CS848
Advanced Topics in Databases
Database Systems on Modern Hardware
Spring 2015

Ken Salem

David R. Cheriton School of Computer Science
University of Waterloo

US Data Center Energy Consumption

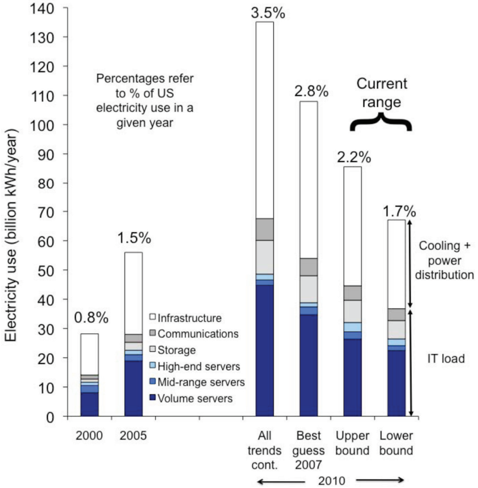

from *Growth in Data center electricity use 2005 to 2010.*, Jonathan Koomey. Analytics Press, Oakland, CA. 2011
`http://www.analyticspress.com/datacenters.html`

US Data Center Energy Consumption

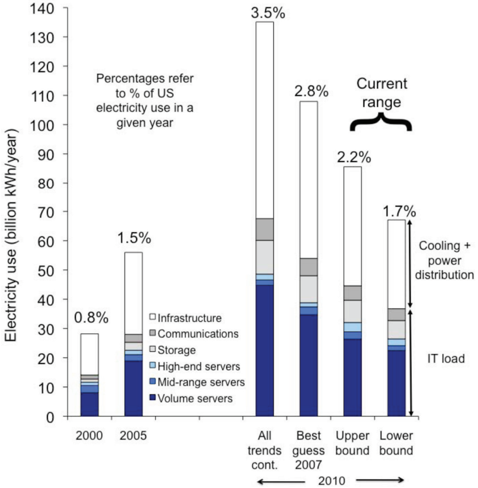

from *Growth in Data center electricity use 2005 to 2010.*, Jonathan Koomey. Analytics Press, Oakland, CA. 2011
http://www.analyticspress.com/datacenters.html

- cooling and distribution double consumption

US Data Center Energy Consumption

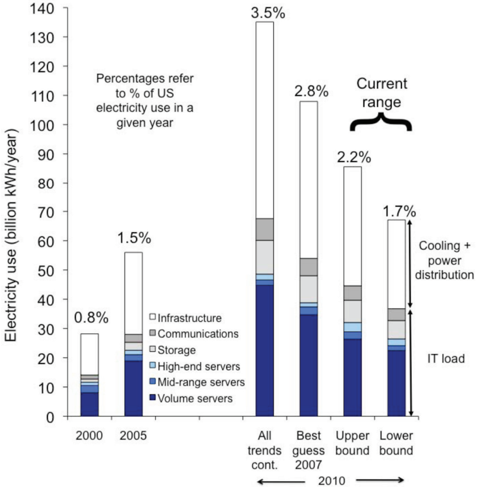

from *Growth in Data center electricity use 2005 to 2010.*, Jonathan Koomey. Analytics Press, Oakland, CA. 2011
http://www.analyticspress.com/datacenters.html

- cooling and distribution double consumption
- 1% reduction ≈ 1 billion kWh/year

US Data Center Energy Consumption

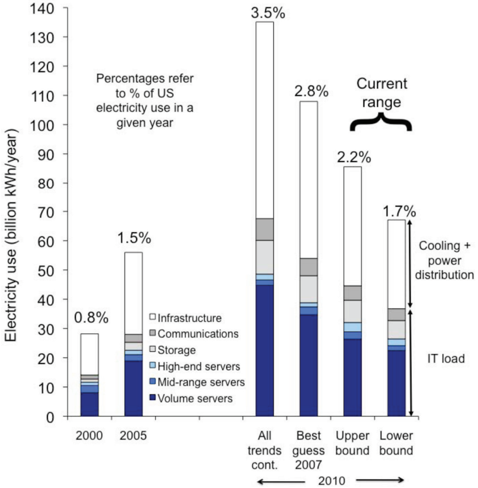

from *Growth in Data center electricity use 2005 to 2010.*, Jonathan Koomey. Analytics Press, Oakland, CA. 2011
http://www.analyticspress.com/datacenters.html

- cooling and distribution double consumption
- 1% reduction ≈ 1 billion kWh/year
- last year, my condo ≈ 4500 kWh

US Data Center Energy Consumption

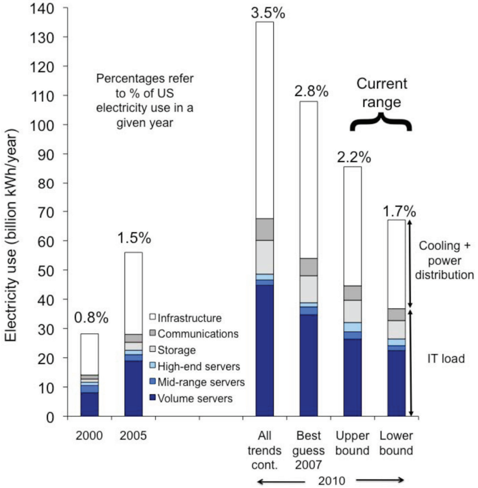

from *Growth in Data center electricity use 2005 to 2010.*, Jonathan Koomey. Analytics Press, Oakland, CA. 2011
http://www.analyticspress.com/datacenters.html

- cooling and distribution double consumption
- 1% reduction ≈ 1 billion kWh/year
- last year, my condo ≈ 4500 kWh
- 1% reduction ≈ 200,000 condos

Energy Efficiency of Cloud Computing

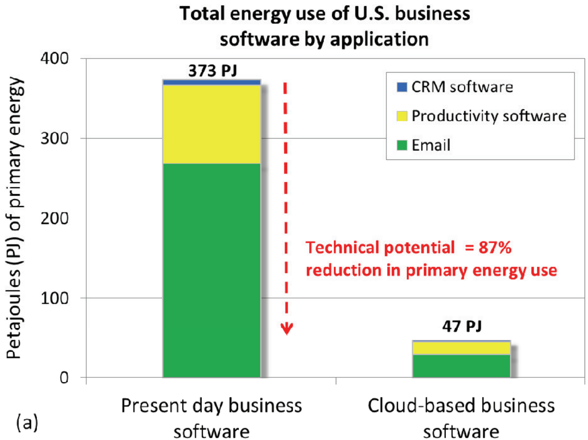

from E. Masanet et al, *The Energy Efficiency Potential of Cloud-Based Software: A U.S. Case Study*, Lawrence Berkeley National Laboratory, June 2013 http://crd.lbl.gov/assets/pubs_presos/ACS/cloud_efficiency_study.pdf

Energy Efficiency of Cloud Computing

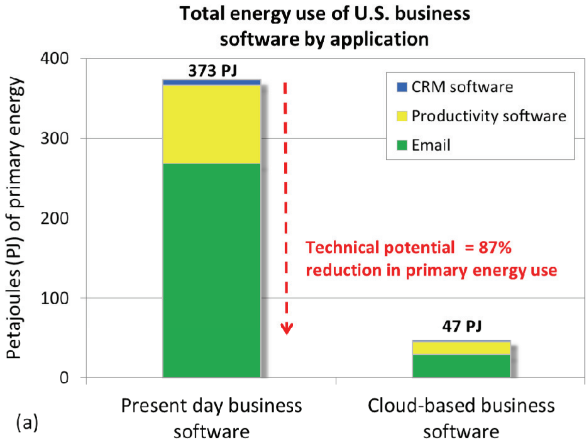

from E. Masanet et al, *The Energy Efficiency Potential of Cloud-Based Software: A U.S. Case Study*, Lawrence Berkeley National Laboratory, June 2013 http://crd.lbl.gov/assets/pubs_presos/ACS/cloud_efficiency_study.pdf

- most hosting still in small server rooms/closets

Energy Efficiency of Cloud Computing

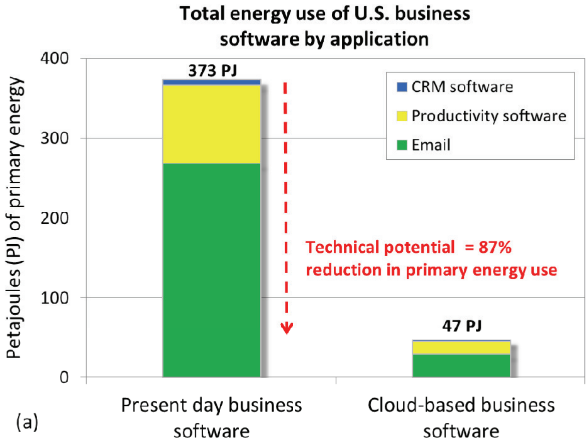

from E. Masanet et al, *The Energy Efficiency Potential of Cloud-Based Software: A U.S. Case Study*,
Lawrence Berkeley National Laboratory, June 2013
`http://crd.lbl.gov/assets/pubs_presos/ACS/cloud_efficiency_study.pdf`

- most hosting still in small server rooms/closets
- step 1: move to cloud

Energy Efficiency of Cloud Computing

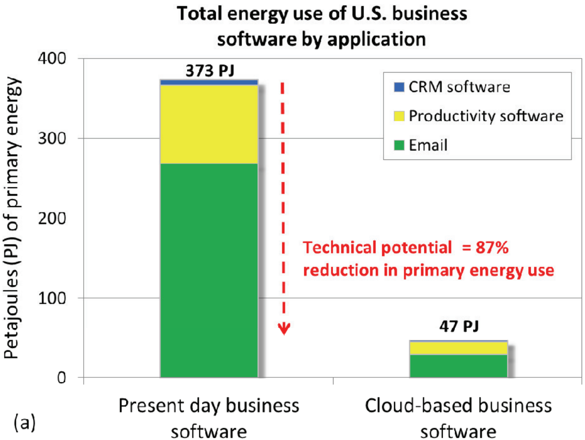

from E. Masanet et al, *The Energy Efficiency Potential of Cloud-Based Software: A U.S. Case Study*,
Lawrence Berkeley National Laboratory, June 2013
http://crd.lbl.gov/assets/pubs_presos/ACS/cloud_efficiency_study.pdf

- most hosting still in small server rooms/closets
- step 1: move to cloud
- step 2: optimize cloud

Cost of Large Data Centers

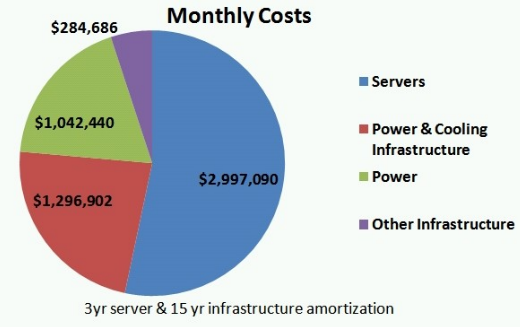

from James Hamilton, *Cost of Power in Large-Scale Data Centers*, Perspectives blog, Nov 2008
`http://perspectives.mvdirona.com/2008/11/cost-of-power-in-large-scale-data-centers/`

Cost of Large Data Centers

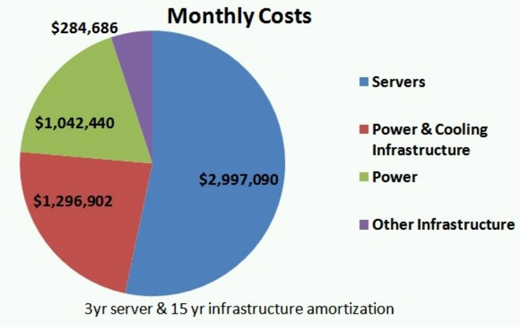

from James Hamilton, *Cost of Power in Large-Scale Data Centers*, Perspectives blog, Nov 2008
`http://perspectives.mvdirona.com/2008/11/cost-of-power-in-large-scale-data-centers/`

- “fully burdened cost of power” ≈ 42%

Cost of Large Data Centers

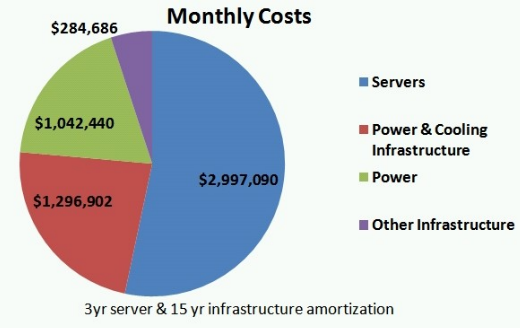

from James Hamilton, *Cost of Power in Large-Scale Data Centers*, Perspectives blog, Nov 2008
`http://perspectives.mvdirona.com/2008/11/cost-of-power-in-large-scale-data-centers/`

- "fully burdened cost of power" ≈ 42%
- (+) server costs decreasing, power cost increasing

Cost of Large Data Centers

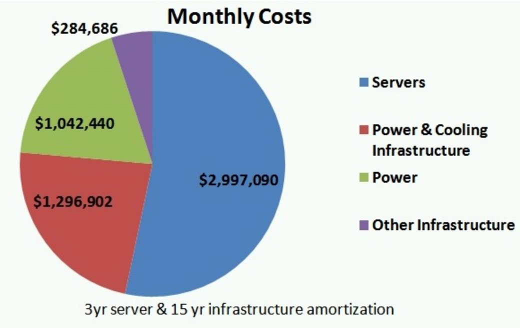

from James Hamilton, *Cost of Power in Large-Scale Data Centers*, Perspectives blog, Nov 2008
`http://perspectives.mvdirona.com/2008/11/cost-of-power-in-large-scale-data-centers/`

- "fully burdened cost of power" ≈ 42%
- (+) server costs decreasing, power cost increasing
- (-) server power efficiency improving

Data Center Server Utilization

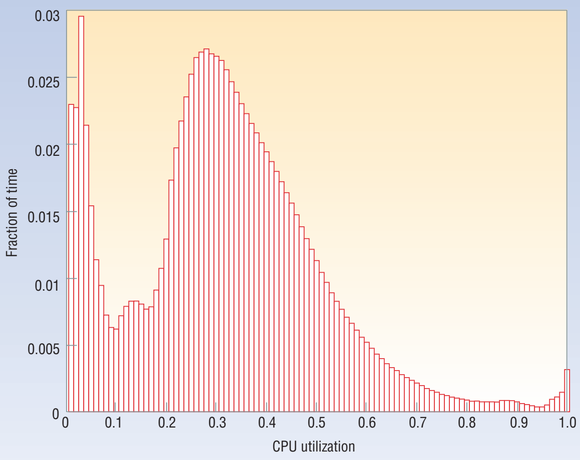

from Barroso and Hölzle, *The Case for Energy-Proportional Computing*, IEEE Computer 40(12), Dec 2007, pp. 33-37

- utilization of > 5000 Google servers over 6 months

Data Center Server Utilization

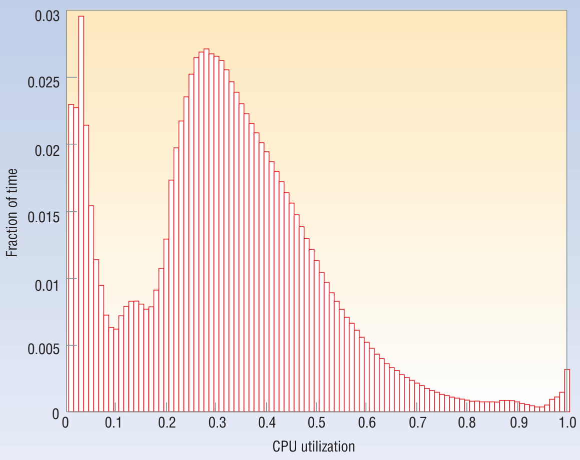

from Barroso and Hölzle, *The Case for Energy-Proportional Computing*, IEEE Computer 40(12), Dec 2007, pp. 33-37

- utilization of > 5000 Google servers over 6 months
- most servers 10%-50%

Data Center Server Utilization

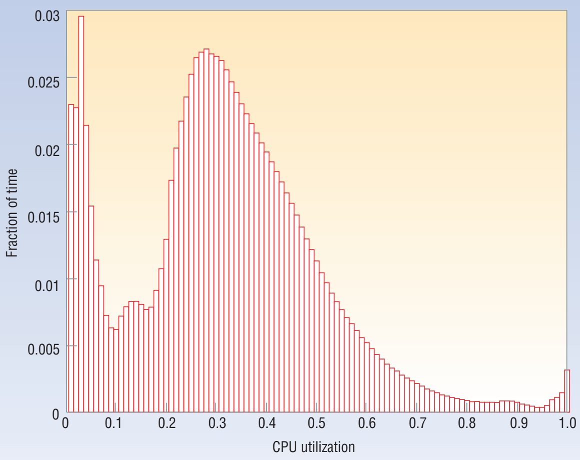

from Barroso and Hölzle, *The Case for Energy-Proportional Computing*, IEEE Computer 40(12), Dec 2007, pp. 33-37

- utilization of > 5000 Google servers over 6 months
- most servers 10%-50%
- full idle unlikely

Power Proportionality

- energy consumption proportional to work done

Power Proportionality

- energy consumption proportional to work done
- SPECpower_ssj2008 benchmark

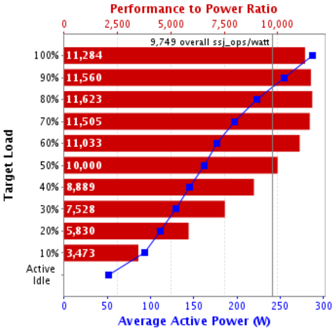

Dell PowerEdge R630
(April 2015)
dual proc/36 cores/64 GB

Power Proportionality

- energy consumption proportional to work done
- SPECpower_ssj2008 benchmark
- power range improving over time?

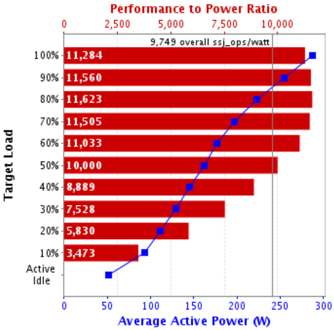

Dell PowerEdge R630
(April 2015)
dual proc/36 cores/64 GB

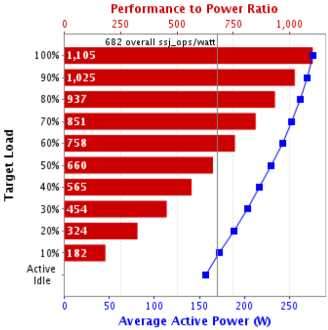

Dell PowerEdge 2950 III
(Dec 2007)
dual proc/8 cores/16 GB

Idle-to-Peak Trend

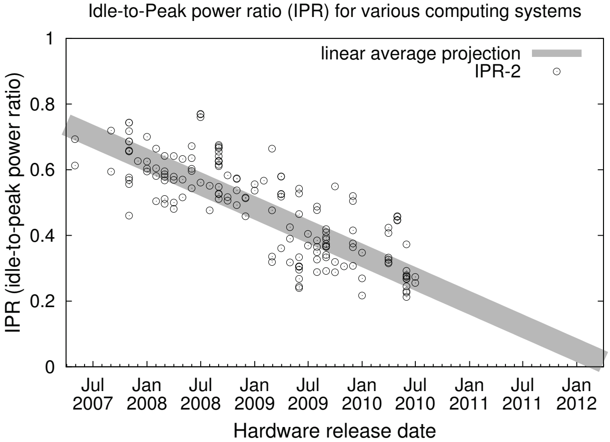

from Varsamopoulos and Gupta
Energy Proportionality and the Future: Metrics and Directions,
In Proc. Int'l Conf. on Parallel Processing Workshops, 2010

Techniques for Energy Efficiency

- dynamic server (de)provisioning
    - adjust number of active servers to load
    - idle or power down unused servers
- frequency and voltage scaling
    - adjust CPU frequency based on workload
    - lower frequency $\Rightarrow$ less power consumed
- energy-aware scheduling
    - choose energy-efficient platform for each workload

Voltage and Frequency Scaling

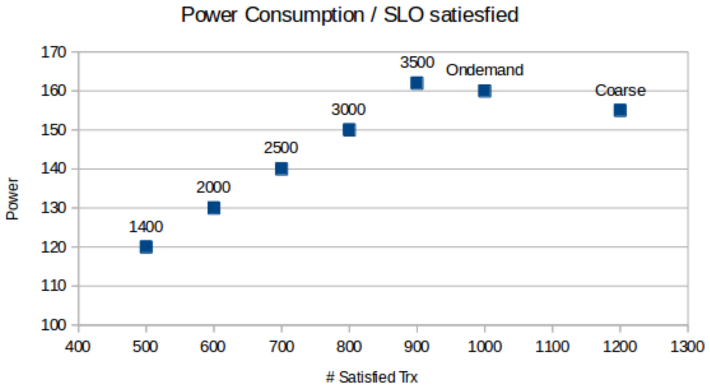

System power consumption vs. TPC-C throughput
in various p-states
Shore-MT, in-memory database

CPU Scaling

$S =$ process feature size ratio, e.g,
32 nm to 22 nm gives $S = 32/22 \approx 1.4$

Dennard scaling

- Δ Quantity $\propto S^2$

Source: M.B. Taylor, *A Landscape of the New Dark Silicon Design Regime*. IEEE Micro 33(5), Aug. 2013, pp. 8-19.

CPU Scaling

$S =$ process feature size ratio, e.g,
32 nm to 22 nm gives $S = 32/22 \approx 1.4$

Dennard scaling

- Δ Quantity $\propto S^2$
- Δ Frequency $\propto S$

Source: M.B. Taylor, *A Landscape of the New Dark Silicon Design Regime*. IEEE Micro 33(5), Aug. 2013, pp. 8-19.

CPU Scaling

$S =$ process feature size ratio, e.g,
32 nm to 22 nm gives $S = 32/22 \approx 1.4$

Dennard scaling

- Δ Quantity $\propto S^2$
- Δ Frequency $\propto S$
- Δ Capacitance $\propto 1/S$

Source: M.B. Taylor, *A Landscape of the New Dark Silicon Design Regime*. IEEE Micro 33(5), Aug. 2013, pp. 8-19.

CPU Scaling

$S =$ process feature size ratio, e.g,
32 nm to 22 nm gives $S = 32/22 \approx 1.4$

Dennard scaling

- Δ Quantity $\propto S^2$
- Δ Frequency $\propto S$
- Δ Capacitance $\propto 1/S$
- Δ Voltage $\propto 1/S^2$

Source: M.B. Taylor, *A Landscape of the New Dark Silicon Design Regime*. IEEE Micro 33(5), Aug. 2013, pp. 8-19.

CPU Scaling

$S =$ process feature size ratio, e.g,
32 nm to 22 nm gives $S = 32/22 \approx 1.4$

Dennard scaling

- Δ Quantity $\propto S^2$
- Δ Frequency $\propto S$
- Δ Capacitance $\propto 1/S$
- Δ Voltage $\propto 1/S^2$
- $\Rightarrow \Delta$ Power $\propto \Delta QFCV^2 = 1$

Source: M.B. Taylor, *A Landscape of the New Dark Silicon Design Regime*. IEEE Micro 33(5), Aug. 2013, pp. 8-19.

CPU Scaling

$S =$ process feature size ratio, e.g,
32 nm to 22 nm gives $S = 32/22 \approx 1.4$

Dennard scaling

- Δ Quantity $\propto S^2$
- Δ Frequency $\propto S$
- Δ Capacitance $\propto 1/S$
- Δ Voltage $\propto 1/S^2$
- $\Rightarrow \Delta$ Power $\propto \Delta QFCV^2 = 1$
- $\Rightarrow \Delta$ Utilization $\propto 1$

Source: M.B. Taylor, *A Landscape of the New Dark Silicon Design Regime*. IEEE Micro 33(5), Aug. 2013, pp. 8-19.

CPU Scaling

$S =$ process feature size ratio, e.g,
32 nm to 22 nm gives $S = 32/22 \approx 1.4$

Dennard scaling

- Δ Quantity $\propto S^2$
- Δ Frequency $\propto S$
- Δ Capacitance $\propto 1/S$
- Δ Voltage $\propto 1/S^2$
- $\Rightarrow \Delta$ Power $\propto \Delta QFCV^2 = 1$
- $\Rightarrow \Delta$ Utilization $\propto 1$

post-Dennard scaling

- Δ Quantity $\propto S^2$
- Δ Frequency $\propto S$
- Δ Capacitance $\propto 1/S$
- Δ Voltage $\propto 1$
- $\Rightarrow \Delta$ Power $\propto \Delta QFCV^2 = S^2$
- $\Rightarrow \Delta$ Utilization $\propto 1/S^2$

Source: M.B. Taylor, *A Landscape of the New Dark Silicon Design Regime*. IEEE Micro 33(5), Aug. 2013, pp. 8-19.

Dark Silicon

- silicon that is not used all the time, or not used at its full frequency
- fixed power envelope limits growth in Q or F or both
  - Denard: QF grows by S^3
  - post-Denard: QF grows by only S

Dark Silicon Example

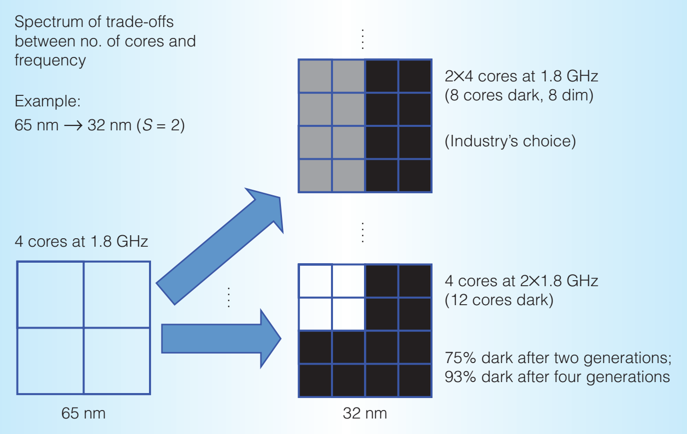

Source: M.B. Taylor,
A Landscape of the New Dark Silicon Design Regime.
IEEE Micro 33(5), Aug. 2013, pp. 8-19.

Responses to Dark Silicon

- smaller chips
- “dim” silicon
  - reduce clock rate, or
  - use more space for low-power functions, e.g., cache,
  - power only part of the time
- functional specialization
  - fast or efficient co-processors
  - execution hops around