

Interconnect-Power Dissipation in a Microprocessor


N. Magen, A. Kolodny, U. Weiser, N. Shamir

Intel corporation®

Technion - Israel Institute of Technology

Interconnect-Power Definition

- **Interconnect-Power** is dynamic power consumption due to interconnect capacitance switching
 - How much power is consumed by Interconnections ?
 - Future generations trends ?
 - How to reduce the interconnect power ?



0.13 μm cross-section, source - Intel

Background


- **Power is becoming a major design issue**
- **Scope: Dynamic power, the majority of power**
- **$P = S A F_i \cdot C_i \cdot V^2 \cdot f$**
- **This work focuses on the capacitance term**

Outline

- **Research methodology**
- **Interconnect Power Analysis**
- **Power-Aware Router Experiment**
- **Interconnect Power Prediction**
- **Summary**

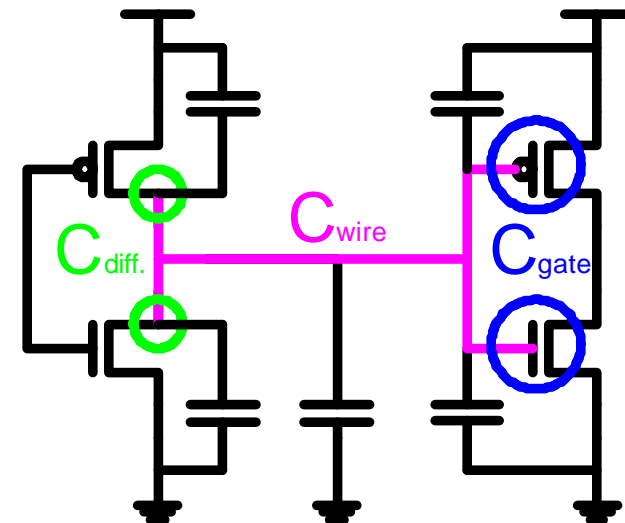

Case study

- **Low-power, state-of-the-art μ -processor**
- **Dynamic switching power analysis**
- **Interconnect attributes:**
 - **Length**
 - **Capacitance**
 - **Fan Out (FO)**
 - **Hierarchy data**
 - **Net type**
 - **Activity factors (AF)**
 - **Miscellaneous.**

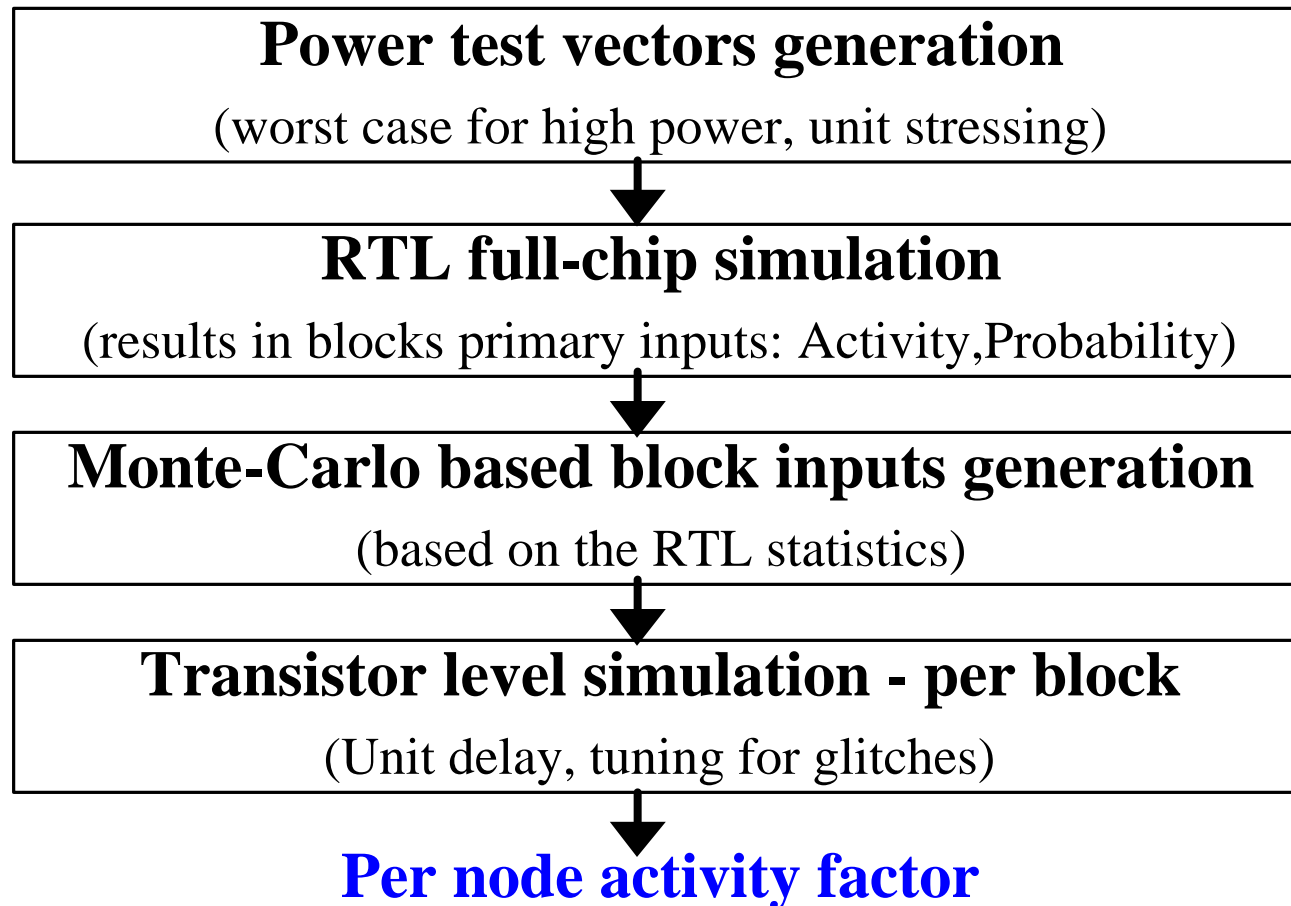


Interconnect Length Model

- **Total** wire length
 - **Stitched** across hierarchies
 - **Summed** over repeaters
-
- **Net** model



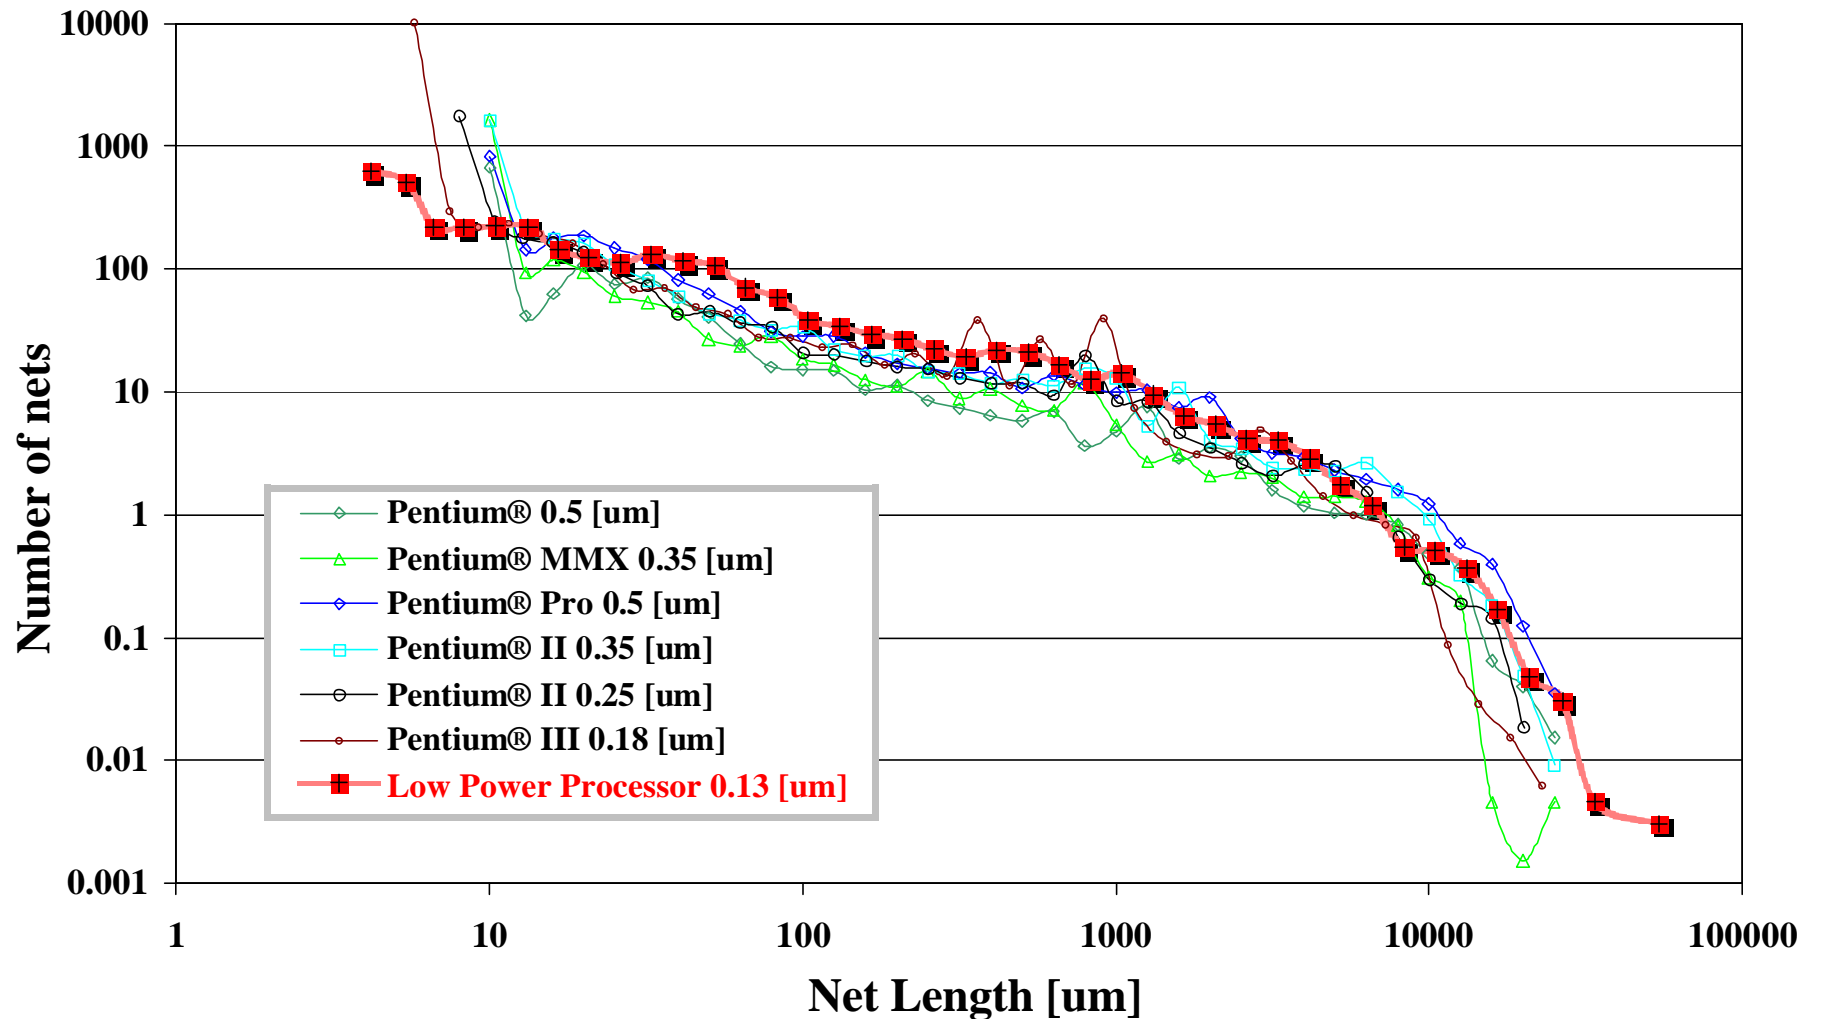
Activity Factors Generation



Outline

- **Research methodology**
- **Interconnect Power Analysis**
- **Power-Aware Router Experiment**
- **Interconnect Power Prediction**
- **Summary**

Interconnect Length Distribution

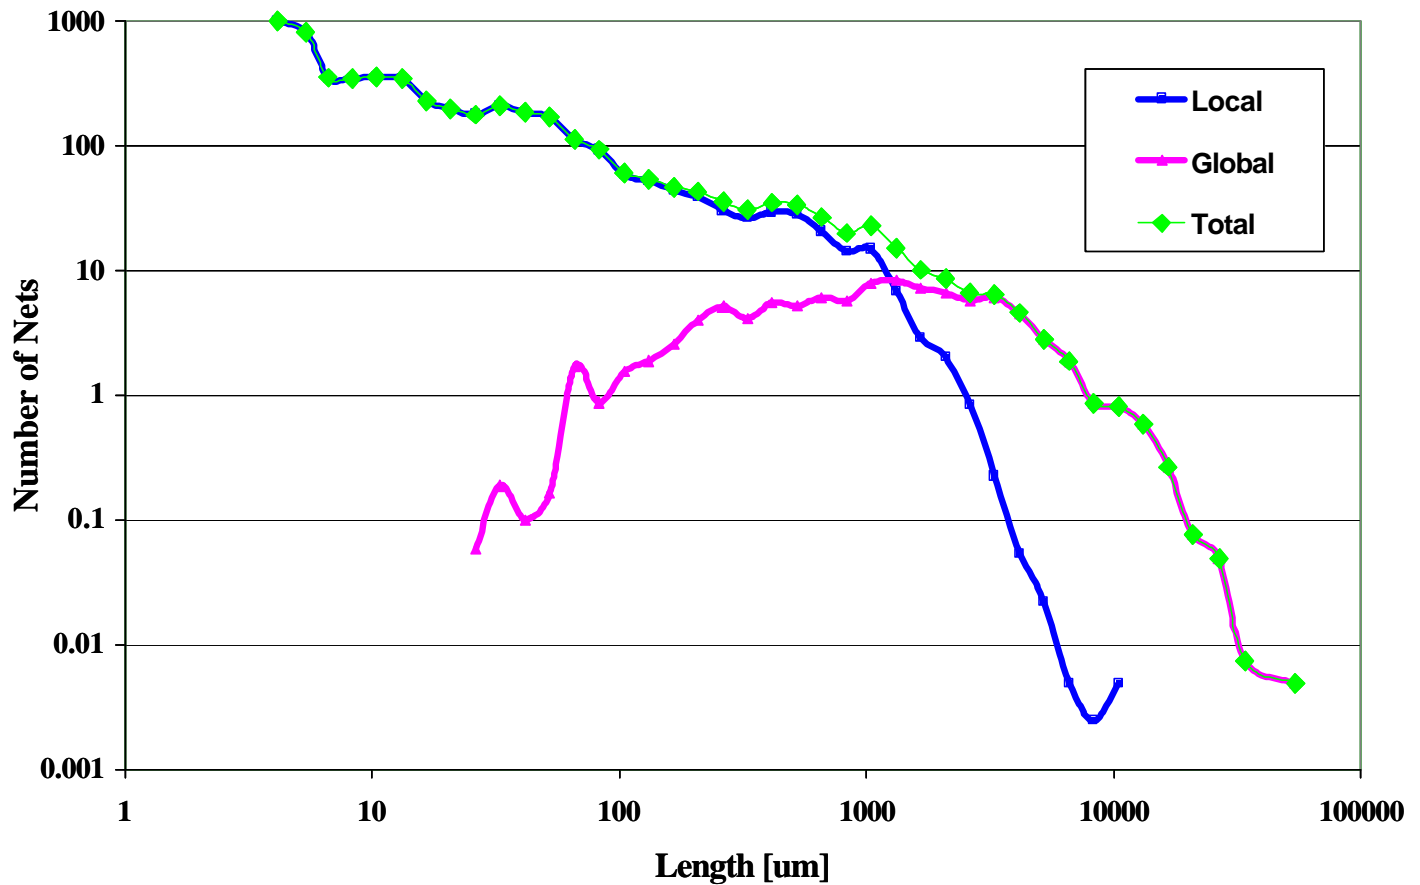


Source: Shekhar Y. Borkar, CRL - Intel

Interconnect Length Distribution

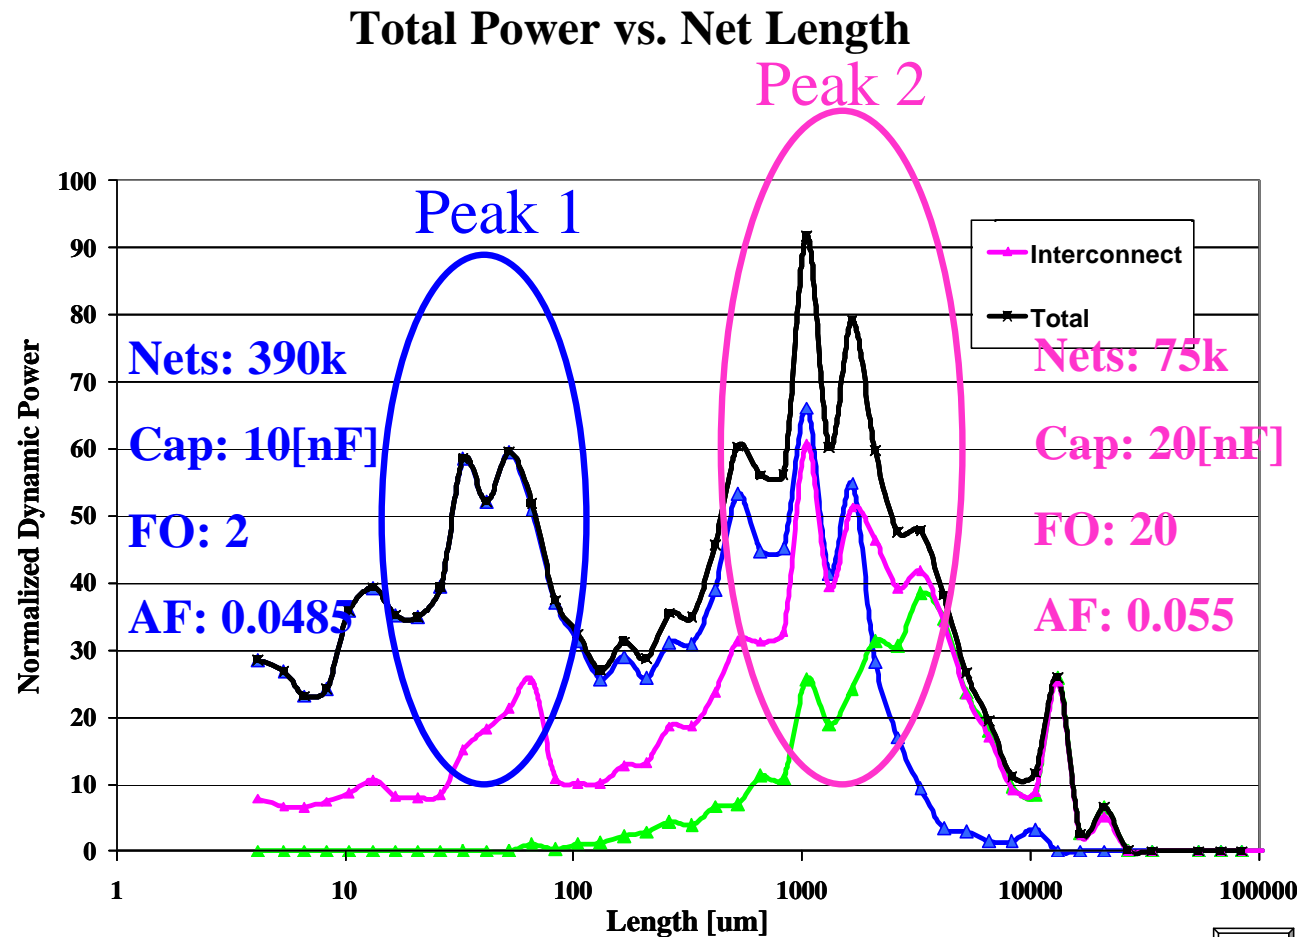
Nets vs. Net Length

- **Log – Log scale**
- **Exponential decrease with length**
- **Global clock – not included**



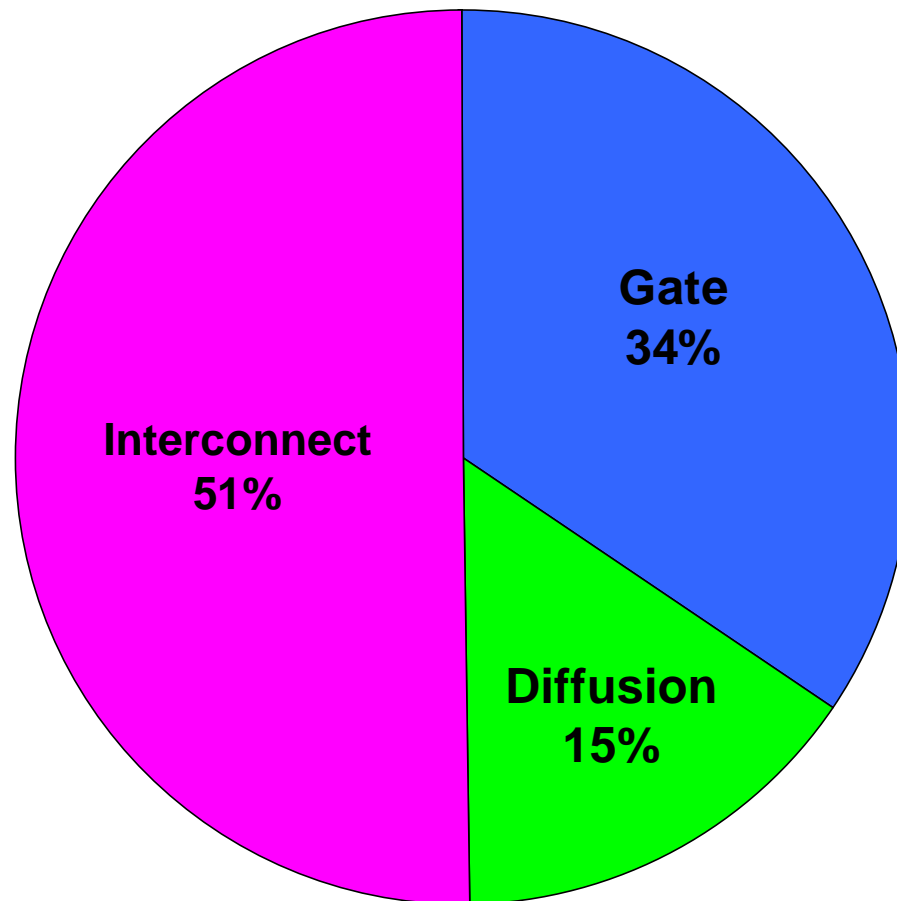
Total Dynamic Power

- **Total Dynamic Power**
- **Global clock – not included**
- **Local nets = 66%**
- **Global nets = 34%**



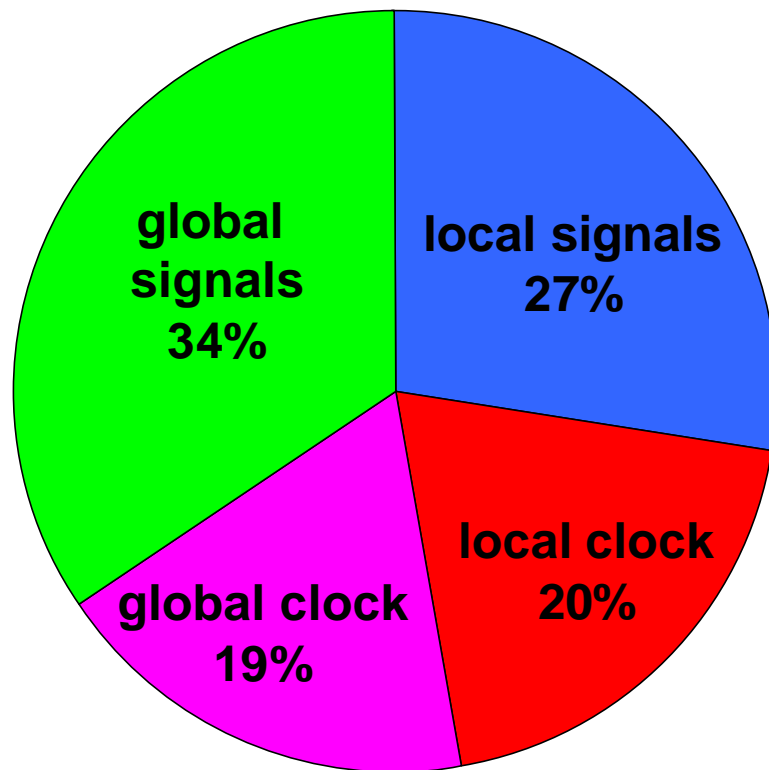
Total Dynamic Power Breakdown

Global clock included




Power Breakdown by Net Types

Global clock included



Interconnect power

(Interconnect only)

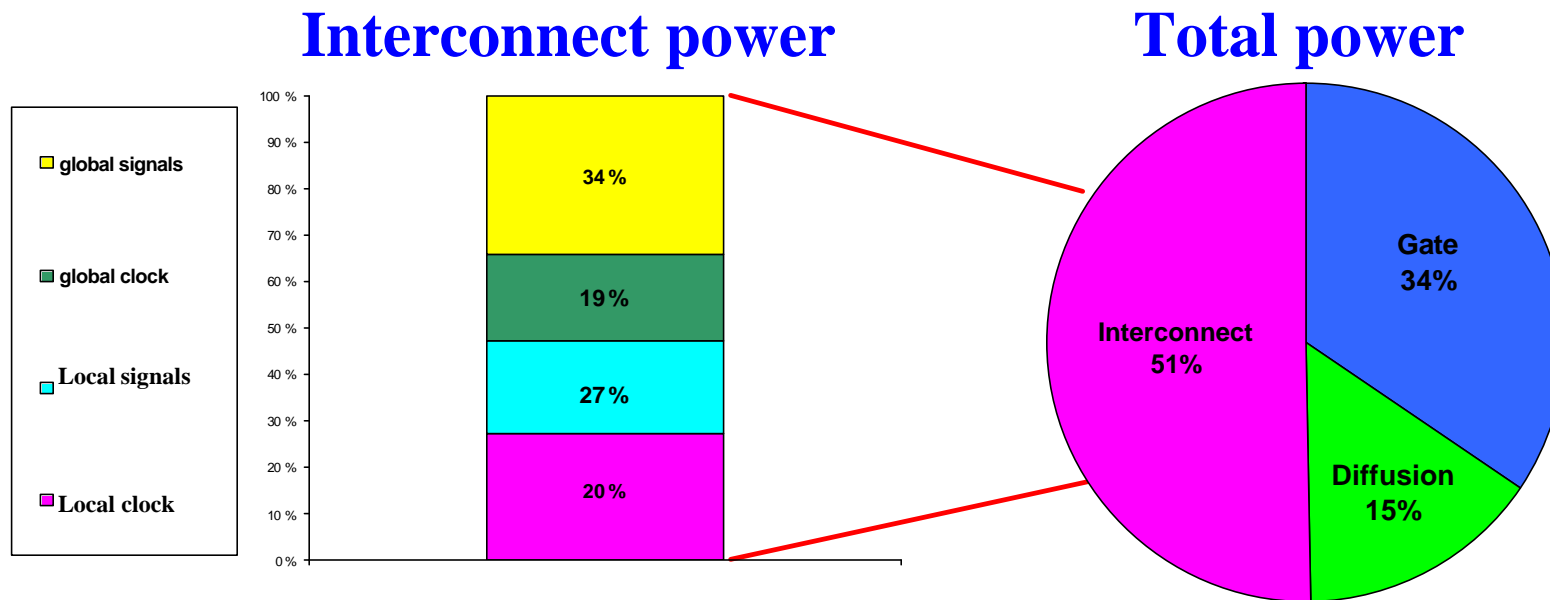


Total power

(Gate, Diffusion and Interconnect)

Interconnect Power Breakdown

- Interconnect consumes 50% of dynamic power
- Clock power ~40% (of Interconnect and total)
- 90% of power consumed by 10% of nets
- Interconnect design is NOT power-aware !
- Predictive model can project the interconnect power.

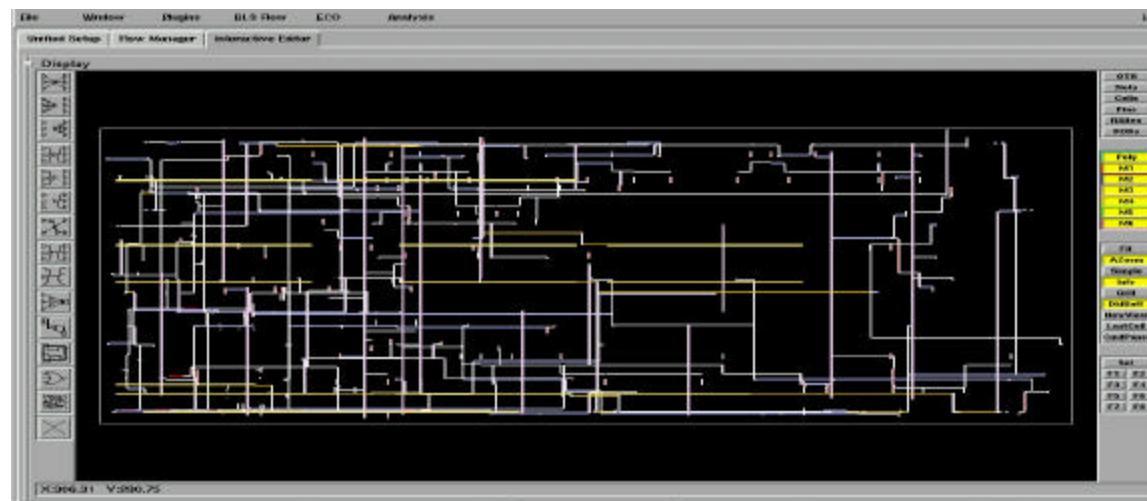


Outline

- **Research methodology**
- **Interconnect Power Analysis**
- **Power-Aware Router Experiment**
- **Interconnect Power Prediction**
- **Summary**

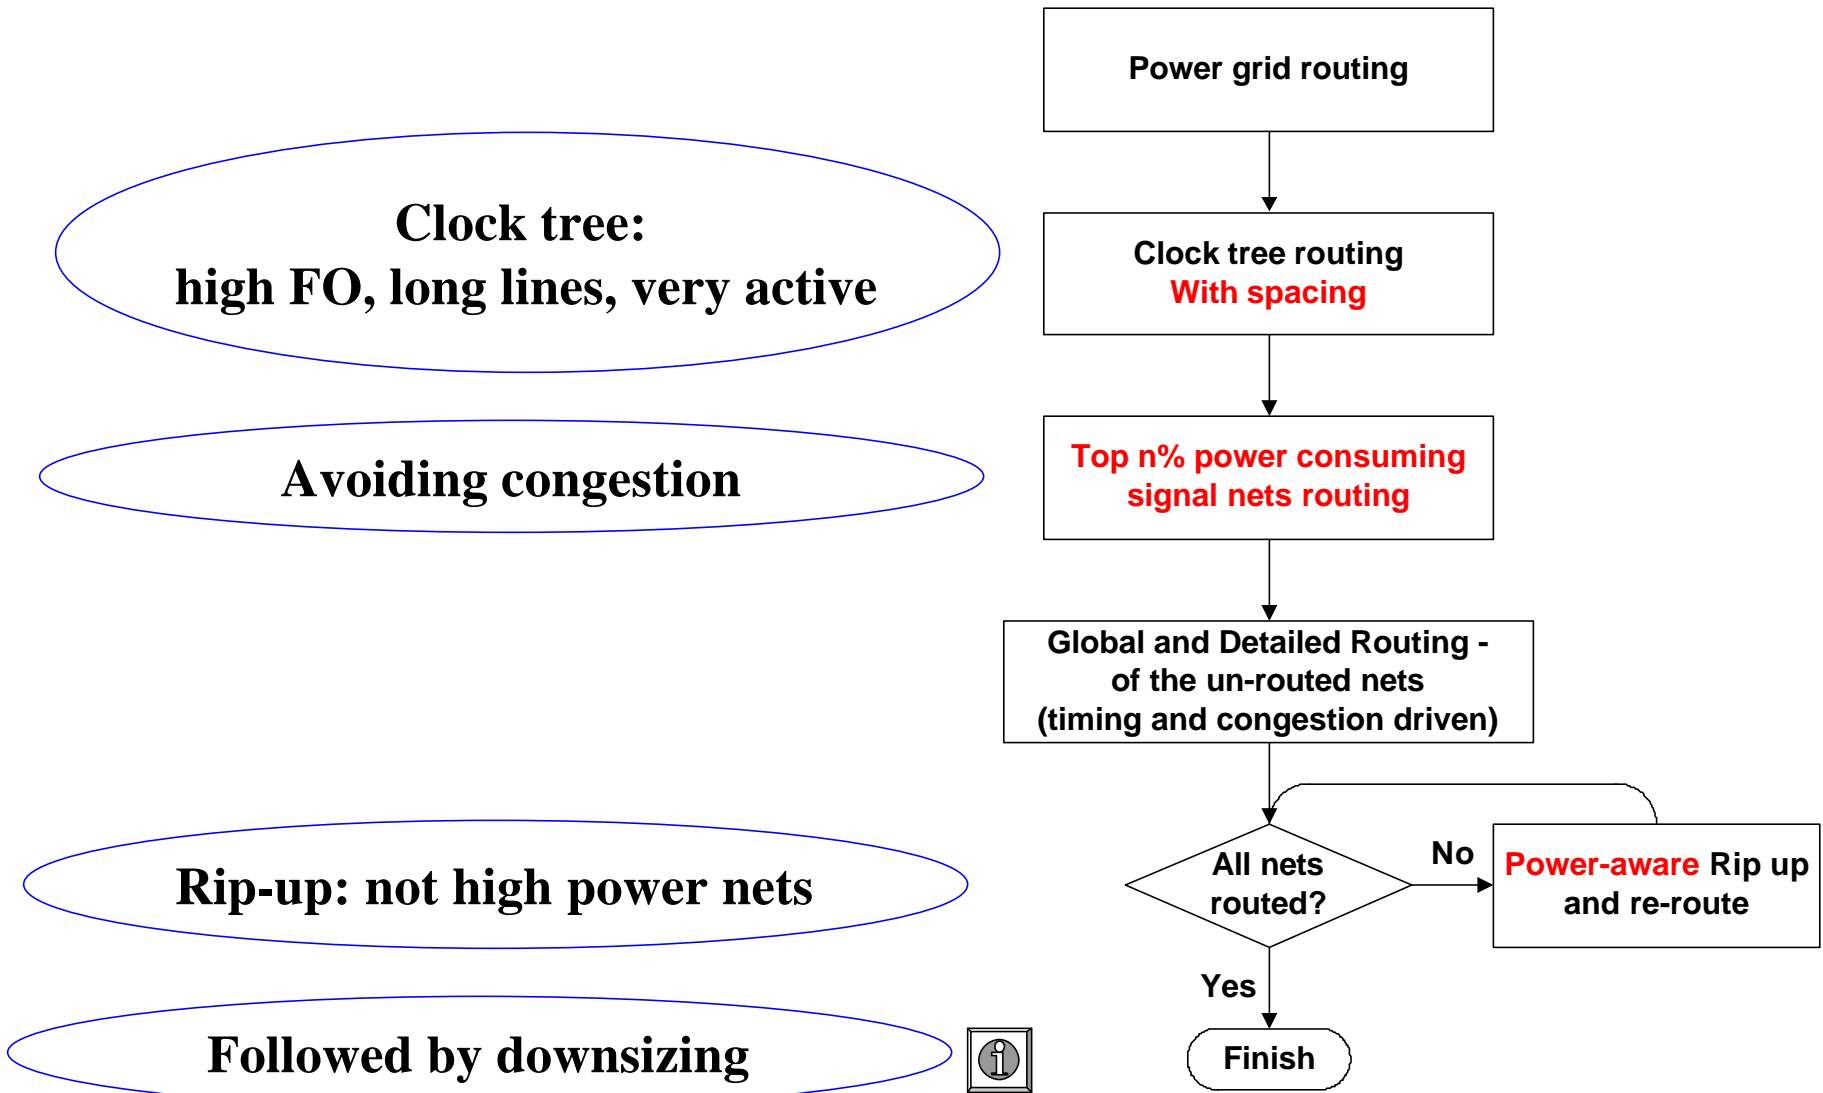
Experiment - Power-Aware Router

- **Routing Experiment optimizing processor's blocks**
 - **Local nodes (clock and signals) consume 66% of dynamic power**
 - **10% of nets consume 90% of power**
 - **Min. spanning trees can save over 20% Interconnect power**
 - **Routing with spacing can save up to 40% Interconnect power**




Small block's local clock network

Power-Aware Router Flow



Results - Power Saving



Average saving results: 14.3% for ASIC blocks ¹

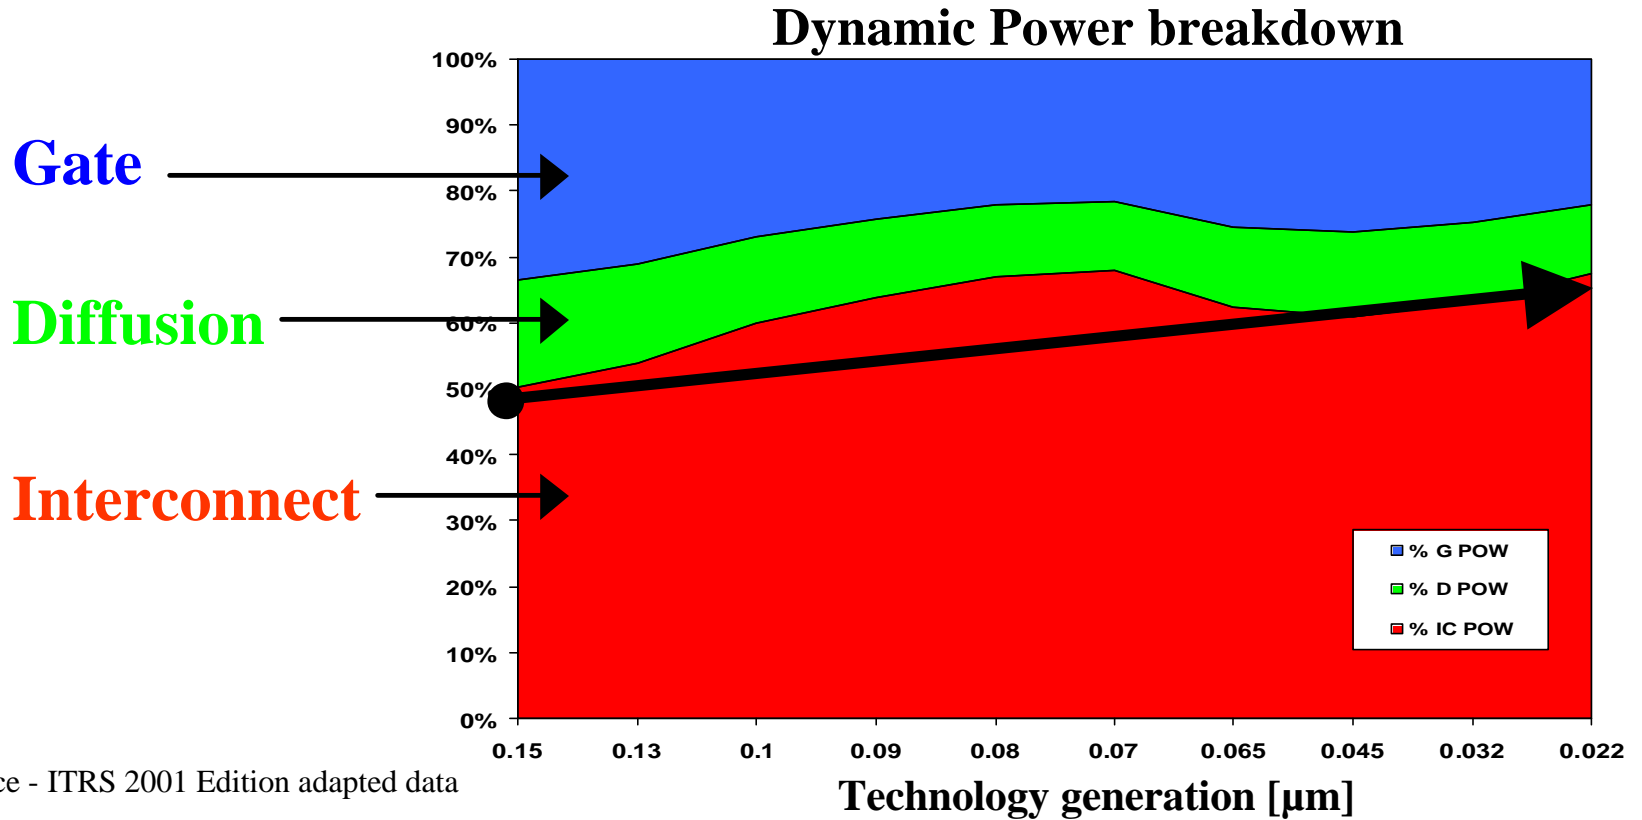
1 - Estimated based on clock interconnect power



Outline

- **Research methodology**
- **Interconnect Power Analysis**
- **Power-Aware Router Experiment**
- **Interconnect Power Prediction**
- **Summary**

Future of Interconnect Power

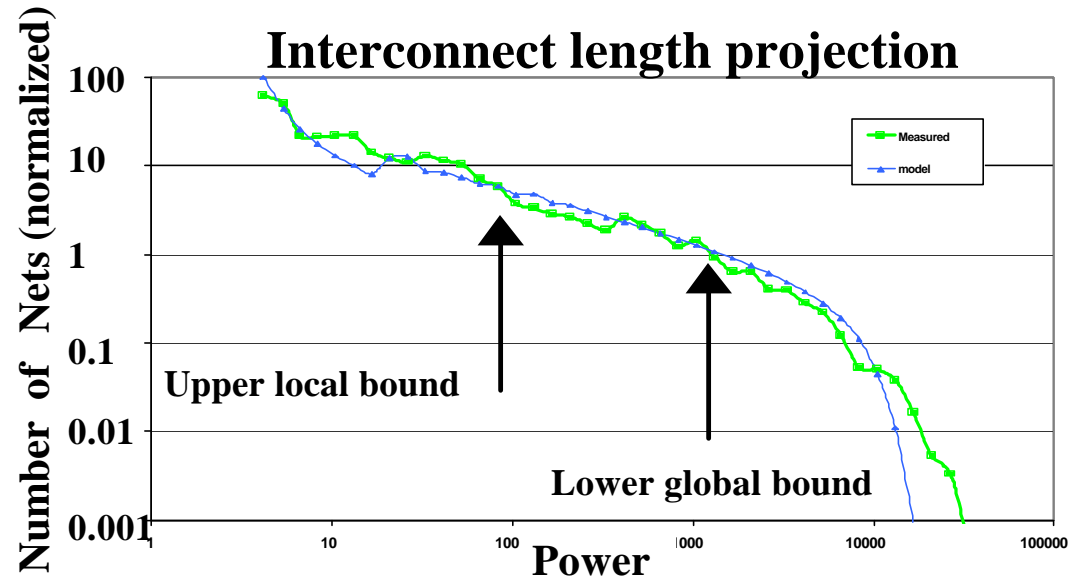


Interconnect power grows to 65%-80% within 5 years !

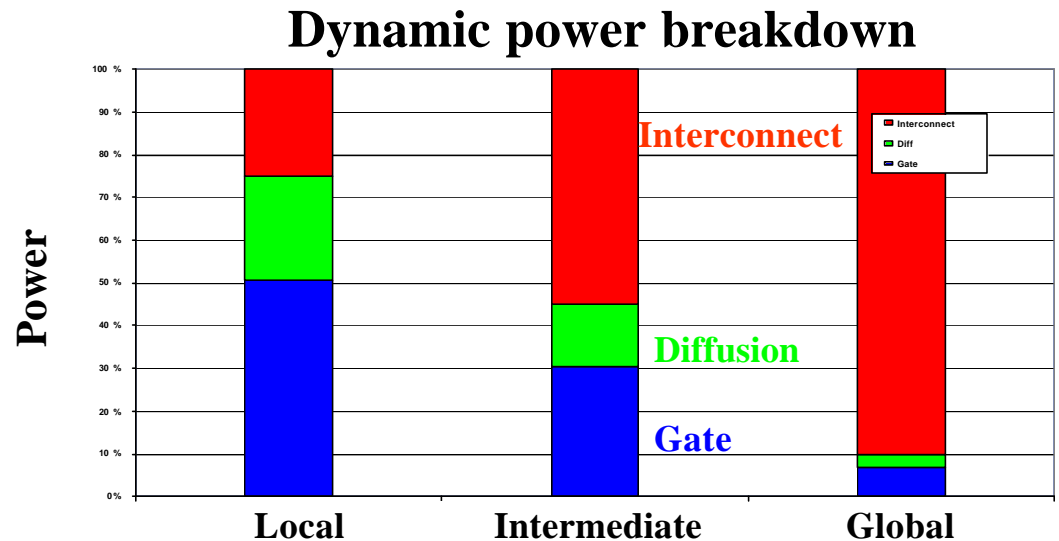
(using optimistic interconnect scaling)

Interconnect Power Prediction

- The number of nets vs. unit length – Modified Davis model

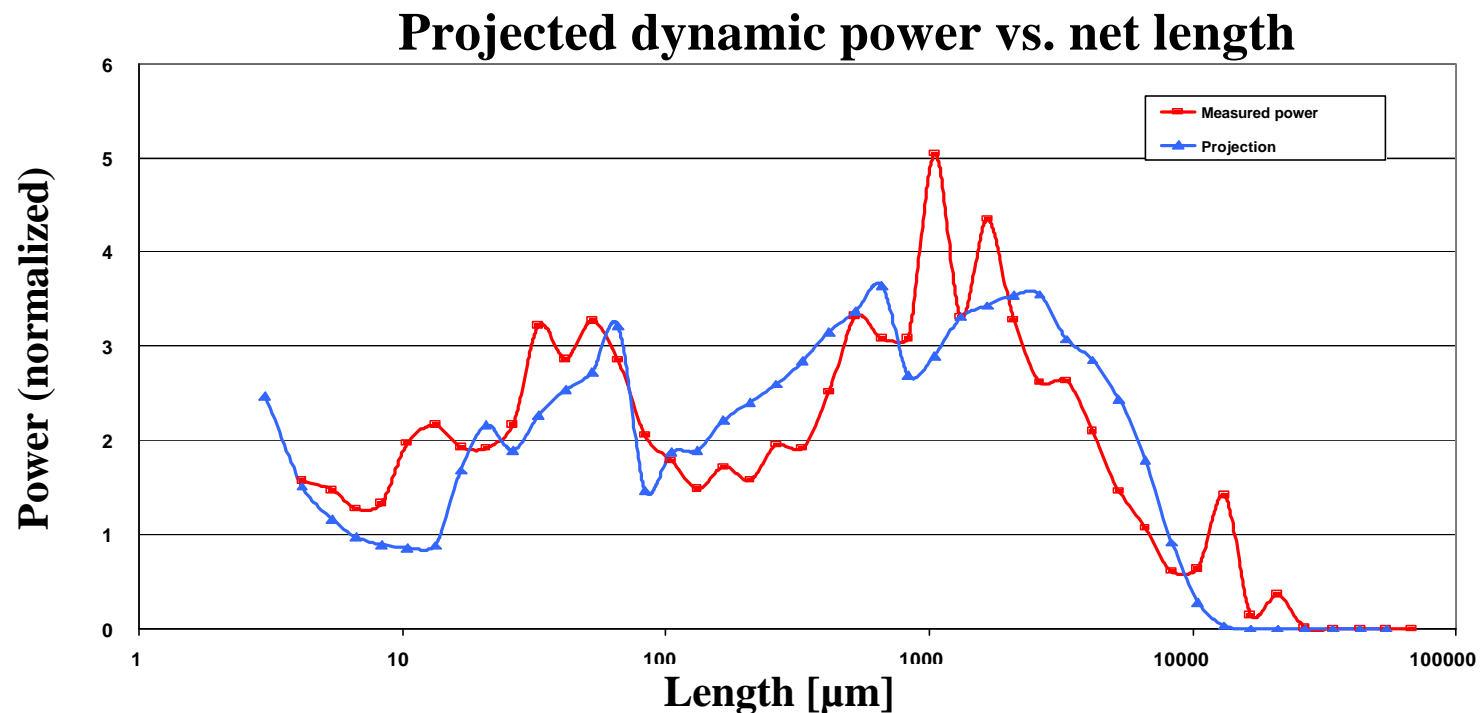


- The dynamic power average breakdown



Interconnect Power Model

- **Multiplication of the number of interconnects with power breakdowns gives:**



The power model matches processor power distribution !

Outline

- **Research methodology**
- **Interconnect Power Analysis**
- **Power-Aware Router Experiment**
- **Interconnect Power Prediction**
- **Summary**

Summary

- **Interconnect is 50% of the dynamic power of processors, and getting worse.**
 - ? **Interconnect power-aware design is recommended**
- **Clock consumes 40% of interconnect power.**
 - ? **Clock interconnect spacing is suggested**
- **Interconnect power is sum of nearly all net lengths and types.**
 - ? **Router level Interconnect power reduction addresses all**
- **Interconnect power has strong dependency on the hierarchy**
 - ? **Per Hierarchy analysis and optimization algorithms**


Future Research

- 1. Interconnect Power characterization and prediction**
- 2. Investigate Interconnect power reduction techniques:**
 - Interconnect-Spacing for power**
 - Interconnect Power-Aware physical design**
 - Aspect Ratio optimization for power**
 - Architectural communication reduction**

Questions ?

BACKUP-Slides

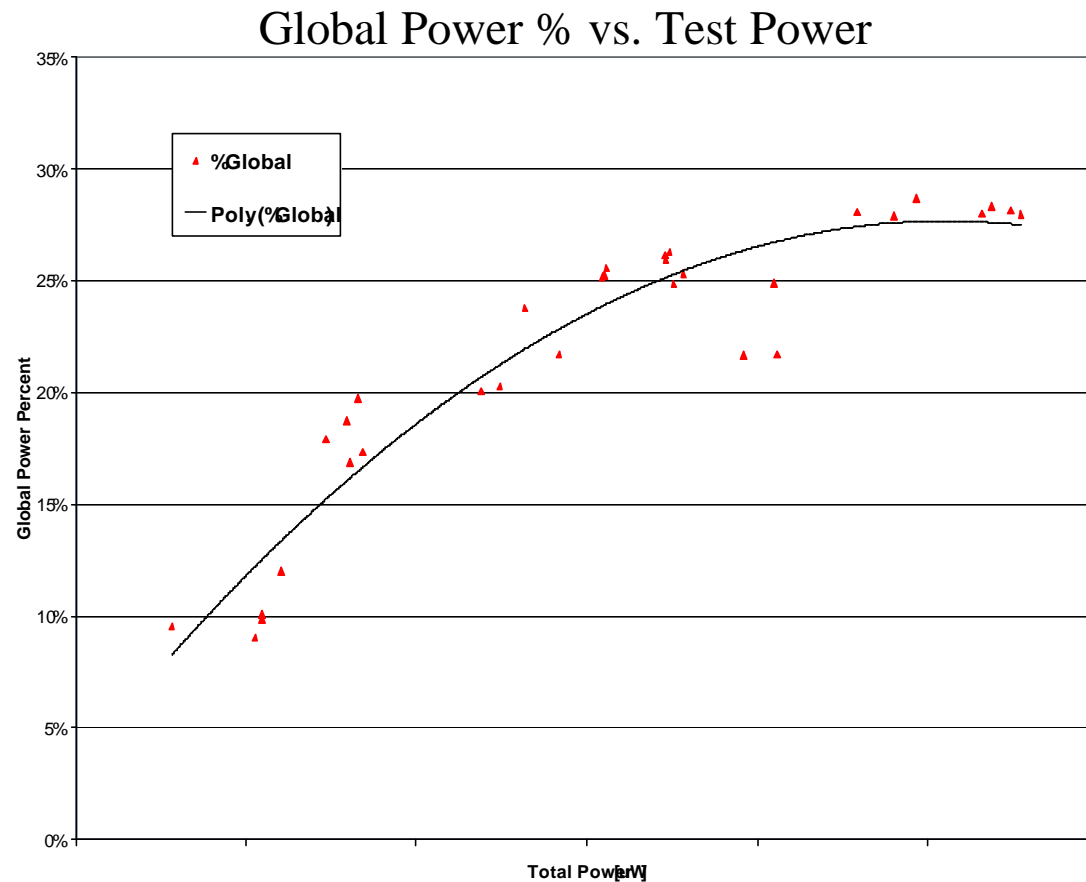
Processor Case Study



- **Analysis subject: Processor, 0.13 [μm]**
- **77 million transistors, die size of 88 [mm^2]**
- **Data sources (AF, Capacitance, Length)**
- **Excluded: L2 cache, global clock, analog units**


Global Communication

- **Global power is important**
- **Global power is mostly IC**
- **For higher power benchmarks – Global power is higher**
- **G-clock excluded**



Benchmark Selection

- **High power test benchmarks**
 - **Worst case design**
 - **Suitable for: thermal design, power grid design**
 - **Average power is a fraction of peak power**
- **Unit stressing benchmarks**
- **Averaging of all high power benchmarks**
 - **High node coverage**



ITC logo

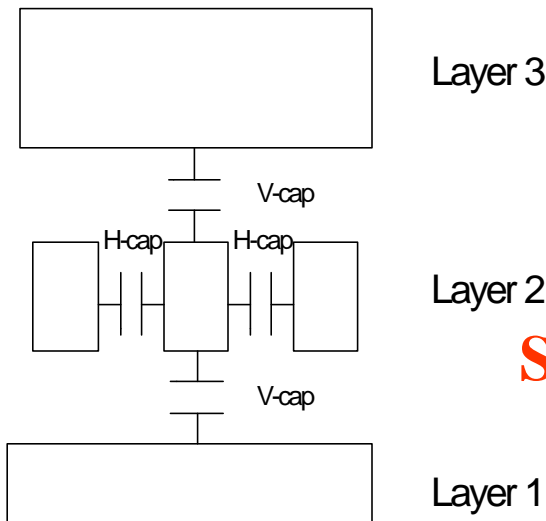
Interconnect Power Implications

- **Interconnect power can be reduced by minimizing switched capacitance:**
 - **Fabrication process (wire parameters)**
 - **Power-driven physical design**
 - **Logic optimization for power**
 - **Architectural interconnect optimization**

Interconnect Capacitance

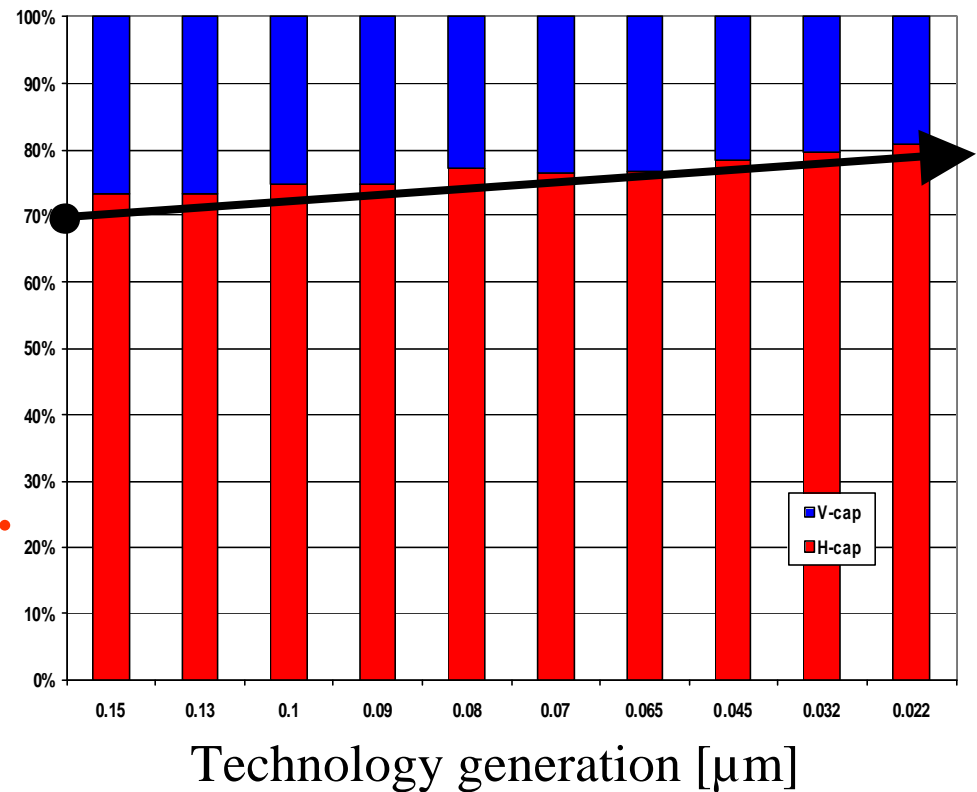
- Side-cap is increasing:
70% to 80%

self-cap.



Side-cap.

Global Capacitance breakdown

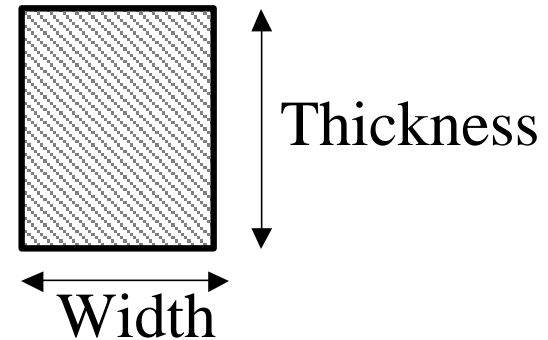


Source - ITRS 2001 Edition adapted data

The majority of interconnect capacitance is side-capacitance !



Fabrication Process – Aspect Ratio (AR)

- **Interconnect AR** = $\frac{\text{Thickness}}{\text{Width}}$



- **Low AR = Low Interconnect power**
- **Low AR = High resistance**
- **Frequency Modeling**

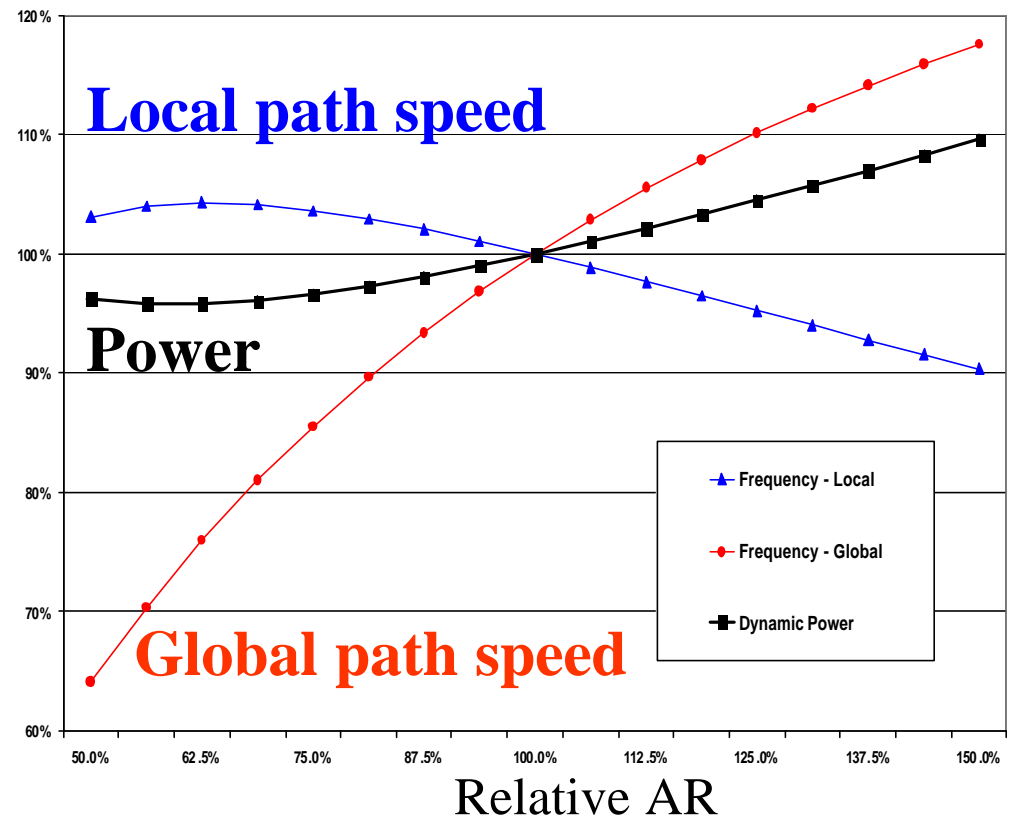
- **Local:** average gate, average IC
- **Global:** optimally buffered global IC



Aspect Ratio – Trade offs

- Power – depends on cap.
- Frequency:
 - Local – gates and IC cap.
 - Global – mostly IC – RC
- Per layer AR optimization !
- Scaling ? more power save, less frequency loss

Freq. And Power vs. Relative AR

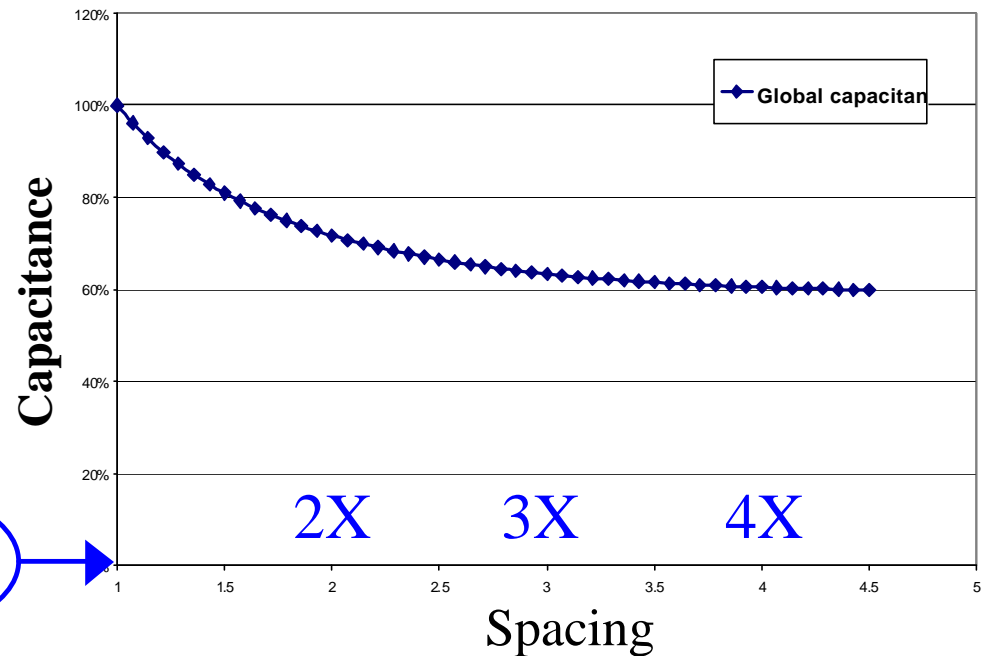


Aspect Ratio optimization can save over 10% of dynamic power !

Physical Design - Spacing

- Spacing can save up to 40%
- About 30% is with double space
- Spacing advantages: scaling, frequency, reliability, noise, easy to modify

0.13 [μm] global IC cap. vs. spacing



Wire spacing can save up to 20% of the dynamic power !



Spacing calculation

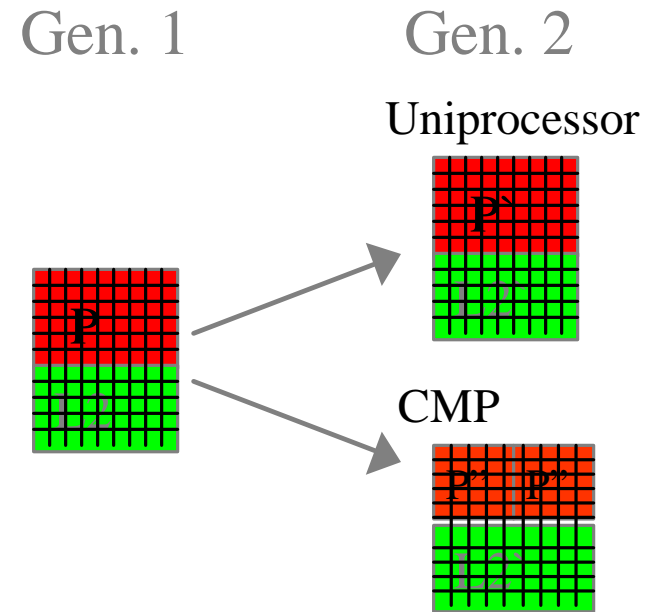
Back of an envelope estimation:

- **10% of Interconnect ? 90% power**
- **X2 spacing = extra 20% wiring**
- **Global clock – not spaced (inductance)**
- **Global clock is 20% of interconnect power**
- **Save: 30% of (90%-20%) = 20%**
- **Interconnect is 50% ? 10% power save**
Expected 20% with downsizing
- **Minor losses - congestion**




μ -Architecture - CMP

- **Comparing two scaling methods, by IC power.**
- **IC - predicted by Rent**
- **L2 - identical, minor**
- **Clock - Identical !**
- **Same average AF.**
- **Result ~5% less dynamic power for CMP**



Power critical vs. Timing critical





Outline

- **Research methodology**
- **Interconnect Power Analysis**
- **Future Trends Analysis**
- **Interconnect Power Implications**
- **Summary**

Interconnect Length Prediction

- **Technology projections - ITRS**
- **Interconnect length predictions:**
 - **ITRS model: 1/3 of the routing space**
- most optimistic
 - **Davis model:**
 - Rent's rule based
 - Predicts number of nets as function of:
the number of gates and complexity factors
- **Models calibrated based on the case study**



Rent's parameters

Rent's rule: $T = k N^r$


T = # of I/O terminals (pins)

N = # of gates

K = avg. I/O's per gate

r = Rent's exponent

can be: $0 < r < 1$, but common -
(simple) $0.5 < r < 0.75$ (complex)



Donath's length estimation model

For the i -th level:

There are 4^i blocks

For each block there are: $k \cdot \left(\frac{N}{4^i}\right)^r$ terminals

Assuming two terminal nets : $\frac{k}{2} \cdot \left(\frac{N}{4^i}\right)^r$ nets

The nets of the $i-1$ level must be subtracted.


$$\text{Nets for level } i : n_i = 4^i \cdot \frac{k}{2} \cdot \left(\frac{N}{4^i}\right)^r - 4^{i-1} \cdot \frac{k}{2} \cdot \left(\frac{N}{4^{i-1}}\right)^r = 4^i \cdot \frac{k}{2} \cdot \left(\frac{N}{4^i}\right)^r \cdot (1 - 4^{r-1})$$

Average interconnection length

The wires can be of two types A and D.

$$L_A = \frac{\sum_{i_A=1}^I \sum_{j_A=1}^I \sum_{i_B=1}^I \sum_{j_B=1}^I [I + i_A - i_B + |j_B - j_A|]}{I^4} = \frac{4}{3} \cdot I - \frac{1}{3I}$$

$$L_D = \frac{\sum_{i_A=1}^I \sum_{j_A=1}^I \sum_{i_B=1}^I \sum_{j_B=1}^I [2I + i_A + j_A - i_B - j_B]}{I^4} = 2 \cdot I$$



Taken from a SLIP 2001 tutorial by Dirk Stroobandt

The average: $r_i = \frac{14 \cdot I}{9} - \frac{2}{9 \cdot I}$



Overall : $\bar{R} = \frac{\sum_{i=1}^I n_i \cdot r_i}{\sum_{i=1}^I n_i}$ equals $\frac{2}{9} \cdot \left(7 \cdot \frac{N^{r-0.5} - 1}{4^{r-0.5} - 1} - \frac{1 - N^{r-1.5}}{1 - 4^{r-1.5}} \right) \cdot \left(\frac{1 - 4^{r-1}}{1 - N^{r-1}} \right)$

Davis Model

- From Rent's rule: $T_r = r \cdot N^P$
- IDF:
$$i(l) = \begin{cases} 1 \leq l \leq \sqrt{N} & : \frac{\mathbf{a} \cdot r}{2} \cdot \Gamma \cdot \left(\frac{l^3}{3} - 2 \cdot \sqrt{N} \cdot l^2 + 2 \cdot N \cdot l \right) \cdot l^{2 \cdot p - 4} \\ \sqrt{N} \leq l \leq 2 \cdot \sqrt{N} & : \frac{\mathbf{a} \cdot r}{6} \cdot \Gamma \cdot (2 \cdot \sqrt{N} \cdot l)^3 \cdot l^{2 \cdot p - 4} \end{cases}$$
- Where: $\mathbf{a} = \frac{FO}{FO+1}$, $\Gamma = \frac{2 \cdot N \cdot (1 - N^{P-1})}{\left(-N^p \cdot \frac{1 + 2 \cdot p - 2^{2 \cdot p - 1}}{p \cdot (2 \cdot p - 1) \cdot (p - 1) \cdot (2 \cdot p - 3)} - \frac{1}{6 \cdot p} + \frac{2 \cdot \sqrt{N}}{2 \cdot p - 1} - \frac{N}{p - 1} \right)}$
- Interconnect total number and length:
 Nets: $I_{total} = \int_1^{2 \cdot \sqrt{N}} i(z) dz$ Length: $L_{total} = \int_1^{2 \cdot \sqrt{N}} i(z) \cdot z \cdot dz$
- Multipoint Length: $L_{multi_terminal} = L_{total} \cdot \mathbf{c}$ where $\mathbf{c} = \frac{4}{FO+3}$

Davis Model - extension

- Constant factor favors shorter nets.
- Short P2P net has higher chance to be a part of a multipoint net.
- Correction factor: $\text{multi-terminal factor}(l) = \frac{\text{number of point to point nets shorter than } l}{\text{total point to point nets}}$
- Length: $I_{\text{multi-terminal}}(l) = \frac{1}{FO} \int_1^l i(z) \text{multi-terminal factor}(z) dz$



RMST - Example


The image displays two side-by-side screenshots of the Tucson fwcntls software interface, showing a circuit layout in the 'Display' window. The left window shows a zoomed-out view, while the right window shows a zoomed-in view of the same circuit. Both windows have a menu bar (File, Window, Plugins, RLS Flow, ECO, Analysis, Help) and a toolbar with icons for various functions. The right window's toolbar has 'AZoom' highlighted. The command window at the bottom of each window shows the command 'select nets late_clkbuf' and its output.

Command: `select nets late_clkbuf`
 % -I- Selected 'late_clkbuf' in nets list. (GuiTclCallbacks)
 % late_clkbuf



Total Dynamic Power


- Total Dynamic Power
- Global clock - not included
- Local nets = 66%
- Global nets = 34%




Local and Global IC

- Local and Global IC are different:
- Number by Length breakdown
- IC breakdown – cap and power
- Fan out
- Metal usage
- AF is similar

Local Power breakdown vs. Net Length




Global Power breakdown vs. Net Length



Benchmarks Comparison

Global Dynamic power vs. Length




High power tests show similar behavior to average SPEC !

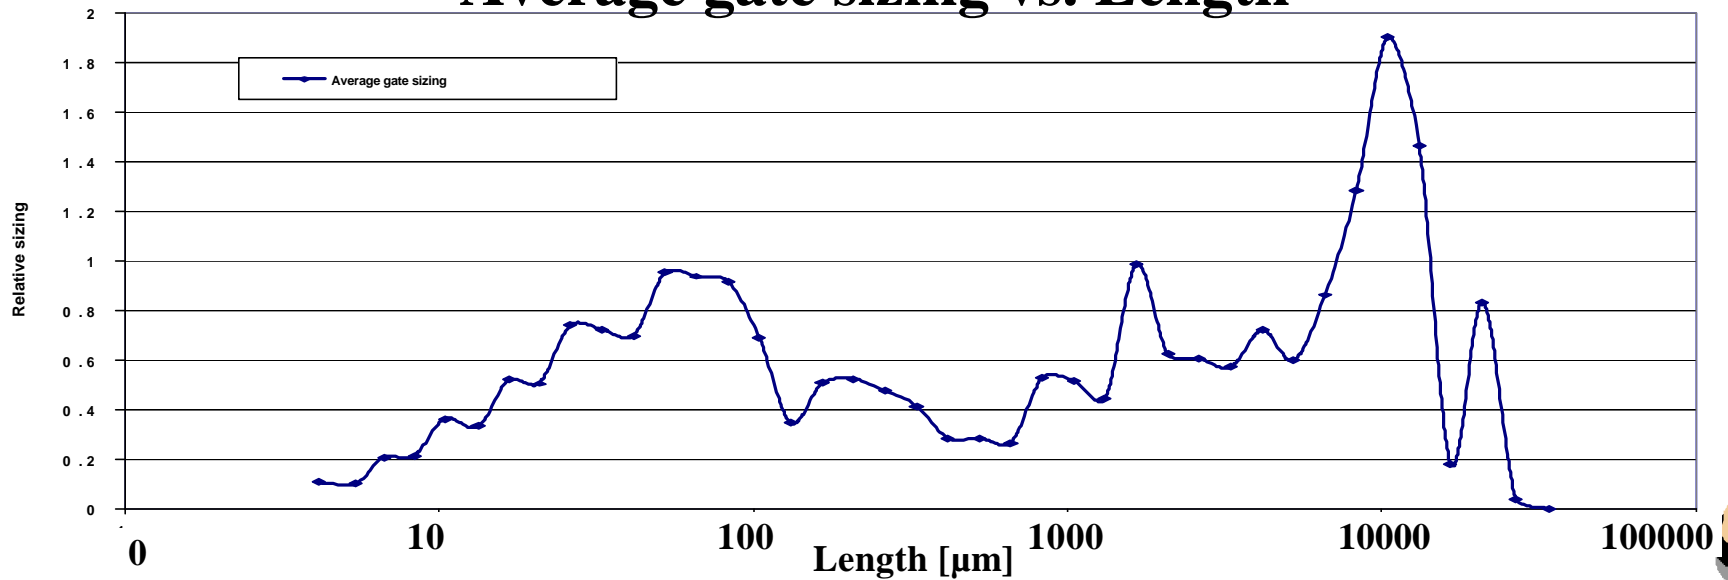


Interconnect Peaks

Total wire length vs. Length



Average gate sizing vs. Length

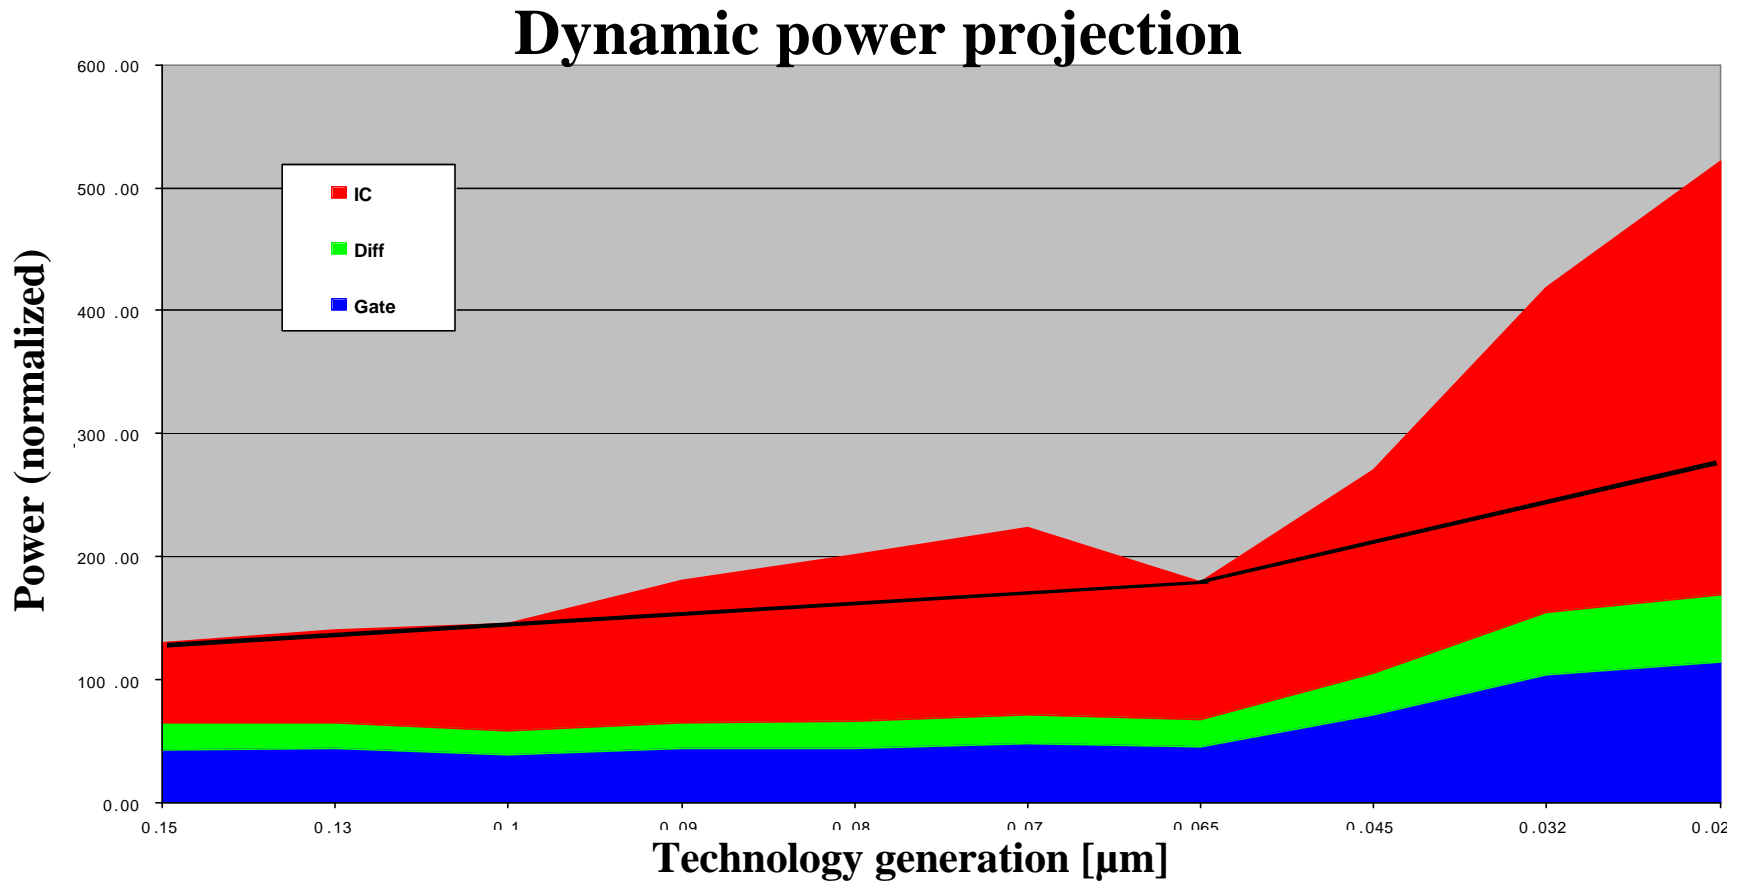


ITRS Power Trends

- The ITRS power projection interconnect power reduction that happens in 2006-2007 is based on:
 1. Aggressive voltage reduction
 2. Low-k dielectric improvements
- The devices capacitance increase by **+30%** (trend **-15%**)
- The combined effect:
 - Interconnect power reduction (relative to voltage)
 - Device power remains constant



Dynamic power - ITRS trend




The Black curve is the ITRS maximum heat removal capabilities



Power-Aware Flow

- The reduced IC cap allows for driver downsizing
- On average it reduced the dynamic power by 1.4 of the IC power saving
- Downsizing is timing verified
- Cells downsizing reduced the total area and leakage by 0.4%
- No signal spacing was applied over 30% unused metal
- Post-layout optimization are possible




FUBS – description

- **A** – medium, randomly picked
- **B** – small, highest clock power
- **C** – small, good potential
- **D** – medium, good potential
- **E** – worse than average

Block Name	Block A	Block B	Block C	Block D	Block E	AVERAGE
Area [μm^2]	138801.6	101274.6	65816.1	164229.1	59766.3	209537.8
Devices	14574	8644	7618	18194	6109	16675
Inactive Nodes	63.66%	98.78%	82.36%	39.22%	35.38%	52.94%
Power [μW]	17170.22	251.15	1786.76	11811.90	6757.11	15373.86
RMST potential power saving	14.3%	17%	22%	29%	4.1%	17%
Clock cap.	11.25%	2.59%	12.75%	13.16%	3.27%	8.01%
Clock power	72.10%	99.99%	96.46%	94.99%	33.84%	60.47%
IC cap.	34.00%	27.70%	38.14%	36.05%	29.86%	34.67%
IC power	28.89%	59.54%	46.74%	48.62%	40.65%	36.83%
Clock IC power	20.19%	59.54%	45.48%	46.26%	16.87%	23.87%
Clock IC length	1.71%	2.34%	2.05%	2.09%	0.74%	3.85%
Relative - Capacitance per Length Unit.	82.23%	113.15%	87.46%	83.74%	85.97%	88.46%



Miller Factor - Power



- Opposite direction switching-

- The current:
$$I_c = \frac{dQ}{dt} = \frac{d(C\Delta V_c)}{dt} = C \frac{d\Delta V_c}{dt}$$

- Energy:
$$E_c = \int_0^T I_c \cdot V_{dd} \cdot dt = \int_0^T C \cdot \frac{d\Delta V_c}{dt} \cdot V_{dd} \cdot dt = C \cdot V_{dd} \int_{-V_{dd}}^{V_{dd}} dv_c = 2 \cdot C \cdot V_{dd}^2$$

- That is 4 times a single switching energy.

Decoupling by Miller factor of '2'.

- Same direction switching => no current.

Decoupling by Miller factor of '0'.

- Average case: Miller factor of '1' suitable for power-average case sum metric.

Routing Model

- **Via blockage:** Layer multiplier = $(1 - \text{blocking fraction})^{\frac{\text{Low layer pitch}}{\text{High layer pitch}}}$
- **Router efficiency: 0.6**
- **Power grid: 20% of routing**
- **Clock grid: 10% of top tier**
- **More accurate than ITRS 2001.**