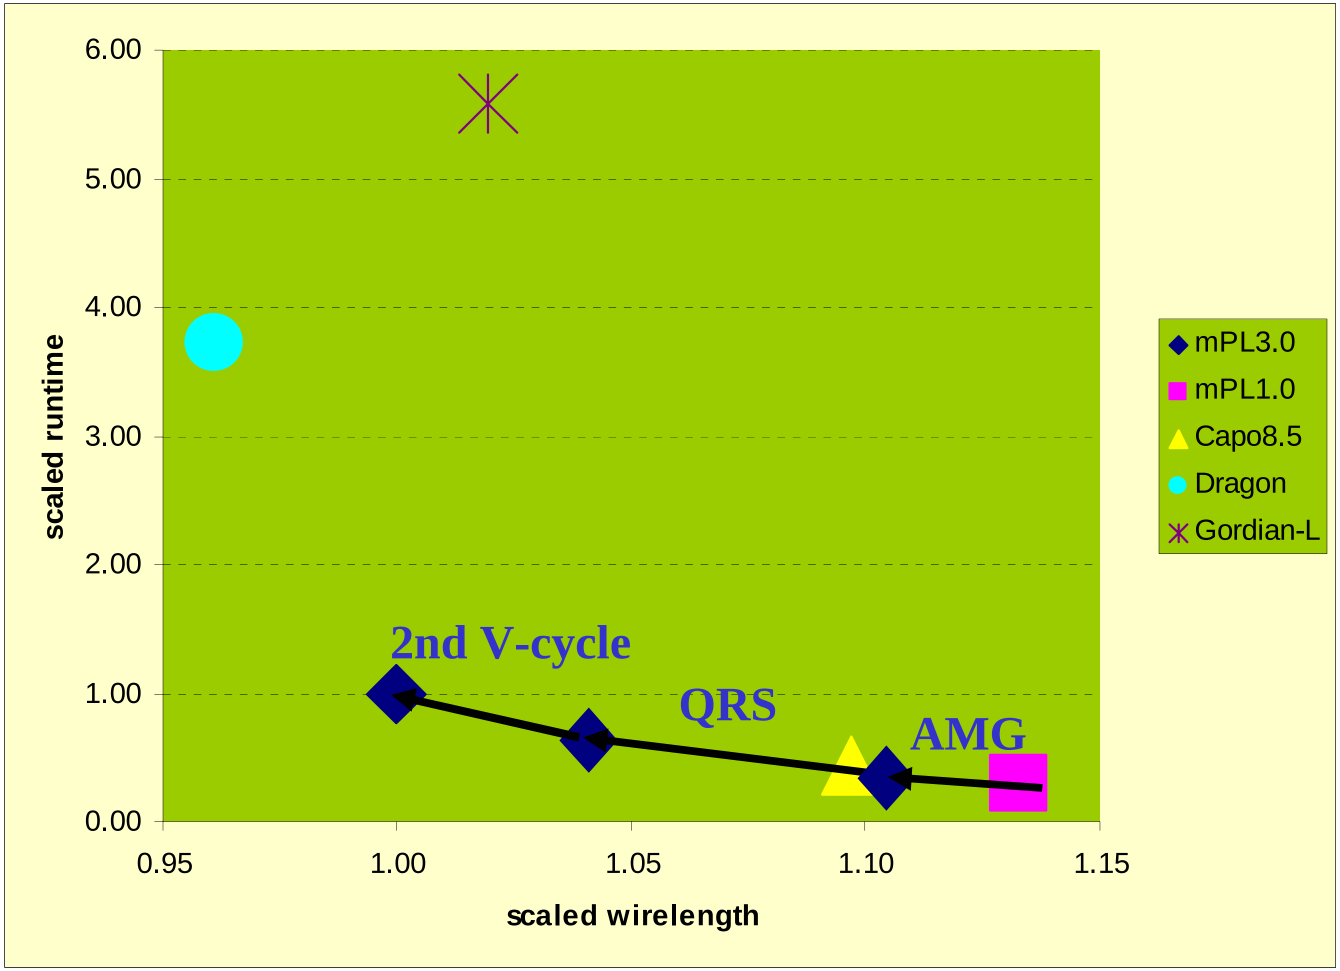

Experimental Results on PEKO and PEKU [C. Chang, J. Cong, M. Romesis, M. Xie, 2003]

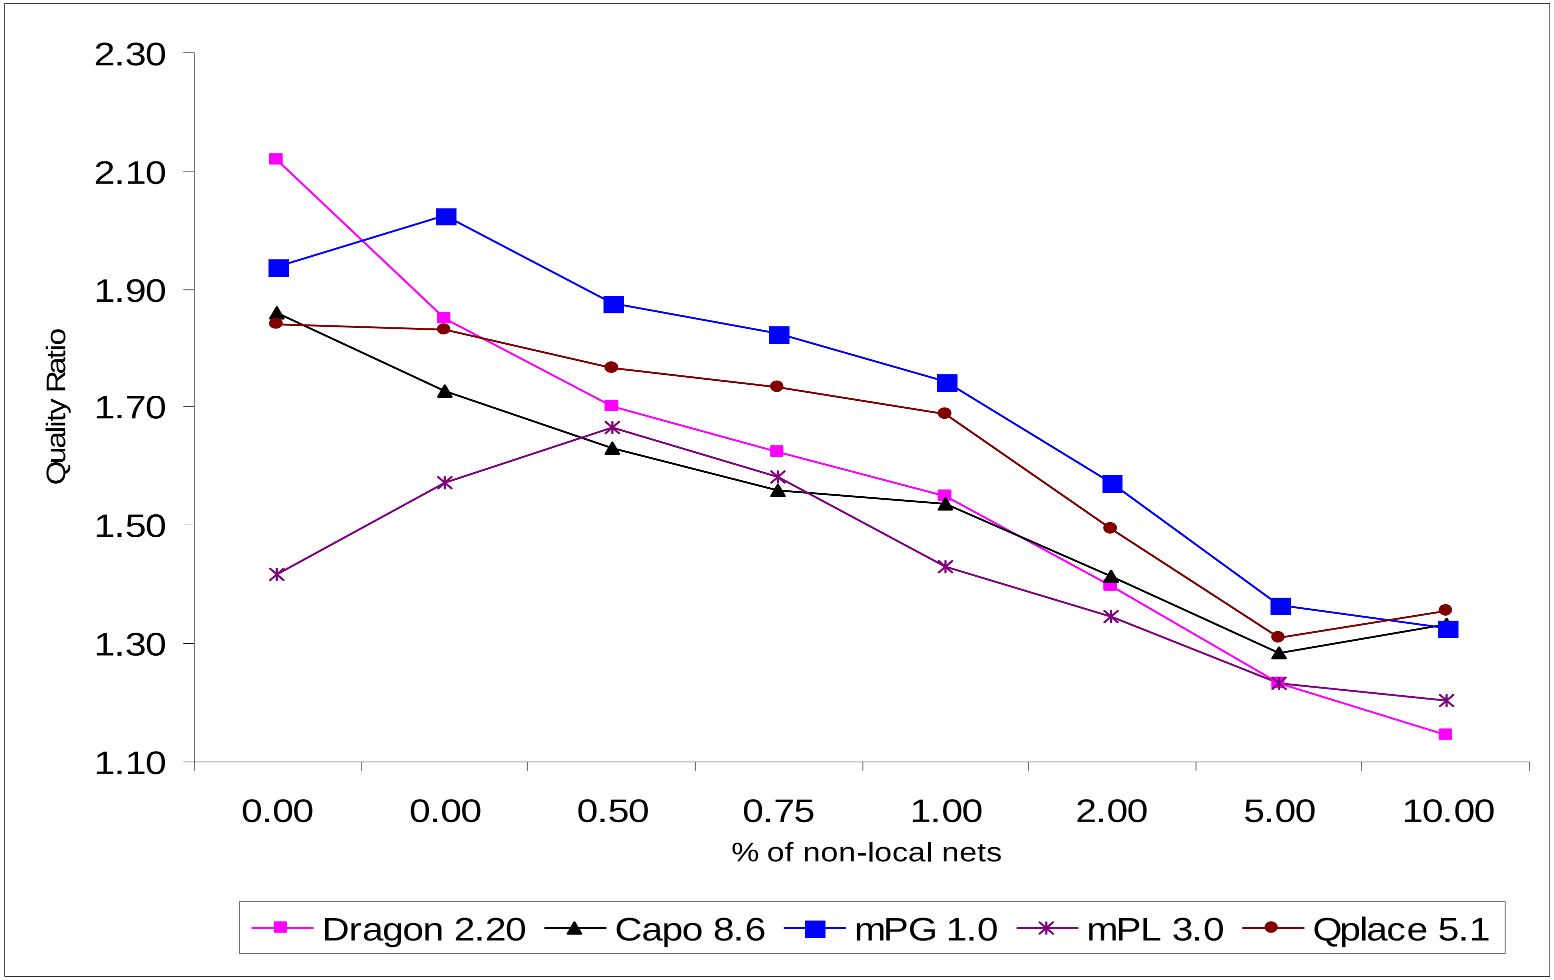

mPL3.0 outperforms other placers on PEKO and keep relative good performance when % of global nets increased

Some Open Issues

- **More complex design objectives and constraints**
  - **clock frequency, routability, heat dissipation, widely varying cell sizes, …**
- **AMG theory/algorithms for optimization on hypergraphs**
  - **control #hyperedges at coarser levels**
  - **continuous vs. discrete relaxations**
- **Multi-level FAS optimization framework**

Relevant Annual Conferences all are associated with IEEE/ACM

- **DAC: Design Automation Conference**
- **ICCAD: International Conference on Computer-Aided Design**
- **ISPD: International Symposium on Physical Design**
- **ASPDAC: Asia South-Pacific Design Automation Conference**

Institute for Pure & Applied Mathematics

IPAM

University of California, Los Angeles

Institute For Pure and Applied Mathematics
University of California, Los Angeles presents a workshop in

Multilevel Optimization in VLSICAD
Computer-Aided Design of Very-Large-Scale Integrated Circuits

December 3—5, 2001

Achi Brandt	Weizmann Institute of Science, Appl. Math. & CS
Jason Cong	UCLA, CS
Ding-zhu Du	University of Minnesota, CS&E
Stephan Hartmann	TU Berlin, Math
Bruce Hendrickson	Sandia National Labs, Parallel CS
Michael Lewis	College of William & Mary, Math
George Karypis	University of Minnesota, CS&E
John Lillis	University of Illinois Chicago, EECS
Stephen Nash	George Mason University, Systems Eng. & Op. Res.
Majid Sarrafzadeh	UCLA, CS
Lieven Vandenberghe	UCLA, EE
Chandu Visweswariah	IBM Research, Design Automation
Chris Walshaw	University of Greenwich, Computing and Math. Sci.
Jacob White	MIT, EECS
Gabriel Wittum	University of Heidelberg, IWR (Sci. Comp.)
Martin Wong	University of Texas at Austin, CS

Members of the organizing committee include Achi Brandt (Weizmann Institute of Science, Israel), Jason Cong (UCLA), and Joseph R. Shinnerl (UCLA)

The program is open to the entire mathematical, computer science, and engineering scientific communities.

Please visit our website for more information, including an online registration form:

http://www.ipam.ucla.edu/programs/mov2001

You can email questions to mov2001@ipam.ucla.edu

kluwer
the language of science

Shopping Cart

About us | Contact | Conference Schedule | Order | Browse by Subject | Customer Service

Search

Title

advanced search

search tips

Products & Services

- Kluwer online
- Journals
- Books
- Looseleafs
- CD-ROMs
- Kluwer Alert

Resources

- Authors
- Libraries

Books » Multilevel Optimization in VLSICAD

Multilevel Optimization in VLSICAD

Add to cart

edited by

Jason Cong

Dept. of Computer Science, University of California, Los Angeles, USA

Joseph R. Shinnerl

Dept. of Computer Science, University of California, Los Angeles, USA

Book Series: COMBINATORIAL OPTIMIZATION : Volume 14

In the last few decades, multiscale algorithms have become a dominant trend in large-scale scientific computation. Researchers have successfully applied these methods to a wide range of simulation and optimization problems. This book gives a general overview of multiscale algorithms; applications to general combinatorial optimization problems such as graph partitioning and the traveling salesman problem; and VLSICAD applications, including circuit partitioning, placement, and VLSI routing. Additional chapters discuss optimization in reconfigurable computing, convergence in multilevel optimization, and model problems with PDE constraints.

Audience: Written at the graduate level, the book is intended for engineers and mathematical and computational scientists studying large-scale optimization in electronic design automation.

Kluwer Academic Publishers, Dordrecht
Hardbound, ISBN 1-4020-1081-8
December 2002 , 320 pp.
EUR 103.00 / USD 98.00 / GBP 66.00

HOME | HELP SECTION | ABOUT US | CONTACT US | CONFERENCE SCHEDULE | ORDERING INFORMATION | BROWSE BY SUBJECT | CUSTOMER SERVICES | SEARCH | PRIVACY POLICY

Book: http://www.wkap.nl/prod/b/1-4020-1081-8