Die weitere Entwicklung

Basis:

- ITRS (International Technology Road Map for Semiconductors)
- Babak Falsafi: Dark Silicon & Its Implications on Server Chip Design, Microsoft Research, Nov. 2010
  Siehe auch *publications* unter http://parsa.epfl.ch/~falsafi/
- Hadi Esmaeilzadeh: Dark Silicon and the End of Multicore Scaling, International Symposium on Computer Architecture (ISCA ’11)

Performanz-Trends

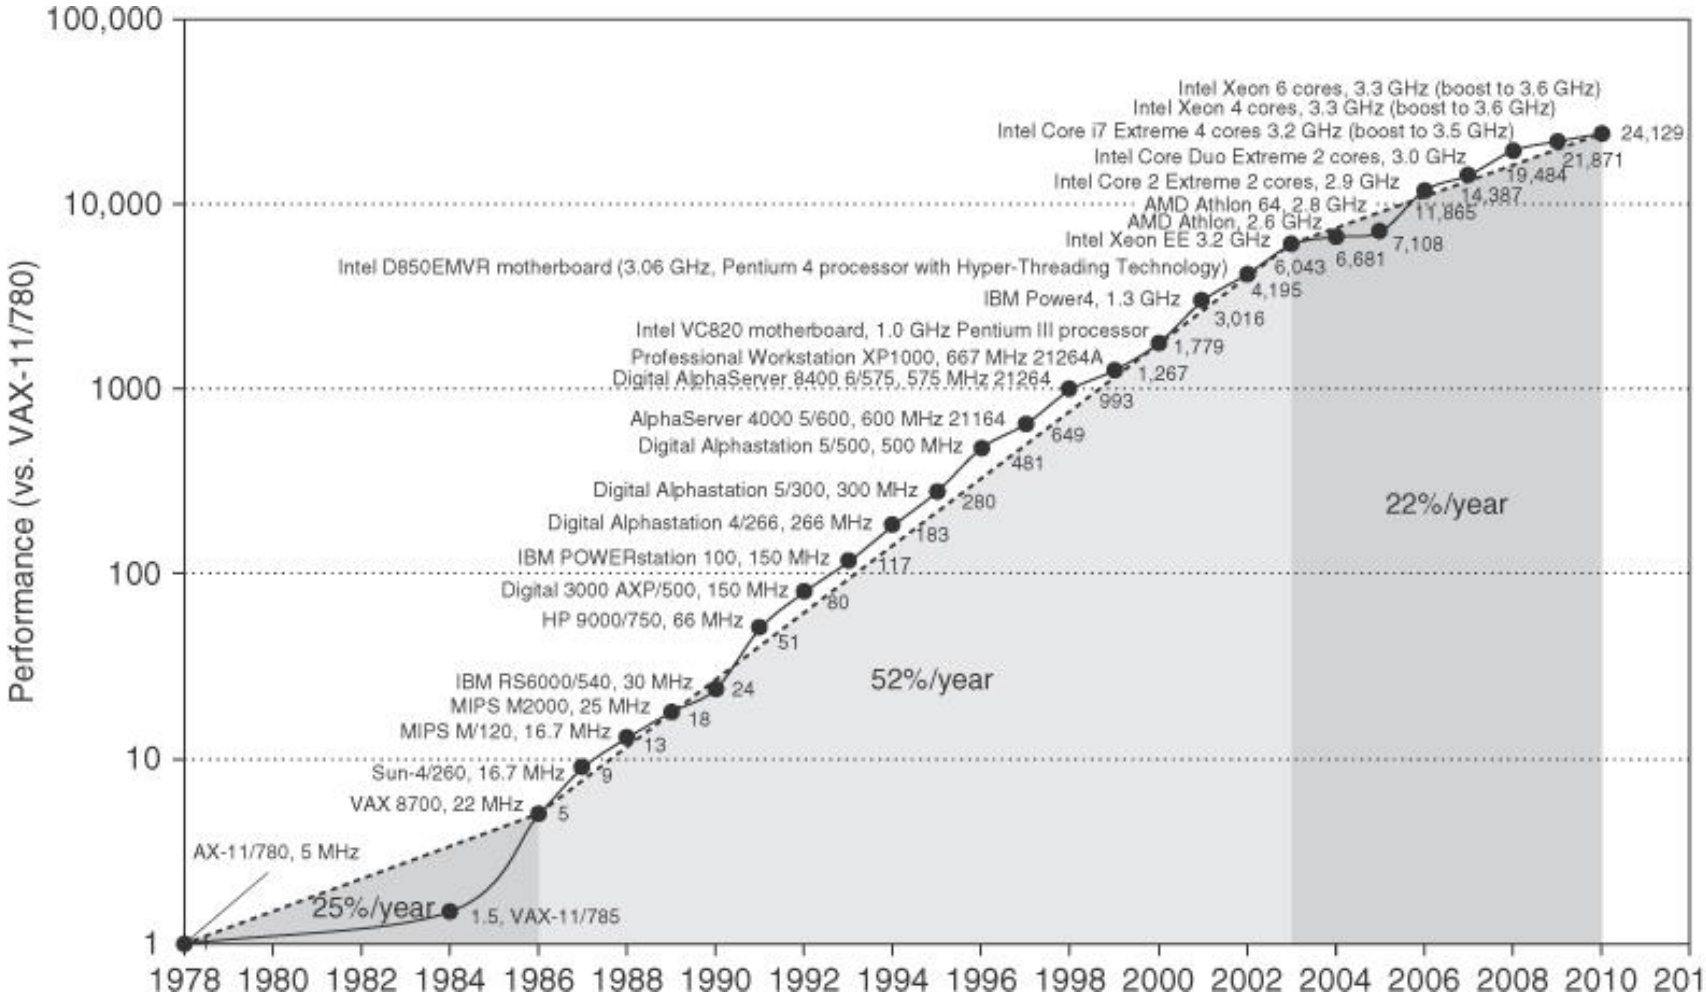

Taktrate-Trends

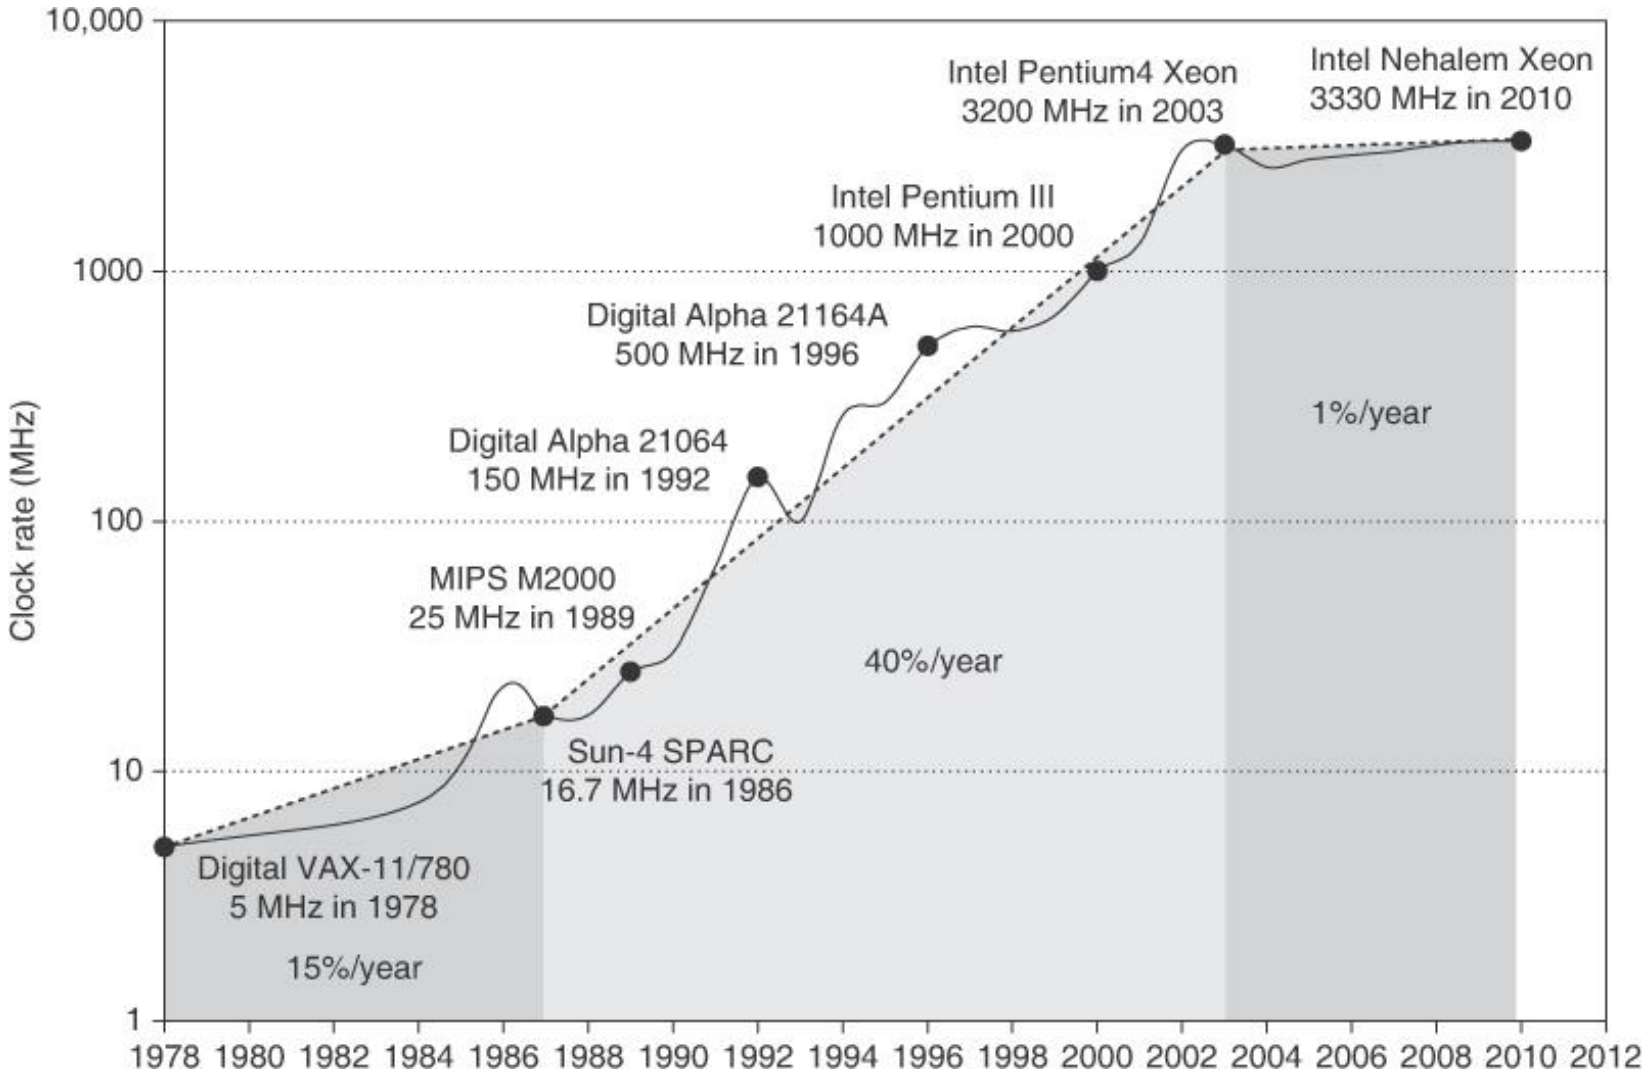

Diskussion in der ITRS zur Entwicklung der Taktraten

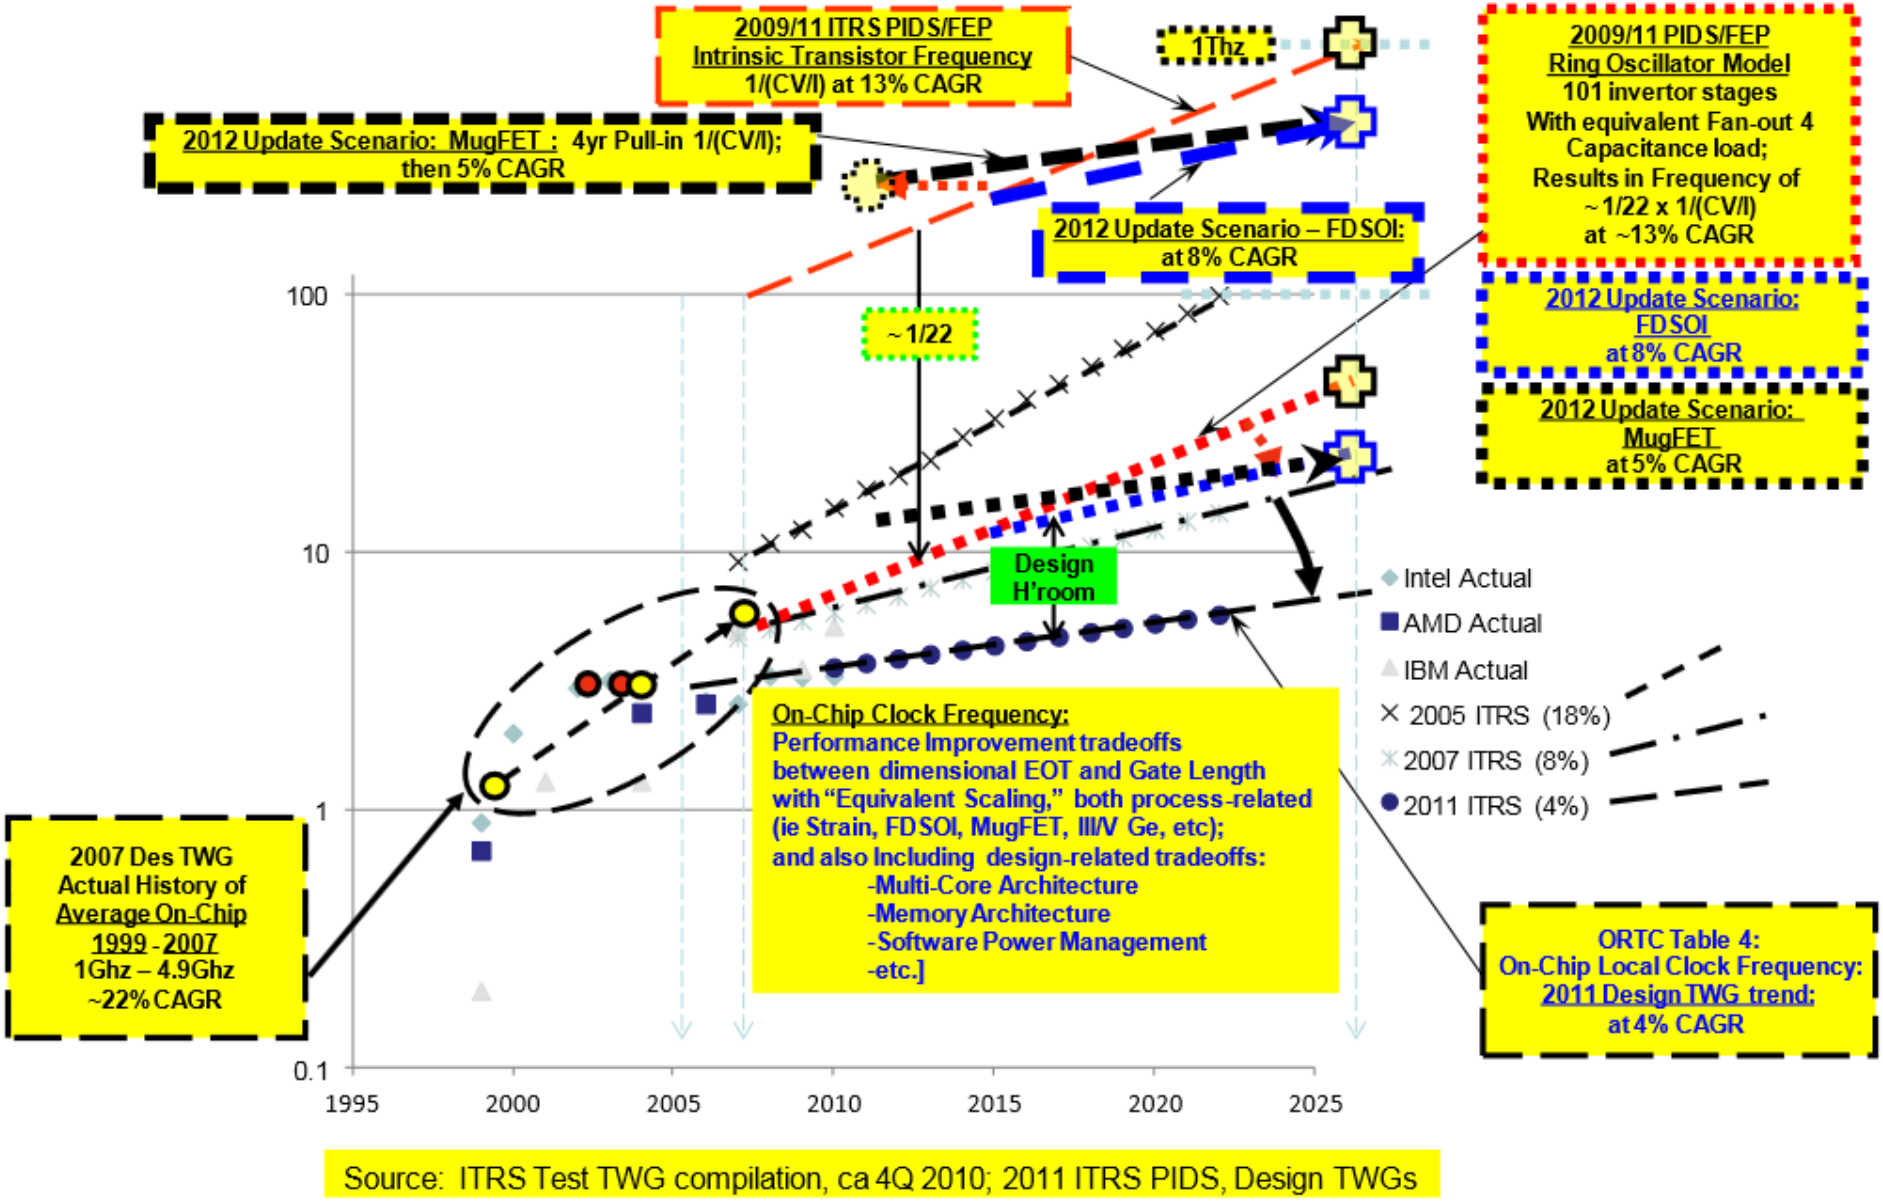

Source: ITRS Test TWG compilation, ca 4Q 2010; 2011 ITRS PIDS, Design TWGs

Trends für Schlüsselparameter gemäß ITRS 2013

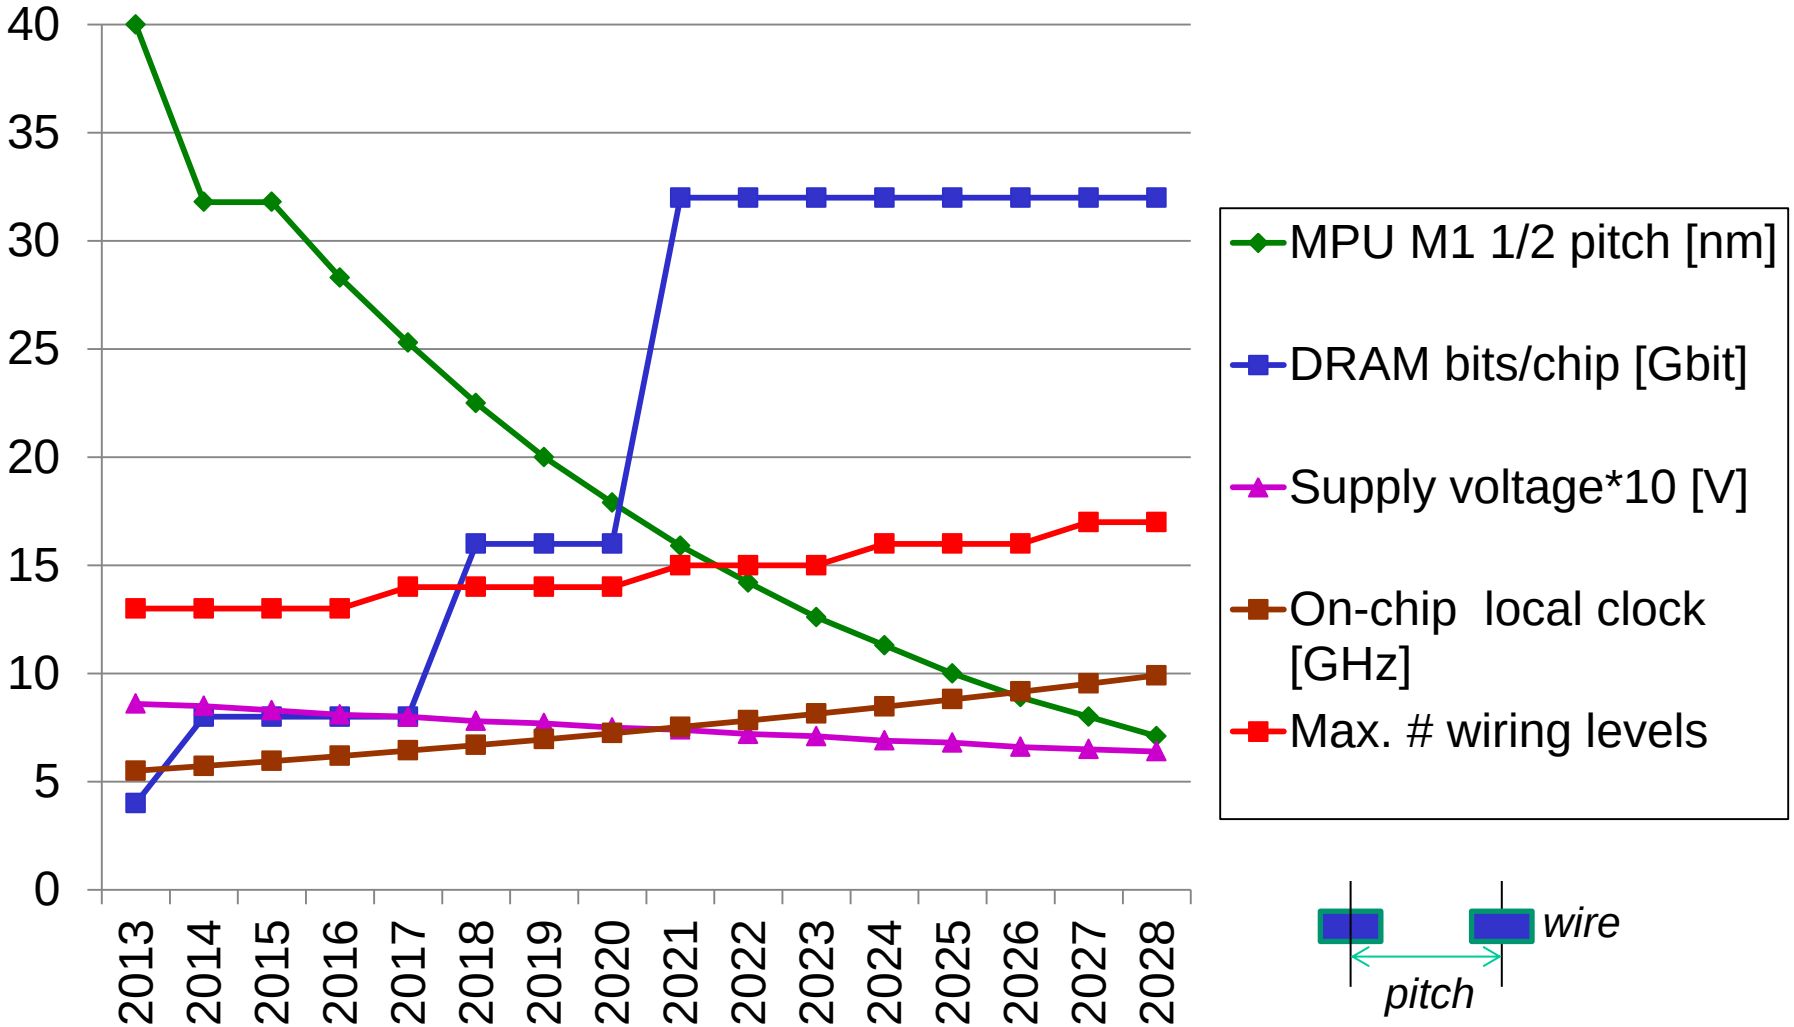

Funktionen pro Chip

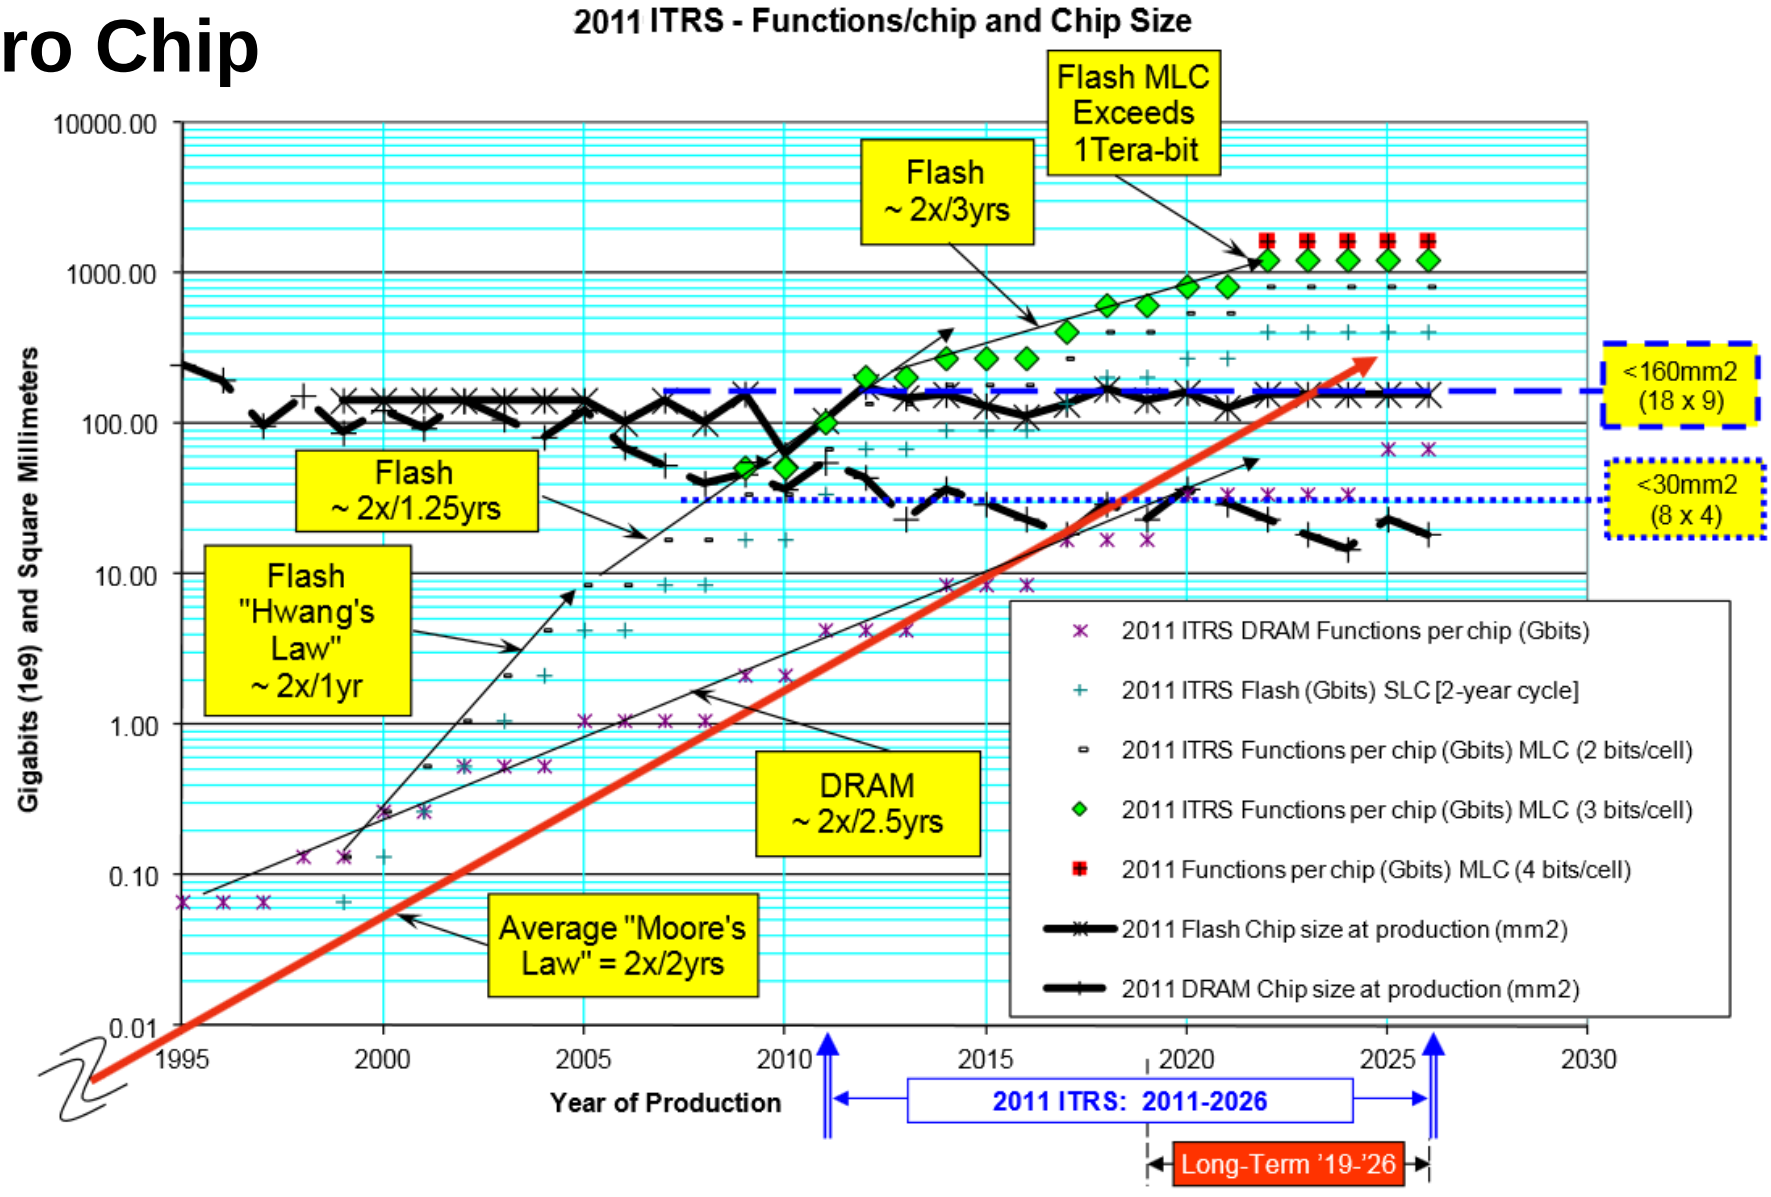

Figure ORTC7 2011 ITRS Product Technology Trends:
Memory Product Functions/Chip and Industry Average "Moore's Law" and Chip Size Trends

[ITRS 2011

Prozessoren für stationäre Systme

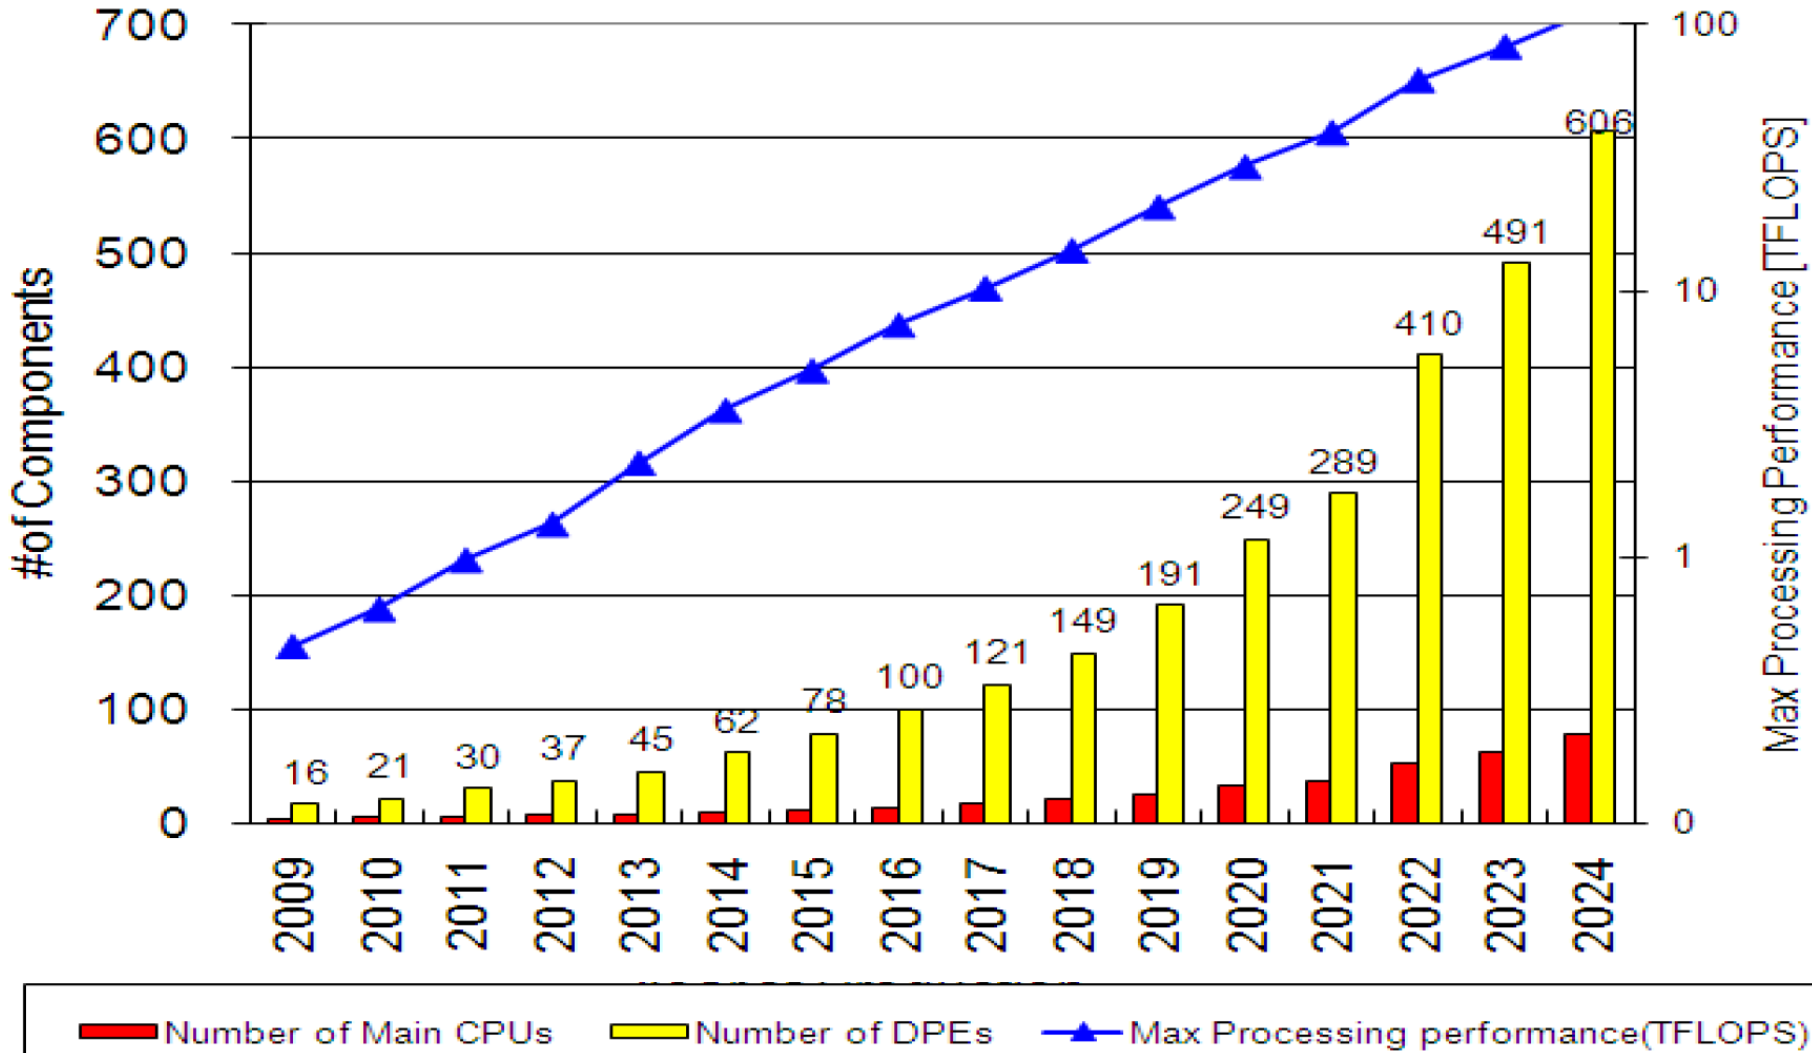

Trends zur Leistungsaufnahme
- Stationäre Systeme -

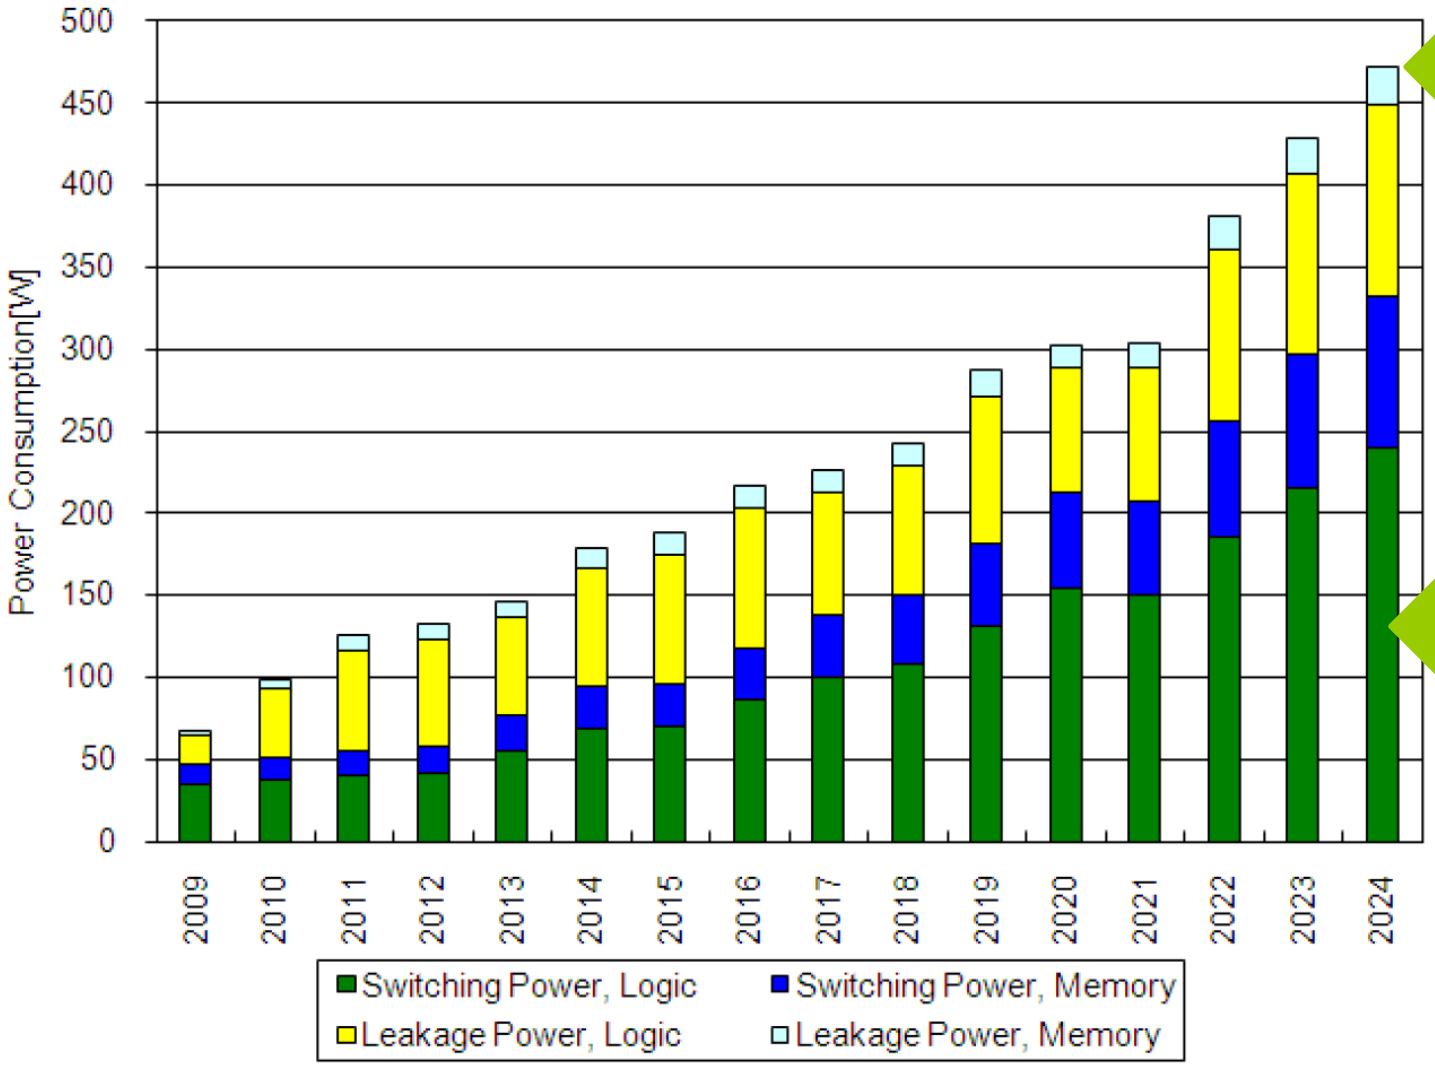

Alptraum (☹) des Umwelt-schutzes

Dominant

Trends zur Leistungsaufnahme
- Mobile Systeme -

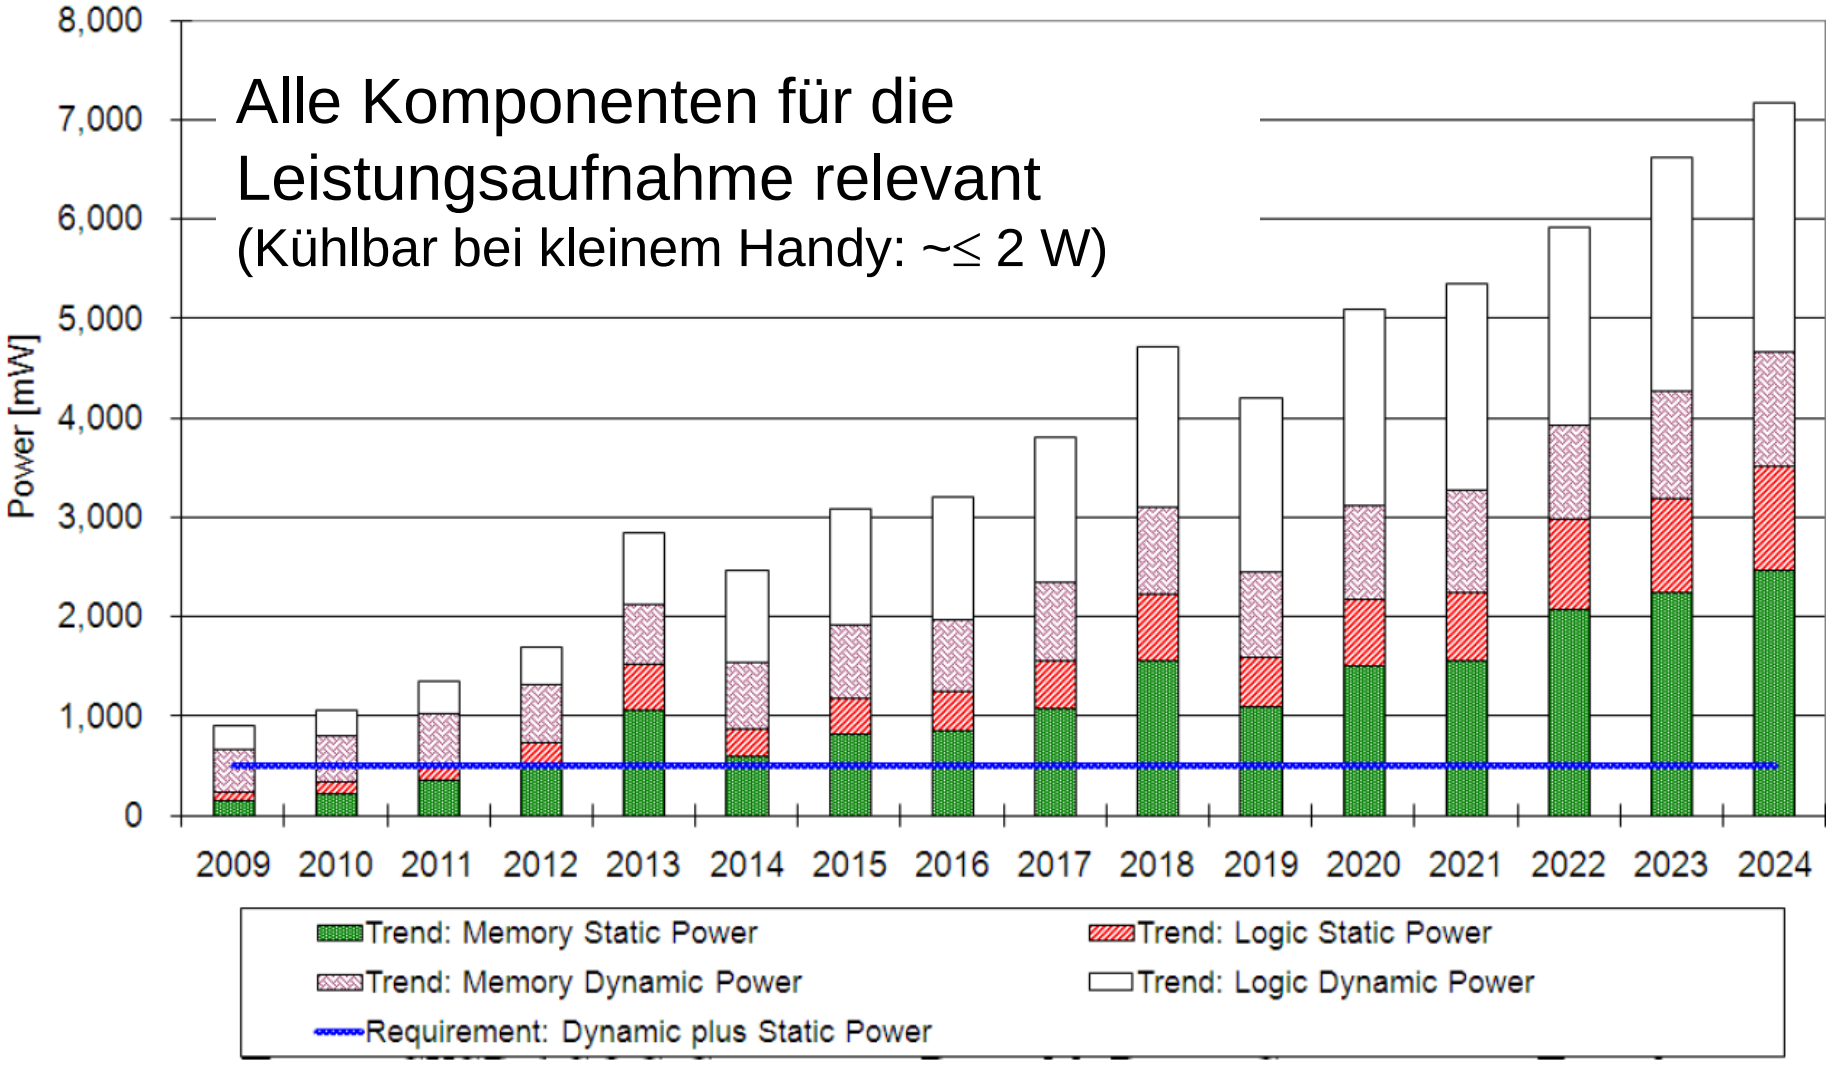

Shift towards Cloud Computing Helps

- Ubiquitous connectivity & access to data
- Consolidate servers ➔ Amortize energy costs

But, the Cloud has hit a wall!

Trends:

- Moore's law continues
- Server density is increasing
- But, voltage scaling has slowed

➔It's too expensive to buy/cool servers

A 1,000m² datacenter is 1.5MW!

(carbon footprint of airlines in 2012)

A few words about our model

Physical char. modeled after Niagara

Area: cores/caches (72% die)

- scaled across tech. nodes

Power:

- Active: projected V_{dd}/ITRS
  - Core=scaled, cache=f(miss), crossbar=f(hops)
- Leakage: projected V_{th}/ITRS, f(area), 62C

Performance:

- Parameters from real server workloads (DB2, Oracle, Apache, Zeus)
- Cache miss rate model (validated)
- CPI model based on miss rate

Caveat: Simple Parallelizable Workloads

Workloads are assumed parallel

- Scaling server workloads is reasonable

CPI model:

- Works well for workloads with low MLP
- OLTP, Web & DSS are mostly memory-latency dependent

Future servers will run a mix of workloads

Area vs. Power Envelope (22nm)

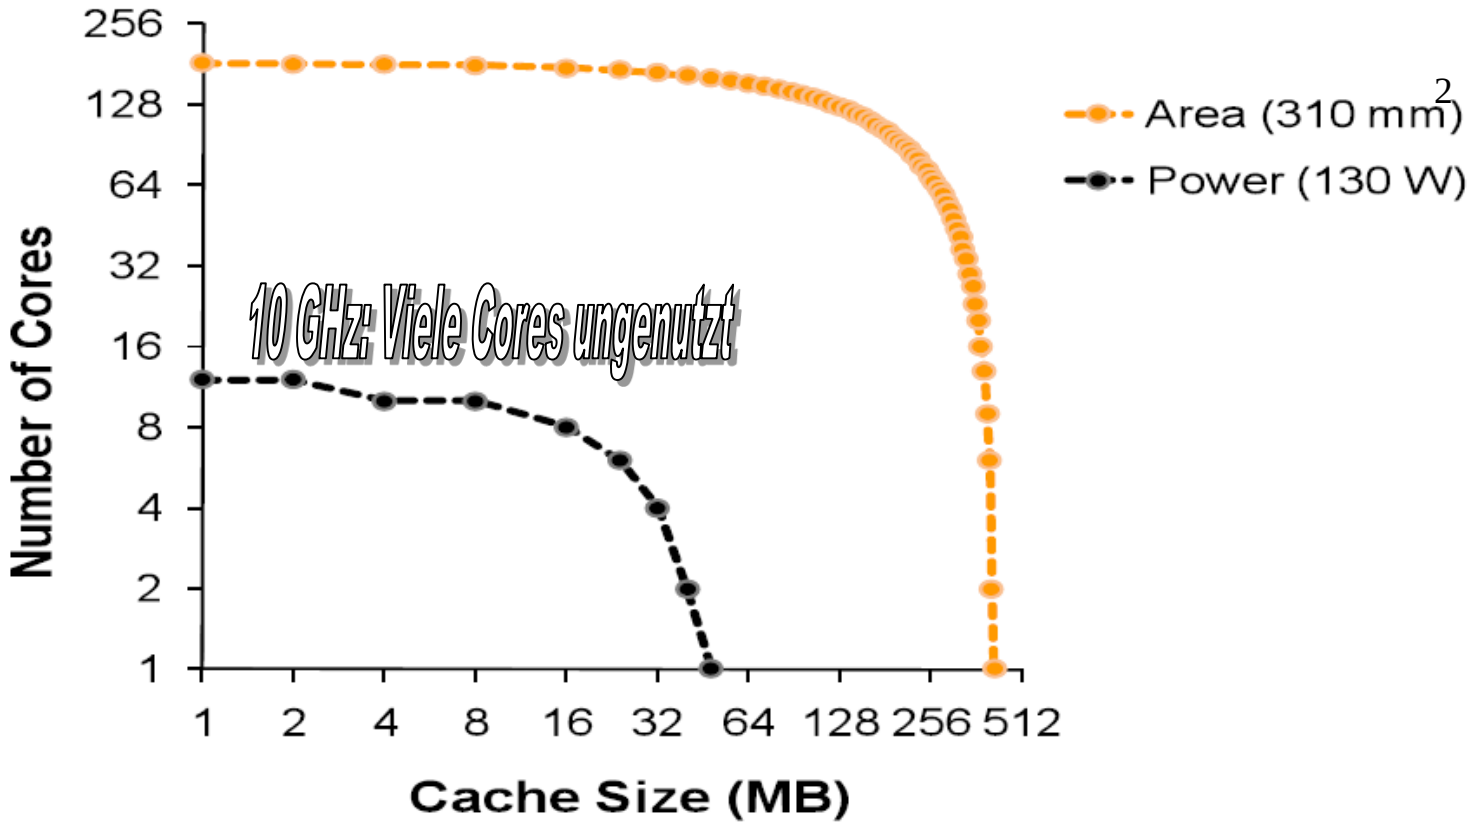

✓ Good news: can fit hundreds of cores

× Can not use them all at highest speed

Of course one could pack more slower cores, cheaper cache

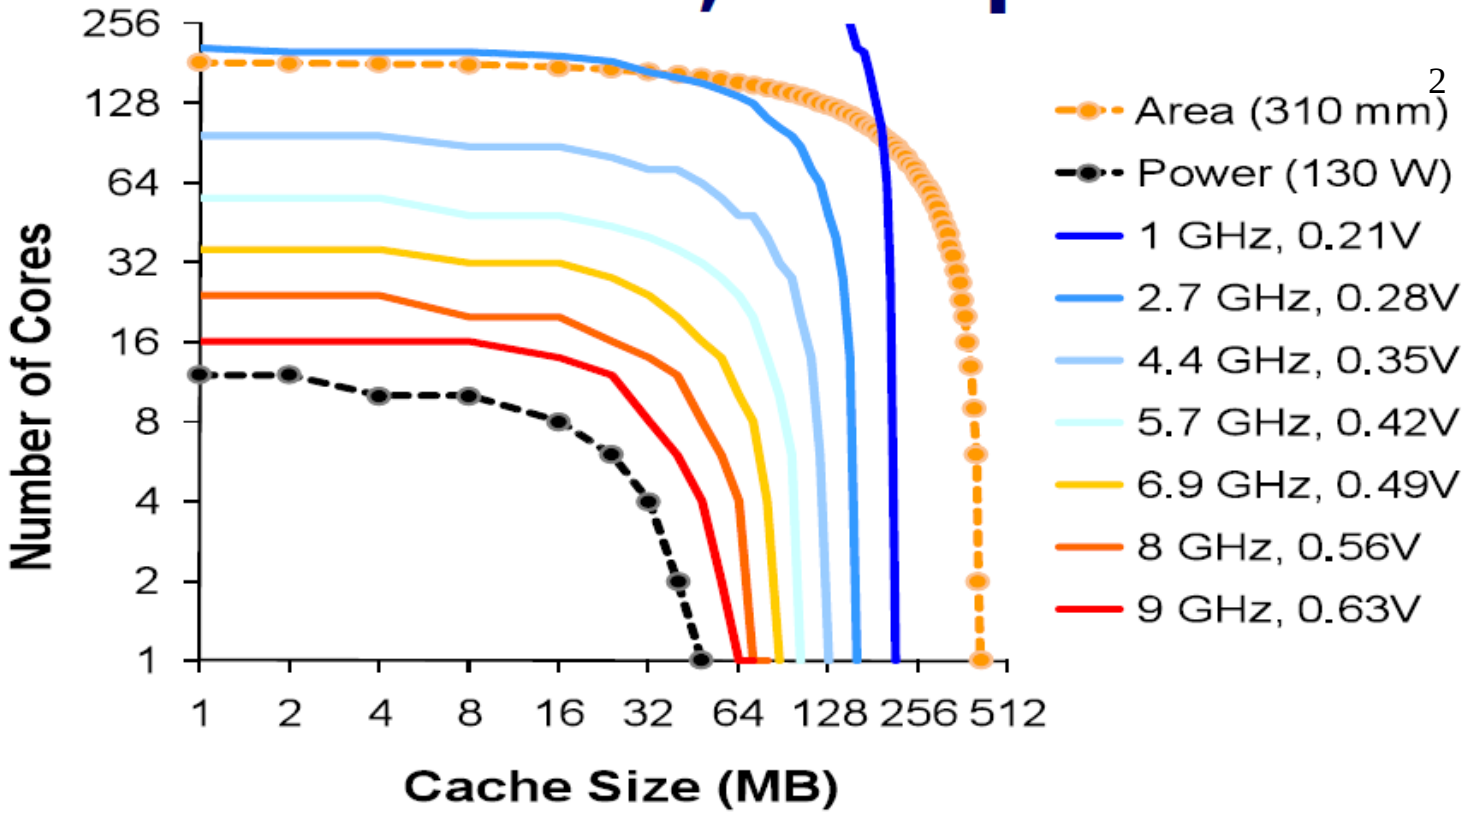

- Result: a performance/power trade-off
- Assuming bandwidth is unlimited

But, limited pin b/w favors fewer cores + more cache

Grenze durch verfügbare Bandbreite (klein bei wenig Cores)

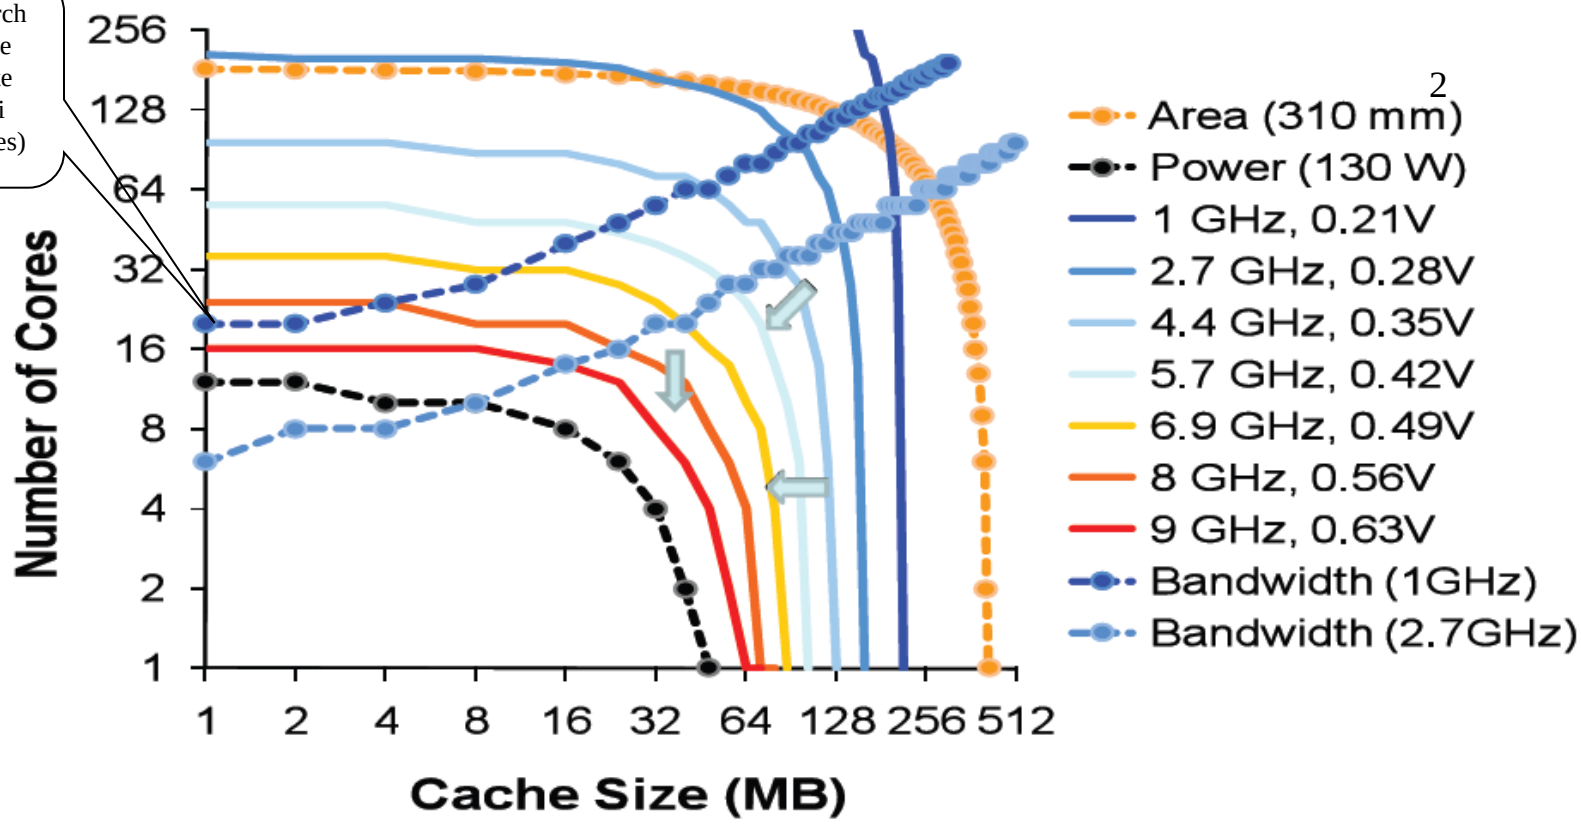

- For clarity, only showing two bandwidth lines
- Where would the best performance be?

Peak Performing with Conventional Memory

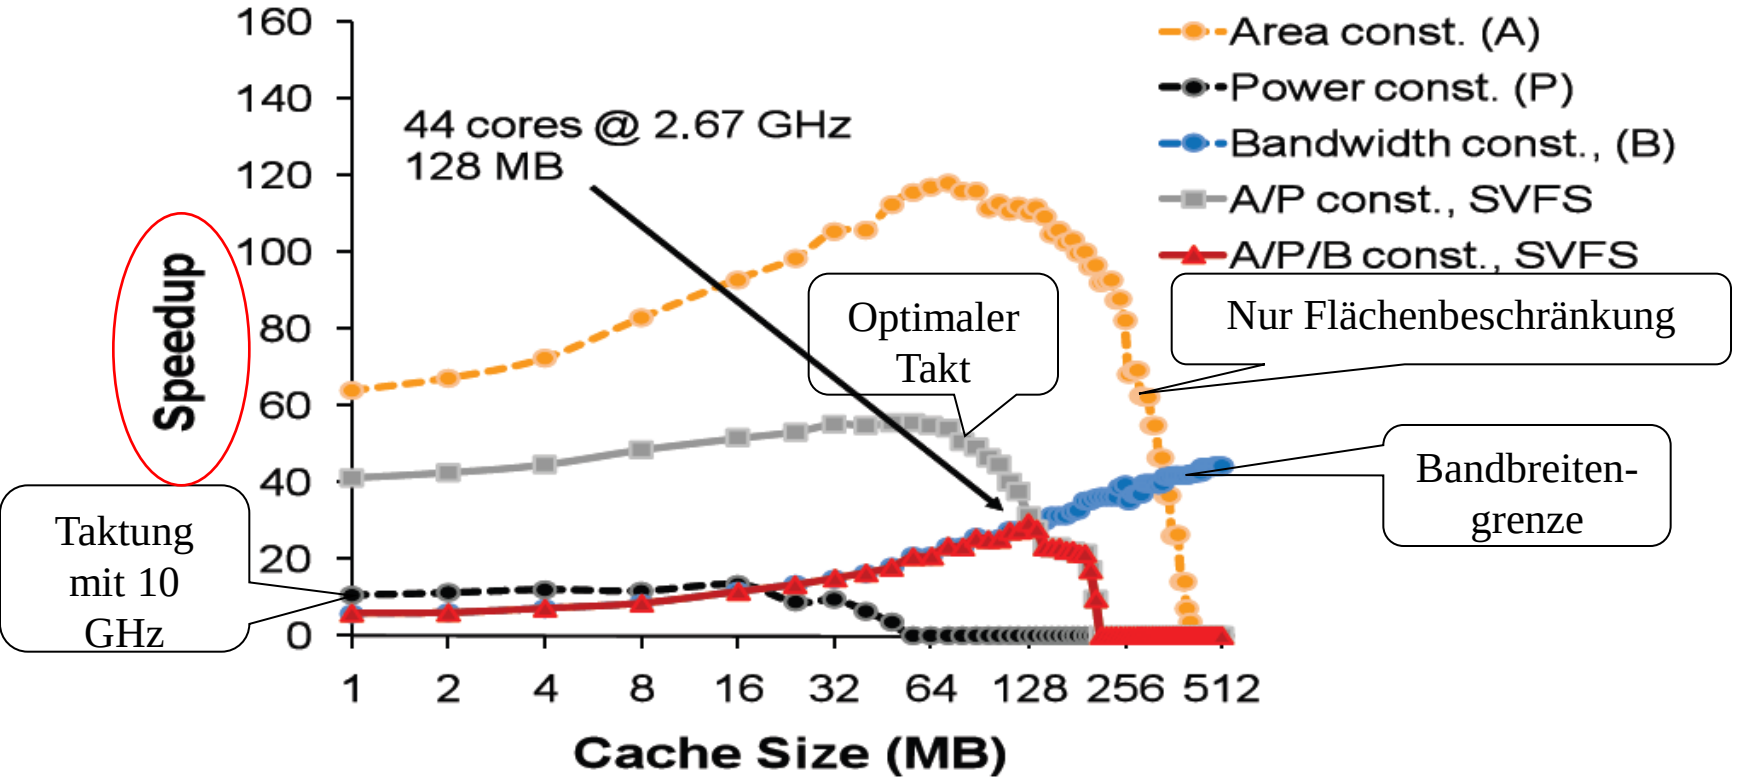

- B/W constrained, then power constrained
- Fewer slower cores, lots of cache

Mitigating B/W Limitations: 3D-stacked Memory

[Loh, ISCA'08]

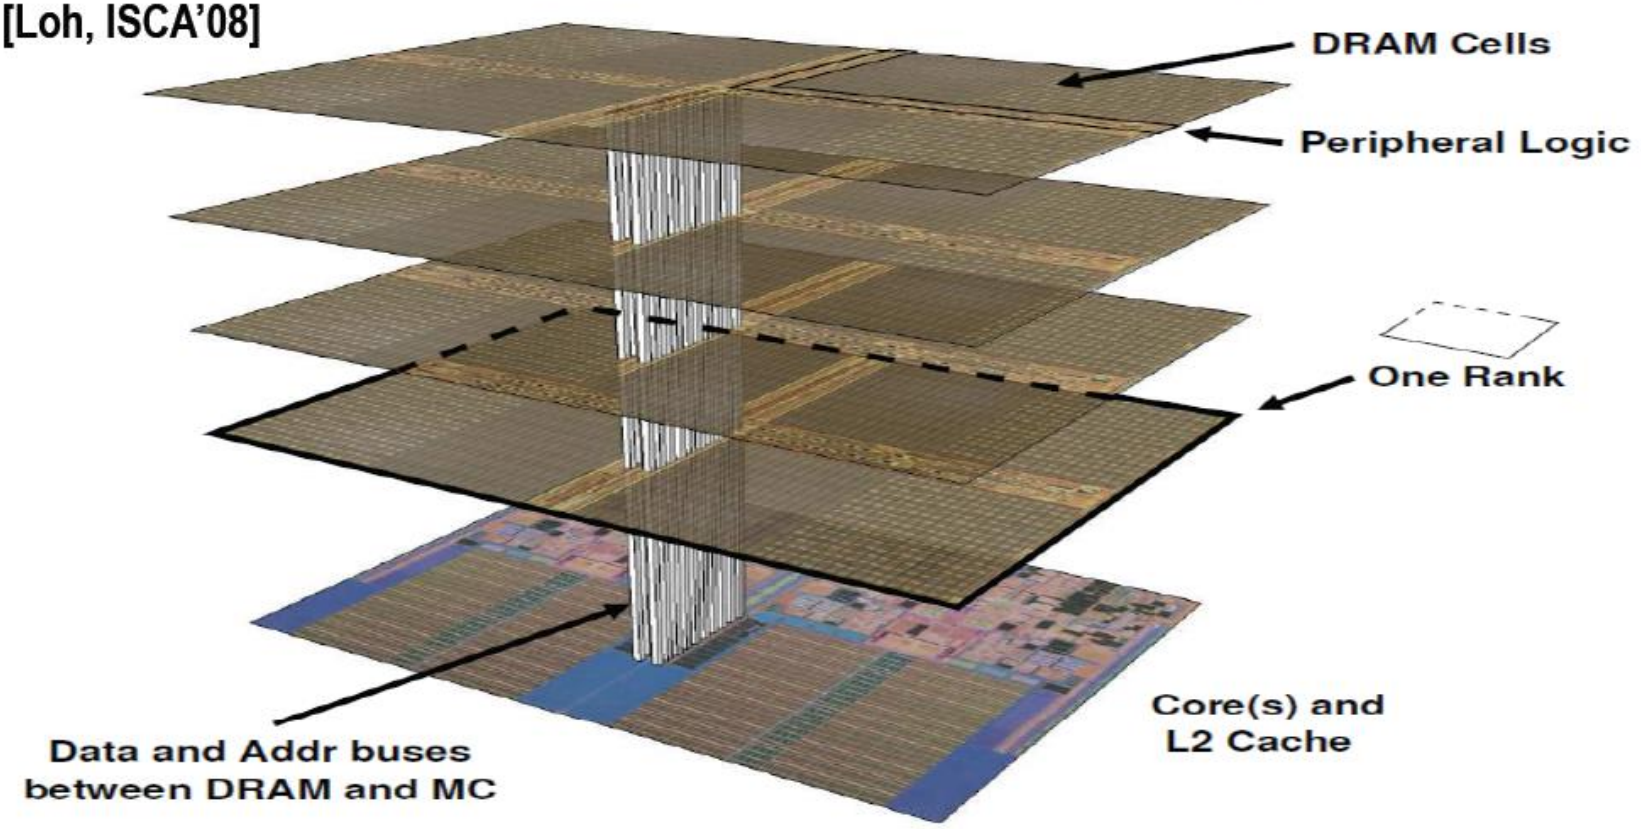

- Delivers TB/sec of bandwidth

Peak Performing w/ 3D-stacked Memory

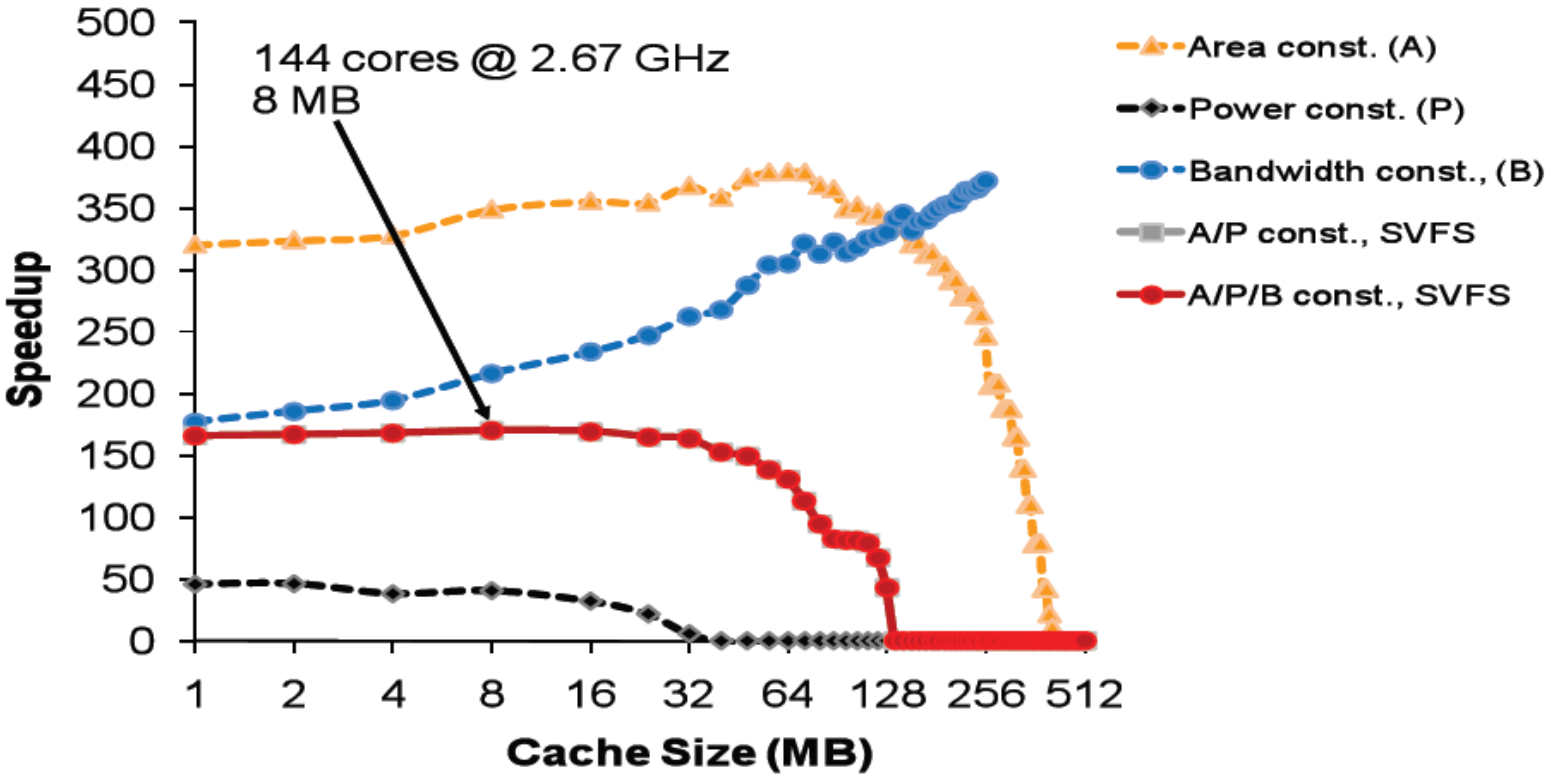

- Only power-constrained
- Virtually eliminates on-chip cache

Long-term: Where to go from here?

1. Redo SW stack
   - Minimize joules/work (algo. down to HW)
   - Program for locality + heterogeneity

2. Pray for technology
   - Energy-scalable silicon devices
   - Emerging nanoscale technologies?

3. Infrastructure technology
   - Renewable/carbon-neutral energy
   - Scalable cooling + power delivery

Short-term Scaling Implications

- Caches are getting huge
  - Need cache architectures to deal with >> MB
  - E.g., Reactive NUCA [ISCA'09]
- Interconnect + cache hierarchy power
  - Need lean on-chip communication/storage
  - Eurocloud chip: ARM+3D [ACLD'10]
- Dark Silicon
  - Specialized processors
  - Use only parts of the chip at a time

Zusammenfassung

- Trend zur Miniaturisierung von Schaltungen hält vermutlich noch einige Jahre an (wie viele?)
- Man trifft neben der ***memory wall*** auf die ***power wall***
- Man kann mehr Transistoren auf den Chips integrieren, aber sie nicht mehr alle gleichzeitig mit Strom versorgen (☞ ***dark silicon***)

- Partielle Abhilfen:
  - *Embedded* Prozessoren (z.B. ARM) als Server?
  - Spezialisierte Prozessoren (z.B. ARM big.LITTLE)
  - 3D-Integration von Speicher?
  - Über die Energieeffizienz von Software nachdenken!