








National Science Foundation
WHERE DISCOVERIES BEGIN




Buffered Clock Tree Sizing for Skew Minimization under Power and Thermal Budgets




Krit Athikulwongse, Xin Zhao, and Sung Kyu Lim
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia, U.S.A.



Outline




- **Introduction**
- **Problem formulation**
- **Algorithm design**
- **Experiments and results**
- **Conclusions**




- **Sequential-Linear-Programming based clock sizing**
 - Low-power and skew minimization [Wang, etc. IEEE Trans. '05]
 - Process-variation-aware clock skew minimization under power budget [Guthaus, etc. DAC'06]
- **Thermal-aware clock network design**
 - Balanced thermal-aware skew under given thermal profiles [Cho, etc. ICCAD'05]
 - Thermal-aware bottom-up merging based on thermal sensitivity [Yu, etc. ISPD'07]
 - Thermal-aware 3D clock routing algorithm design [Minz, etc. ASPDAC'08]


Challenges



- **Wire/buffer sizing impacts on delay / thermal / power**
 - Impact of sizing on thermal variations can not be ignored
- **Thermal profiles are changed during optimization**
- **Consider power and thermal budget simultaneously**




Contributions



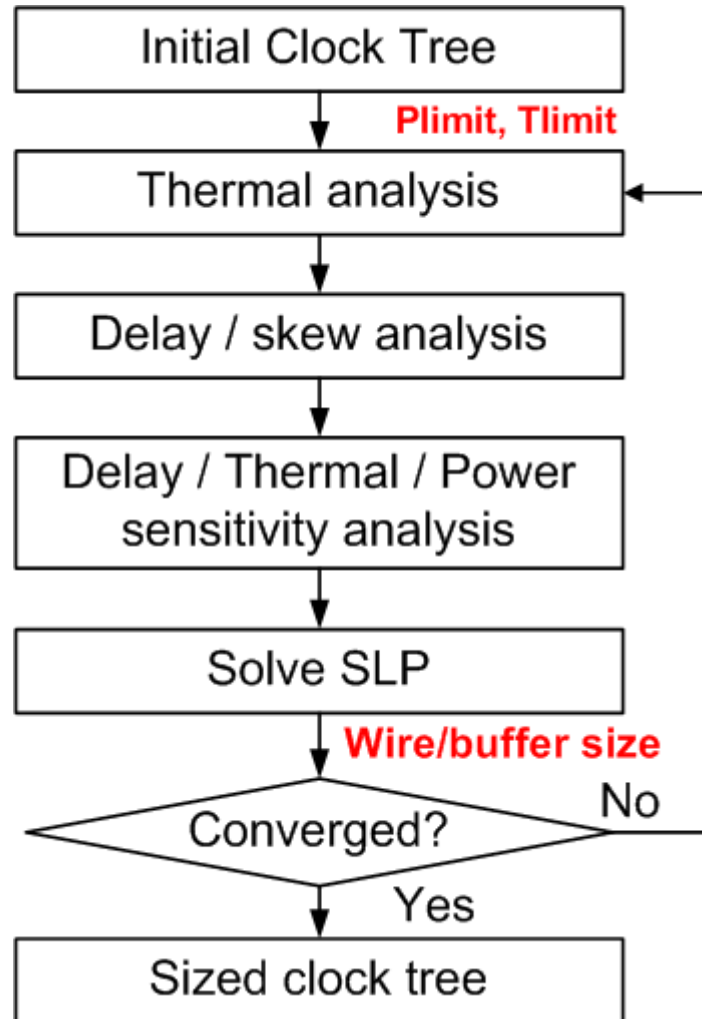

- **Thermal-Aware Sequential-Linear-Programming (TA-SLP)**
 - Thermal-aware clock skew minimization under power and thermal budgets
- **Single and multiple thermal profiles**
- **Thermal-aware skew or skew range minimization**
 - Minimize skew value under single steady-state thermal profile
 - Minimize skew range under multiple thermal profiles
- **Loosely-Thermal-Aware-Sequential-Linear-Programming (LTA-SLP)**
 - A speedup scheme to update delay, thermal and power sensitivity

Problem formulation: Thermal-aware clock sizing




- **Input:**
 - Non-sized buffered clock tree
 - Thermal profile (single or multiple)
 - Power and thermal budgets
- **Output:**
 - A sized-clock tree: wire widths and buffer sizes are determined
- **Object:**
 - Minimize thermal-aware clock skew (skew value or skew range)
- **Constraint:**
 - Power budget
 - Temperature budget

Overview: TA-SLP sizing flow



Thermal analysis



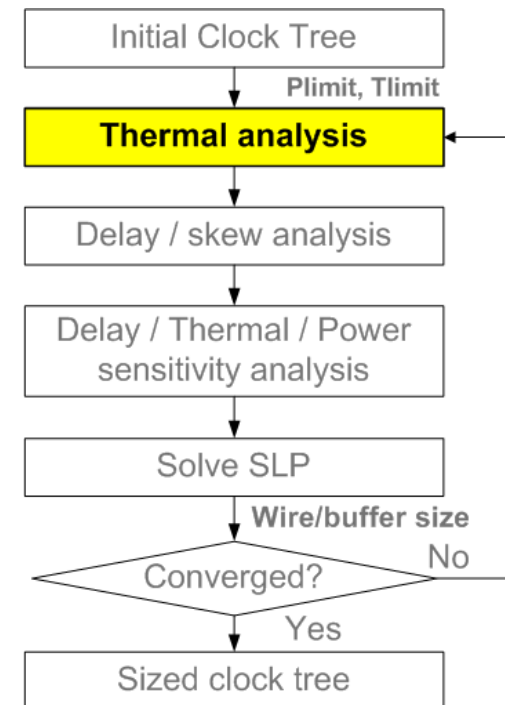
- **Power consumption**
 - Underlying power: Substrate devices
 - Clock power: Clock wires and buffers
- **Compact substrate thermal model***

$$\mathbf{t} = \mathbf{R} \times \mathbf{p}$$


$$R_{ij} = \partial T_i / \partial p_j$$

- **n x n thermal tiles**
- **R: thermal resistance matrix**
- **t, p: temperature and power vectors**

*C.-H. Tsai and S.-M. Kang. Cell-level placement for improving substrate thermal distribution. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 19(2):253–266, Feb. 2000.



Thermal-aware delay models




- **Thermal impact on wire model**
 - Temperature-dependent resistance


$$R(T) = \frac{\rho_0 \cdot l}{w \cdot t} \times (1 + \beta_0 \cdot T)$$

- **Thermal impact on buffer model**
 - Driving resistance
 - linearly depends on temperature within a small region.
 - Intrinsic delay with buffer size S and temperature T
 - 10x to 70x buffer size, at 60 to 110 °C, $\alpha = 0.005$, $\tau = 0.00013$, mean absolute error = 0.04ps.

$$t_d(T, S) = t_{d_0}(S) \times (1 + \beta_{in}(S) \cdot (T - T_0)),$$
$$\beta_{in}(S) = \alpha - S \cdot \tau,$$



Sensitivity computation



For a wire/buffer j , with w_j^0 changes to $w_j^0 + \epsilon_j$:

$$\hat{\mathbf{w}}_j^0 = [w_1^0 \ w_2^0 \ \cdots \ w_j^0 + \epsilon_j \ \cdots \ w_m^0]^T$$


- **G: Delay sensitivity matrix**
 - Delay variation caused by sizing

$$G_{ij} = \left. \frac{\partial d_i}{\partial w_j} \right|_{w_j = w_j^0} \approx [d_i(\hat{\mathbf{t}}_j^0, \hat{\mathbf{w}}_j^0) - d_i(\mathbf{t}^0, \mathbf{w}^0)] / \epsilon_j$$


$$d_i^1 = d_i^0 + \sum_{\forall j} G_{ij} \cdot \delta_j$$

Where $i \in$ sinks, $j \in$ wires/buffers

$$\hat{\mathbf{t}}_j^0 = \mathbf{t}(\hat{\mathbf{w}}_j^0)$$



Sensitivity computation (cont.)



- **Γ : Thermal sensitivity matrix**

- Average thermal variation of thermal tiles caused by sizing

$$\Gamma_{ij} = \left. \frac{\partial t_i}{\partial w_j} \right|_{w_j=w_j^0} \approx [t_i(\widehat{\mathbf{w}}_j^0) - t_i(\mathbf{w}^0)]/\epsilon_j$$

$$t_i^1 = t_i^0 + \sum_{\forall j} \Gamma_{ij} \cdot \delta_j$$

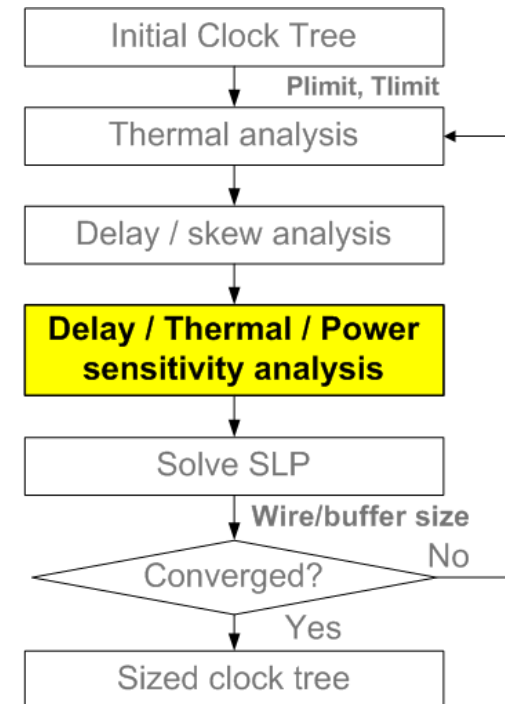
Where $i \in \text{tiles}$, $j \in \text{wires/buffers}$

- **β : Power sensitivity vector**


- Total power variation caused by sizing

$$\beta_j = \left. \frac{\partial p}{\partial w_j} \right|_{w_j=w_j^0} \approx [p(\widehat{\mathbf{w}}_j^0) - p(\mathbf{w}^0)]/\epsilon_j$$

Where $j \in \text{wires/buffers}$



Linear sub-problem in TA-SLP



- Given the current wire/buffer size w^0 , decide the size change δ

Minimize

$$d_{\max} - d_{\min}$$

Subject to

$$\mathbf{d}^1 = \mathbf{d}^0 + \mathbf{G} \cdot \delta$$

$$\mathbf{t}^1 = \mathbf{t}^0 + \mathbf{\Gamma} \cdot \delta$$

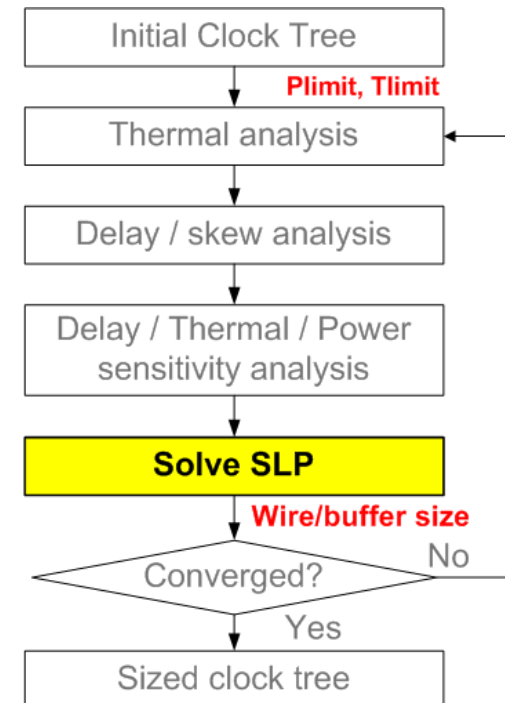
$$p^1 = p^0 + \beta^T \cdot \delta$$

$$d_i^1 \geq d_{\min}, \quad \forall i \in \text{sinks}$$


$$d_i^1 \leq d_{\max}, \quad \forall i \in \text{sinks}$$

$$t_j^1 \leq T_{\text{limit}}, \quad \forall j \in \text{tiles}$$

$$p^1 \leq P_{\text{limit}}$$



Wire/buffer size update and control



New size after sizing:

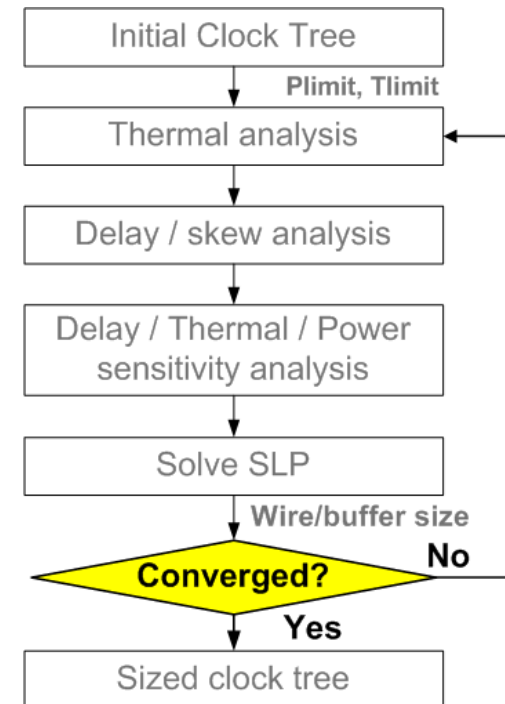
$$w_j^1 = w_j^0 + \delta_j$$

Restrict to:


$$w_j^1 \in [\max(L_j, w_j^0 - \epsilon_j), \min(U_j, w_j^0 + \epsilon_j)]$$

ϵ_j : size perturbation for

In each linear sub-problem, the variable δ_j for each wire/buffer is restricted within a small range, which provides the approximated linear range for the delay gradient




Extension: multi-thermal profiles



- **Single thermal profile -> multiple thermal profiles**
- **Minimization: thermal-aware skew -> thermal-aware skew range**
- **Our strategy:**
 - **Single average thermal profile:**
Average thermal profile among the multiple non-uniform profiles
 - **Thermal-aware clock sizing under the single average thermal profile, and then evaluate it using multiple profiles**

Loosely-Thermal-Aware-Sequential-Linear-Programming




- **Motivation:**
 - TA-SLP spends most of time solving LP sub-problems
 - Density of G matrix affects the runtime
- **A speed-up scheme: LTA-SLP**
 - Keep thermal profile the same in G matrix calculation
 - Update the thermal profile in each iteration

$$G_{ij} = \left. \frac{\partial d_i}{\partial w_j} \right|_{w_j=w_j^0} \approx [d_i(\hat{\mathbf{t}}_j^0, \hat{\mathbf{w}}_j^0) - d_i(\mathbf{t}^0, \mathbf{w}^0)] / \epsilon_j$$

\downarrow
 $d_i(\mathbf{t}^0, \hat{\mathbf{w}}_j^0)$

Detail experiment settings



- **65nm technology:**

- **Wire:**

- $R = 0.15 \Omega/\mu\text{m}$, $C = 0.2 \text{ fF}/\mu\text{m}$, at 0°C , minimum size

- Size: $0.24\mu\text{m} - 0.96\mu\text{m}$

- **Buffer:**

- $R_d = 4.7\text{k}\Omega$, $C_L = 0.47\text{fF}$, $t_d = 17.4\text{ps}$, at 0°C , minimum size

- Size: 12X – 64X


- **Frequency: 5GHz, V_{dd} : 1.2V**

- **Thermal profile:**

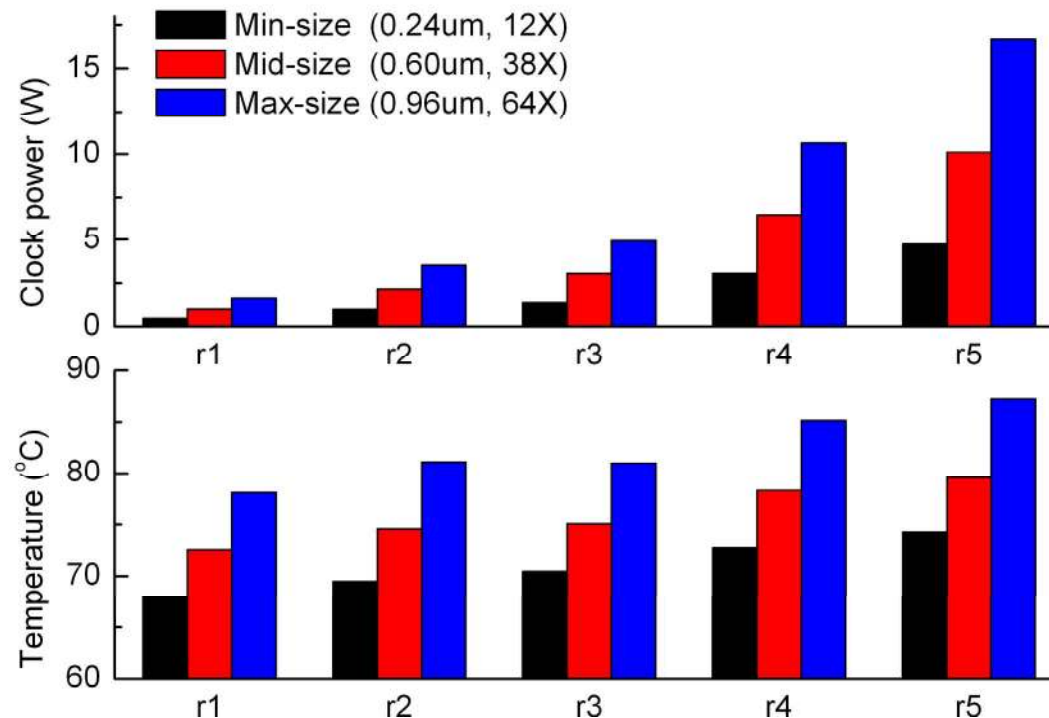
- **Multiple non-uniform: 10 power / thermal profiles**

| Circuits | # Sinks |
|----------|---------|
| r1 | 267 |
| r2 | 598 |
| r3 | 862 |
| r4 | 1903 |
| r5 | 3101 |


Impact of sizing on power and temperature



- Wire and buffer sizing greatly impacts power and temperature.
- Three cases: min-size, mid-size and max-size
 - Total power increases by 4 times
 - Temperature increases by 10 °C to 13 °C.



Impact of initial sizing




- Three initial sizing: min-size, mid-size and max-size
- The single average thermal profile out of 10 thermal profiles
- TA-SLP reduces thermal-aware clock skew to near zero, regardless of the initial solution

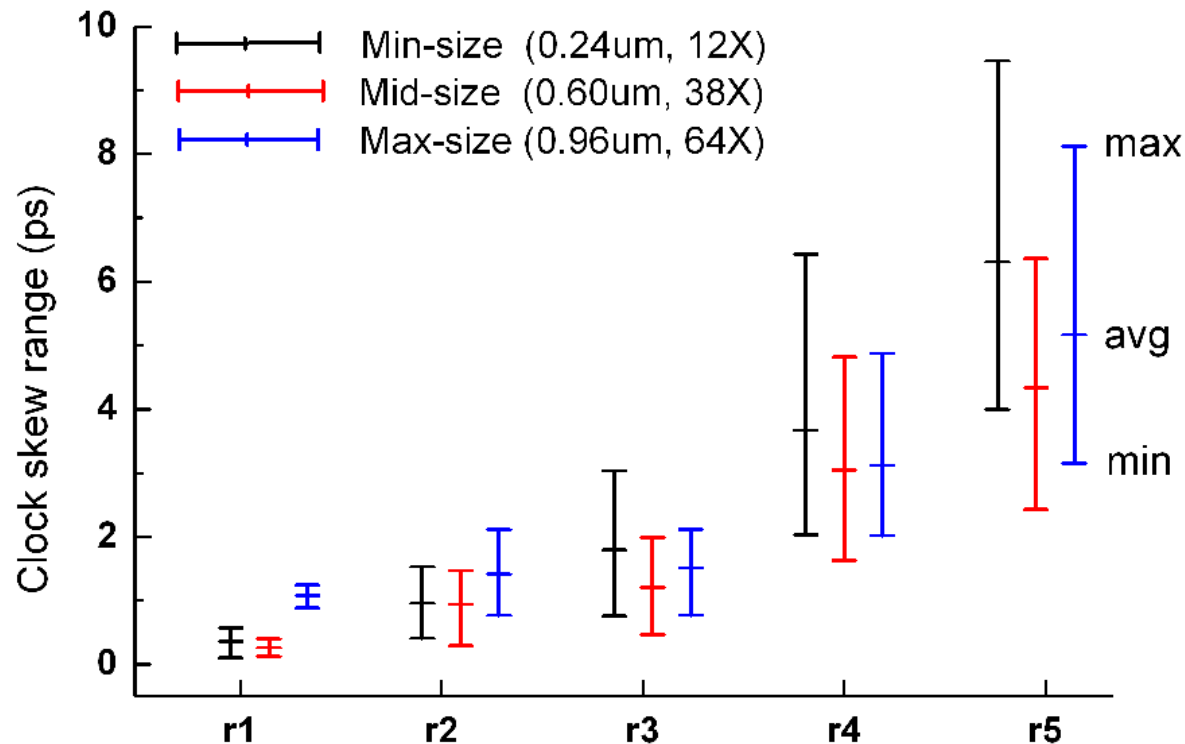
| Ckt. | Min-size 0.24um 12X | Mid-size 0.60um 38X | Max-size 0.96um 64X |
|------|---------------------------|---------------------------|---------------------------|
| r1 | 0.00 | 0.00 | 0.77 |
| r2 | 0.00 | 0.00 | 0.44 |
| r3 | 0.01 | 0.00 | 0.31 |
| r4 | 0.01 | 0.01 | 0.42 |
| r5 | 0.01 | 0.01 | 0.44 |

Skew value based on single average thermal profile (ps)


Impact of initial sizing (cont.)



- **Clock skew range based on 10 thermal profiles**
 - Mid-size case is slightly better than max-size case
 - Min-size case is the worst
- **We choose mid-size as initial sizing**



Comparison: NTA-SLP vs TA-SLP




- **NTA-SLP: Non-thermal-aware SLP -based sizing**
 - Minimize skew under power budget only, which resembles an existing work*

| | P _{total} (W) | | | T _{max} (°C) | | |
|----|--------------------------|---------|--------|-------------------------|---------|--------|
| | Budget | NTA-SLP | TA-SLP | Budget | NTA-SLP | TA-SLP |
| r1 | 2.87 | 2.86 | 2.86 | 74.18 | 71.16 | 71.12 |
| r2 | 6.07 | 6.05 | 6.04 | 77.07 | 72.98 | 72.91 |
| r3 | 8.65 | 8.60 | 8.59 | 76.96 | 73.62 | 73.60 |
| r4 | 18.32 | 18.25 | 18.23 | 80.88 | 76.59 | 76.59 |
| r5 | 28.61 | 28.46 | 28.52 | 82.85 | 78.00 | 78.03 |

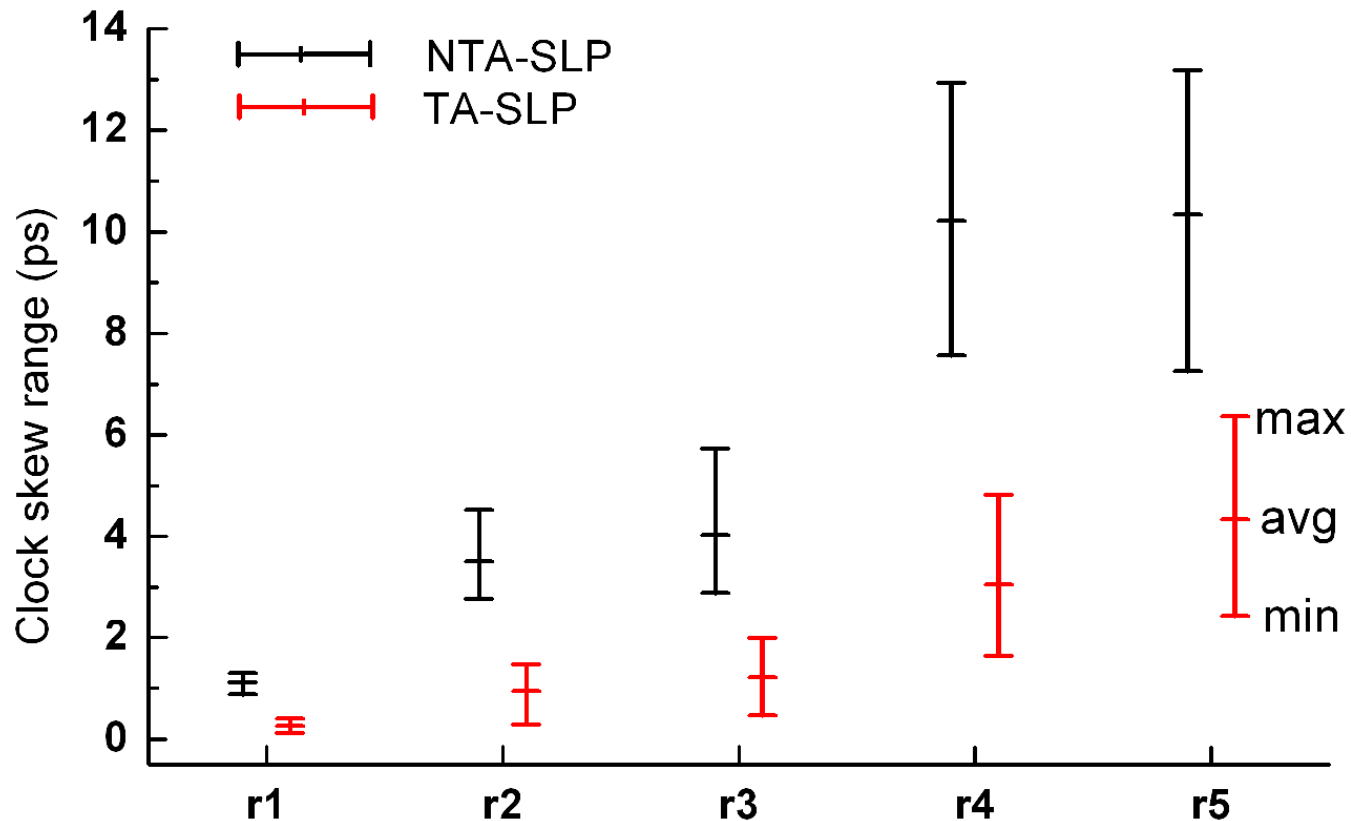
*M. R. Guthaus, D. Sylvester, and R. B. Brown. Clock buffer and wire sizing using sequential programming. In *Proc. ACM Design Automation Conf.*, pages 1041–1046, San Francisco, CA, July 24–28 2006.

Comparison: NTA-SLP vs TA-SLP (cont.)




- **Skew reduction:**


- Avg. skew saving: 59%-77%
- Max skew saving: 52%-69%




Impact of multiple thermal profiles



- TA-SLP with mid-size initial solution
- Under 10, 100, 1000 thermal profiles:
 - Average skew stays stable
 - Skew range increase with more thermal profiles



Runtime




- Comparisons among NTA-SLP, TA-SLP, and LTA-SLP


| Ckt. | NTA-SLP | | TA-SLP | | LTA-SLP | |
|------|---------|---------|--------|---------|---------|---------|
| | # itr | CPU/itr | # itr | CPU/itr | # itr | CPU/itr |
| r1 | 12 | 2.0 | 12 | 12.0 | 17 | 11.4 |
| r2 | 12 | 8.8 | 13 | 43.8 | 20 | 29.3 |
| r3 | 14 | 18.4 | 15 | 84.0 | 16 | 46.6 |
| r4 | 15 | 75.9 | 13 | 692.2 | 14 | 139.9 |
| r5 | 15 | 184.5 | 15 | 2190.5 | 16 | 292.8 |

- Density of G matrix:
 - r5: G matrix in TA-SLP has 23.6 million non-zero elements
 - r5: G matrix in LTA-SLP has 2.0 million non-zero elements

Conclusions



- **Proposed thermal-aware sequential-linear-programming (TA-SLP) based clock sizing**
- **Addressed the impact of clock sizing on power, thermal and delay**
- **Minimized thermal-aware clock skew under power and thermal budgets**
- **Narrowed clock skew range under multiple thermal profiles**
- **Proposed loosely-thermal-aware sequential-linear-programming (LTA-SLP)**



Thank you