A Tale of Two Nets: Studies in Wirelength Progression in Physical Design

Andrew B. Kahng   Sherief Reda
CSE Department
University of CA, San Diego

This talk addresses two main questions

- Are our wirelength optimizations consistent across the physical design flow?

- How can we measure and report similarity of different placements?

Outline

- Motivation
- Definitions of consistency and similarity
- Empirical findings:
  - Wirelength estimates vs. HPWL
  - HPWL vs. Steiner length
  - Steiner length vs. routed length
  - The big picture
- Implications to physical design

The length of a net can have many values during the physical design flow

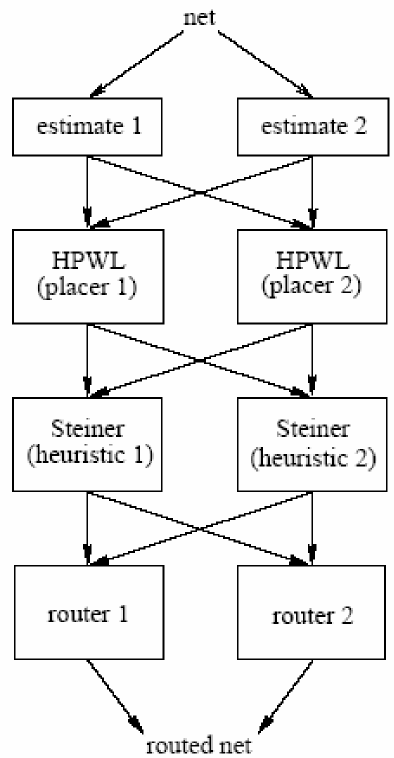

Technical inability to find optimal solutions

$\rightarrow$ Divide physical design into a number of suboptimal steps

$\rightarrow$ Use a heuristic (at each step) to optimize a *compromised* objective instead of the ideal objective

1. How reasonable is it to optimize for a compromised objective?

- Logic synthesis optimizes estimates (WLM, fanout) instead of routed wirelength
- Placement optimizes HPWL (or WWL) instead of routed wirelength
- Steiner tree construction methods optimizes individual nets
- Global routing versus detailed routing

2. How similar are our results if we use different tools / metrics?

- How different are results of the same tools?
- How different are results of different tools?
- Does it matter which estimate to use?

→ Similarity increases tool interoperability and supports incremental changes

Outline

- Motivation
- Definitions of consistency and similarity
- Empirical findings:
  - Wirelength estimates vs. HPWL
  - HPWL vs. Steiner length
  - Steiner length vs. routed length
  - The big picture
- Implications to physical design

Ideal measuring of consistency is not technically feasible

- Place a design optimally according to HPWL and then route it optimally $\Rightarrow WL_1$
- Simultaneously place and route a design optimally according to routed WL $\Rightarrow WL_2$
- Consistency: How far is WL_1 from WL_2?

We indirectly measure consistency by tracing lengths of pairs of nets

- Pick a random pair of nets *(i, j)* with the same number of pins
- Trace their lengths during physical design

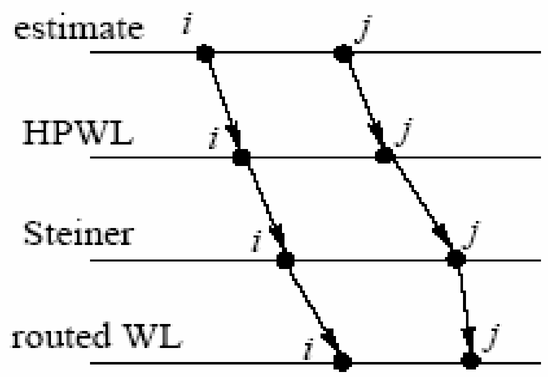

Consistent optimization

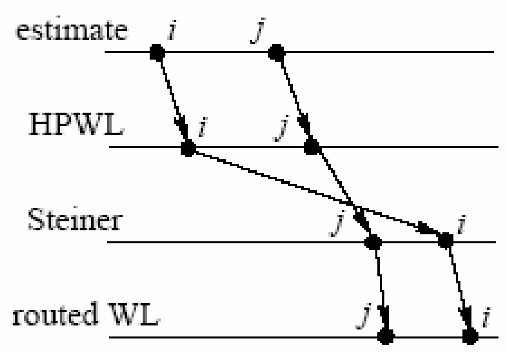

Inconsistent optimization

<u>Definition:</u> If *a* and *b* are two stages: a pair of nets are *consistent* if and only if $l_a(i) \geq l_a(j) \Leftrightarrow l_b(i) \geq l_b(j)$

Consistency = fraction of pairs of nets that their relative lengths do not change between two physical design stages

Achieving a consistency of 50% is trivial

- Given a netlist and its placement:
  - Let E be any set of estimate values randomly assigned to the nets
  - Let P be the lengths of the nets from the given placement
- Tracing relative lengths between E and P

  $\Rightarrow$ 50% of the time, pairs of nets would have the same relative lengths in both E and P

How much increase in consistency we can get if some tolerance is allowed?

Definition:

If *a* and *b* are two stages: a pair of nets i and j is *consistent* with some tolerance *tol* if and only if

$$l_a(i) \geq l_a(j) \Leftrightarrow l_b(i) \geq l_b(j) \quad \text{or} \quad \frac{|l_b(i) - l_b(j)|}{\max(l_b(i), l_b(j))} \leq tol$$

$\Rightarrow$ Consistency with tolerance allows us to tell whether small changes in wirelength are responsible for inconsistency

Placement is a means towards an end; similarity should be based on just WL

Given two outputs *q, p* (placements / routings) and some net *i*: compare the length of *i, l(i),* between *p* and *q*

Similarity $$s_{p,q}(i) = \frac{\min(l_p(i), l_q(i))}{\max(l_p(i), l_q(i))}$$

Our notion of similarity satisfies three properties: [Alpert et al. ASPDAC'05]

Reflexive: $s_{p,p}(i) = 1$

Symmetry: $s_{p,q}(i) = s_{q,p}(i)$

Triangle Inequality: $s_{p,q}(i) s_{q,t}(i) \le s_{p,t}(i)$

Outline

- Motivation
- Definitions of consistency and similarity
- Empirical findings:
  - Wirelength estimates vs. HPWL
  - HPWL vs. Steiner length
  - Steiner length vs. routed length
  - The big picture
- Implications to physical design

Let’s start measuring how consistent and similar our optimizations are

- Estimators: Mutual Contraction (MC) and Intrinsic Shortest Path Length (ISPL)
- Placers: Capo9.3 and APlace2.0
- Steiner tree heuristics: FLUTE
- Routers: Cadence NanoRoute (V 4.10)
- Benchmarks: IBMV2 benchmarks (easy suite

Study Plan

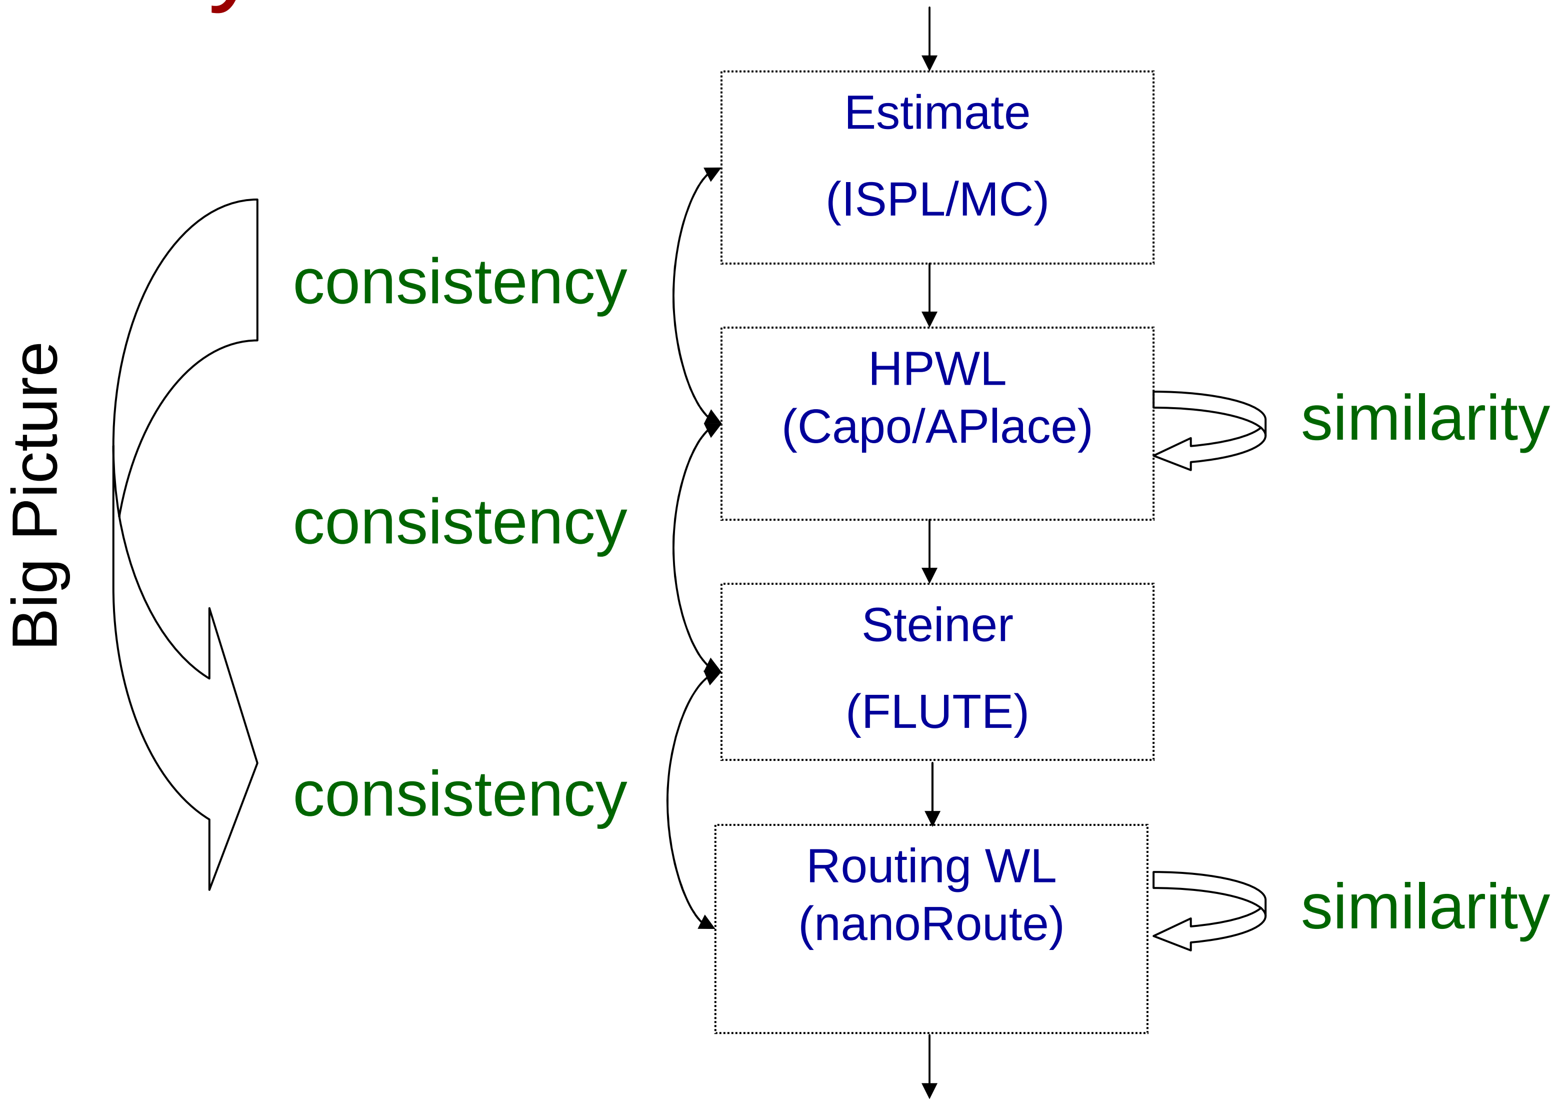
Big Picture
consistency
consistency
consistency
Estimate
(ISPL/MC)
HPWL
(Capo/APlace)
Steiner
(FLUTE)
Routing WL
(nanoRoute)
similarity
similarity

A grading system assigns grades to various consistency/similarity ranges

Consistency / Similarity range	Grade
90%-100%	Excellent
80%-90%	Very good
70%-80%	Good
60%-70%	Average
50%-60%	Poor

Consistency between estimates and HPWL is average/poor

MC versus HPWL

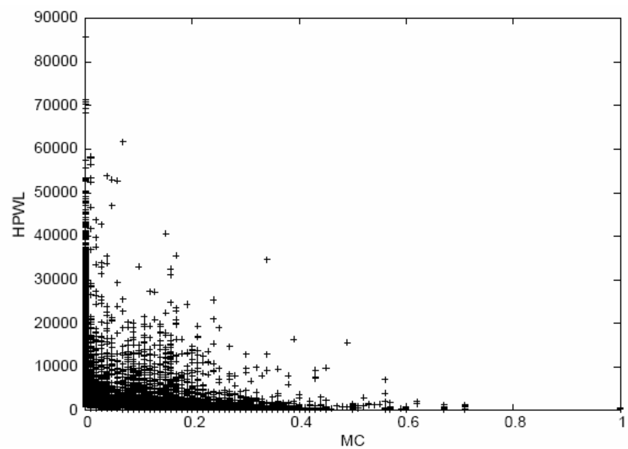

ISPL versus HPWL

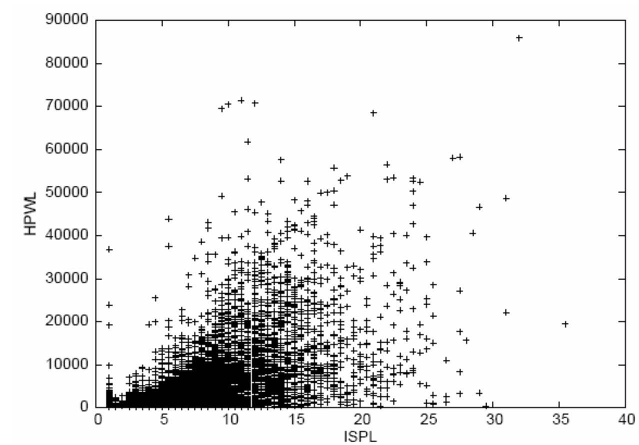

Using Capo

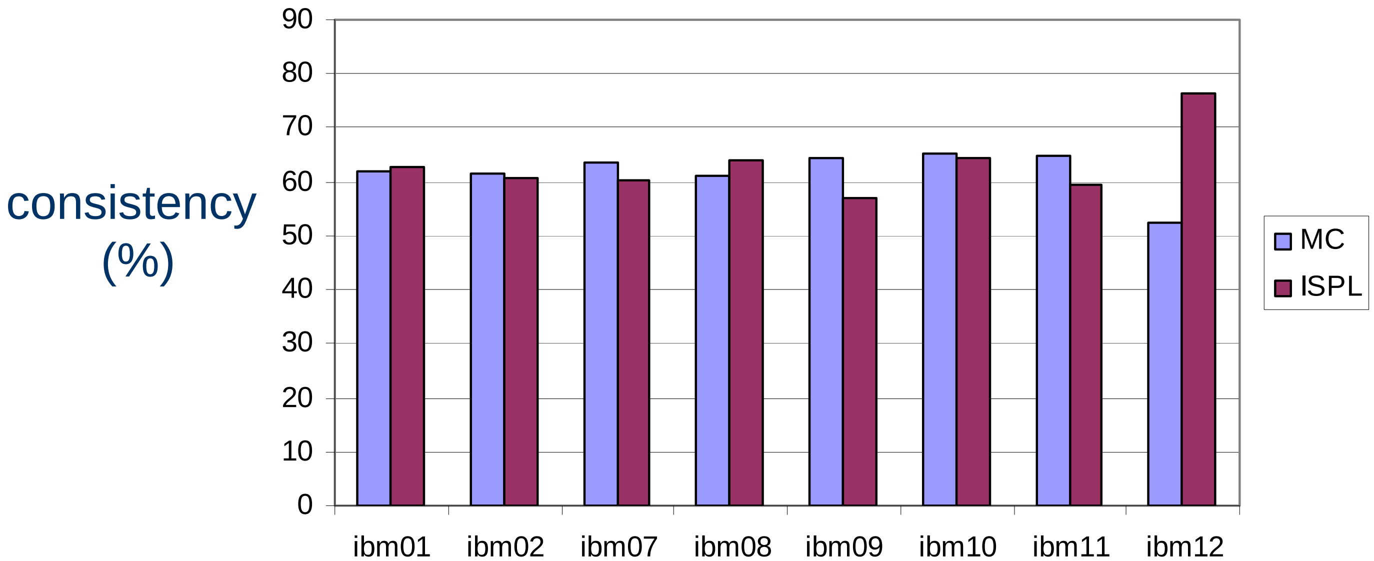

Consistency between estimates increases with increased tolerance

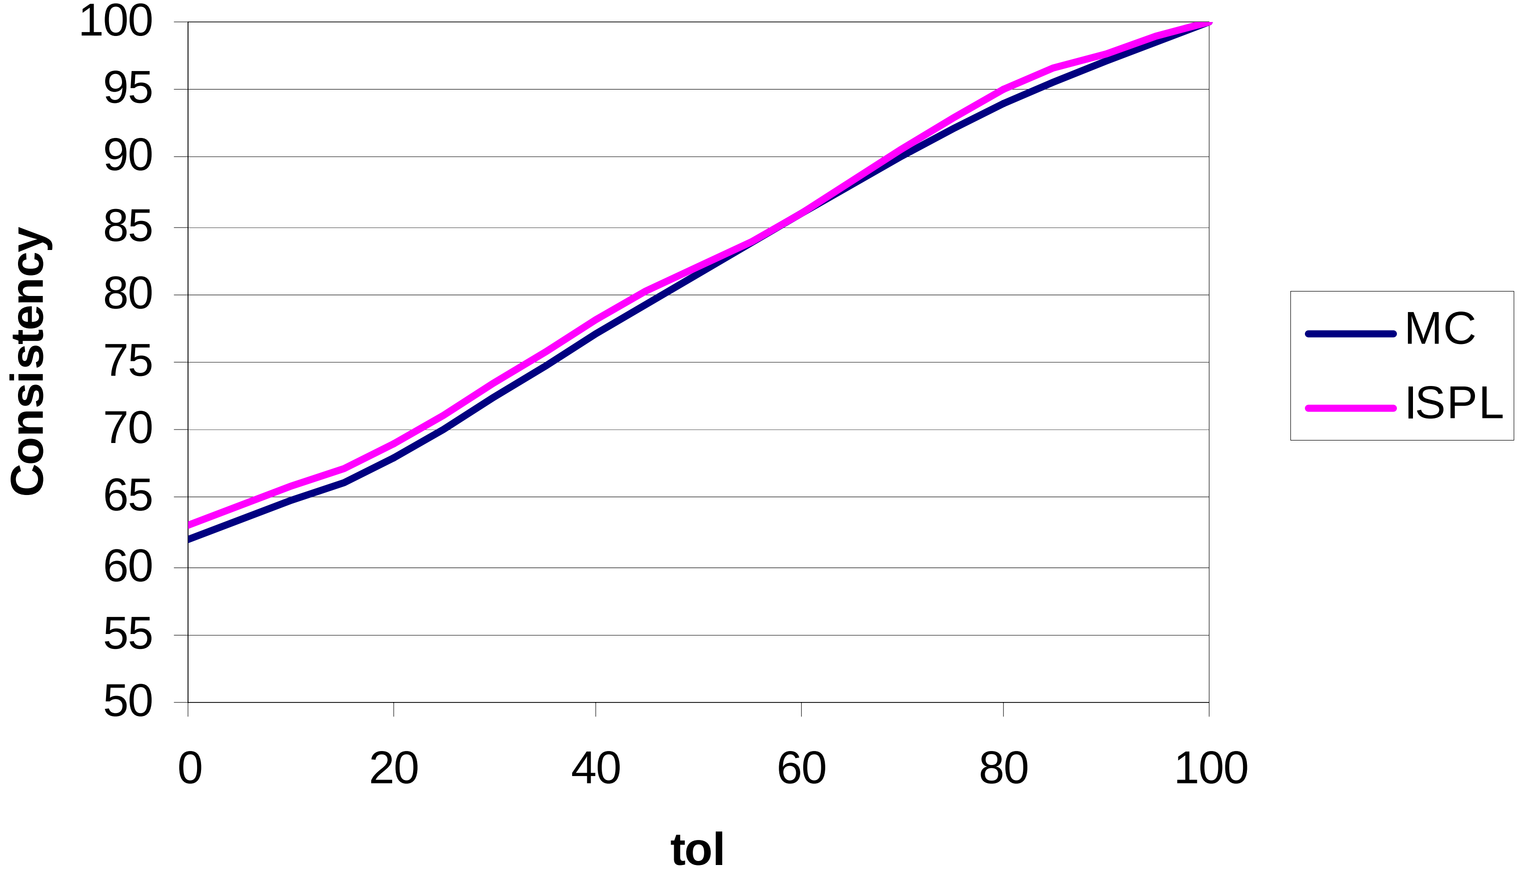

Consistency between HPWL and Steiner length is excellent

circuit	All nets Consistency	Degree ≥ 4 nets
ibm01	99.81%	96.24%
ibm02	99.72%	96.83%
ibm07	99.91%	97.60%
ibm08	99.89%	97.32%
ibm09	99.9%	96.60%
ibm10	99.82%	97.35%
ibm11	99.86%	97.27%
ibm12	70.21%	75.35%

Consistency between Steiner length and routed wirelength is very good

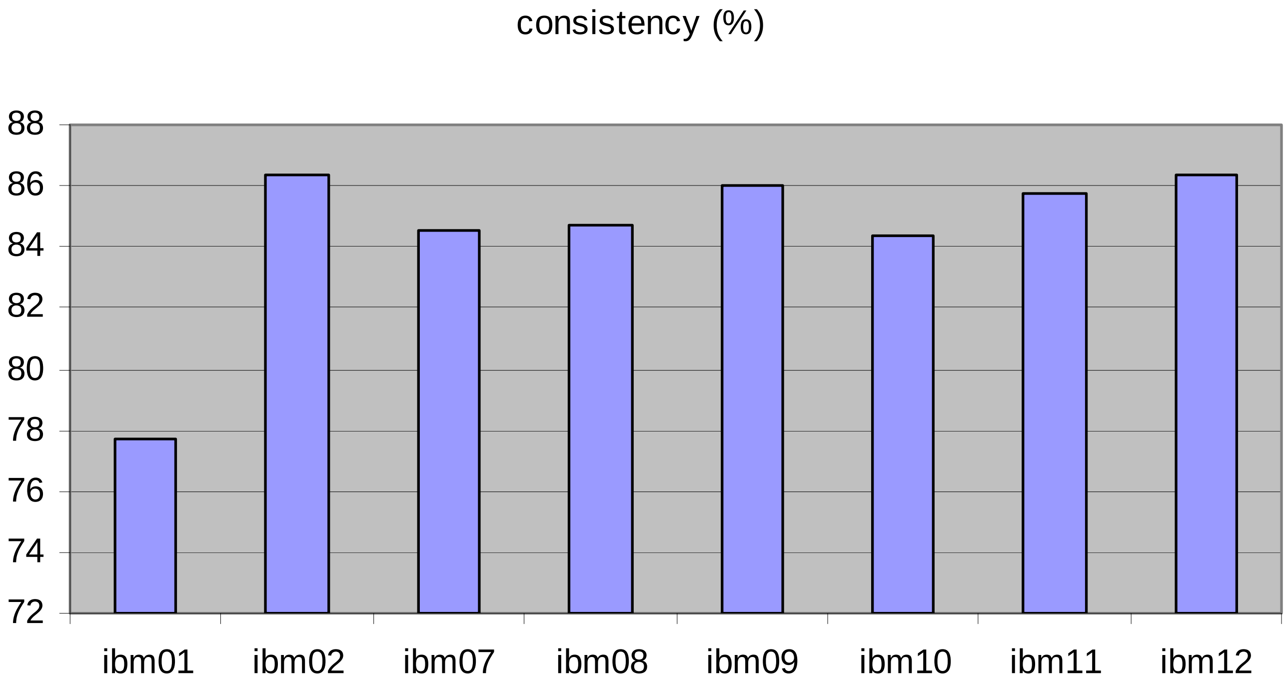

Adding tolerance improves consistency between Steiner length and routed WL

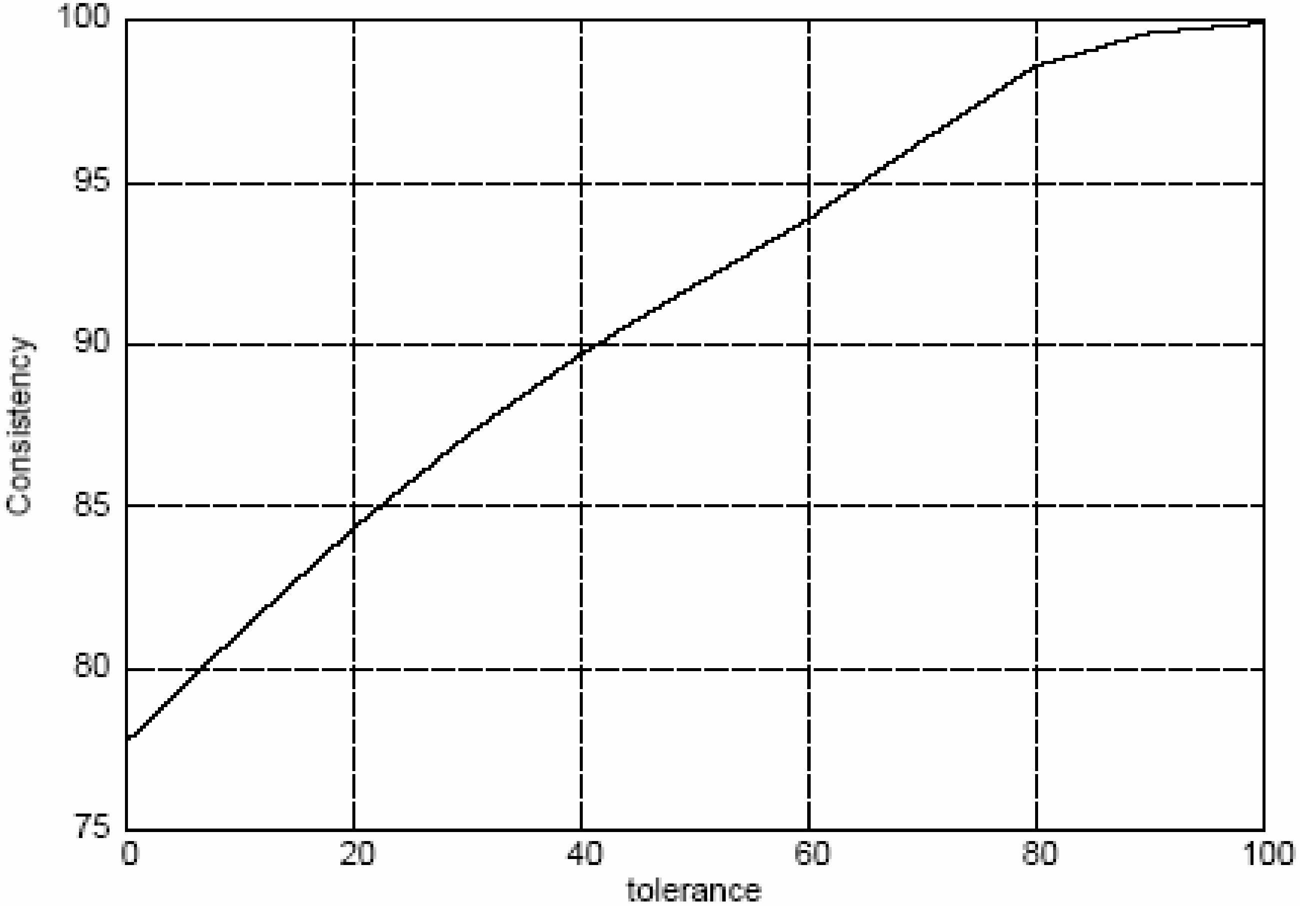

For ibm01 using Capo's placement

How good are Weighted Wirelength Objectives?

From Capo's placement:

- Pick two nets *i* and *j* with the same HPWL and number of pins (≥4)
- Calculate Steiner length *t(i)* and *t(j)* for both *i* and *j*
- Ideally, according to WWL, t(i) = t(j)
- Define skew=*max(t(i), t(j))/min(t(i), t(j))*

Let's look at the fraction of nets *p(x)* with skew between [1…x]

Skew in Steiner lengths between nets with the same HPWL and #pins

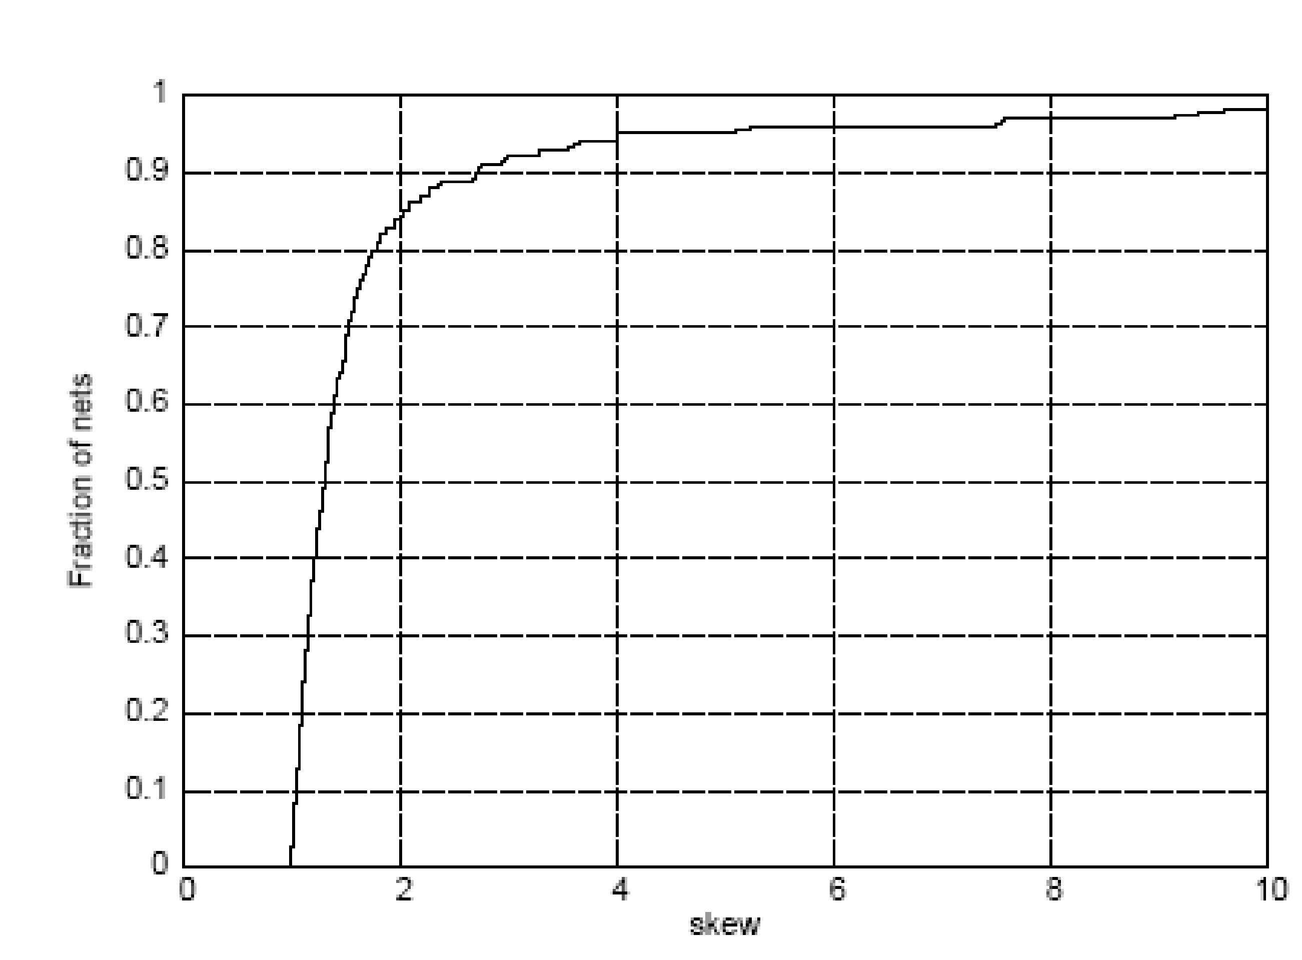

Most nets ~85% are just within skew < 2

The big picture: consistency across more than two consecutive stages

	HPWL	Steiner	Routed WL
Estimate	Poor/average	Poor	Poor
HPWL	-	Excellent	Very good / good
Steiner	-	-	Very good

Placements of the same placer or different placers share little similarity

	APlace vs Capo	Capo vs Capo	APlace vs APlace
ibm01	64.92 %	64.61%	71.17%
ibm02	63.06%	65.10%	79.34%
ibm07	60.28%	59.79%	68.95%
Ibm08	62.38%	61.93%	65.31%
Ibm09	58.74%	61.54%	62.46%
ibm10	63.27%	64.49%	67.56%
ibm11	59.90%	60.95%	64.83%
ibm12	61.26%	61.25%	69.83%

Similarity between different routing results of different placers is poor

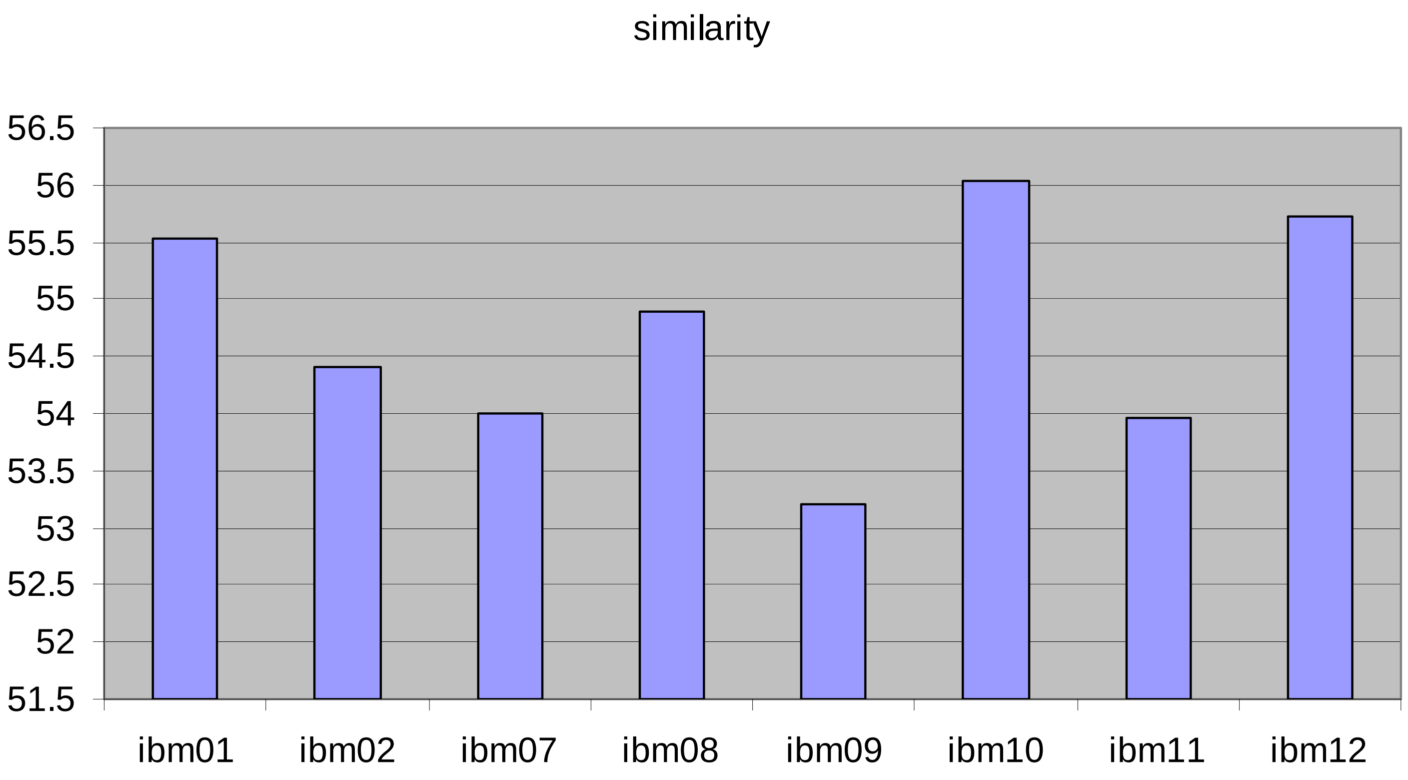

Outline

- Motivation
- Definitions of consistency and similarity
- Empirical findings:
  - Wirelength estimates vs. HPWL
  - HPWL vs. Steiner length
  - Steiner length vs. routed length
  - The big picture
- Implications to physical design

Implications for *a priori* WL estimators

- Do not trust individual *a priori* WL estimators with every net $\Rightarrow$ have poor consistency with WL
- Rely on a priori estimators to predict *long* interconnects [KahngR05 - ICCAD'05]
- Aggregating individual estimates smooths out variations and can successfully predict total WL [KahngR05 - ICCAD'05]

$\rightarrow$ More work needed in the area of *a priori* wirelength estimation

Implications for placers

- Continue using HPWL as an optimization objective
  - excellent consistency with Steiner length
- Find better ways to integrate routing into placement
- Find ways to make placements more similar (at least from the same placer)

Implications for Steiner length objectives and routers

- WWL sometimes has slight discrepancies
- Focusing on individual Steiner lengths as a way for improving routing results might not be the best strategy
- Integration of placement and global routing must be the standard framework

Questions / Answers

Definitions

- *A priori:* predicting HPWL before placement just by looking at the netlist
- Stability = similarity (as defined in slide 12)
- Consistency (defined in slides 9-10)
- WWL = weight factors for HPWL based on lookup tables indexed by the number of pins and HPWL of nets. They approximate Steiner length

Questions and Answers

- How do you explain that APlace shows so little similarity?
- APlace uses an initial random placement around the center to establish the gradients. This placement changes when the random seed changes
- How do you explain that estimates give such poor results?
- We have found relatively little similarity (~65%) between same-placer results. How can we expect something better from estimates?
- Why do you think GI can be well predicted with estimators?
- ISPL in particular can well detect GIs, because these interconnects tend to be quite structurally separated at the netlist

Questions and Answers

- What are the unexpected data found in your studies?
- We are surprised by how there is very poor consistency between a priori estimates and routed WL. We are surprised by excellent consistency between HPWL and Steiner length. I would say our focus should be improving the consistency between Steiner length and routing WL.

Questions and Answers

- If this would have been a “more perfect” paper, what would we have seen in it?
- We would like to see comparisons of different routers and of results of the same router. Perhaps we would also like to see how consistency between HPWL and routing changes if whitespace (WS) changes. WS increase is likely to improve consistency.
- We are also interested in seeing what would be the similarity of placements produced from industrial placers, since these should tend to be more stable especially to support ECO changes.
- We are also looking to study the consistency between Steiner WL and routed WL of a Steiner-tree driven placer