Announced Title:

Energy Efficiency Metrics for Network Equipment

e.g Energy Efficient Ethernet

Actual Title(s)

What is Happening in Energy Efficient Ethernet (P802.3az)

AND

The Challenge of Energy Efficiency Metrics for Network Equipment

Presented by:

Geoff Thompson

- Strategic Standards/Nortel Networks
- Member emeritus, IEEE 802 Exec Comm.
- Past Chair & active member, IEEE 802.3
- Member NESCOM, RAC
- E-mail: thompson@ieee.org
- Telephone: Seldom there, send e-mail

First:

What is Happening in:

Energy Efficient Ethernet (P802.3az)

P802.3az Energy Efficient Ethernet

- Is an approved IEEE Standards project
  - Approved September, 2007
- Is a Task Force of IEEE 802.3 Ethernet Working Group.
- The PAR is for an amendment to IEEE Std 802.3

Thus (as an 802.3 project)

The scope of the project is limited to:

- Items appropriately within the scope of 802.3
- Interoperability specifications
- NOT implementations

Further…

We have chosen the scope of the project to be limited to:

- Low hanging fruit
- Twisted pair copper interfaces

Where does the system power go?

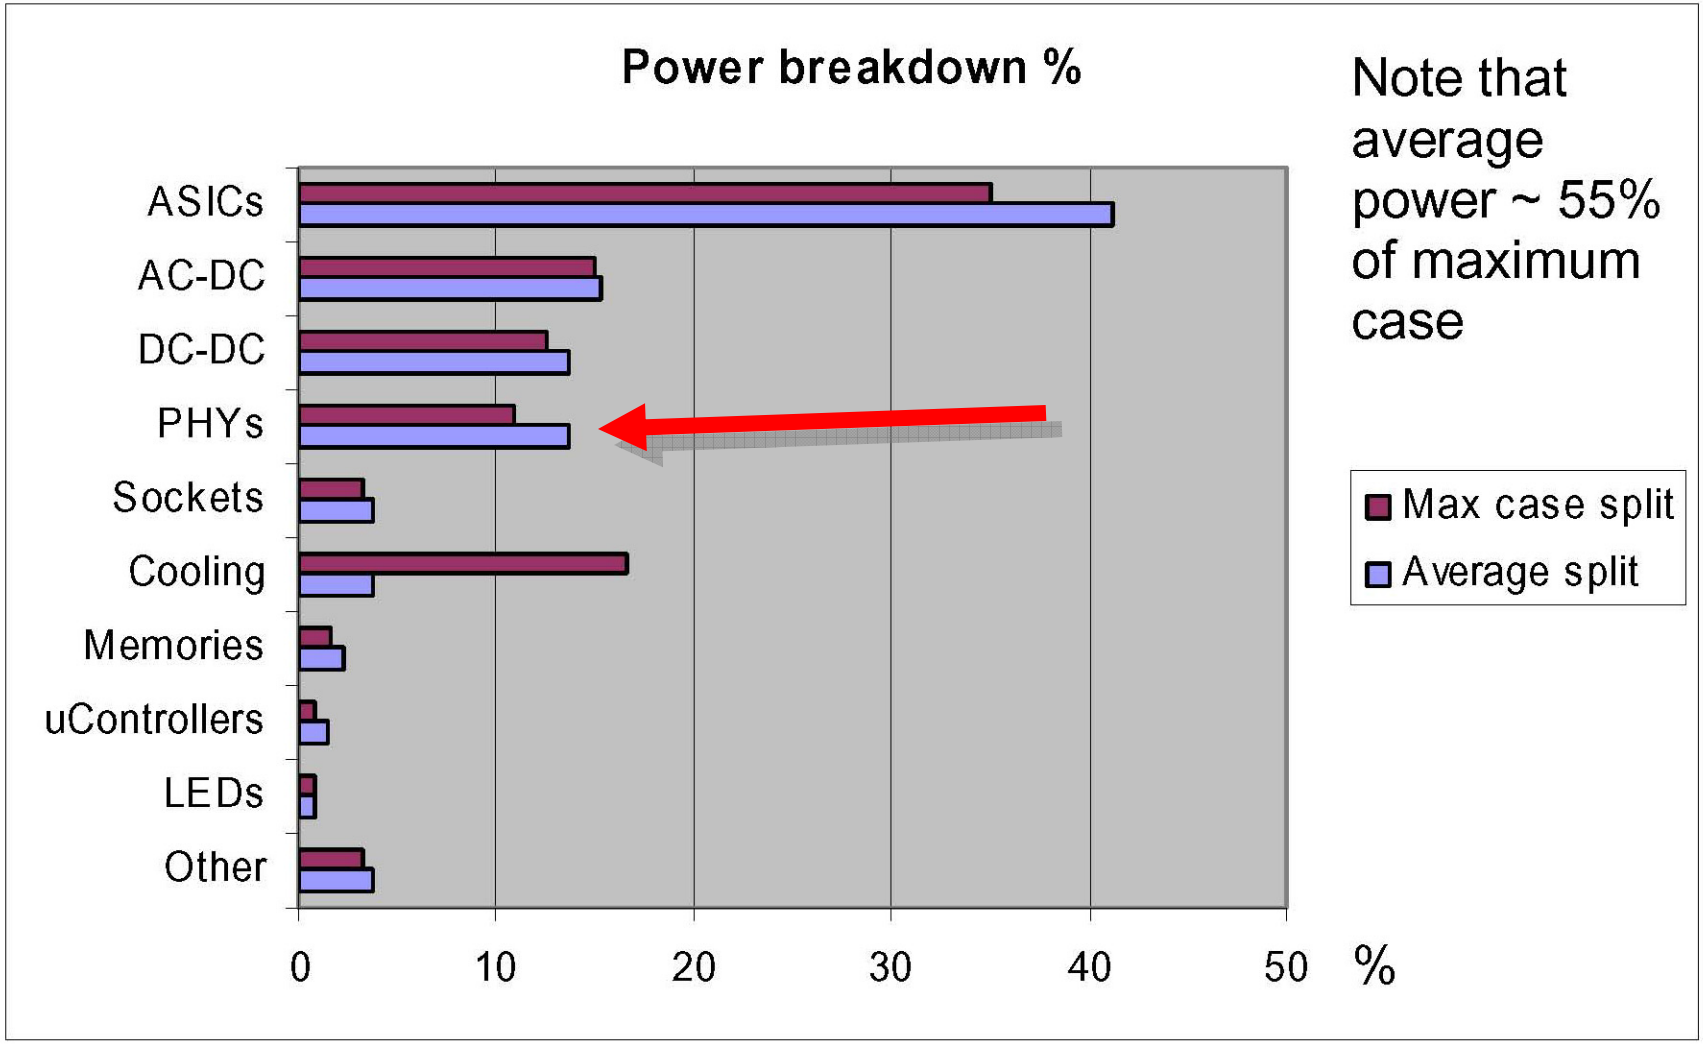

This slide "borrowed" from presentation by Hugh Barass/Cisco

Focus Areas

- Data Center Copper
  - Code speak for 10GBASE-T
- Low power between packets, 100M to 10G
  - Current code sets are constant power
  - i.e. idle consumes as much power as 1s & 0s
  - Keeps latencies low
- Improved voltage efficiency for 10BASE-T
  - Std was done when 5 volt was the norm

Problem areas

- Bias in participation towards PHY vendors
  - Tilts effort towards: Fix "problem" with a new, high margin chip
  - Focuses effort towards 10GBASE-T problem
  - Self-contained localized solutions
- Layering of Communication standards
  (ISO 7 Layer Model & derivations)
- Established dependence on unspecified but usual behavior
  - e.g. Maintaining TCP Connections

Problem areas (2)

- Changing requirements for "unspecified behavior"
  - e.g. AV Bridging, time keeping, synchronous applications
  - e.g. Latency dependencies (Data Center Ethernet)

Approaches

- Auto-Negotiation (rejected, too slow, drops link)
- Rapid PHY Switching (rejected, requires predictive logic or big buffers)
- Subset PHYs (variant of RPS), (Rejected, RPS problems + too many new signals)
- Low power IDLE + refresh (current dir.)
- Low voltage 10BASE-T (current direction)

Still to be addressed

The More System Oriented Aspects

- Any kind of “WAIT” signals
- Explicit speed shifts/speed locks
- Power up/down facilities

Issues

- Any kind of “WAIT” signals
- Explicit speed shifts/speed locks
- Power up/down facilities

Part 2

The Challenge of Energy Efficiency Metrics for Network Equipment

What is P802.3az doing?

- Not much
  - Most discussions have not been rigorously quanitative.
  - Implementation vs standards are ultimately the issues

Central Issues

- Simple power per bit across the network is not the issue
  - Delay is an issue and cost (Packets transmitted at speeds lower than their generation speed “take longer” to get there).
  - Transit variability can mess things up.
  - Newer applications tend to be more sensitive than traditional “best effort” data service.
  - Capital/OpEx shows up too.
- One dimensional metric is not the answer

Even the one dimensional metric is not that simple

- IF power consumption is traffic dependent
  (not currently the case)
- THEN an agreed upon traffic model is needed
  (much like EPA City/Hiway driving cycles)
- Traffic model SHOULD have a reasonable relationship to real network traffic patterns
- At best, THE NUMBER for most networks
  (power cycled networks excluded) will be of the form:
  Watts/hour of up time (per unit of peak capacity?)
  PLUS
  n microwatts per kilobit transmitted.

Where does the system power go?

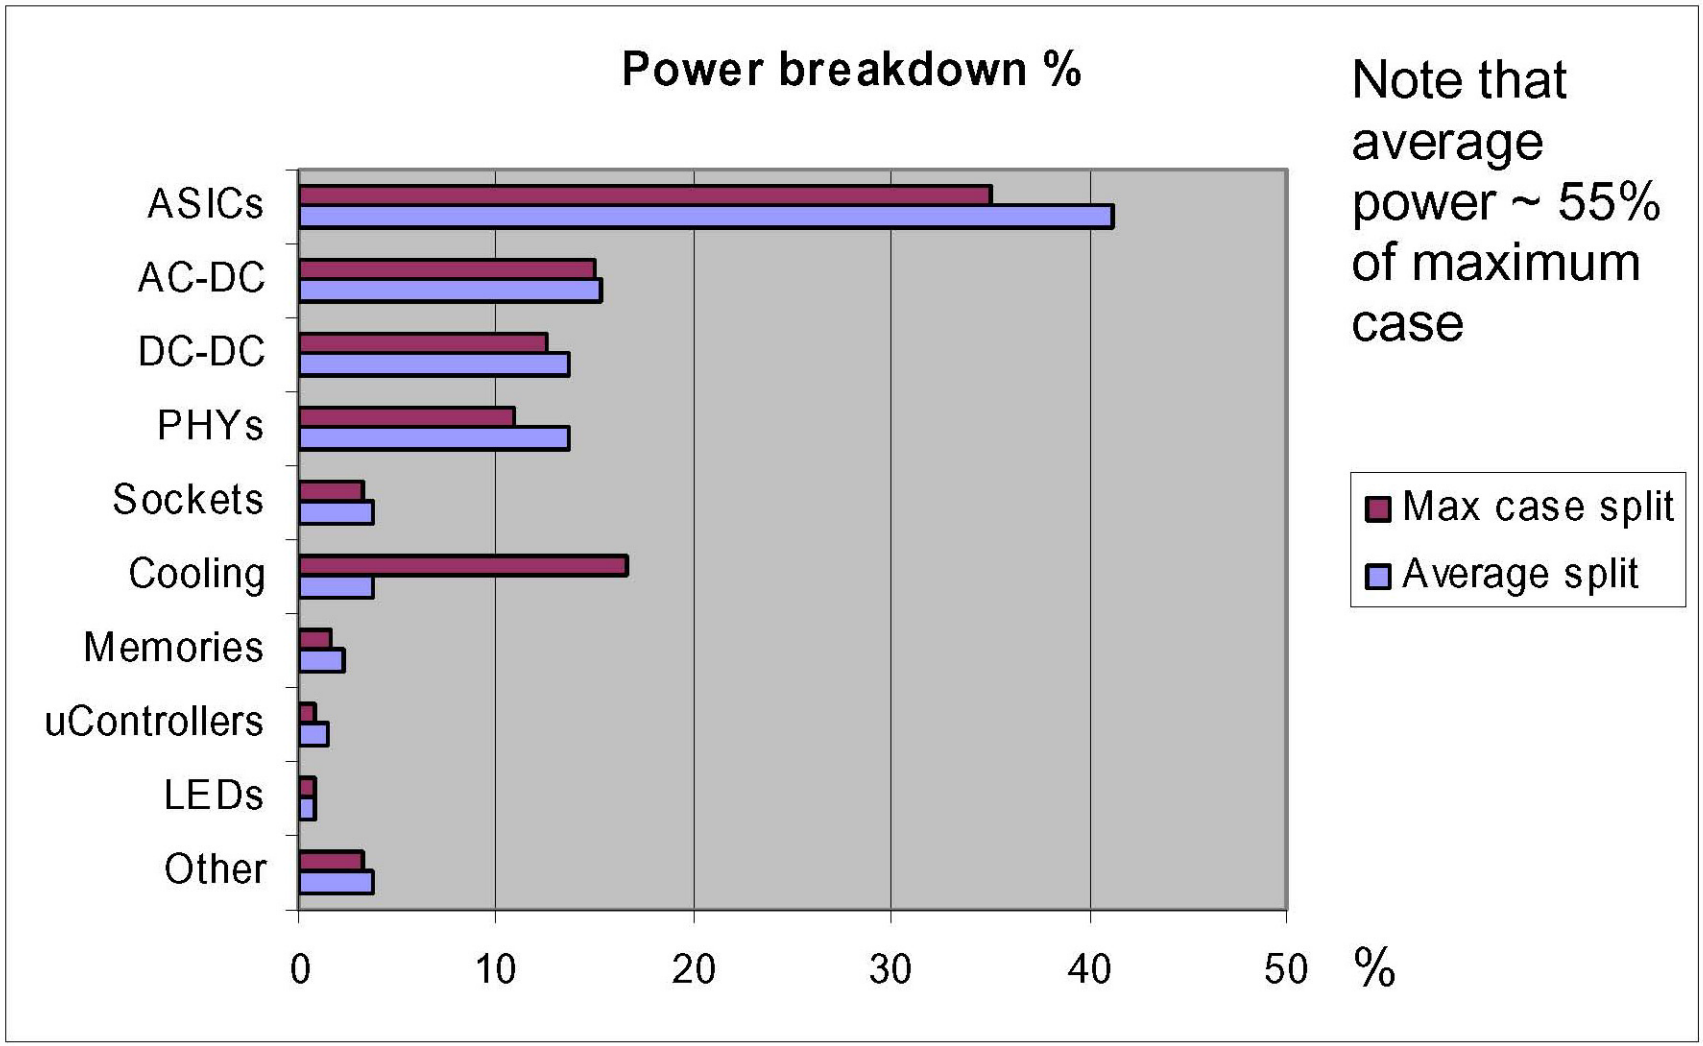

This slide "borrowed" from presentation by Hugh Barass/Cisco

CONCLUSION

MOST OF THE POWER DISSIPATION IS GOING TO BE FROM CONSTANT DISSIPATION PORTIONS OF THE SYSTEM

However,

There are SO MANY Ethernet copper ports sold that a small power reduction is worth while on a national scale.