Large-Scale Circuit Placement: Gap and Progress

EDPS Workshop; April 6–7, 2005; Monterey, CA.

Tony F. Chan
Kenton Sze

Mathematics Department, UCLA
Los Angeles, CA 90095-1555
{chan,nksze}@math.ucla.edu

Jason Cong
Joseph Shinnerl
Min Xie
Computer Science Department, UCLA
Los Angeles, CA 90095-1596
{cong,shinnerl,xie}@cs.ucla.edu

The presentation reviews results from some recent optimality studies and recent progress in the UCLA mPL placer.

1 Placement Examples with Known Optima (PEKO)

Recently, the PEKO studies have given evidence of how far computed placement solutions are from optimal and how much the deviation from optimality is likely to grow with respect to problem size. The original PEKO examples have all cells of equal size. Each PEKO is constructed to match the net-degree distribution of a given benchmark exactly, but with all nets defined to support a placement in which each net is purely local and has the minimum possible half-perimeter wirelength.

Four state-of-the-art placers from academia including Dragon [12], Capo [3], mPL [4], mPG [6], and one industrial placer, QPlace [2] were compared on PEKO. Overall, their computed wirelengths were 1.59 to 2.40 times the optimal in the worst cases and are 1.43 to 2.12 times the optimal on average. As for scalability, the average solution quality of each tool deteriorates by an additional 9% to 17% when the problem size increases by a factor of 10.

The PEKO studies present strong evidence that the quality gap between globally optimal and computed solutions remains large. Although recent work suggests that the gap is shrinking fast on the smaller benchmarks ($10^4 \leq N \leq 10^5$), with specialized approaches [11, 13] achieving quality ratios as low as 10%, the gap remains large on larger test cases. Internal unpublished studies show that simple perturbations of the PEKO test cases, e.g., aggregating large sets of standard cells into large macro blocks, create cases for which the quality gap may exceed 8×.

In addition to quantifying suboptimality, these benchmarks provide clues useful in targeting algorithm weaknesses. Recent work by Ono and Madden [11] examines placements computed by existing tools on PEKO and shows that displacements of cells from known optimal locations may all be small, even when corresponding total wirelength is more than 30% longer than optimal. While this result might be seen as indicating a need for improved algorithms for detailed placement, we will show displacements whose global regularity suggests otherwise.

2 Progress in UCLA mPL5

The mPL5 framework [5] makes several improvements to the Kraftwerk framework for generic force-directed placement [7] and embeds this improved engine in a multilevel-optimization flow. A bounding-box weighted half-perimeter wirelength objective is approximated by the log-sum-exp model [10, 9], in which $W(\mathbf{x}, \mathbf{y}) =$

$$\gamma \sum_{\text{nets } e \in E} \left(\log \sum_{\text{nodes } v_k \in e} \exp(x_k/\gamma) + \log \sum_{v_k \in e} \exp(-x_k/\gamma) \; + \; \log \sum_{v_k \in e} \exp(y_k/\gamma) + \log \sum_{v_k \in e} \exp(-y_k/\gamma) \right),$$

where $\mathbf{x}$ and $\mathbf{y}$ denote vectors of cell x- and y-coordinates. The smaller the parameter γ, the more accurate the approximation. For simplicity, restrict attention to area density constraints only. These are explicitly imposed separately in each rectangular bin of a uniform grid laid over the placement region. Let d_j denote the cell-area density of bin B_j and let K denote total cell area divided by the total placement area. The

area-density constraints are initially expressed simply as $D_j = K$ over all bins B_j. Viewing the D_j as a discretization of a continuous density function $d(x, y)$ defined at points $(x, y) \in R$, these constraints are smoothed by approximating d by the solution ψ to the Helmholtz equation

$$\begin{cases} \Delta\psi(x,y) - \epsilon\psi(x,y) = d(x,y), & (x,y) \in R \\ \frac{\partial\psi}{\partial\nu} = 0, & (x,y) \in \partial R \end{cases} \tag{1}$$

where $\epsilon > 0$, ν is the outer unit normal, ∂R is the boundary of the placement region R, $d(x, y)$ is the continuous density function at a point $(x, y) \in R$, and Δ is the Laplacian operator $\Delta \equiv \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$. The smoothing operator $\Delta_\epsilon^{-1} d(x, y)$ defined by solving (1) is well defined, because (1) has a unique solution for any $\epsilon > 0$. As the solution of (1) has two more derivatives [8, e.g.] than $d(x, y)$, ψ is a smoothed version of d. Discretized versions of (1) can be solved rapidly by fast numerical multilevel methods. Recasting the density constraints as a discretization of ψ gives the nonlinear programming problem

$$\min\ W(\mathbf{x}, \mathbf{y}) \qquad \text{s.t.} \quad \psi_j = -K/\epsilon, \quad 1 \le i \le m, 1 \le j \le n,$$

where the ψ_j are obtained by solving (1) with the discretization defined by the given bin grid. Interpolation from the adjacent coarser level defines a starting point. This nonlinear-programming problem is solved by the Uzawa iterative algorithm [1], which does not require second derivatives or large linear-system solves.

Experiments indicate that mPL5's results are extremely competitive in both total wirelength and speed and scalability.

References

[1] K. Arrow, L. Huriwicz, and H. Uzawa. *Studies in Nonlinear Programming*. Stanford University Press, Stanford, CA, 1958.

[2] Cadence Design Systems Inc. Envisia ultra placer reference. In *http://www.cadence.com*, QPlace version 5.1.55, compiled on 10/25/1999.

[3] A.E. Caldwell, A.B. Kahng, and I.L. Markov. Can recursive bisection produce routable placements? In *Proc. 37th IEEE/ACM Design Automation Conf.*, pages 477–482, 2000.

[4] T.F. Chan, J. Cong, T. Kong, and J. Shinnerl. Multilevel optimization for large-scale circuit placement. In *Proc. IEEE International Conference on Computer Aided Design*, pages 171–176, San Jose, CA, Nov 2000.

[5] T.F. Chan, J. Cong, and K. Sze. Multilevel generalized force-directed method for circuit placement. In *Proc. Int'l Symp. on Phys. Design*, 2005.

[6] C.C. Chang, J. Cong, Z. Pan, and X. Yuan. Physical hierarchy generation with routing congestion control. In *Proc. ACM International Symposium on Physical Design*, pages 36–41, San Diego, CA, Apr 2002.

[7] H. Eisenmann and F.M. Johannes. Generic global placement and floorplanning. In *Proc. 35th ACM/IEEE Design Automation Conference*, pages 269–274, 1998.

[8] L. C. Evans. *Partial Diferential Equations*. American Mathematical Society, 2002.

[9] A.B. Kahng and Q. Wang. Implementation and extensibility of an analytic placer. In *Proc. Int'l Symp. on Phys. Design*, pages 18–25, 2004.

[10] W. Naylor et al. Nonlinear optimization system and method for wire length and delay optimization for an automatic electric circuit placer, Oct. 2001.

[11] S. Ono and P.H. Madden. On structure and suboptimality in placement. In *Asia South Pacific Design Automation Conf.*, Jan 2005.

[12] M. Wang, X. Yang, and M. Sarrafzadeh. Dragon2000: Standard-cell placement tool for large industry circuits. In *Proc. International Conference on Computer-Aided Design*, pages 260–264, 2000.

[13] Q. Wang, D. Jariwala, and J. Lillis. A study of tighter lower bounds in LP relaxation based placement. In *Proc. Great Lakes Symposium on VLSI*, Chicago, IL, Apr 2005.