MEPTEC Thermal Management Symposium, February, 2006

Sun microsystems

Thermal Challenges for SPARC Based Microprocessors

Bidyut Sen & Jim Jones
Semiconductor Packaging
Sun Microsystems

SP

Make Money
Grow
Re-enlist Champions
Leverage Our Partners
Simplify Our Business

Agenda

- Thermal Roadmap
- Issues in Various Areas
- CMT
- Solution Space

Moore's Law

What is the best use of these transistors?

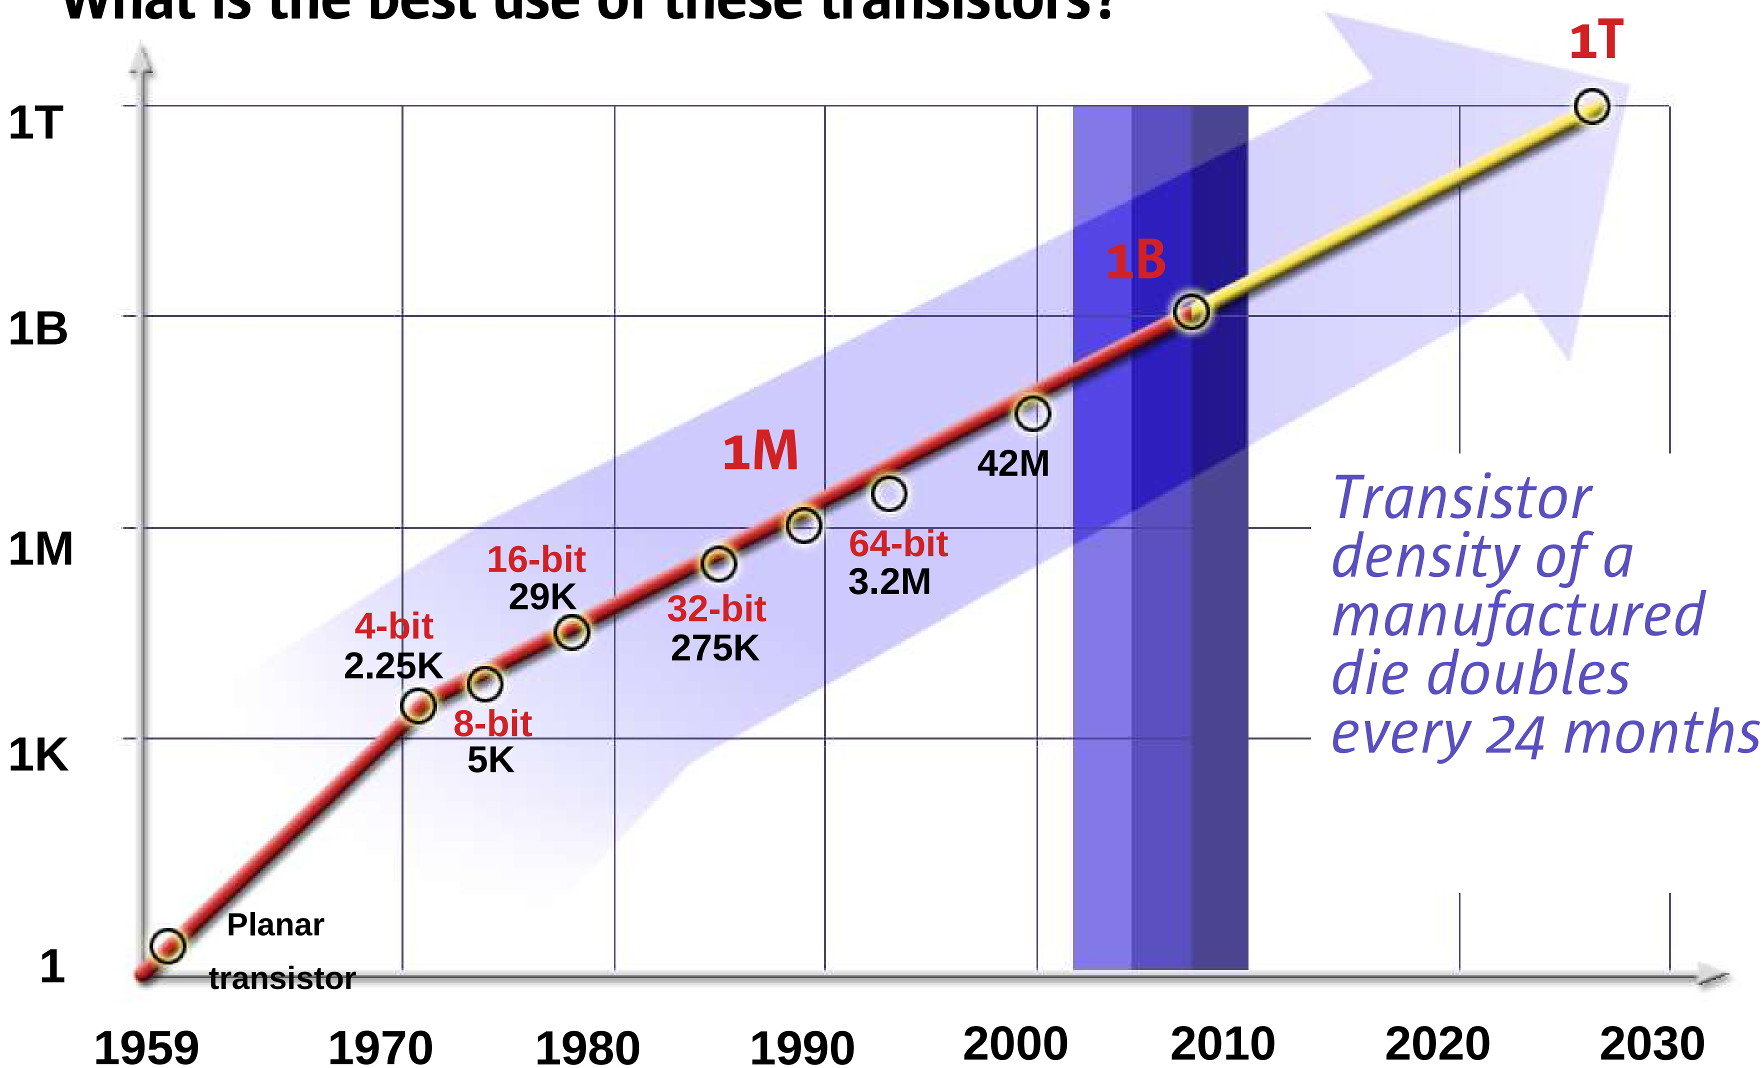

Microprocessor Thermal Dissipation Roadmap

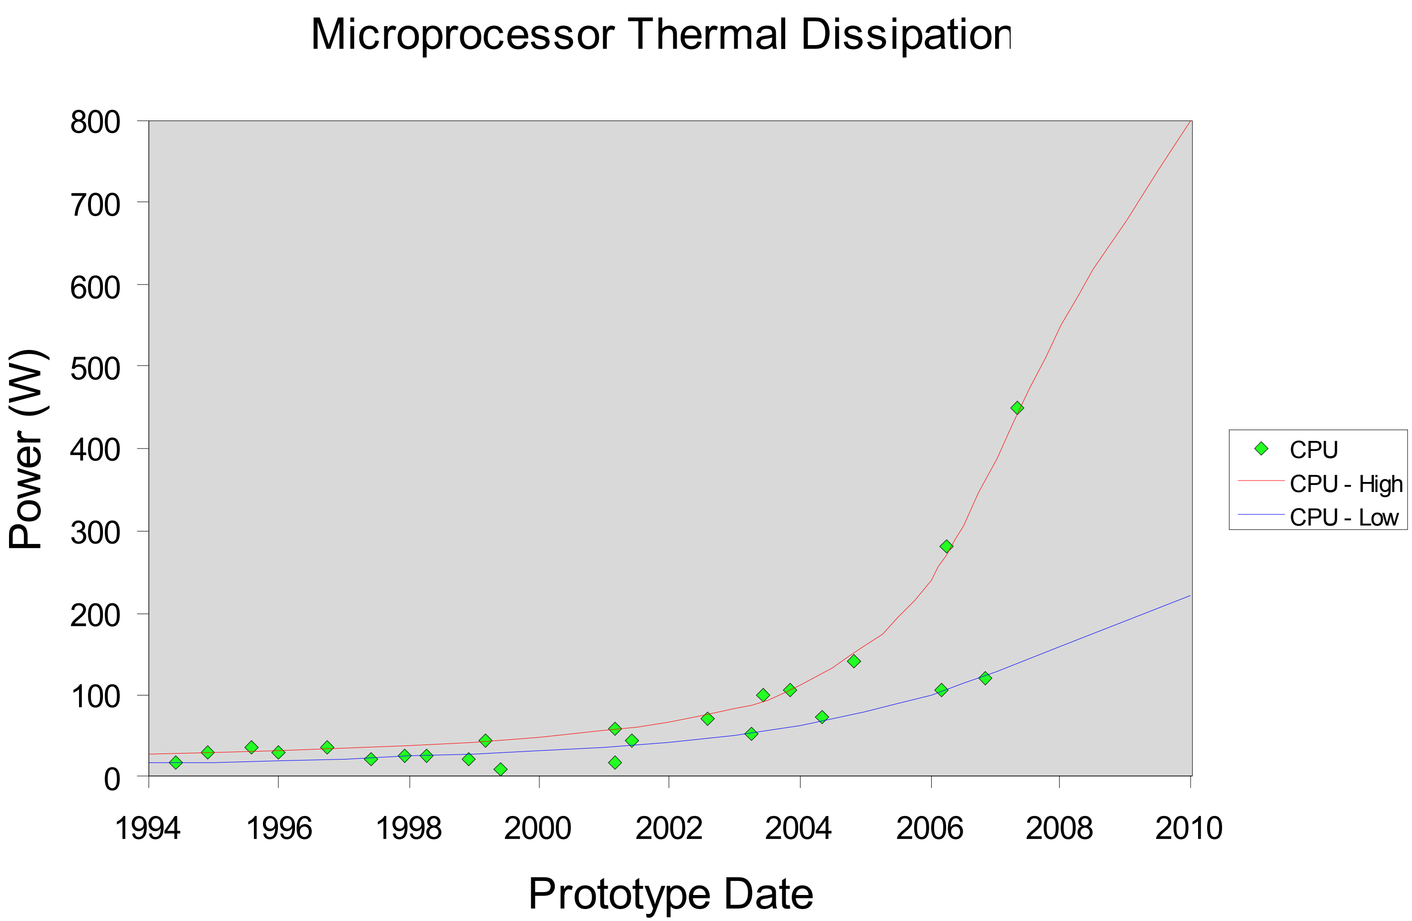

Industry System Thermal Power Density

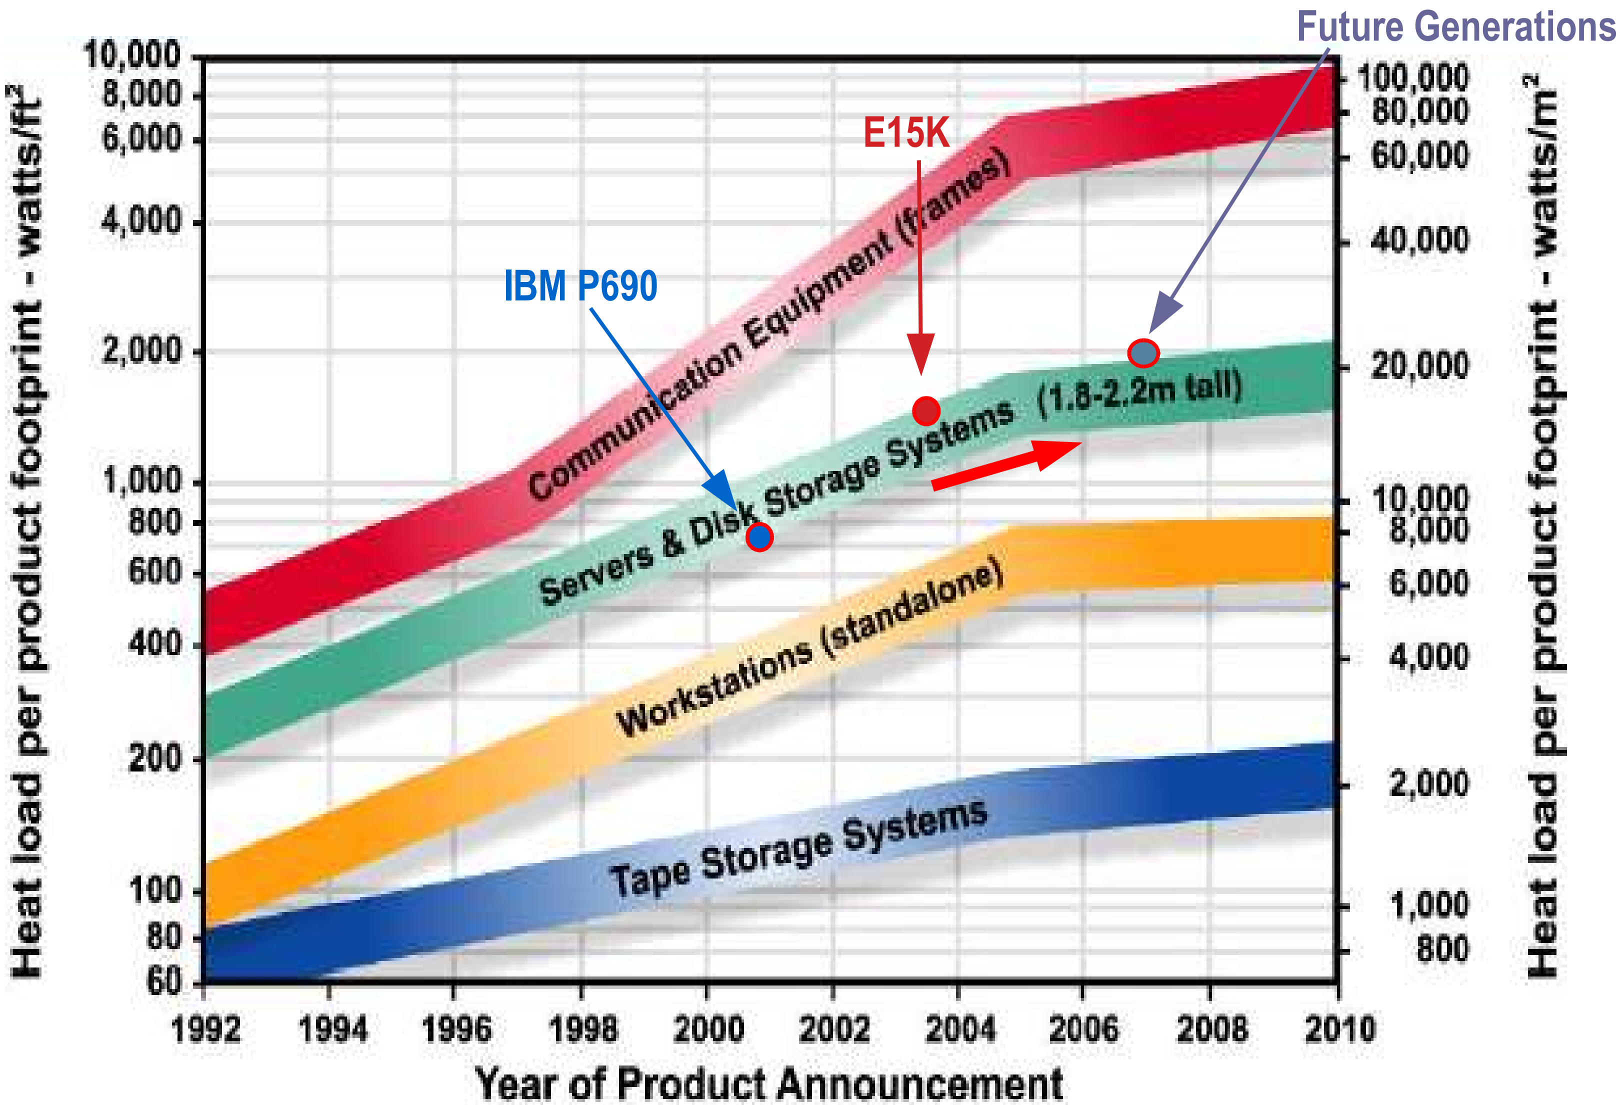

Processor Trends

Spending transistors on complex high frequency processors is showing diminishing returns....

Frequency

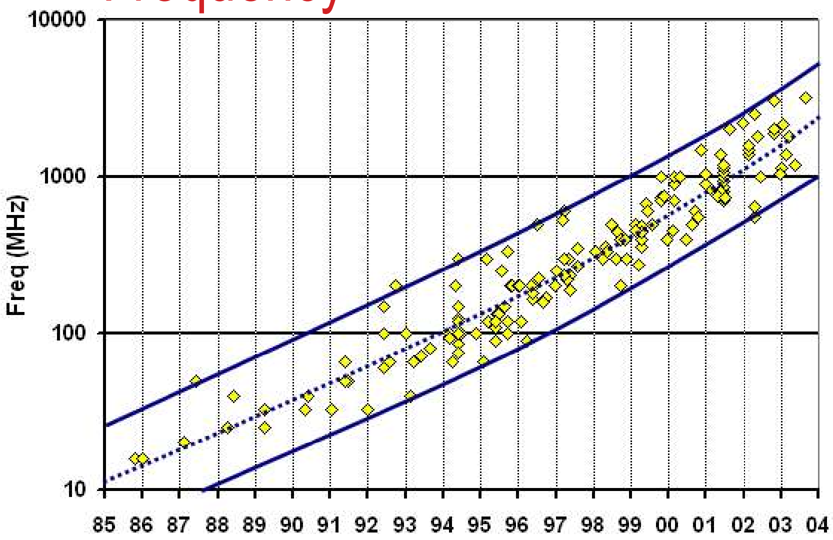

Performance / Power Density

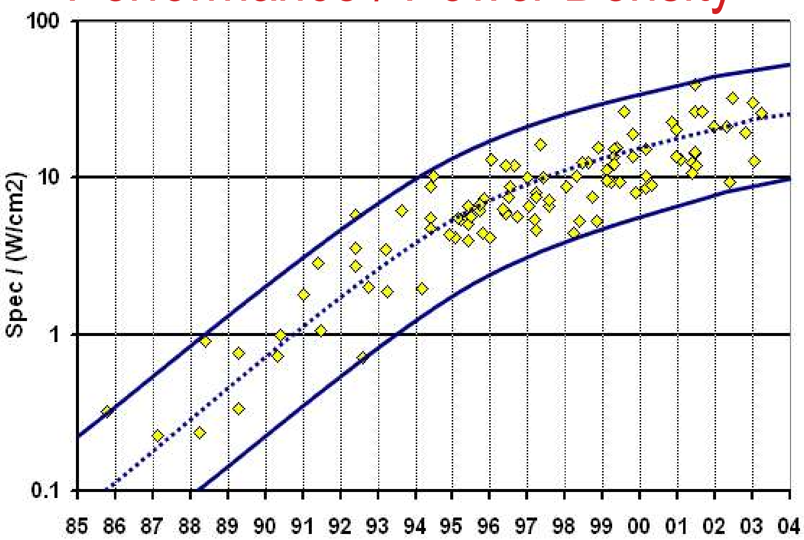

Issues:

- Memory Latency
- Design Complexity
- High power & hot chips

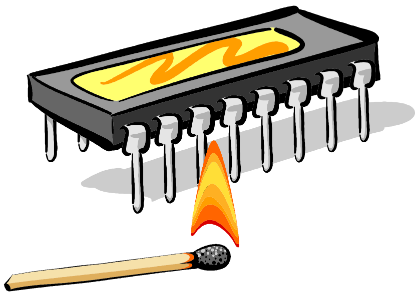

Total Chip Dynamic and Static Power Dissipation Trends Based on the ITRS — Michigan: Kim et al., 2003

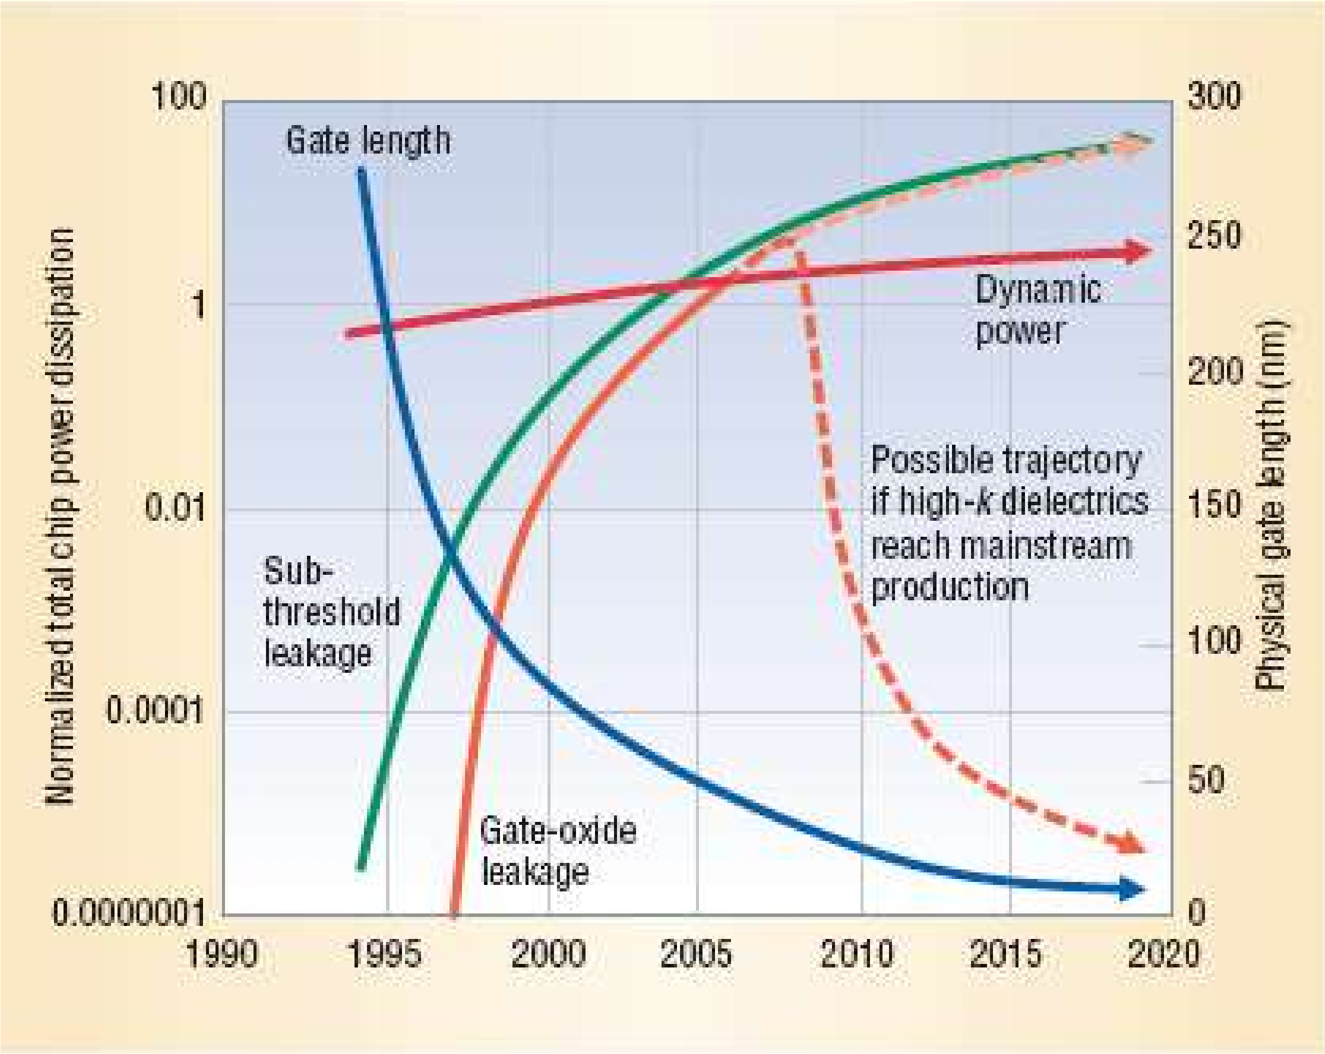

Real End Use Condition - Ranch

Ranch Machine – CPU Temperature

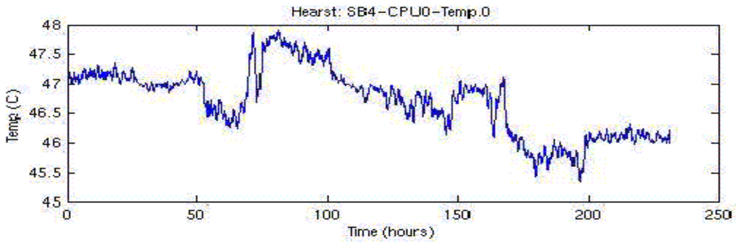

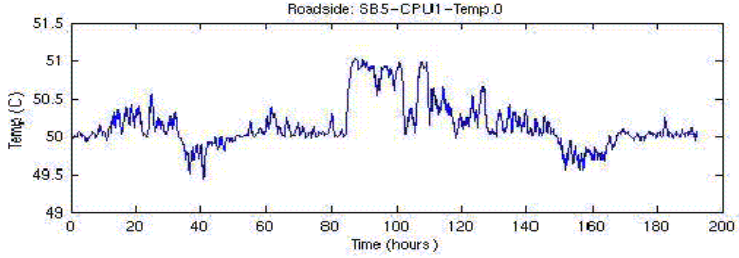

< 3C fluctuations

Idle at 45C or 49C

First Forced Mini-Cycle Results (Temperature)

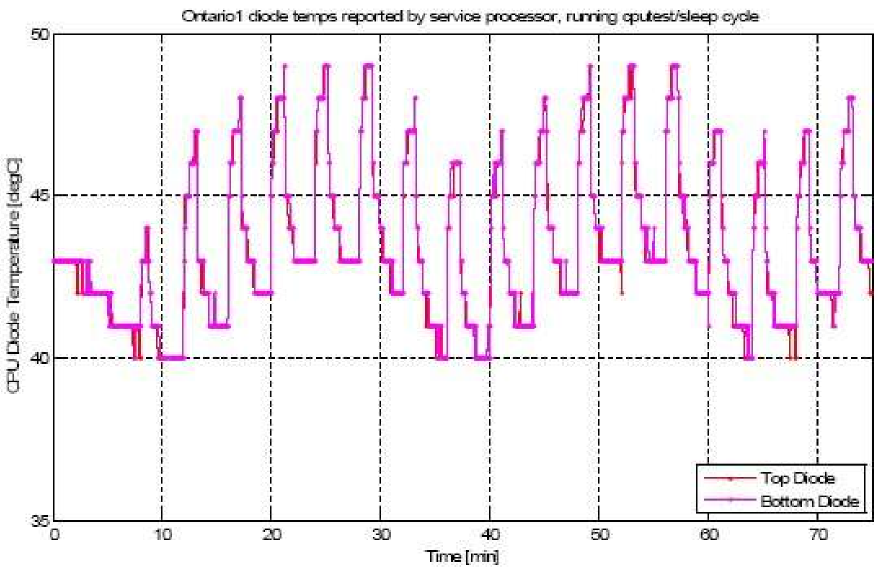

- Temperature swing is ~ 9 degrees, including the low-frequency component.

Courtesy of Kenny Gross & Kalyan Vaidyanathan and Leesa Noujeim

ron.zhang@sun.com

Determine Equivalent ATC based on Used Condition

Assume

Solder Damage by Power Cycles = Solder Damage by ATC

Then

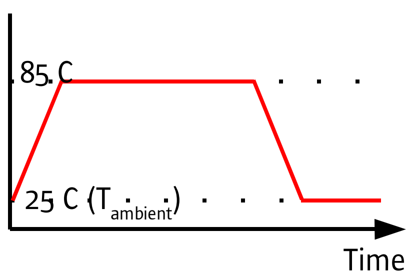

+

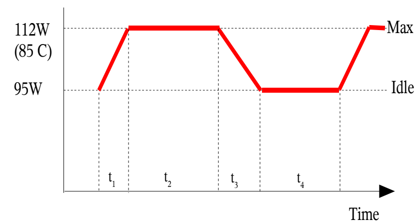

=

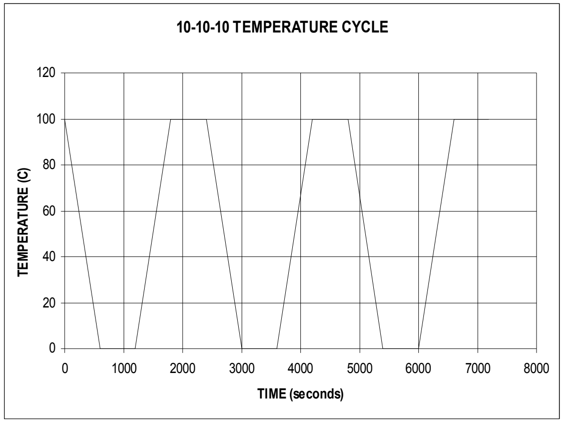

$$(d_{\text{big cycle}} \times N_{\text{big cycle}} + d_{\text{mini cycle}} \times N_{\text{mini cycle}}) = d_{ATC} \times N_{ATC}$$

d – Damage per cycle
N – Number of cycles

ron.zhang@sun.com

CMT Processor Advantages

Addresses Bottlenecks - Memory, Power

- Simpler cores operate at lower frequencies while maintaining processor performance/throughput
- No extra transistors (and power) for multiple-issue or out-of-order execution
- Time-share a simpler pipeline
- Share under-utilized resources (crossbar switch): Memory and I/O controllers, Second-level caches
- Improved power efficiency & Die Yields
- Active power & temperature control by scheduling or idling threads & cores

CMT – Multiple Multithreaded Cores

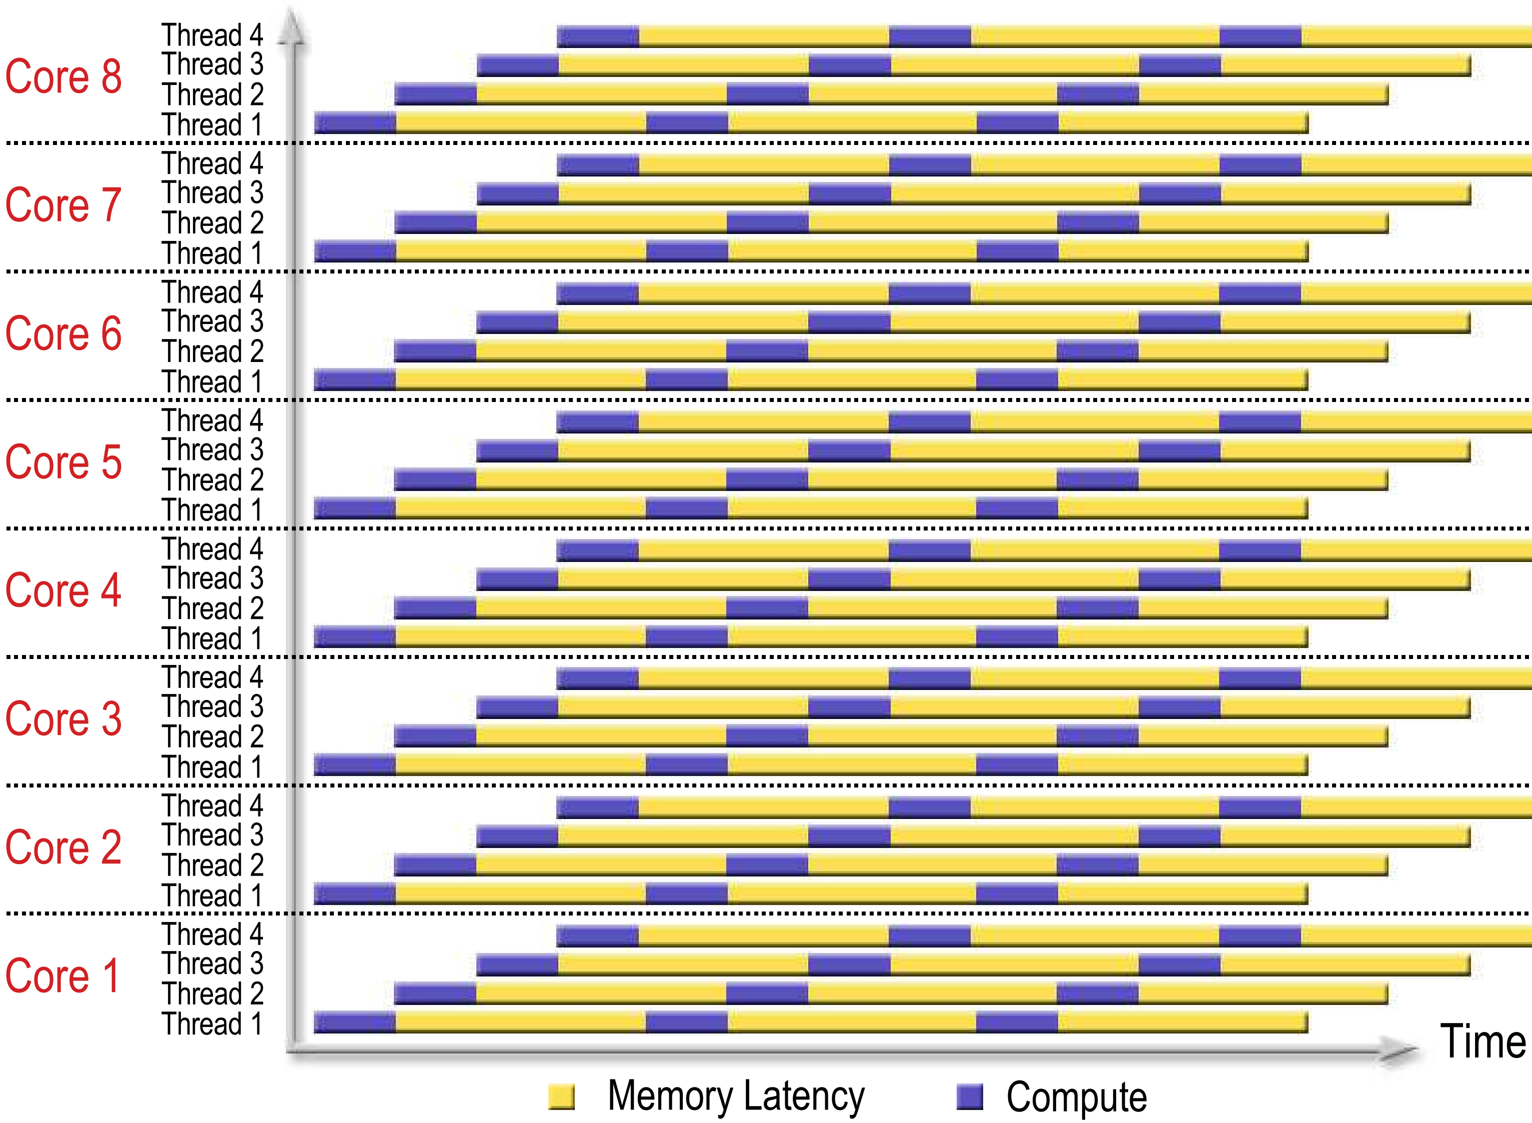

Processor Power Density Trends

CMT Decreases Power Density without Sacrificing Performance

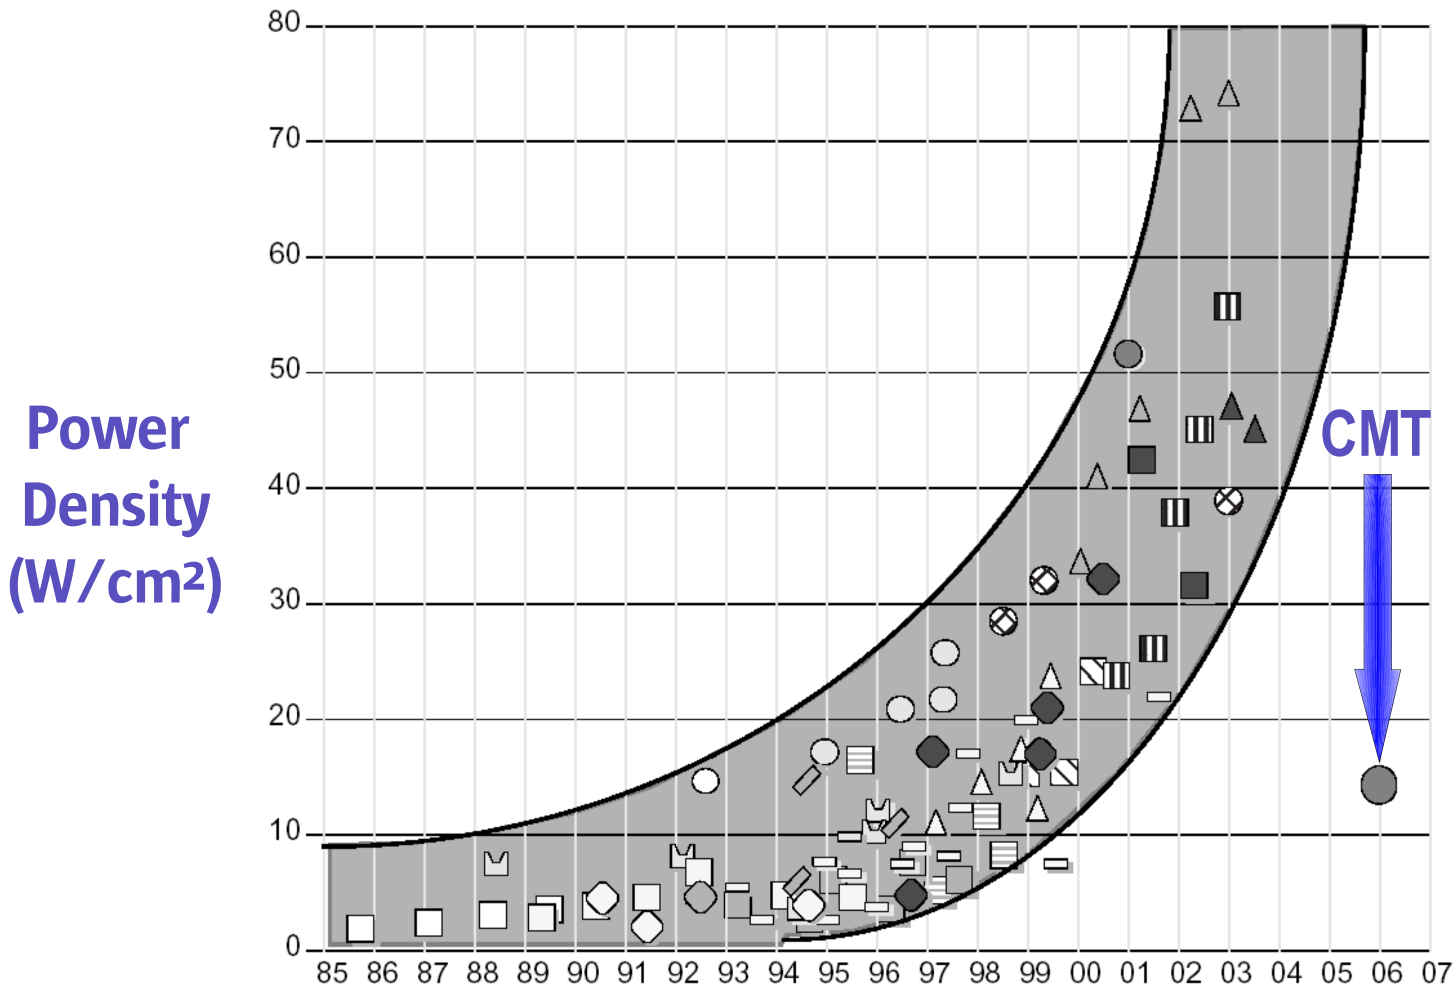

CMT Processor Thermal Advantage

"Cool Threads"

- Cool – improves performance, power & reliability
- Uniform – peak power close to average power, improves thermal management, clock distribution & reliability

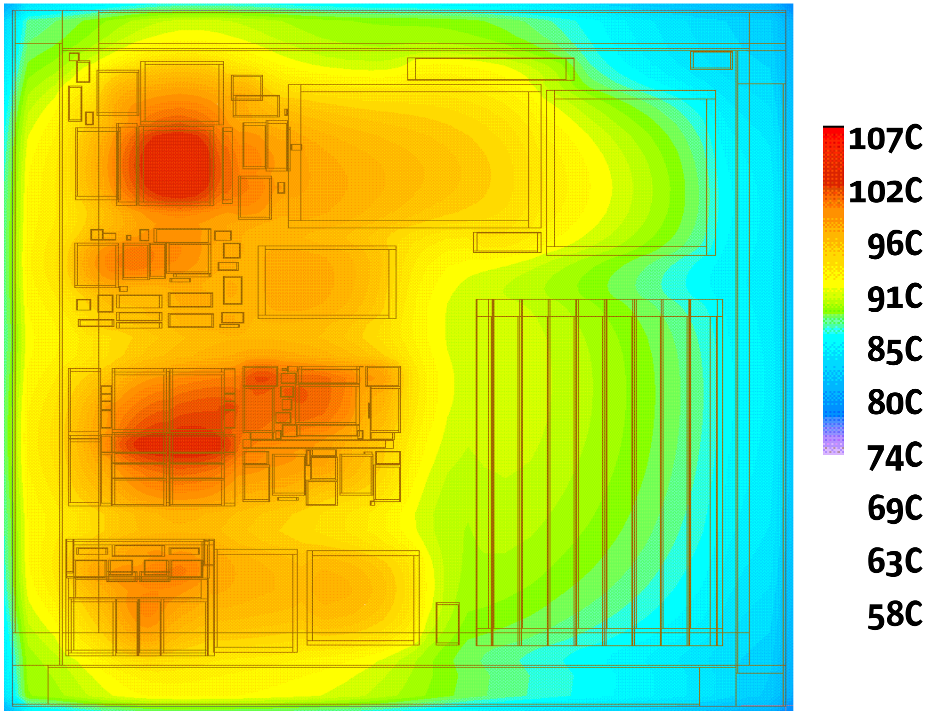

Single-core Processor

(size not to scale)

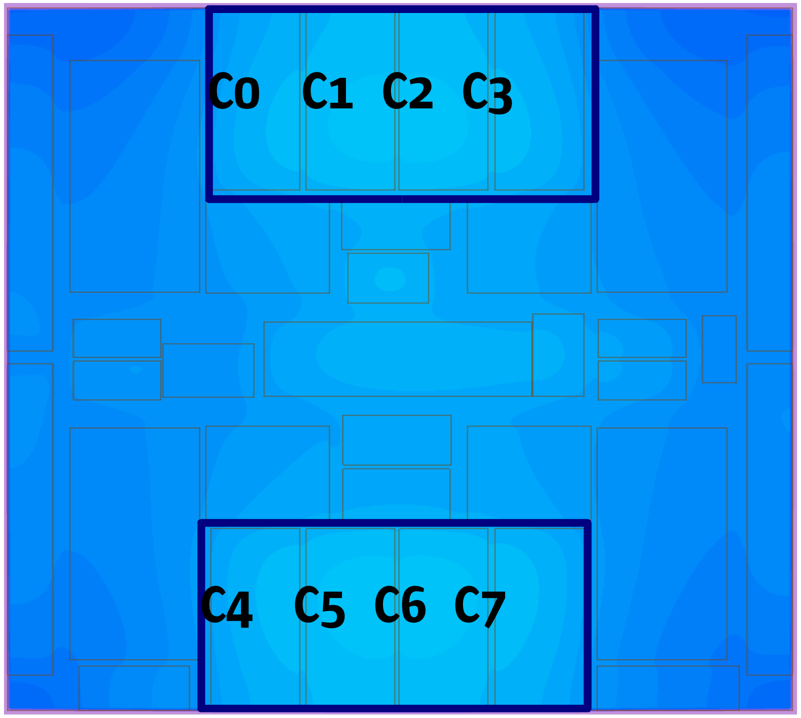

CMT Processor

Cooling Hierarchy

- Package
- Board Level
  - TIM2
  - Heat Sink
  - Advacned Cooling Solutions
- Systems
- Facilities

Concept of "Thermal Resistance"

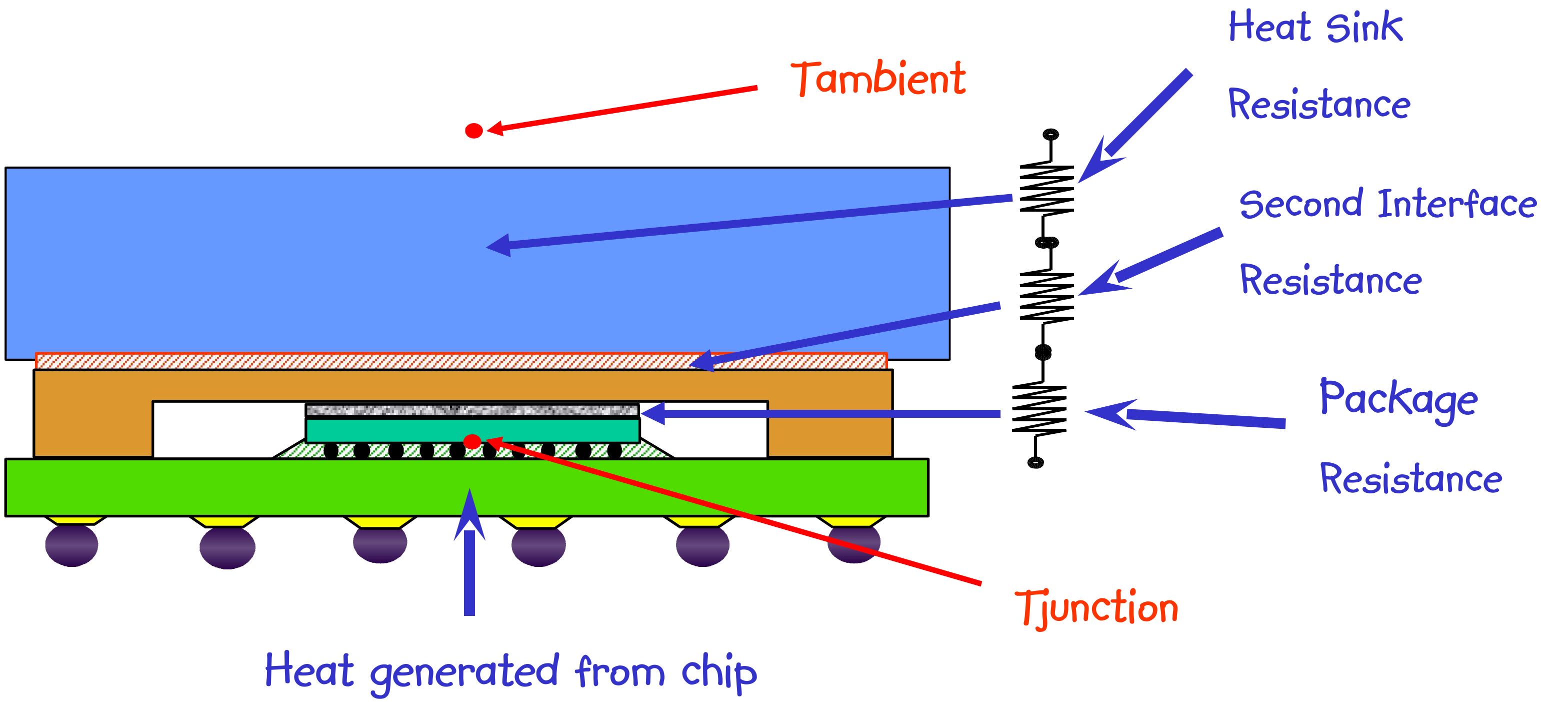

Solution Development Trends Thermal Budget Breakdown

Future Products

Current Product

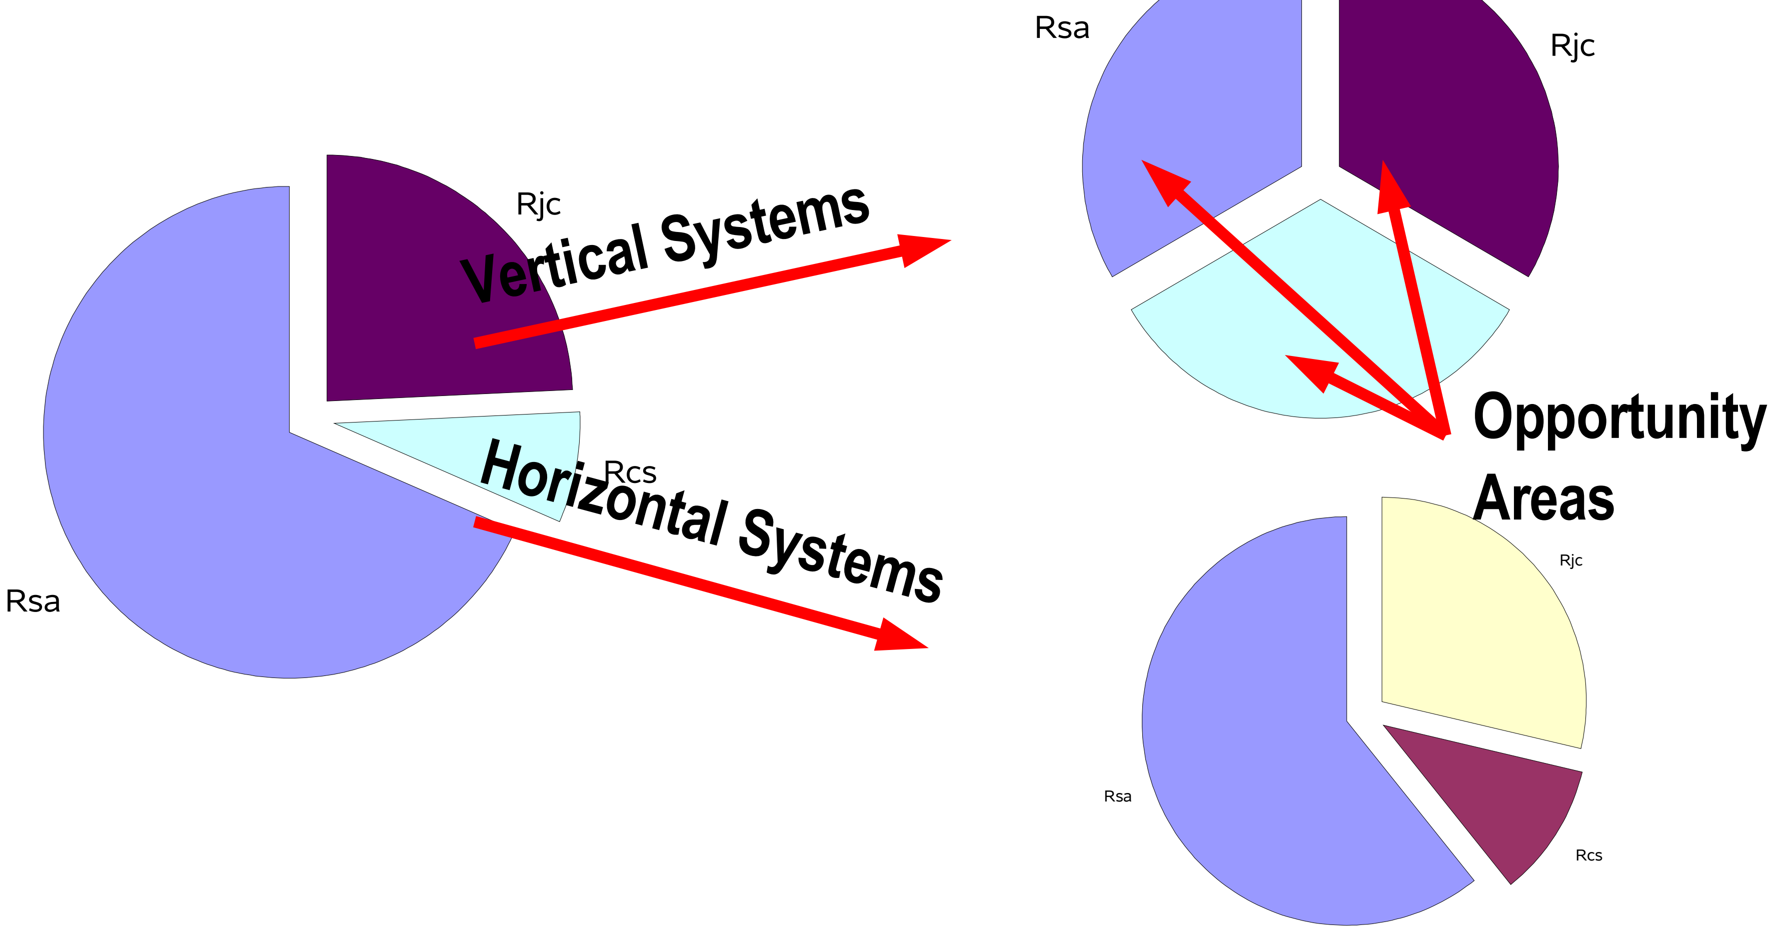

Selected DOE Solution Space

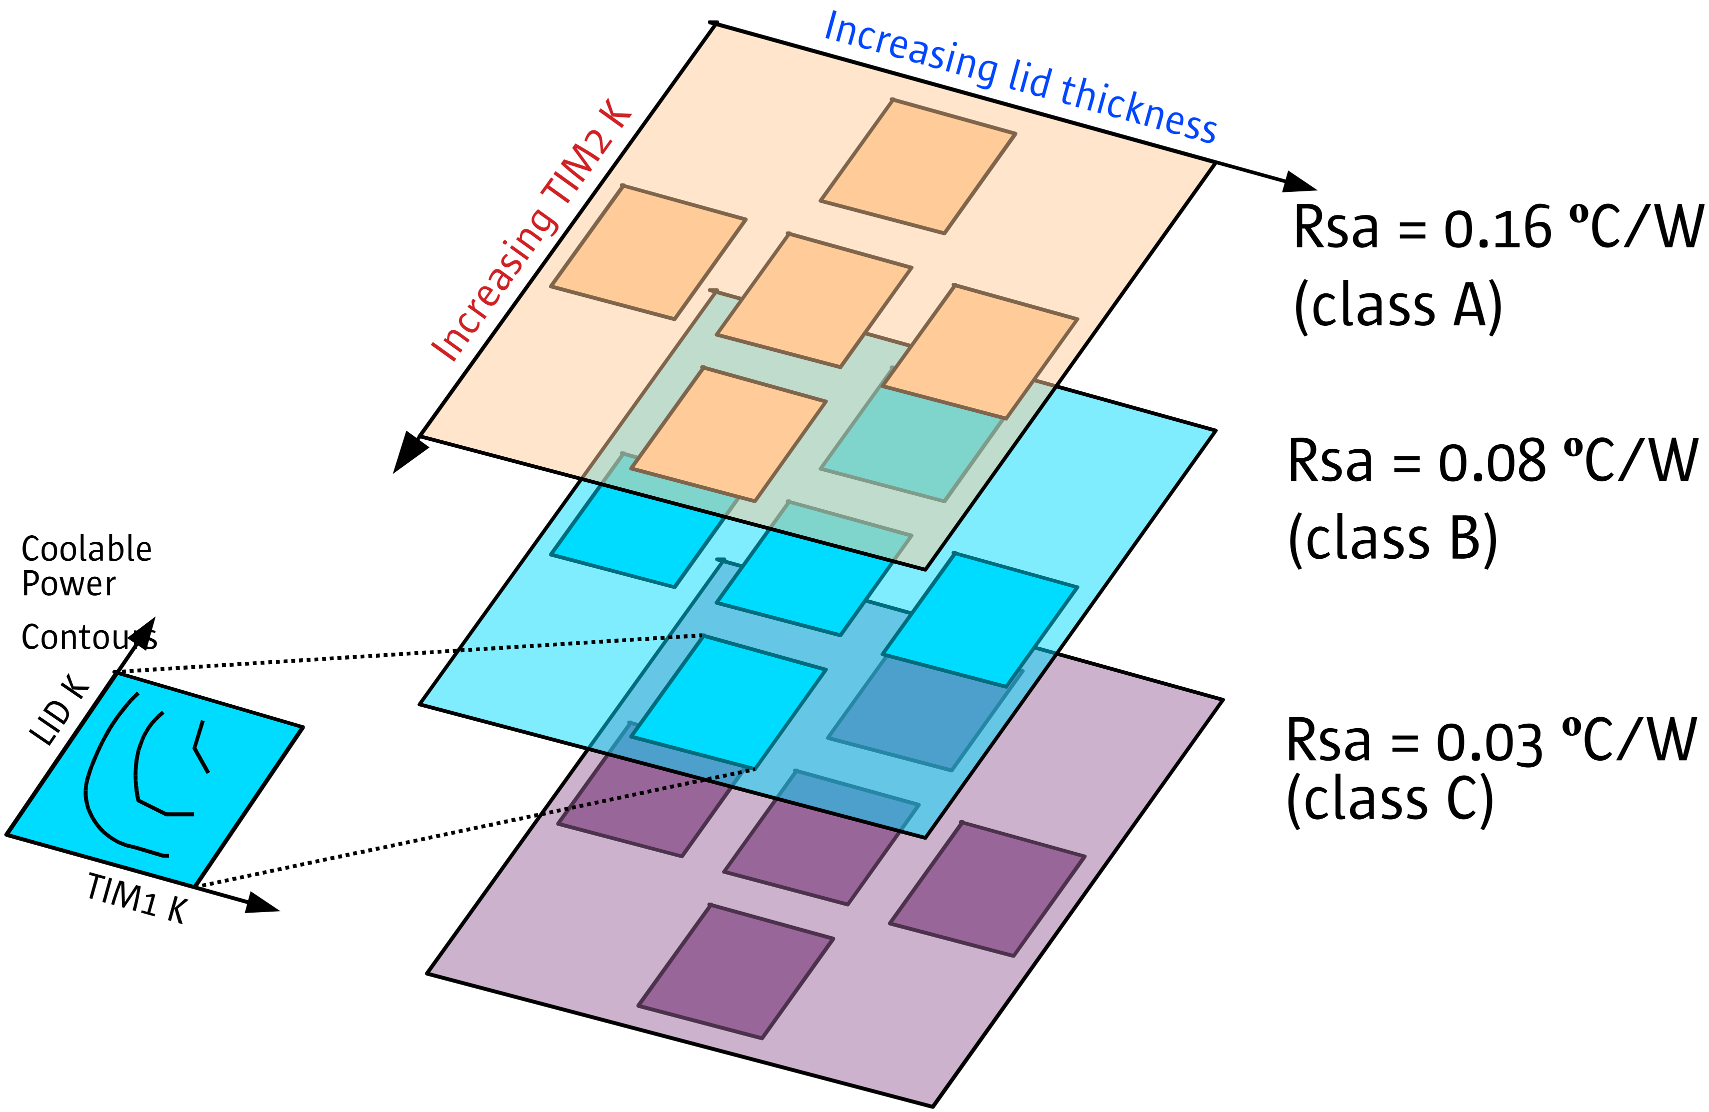

DOE Study Conclusions

- The DOE analysis reveals different contributions to the coolable power gains from the considered **design variables**
  - Within the package – **TIM1** effective thermal conductivity has the most impact, followed by lid conductivity
  - Outside the package – the **cooling solution, Rsa**, has the most impact
  - **TIM2** has a small impact up to k ≈ 10 W/mK; and even that only with highly efficient cooling solutions
  - Coolable power is largely independent of **lid thickness** for TIM2 k below ~10 W/mK. For TIM2 k above ~10 W/mK, increasing lid thickness could even be detrimental at CuW-like lid conductivity; this trend reverses at higher lid conductivity (like ScD)
- Improving TIM1/lid/TIM2 conductivity, alone with less efficient cooling solutions (Rsa ≈ 0.5 C/W), produces limited coolable power gains. To achieve the full potential, these changes must be made in conjunction with improvements in the cooling solution.

Thermal Solution Development - Rjc Improvements (Lid)

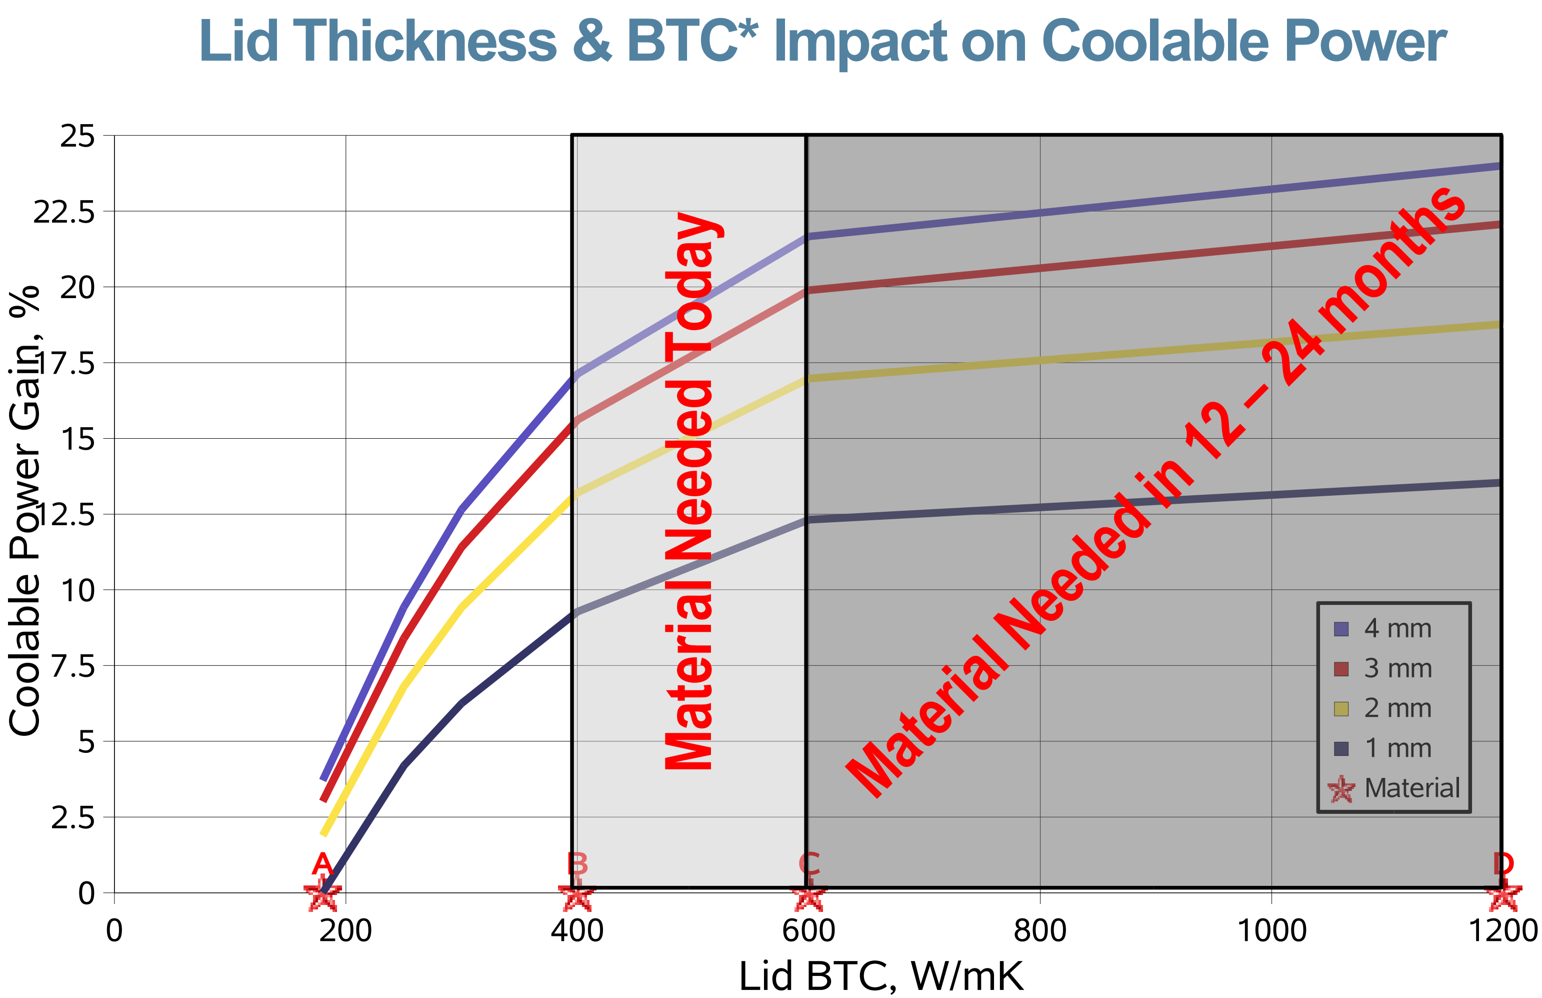

*BTC – Bulk Thermal Conductivity

Thermal Solution Development - Rjc Improvements (TIM1's)

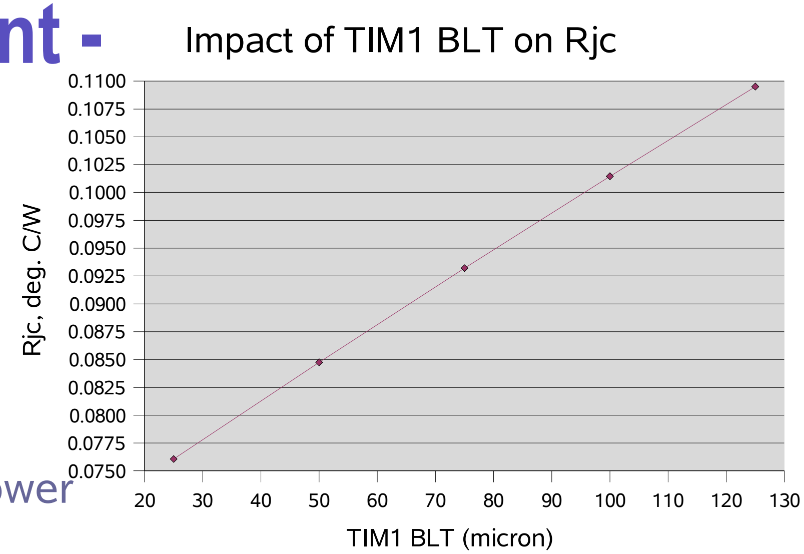

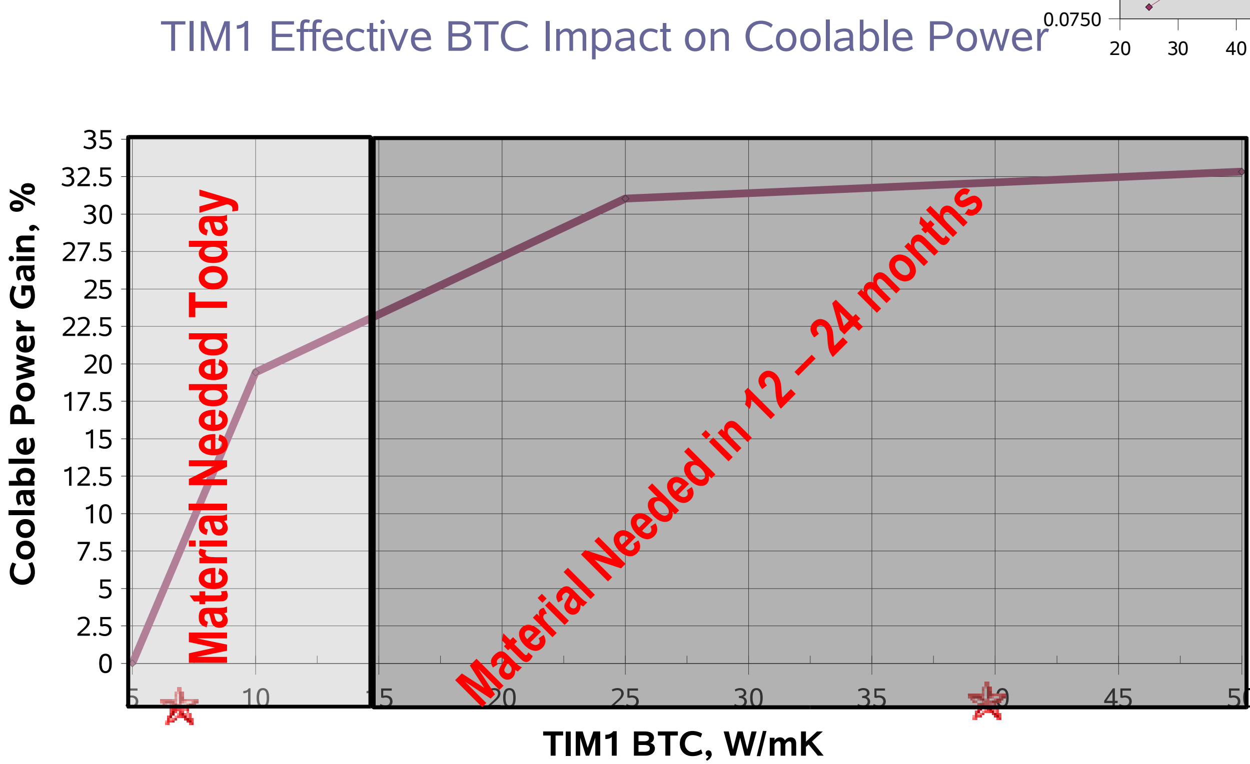

*BTC – Bulk Thermal Conductivity

Impact of TIM2 BTC* on Rjc

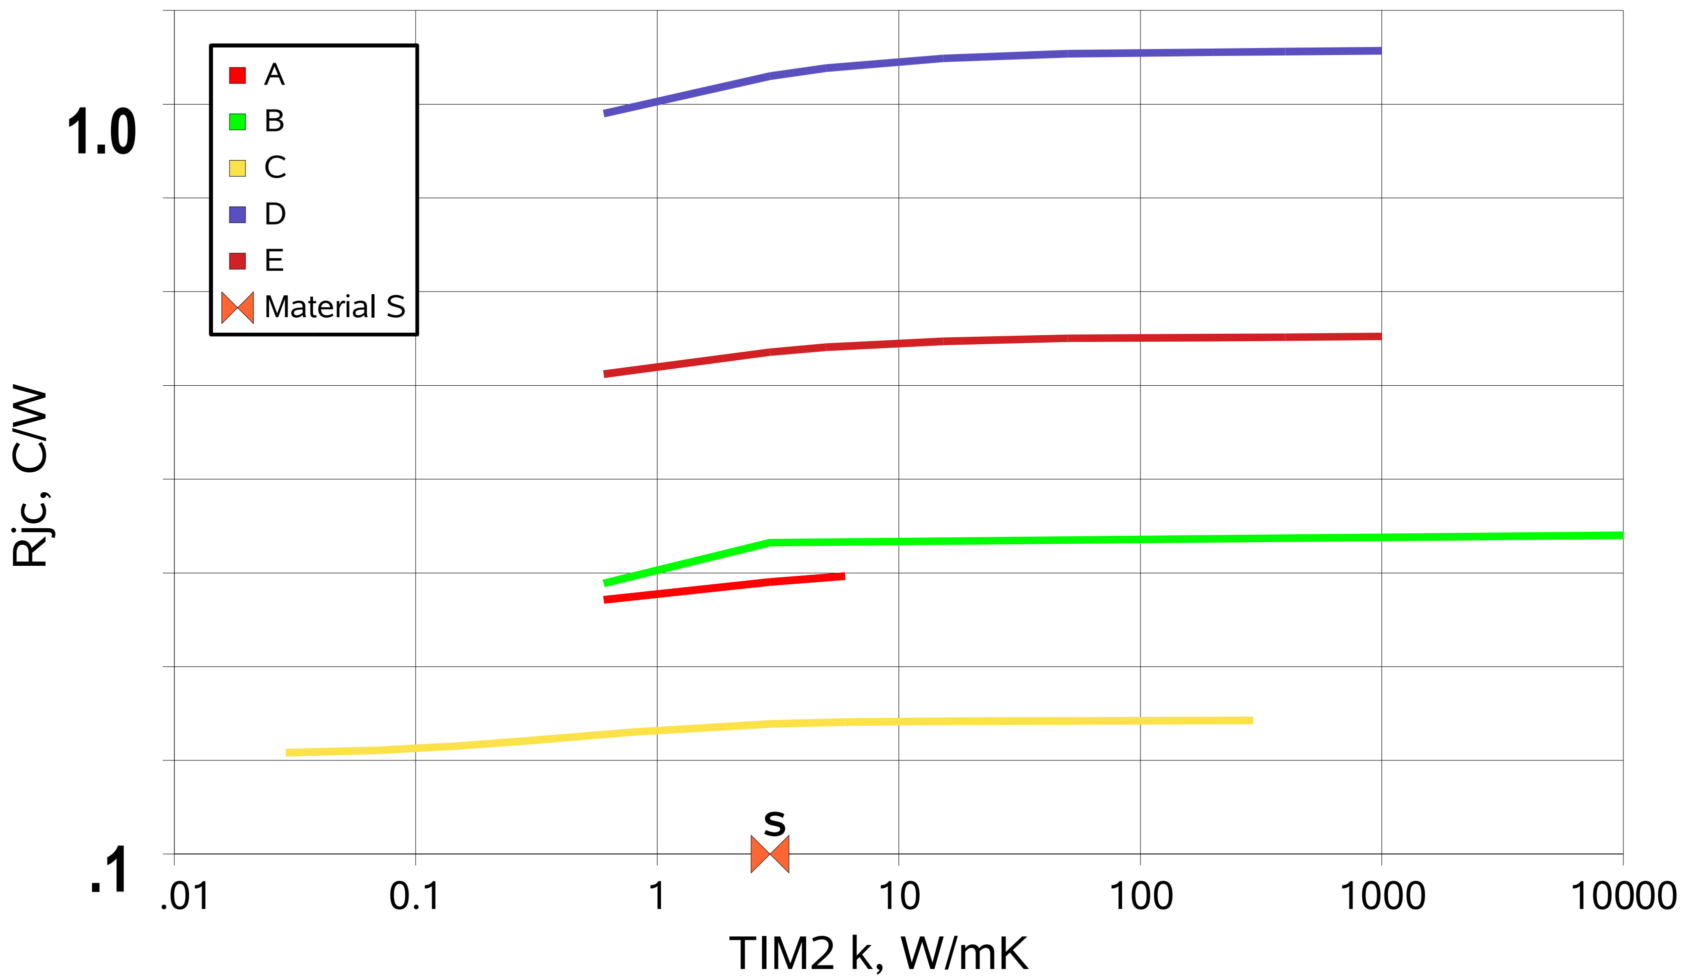

*BTC – Bulk Thermal Conductivity

Impact of TIM2 BTC* on Die Thermal Gradient

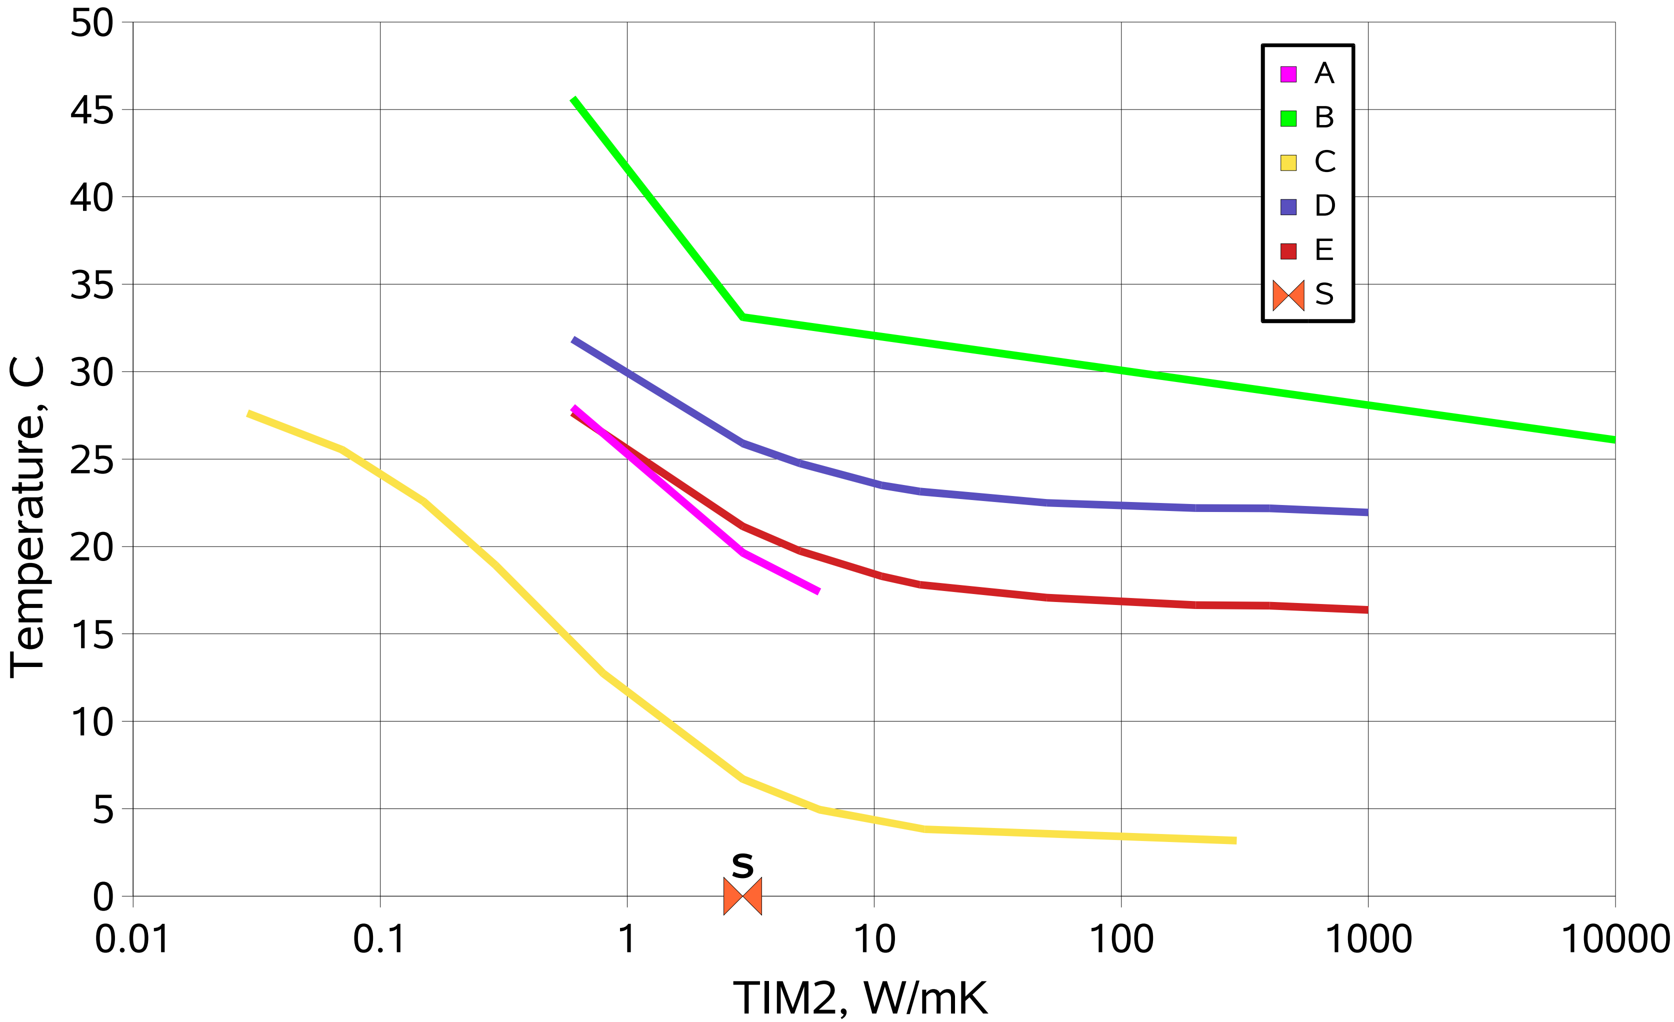

*BTC – Bulk Thermal Conductivity

Cooling Solution Options

- BLT Improvements
  - > Thinner BLT
- Organic TIMs
  - > Higher effective TC
  - > Material investigation
- Inorganic TIM
  - > Materials Availble
  - > Process validation
- Advanced Lid Materials
  - > Issues with supply
  - > unproven reliability

bidyut.sen@sun.com

Cooling Solution Options

- TEC
  - > Early stages of evaluation – potential is uncertain
- Lidless
  - > Heat Spreading with Air Cooling
  - > Handling Issues
- Composite Die
  - > Thinner silicon with diamond-like layer
  - > Needs Development

bidyut.sen@sun.com

Cooling Solution Options

- **Liquid Cooling**
  - > Direct and Indirect Attach
  - > Liquid Cold Plate
  - > System Level
  - > Data Center Level
  - > Similar complexity to refrigeration with much shorter usable life and small performance gain over air
- **Refrigeration**
  - > Infrastructure
  - > Reliability

Microprocessor Power Projection

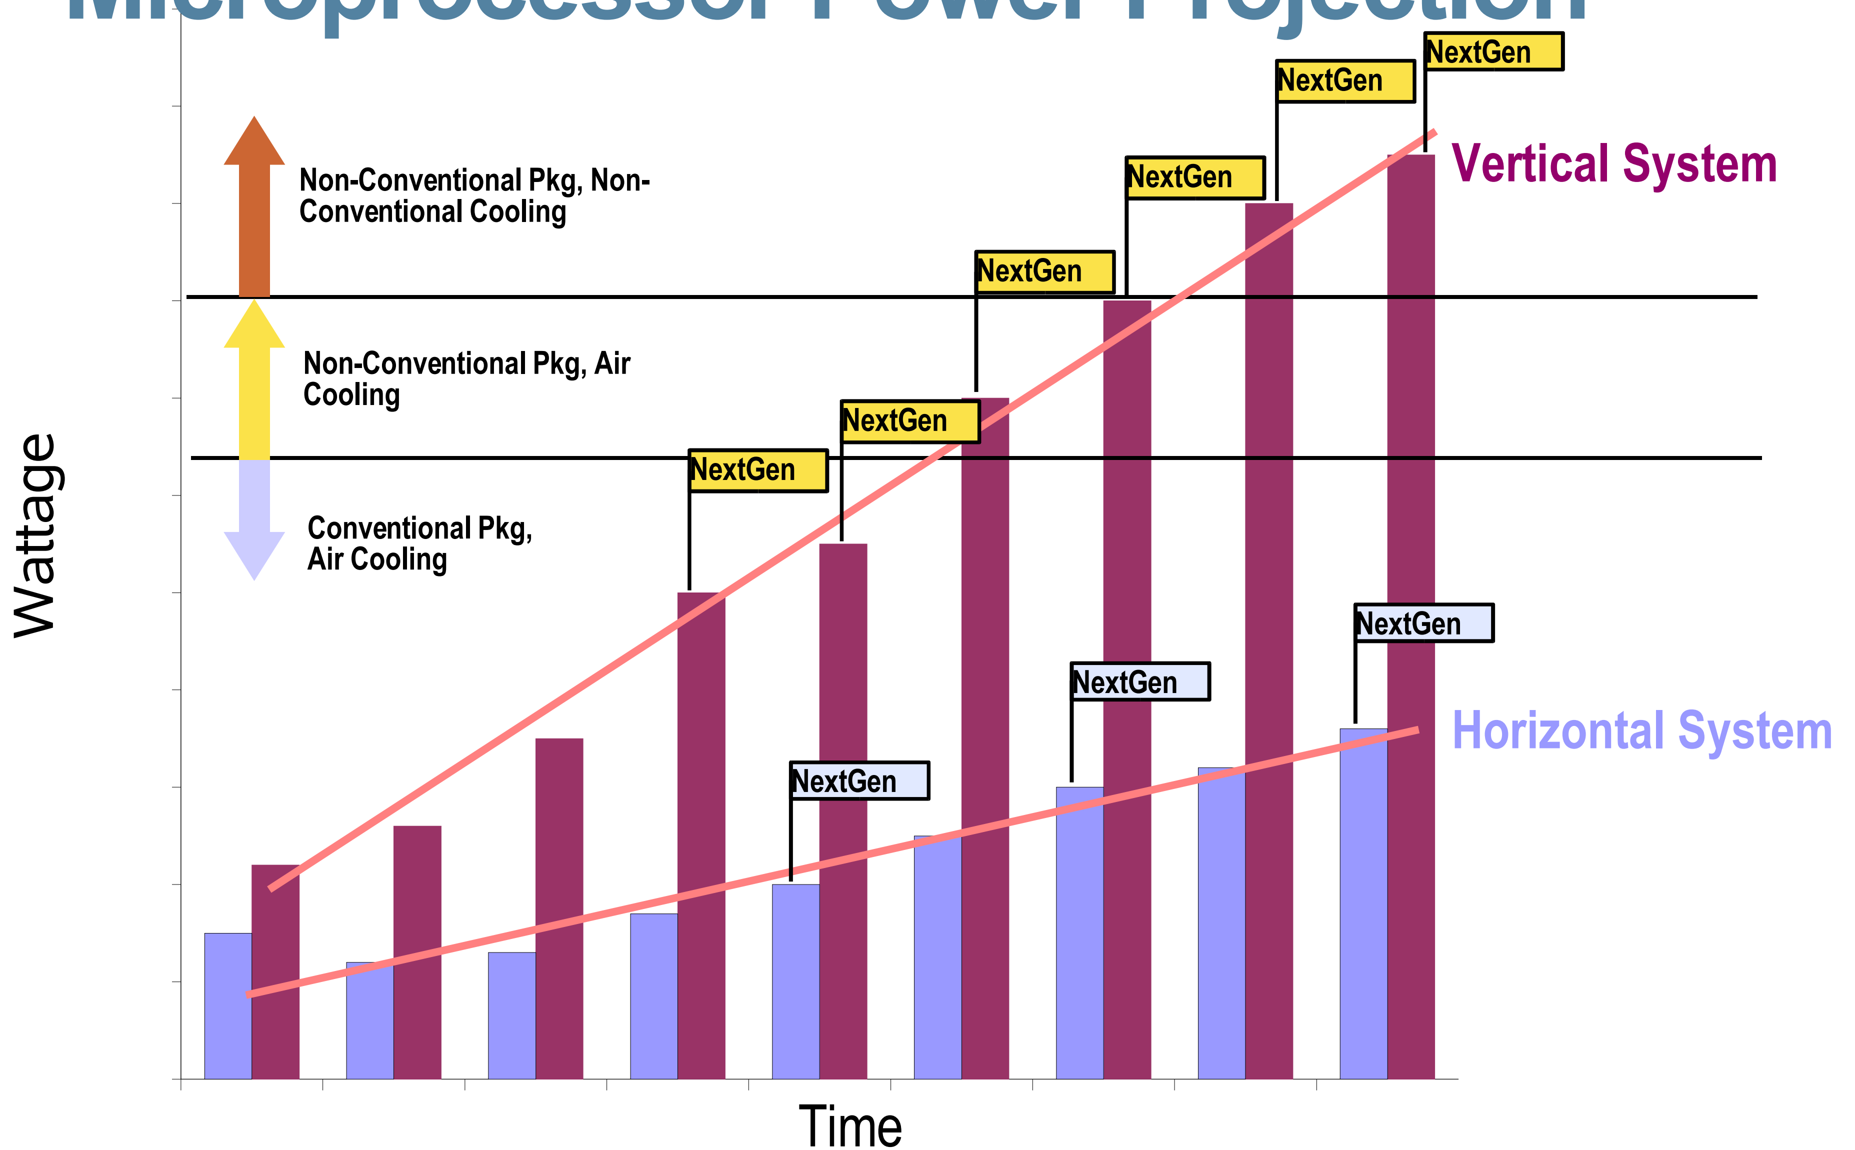

Cooling Options

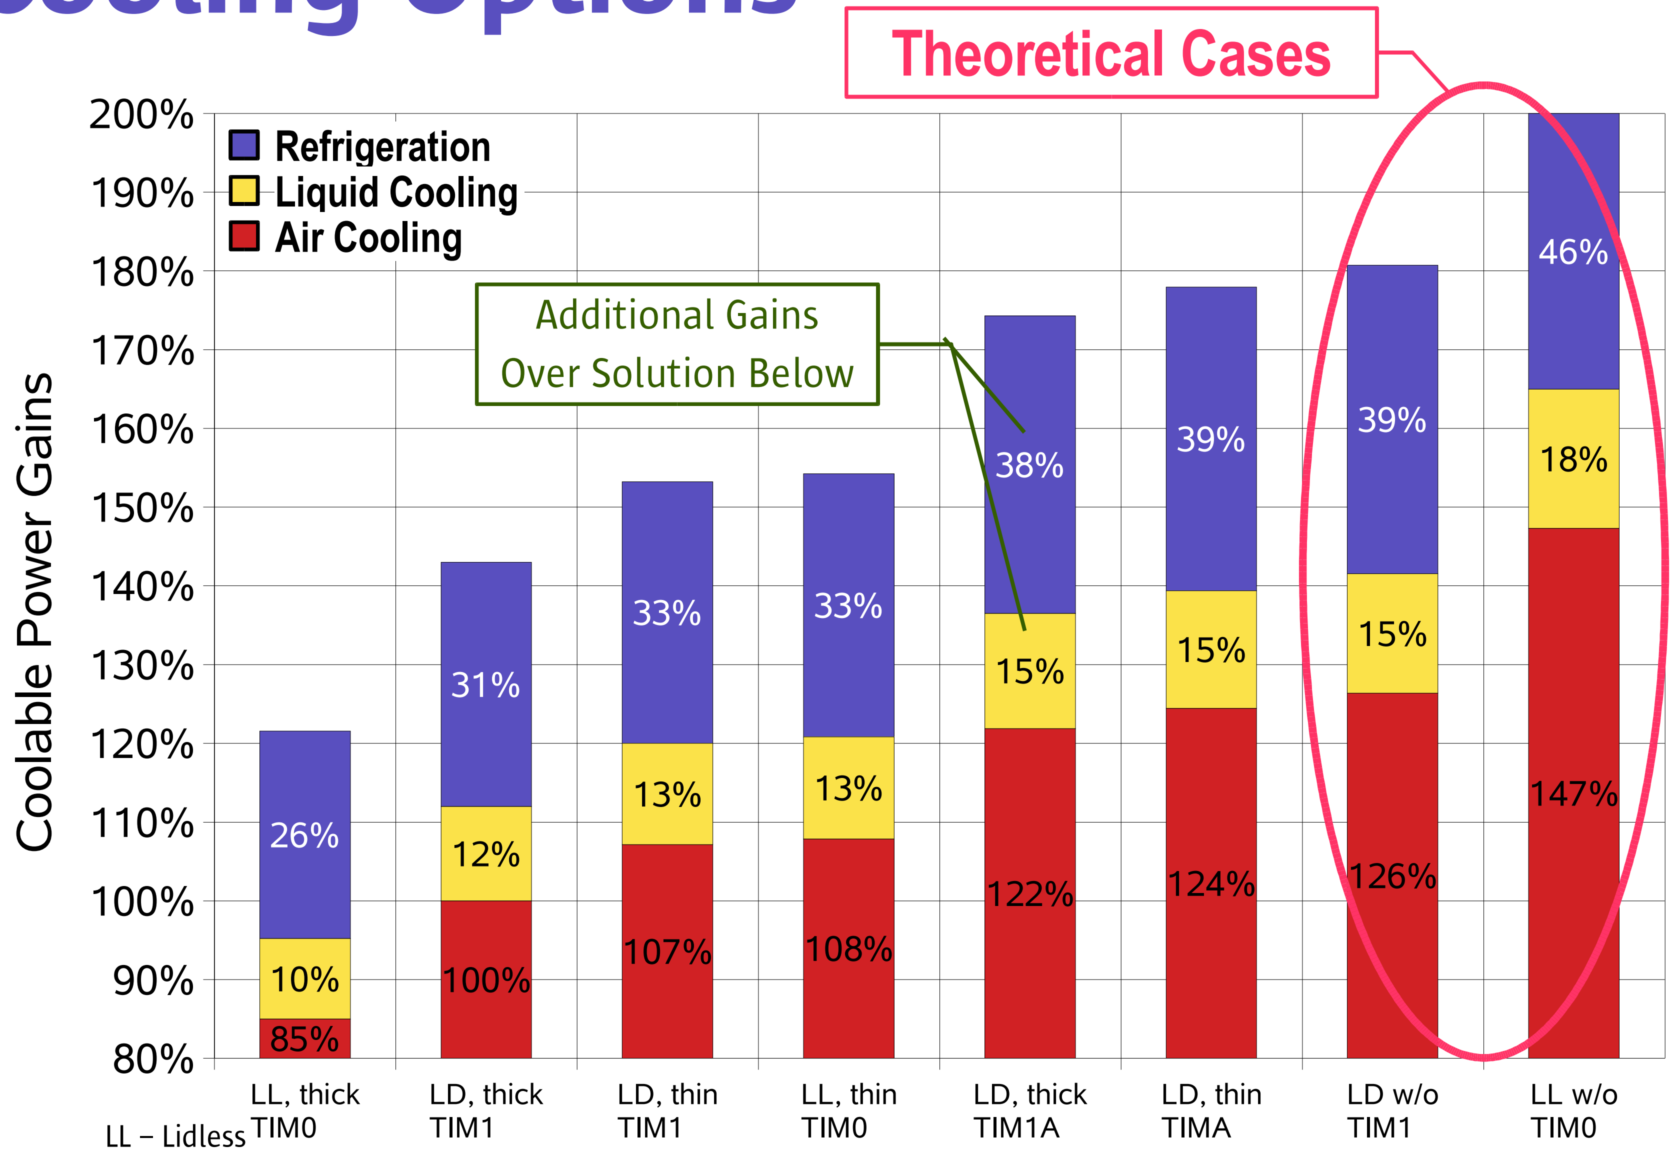

LL – Lidless

LD – Lidded

Thermal Solution Development - Rsa Improvements

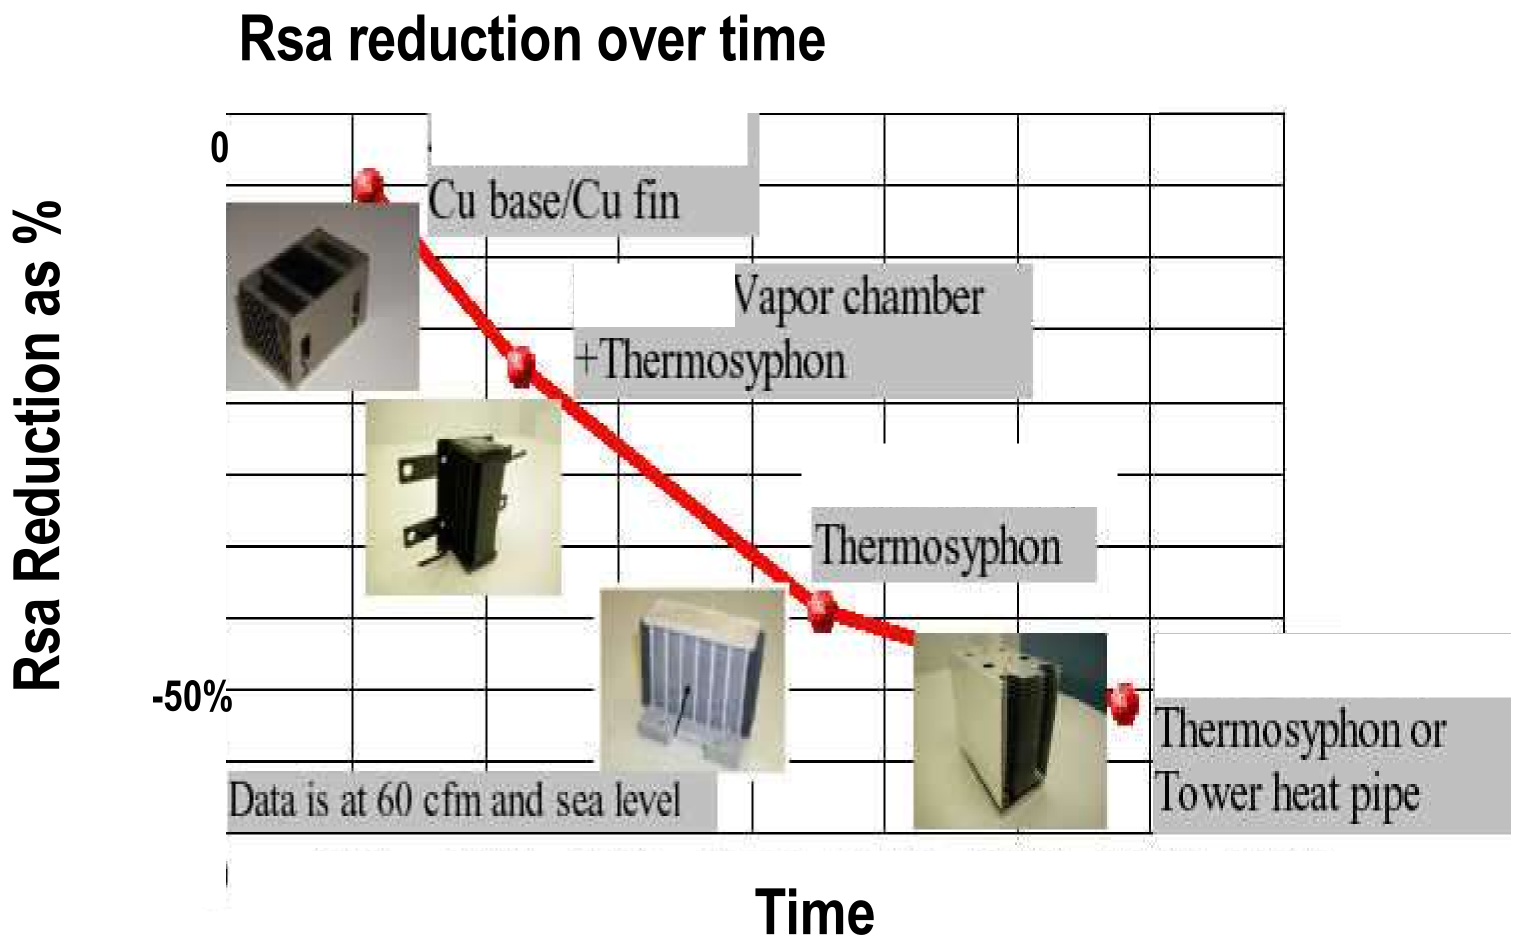

source: Guoping Xu

Chip, compressor, and total power

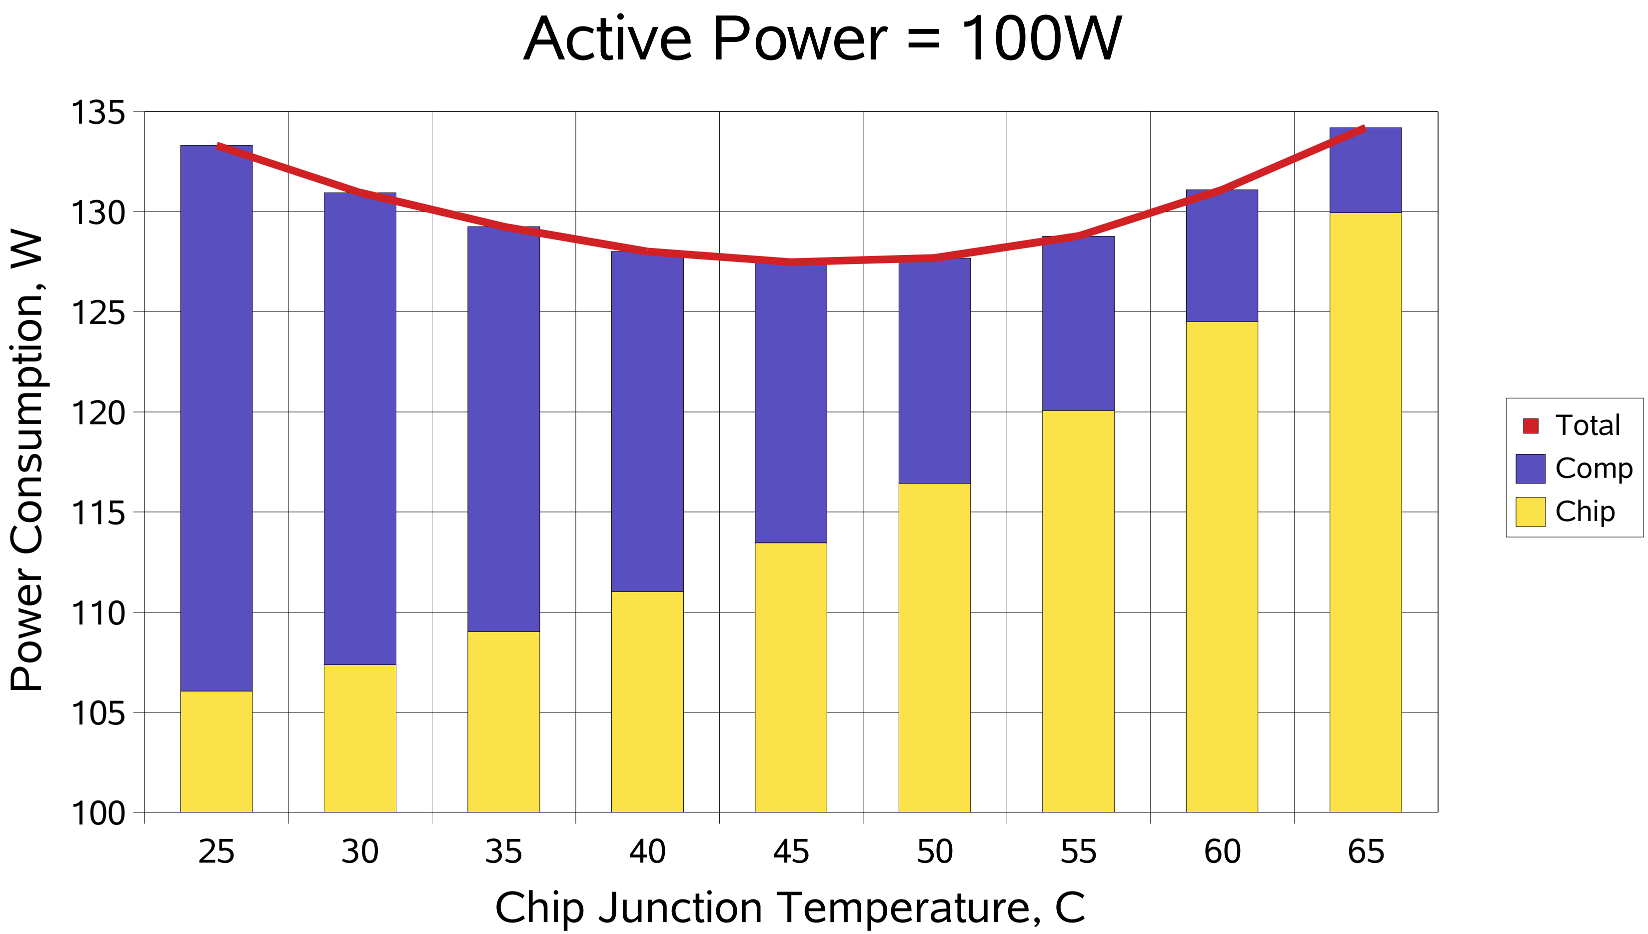

Sun microsystems

Data Center Cooling Options

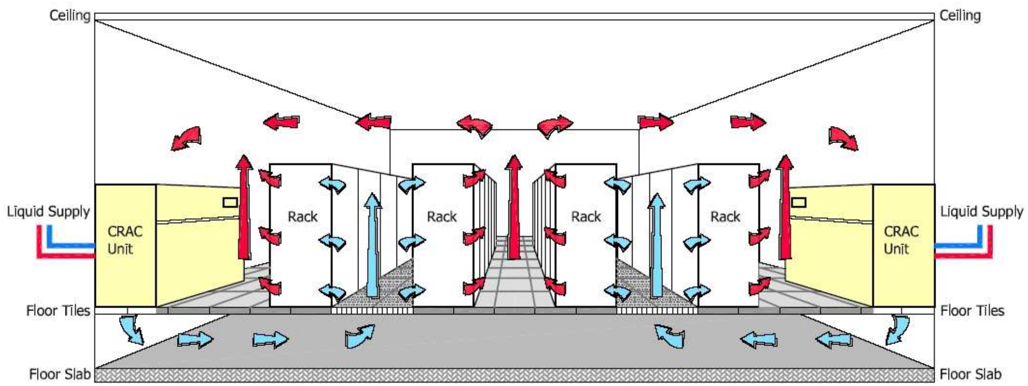

Past

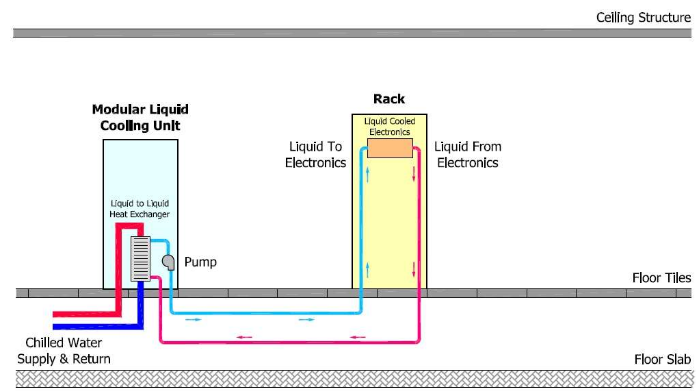

Future

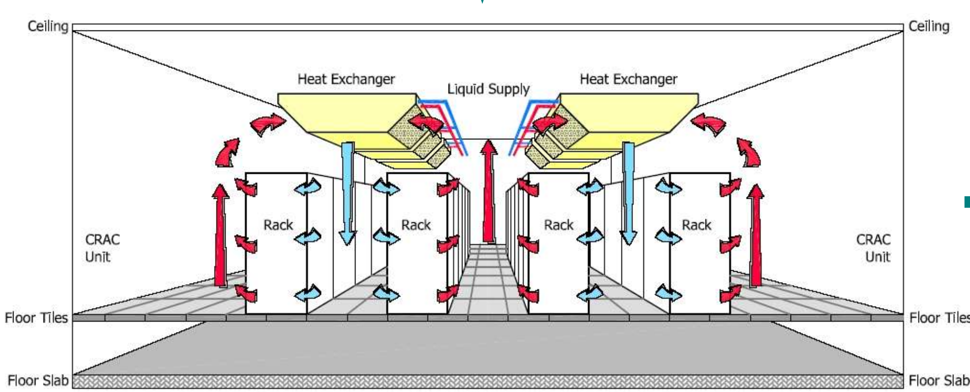

Today

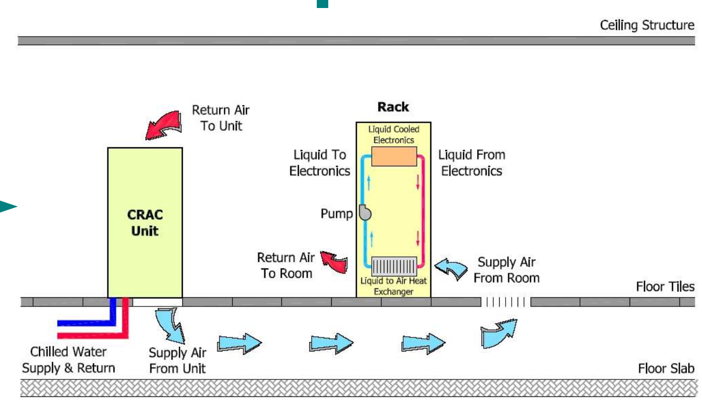

Tomorrow

bidyut.sen@sun.com

FLOORPLAN

What is Junction-Case Thermal Resistance (R_{jc})?

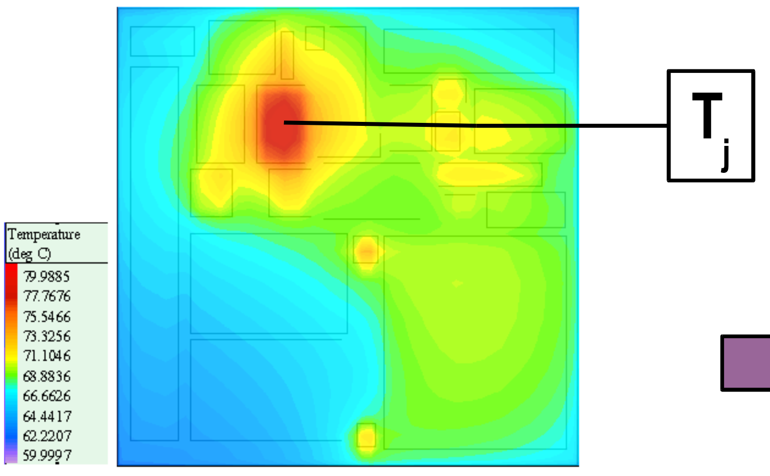

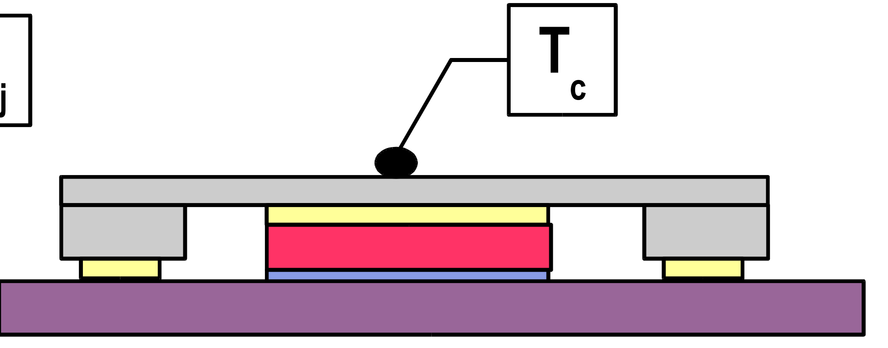

- R_{jc} is an industry-wide accepted measure of package thermal performance

$$R_{jc} = \frac{(T_j - T_c)}{Q}$$

Thermal Resistance with Floorplan

- PkgA: 15 mm x 15 mm, 120 W : a bulk power density of 0.53 W/mm^2 as before
- PkgC: 10 mm x 10 mm, 25 W: a bulk power density of 0.25 W/mm^2, lower than that of PkgA

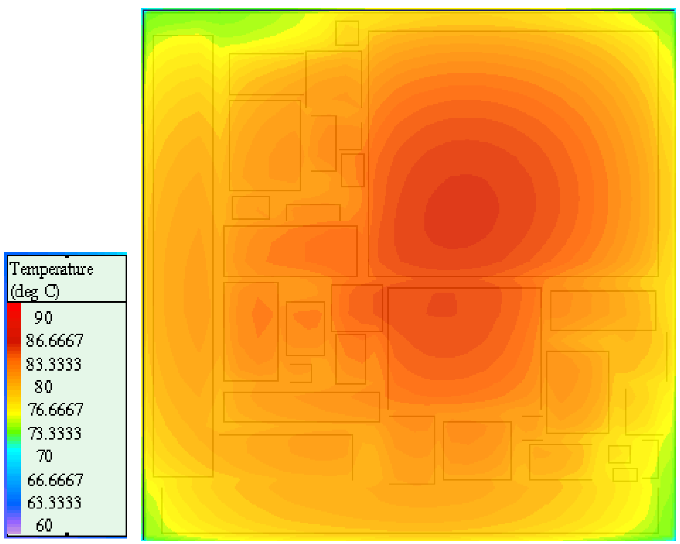

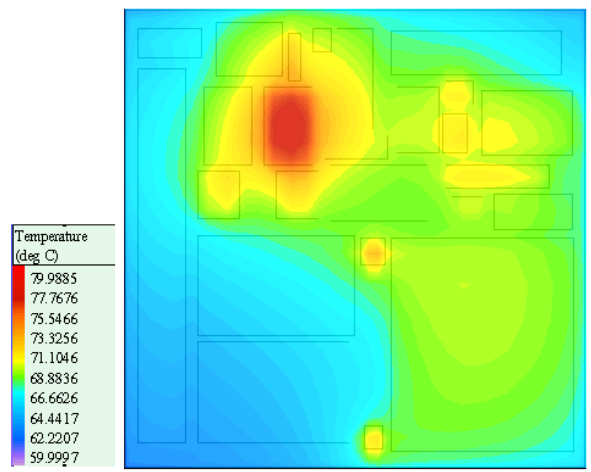

- PkgC resistance is 0.59 deg. C/W, higher than that of PkgA !

Summary

- Roadmap suggests continuation of increase of Chip Power unless a Silicon technology change happens
- Several Performance and Reliability Level Problems needs to be addressed
- Architecturally, Chip Multithreading (CMT) and Throughput Computing should help with Power mangement
- New Developments on Packaging and System Level Cooling is necessary