LET'S CREATE IT

Multicore for 4G: 3GPP versus ITRS

3GPP A GLOBAL INITIATIVE

MPSoC 2011, Beaune, July 6

Kees van Berkel

professor @ TU/e

fellow @ ST ERICSSON

ATRENE

Mobile internet: cost in US$/GB

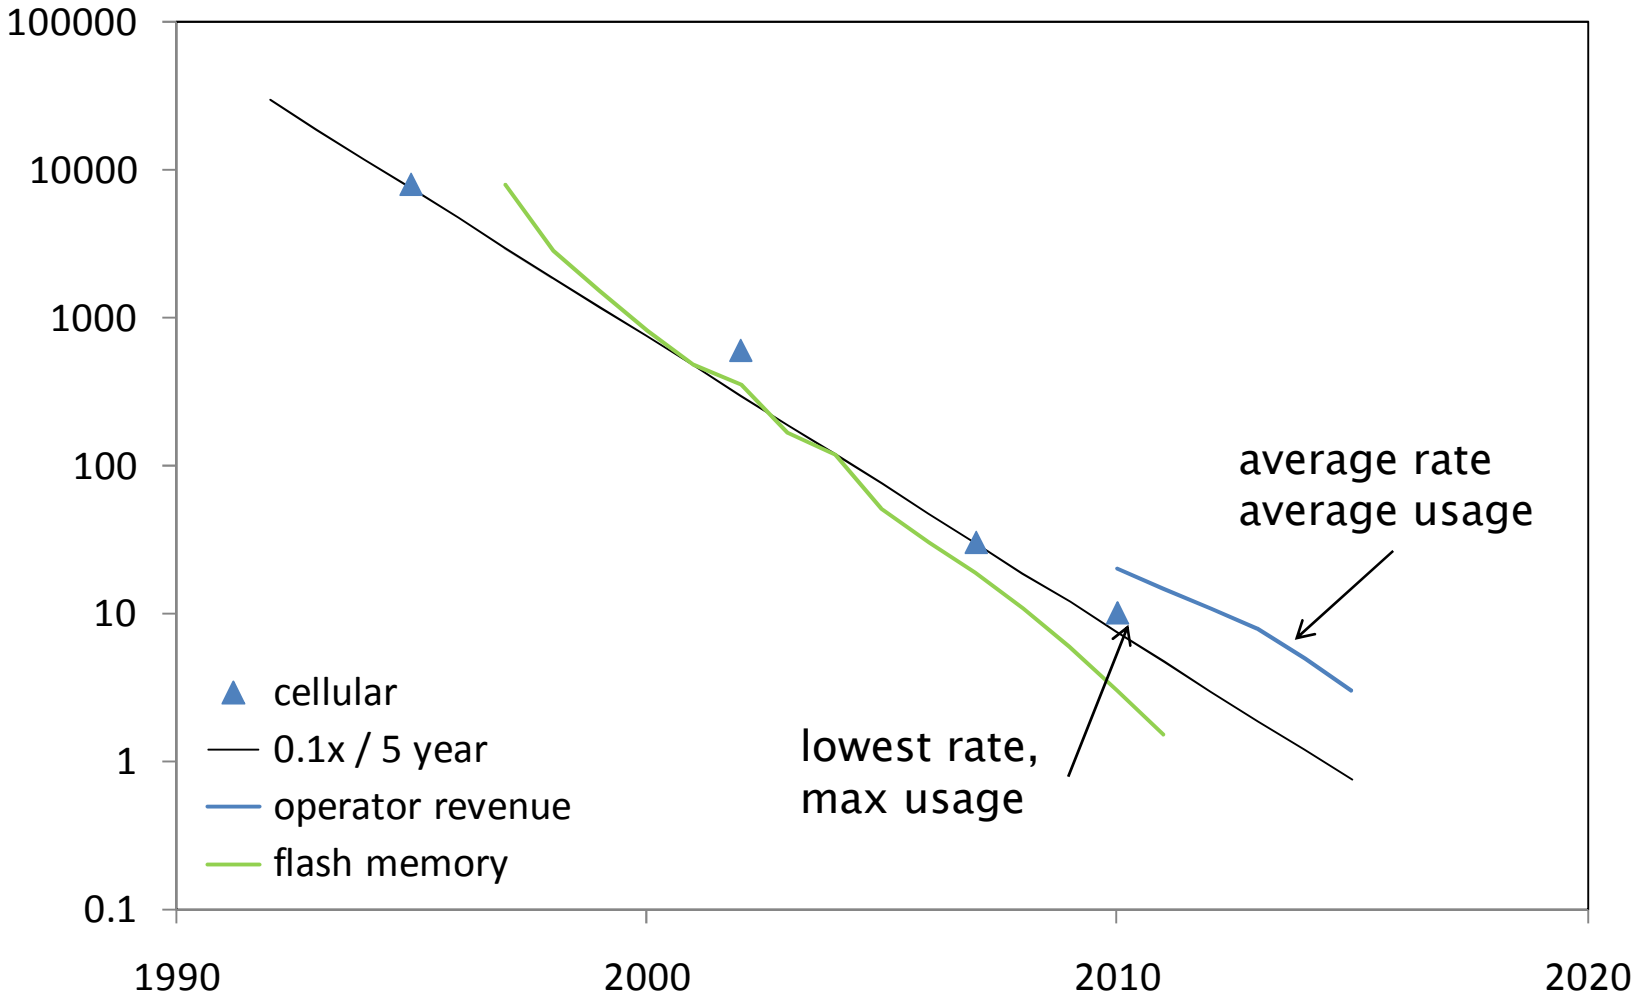

Mobile internet is fueled by a steady decrease in cost per GB:
0.1× /5 years, ...

Cellular downlink [Mbit/sec]

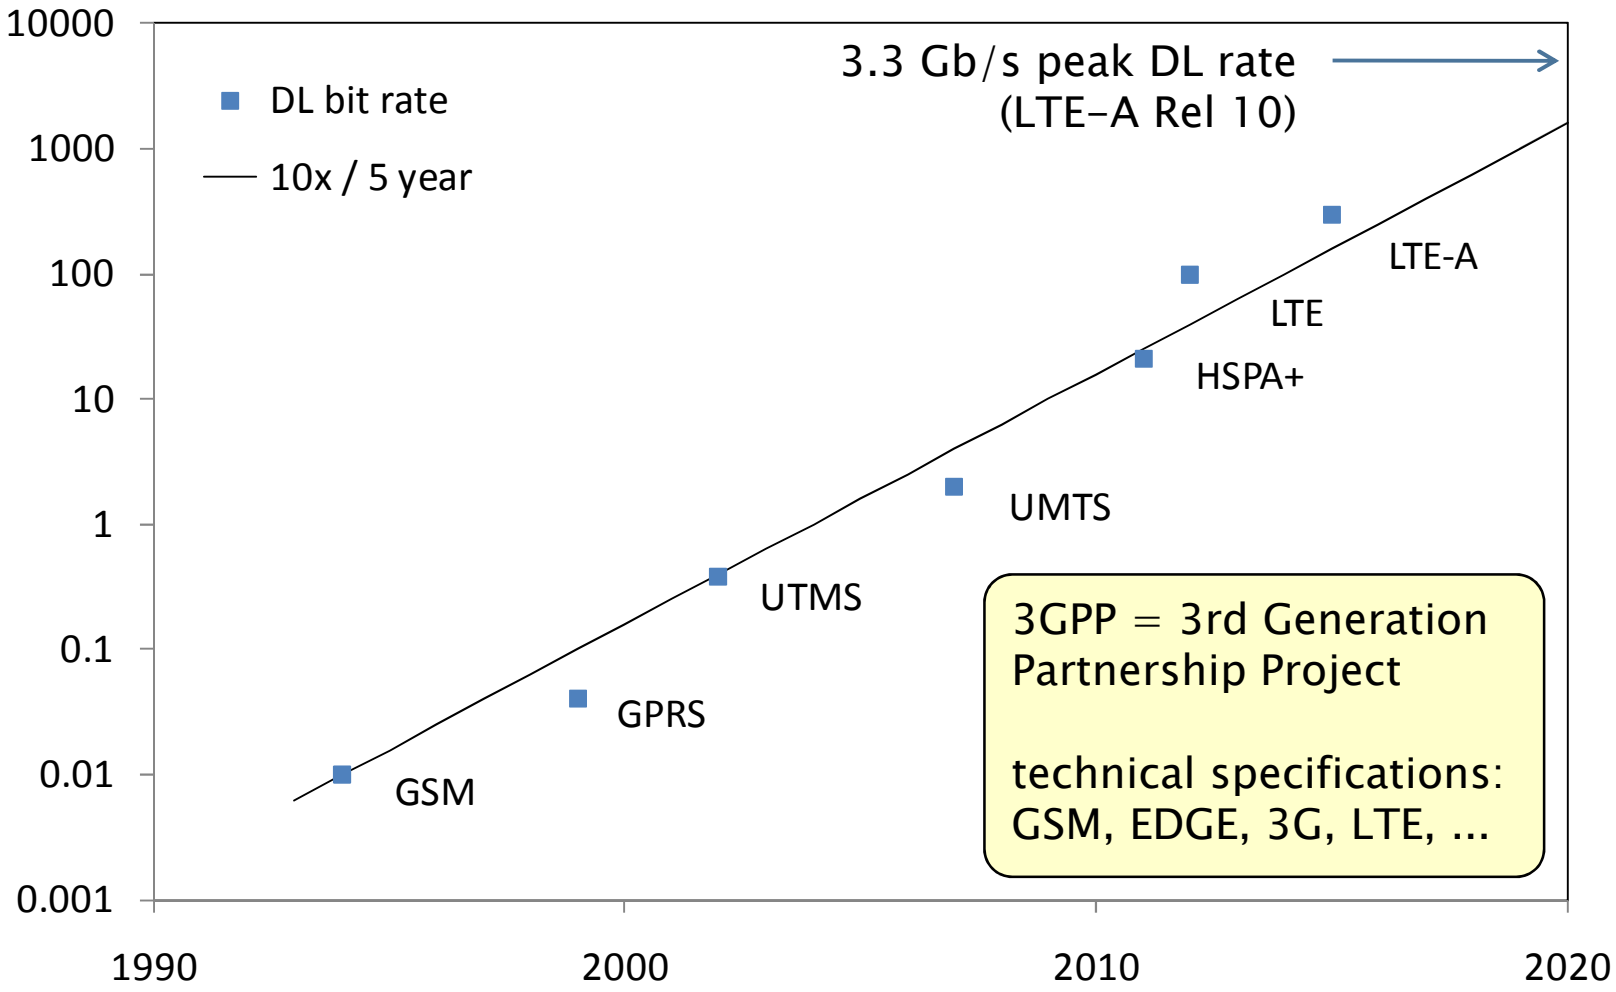

... and by a matching steady increase in (peak) downlink data rate of 10× /5 years, ...

CMOS feature size [nm]

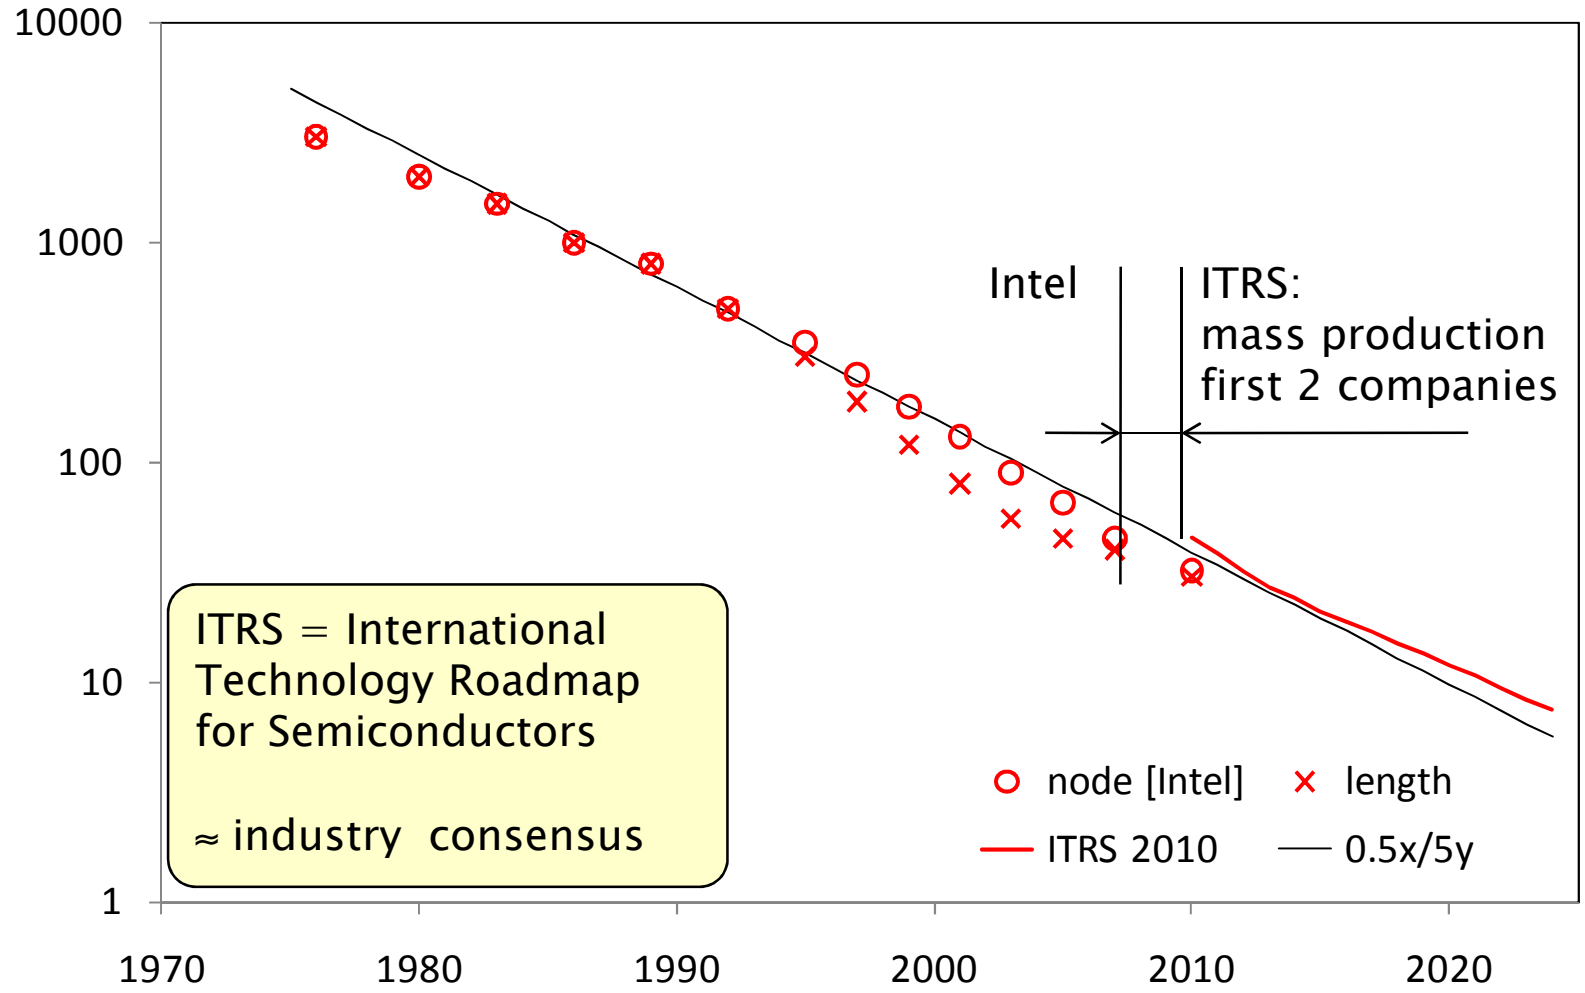

... enabled by a steady decrease in (CMOS) feature size of "only" 0.5× /5 years, ...

3GPP versus ITRS

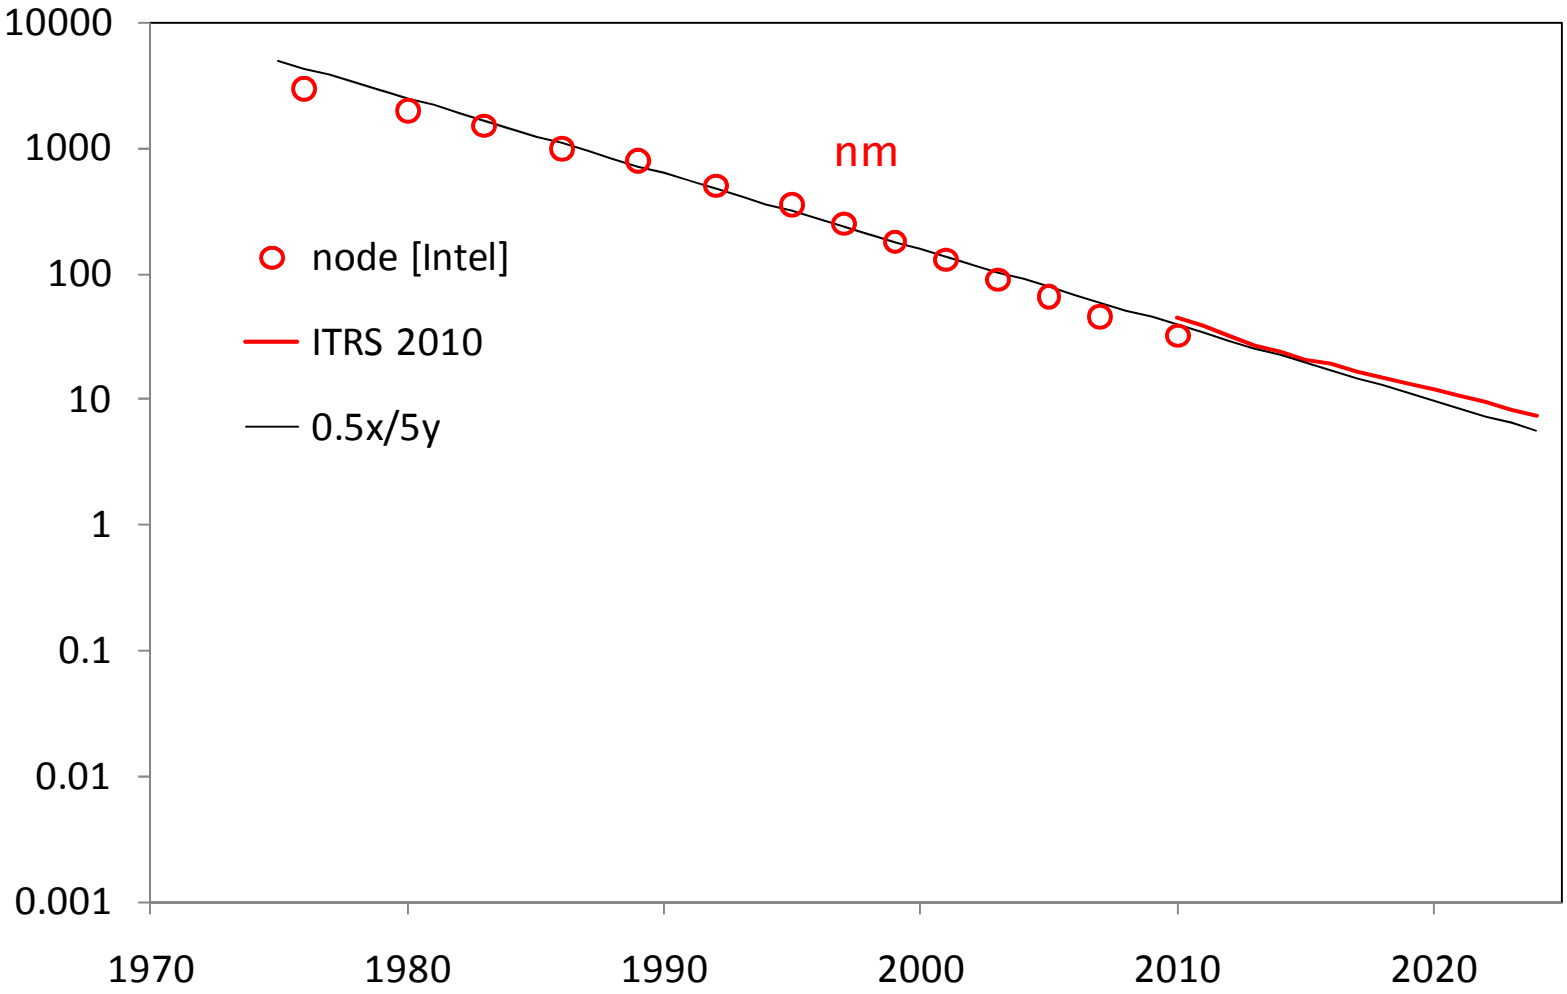

3GPP A GLOBAL INITIATIVE versus ITRS

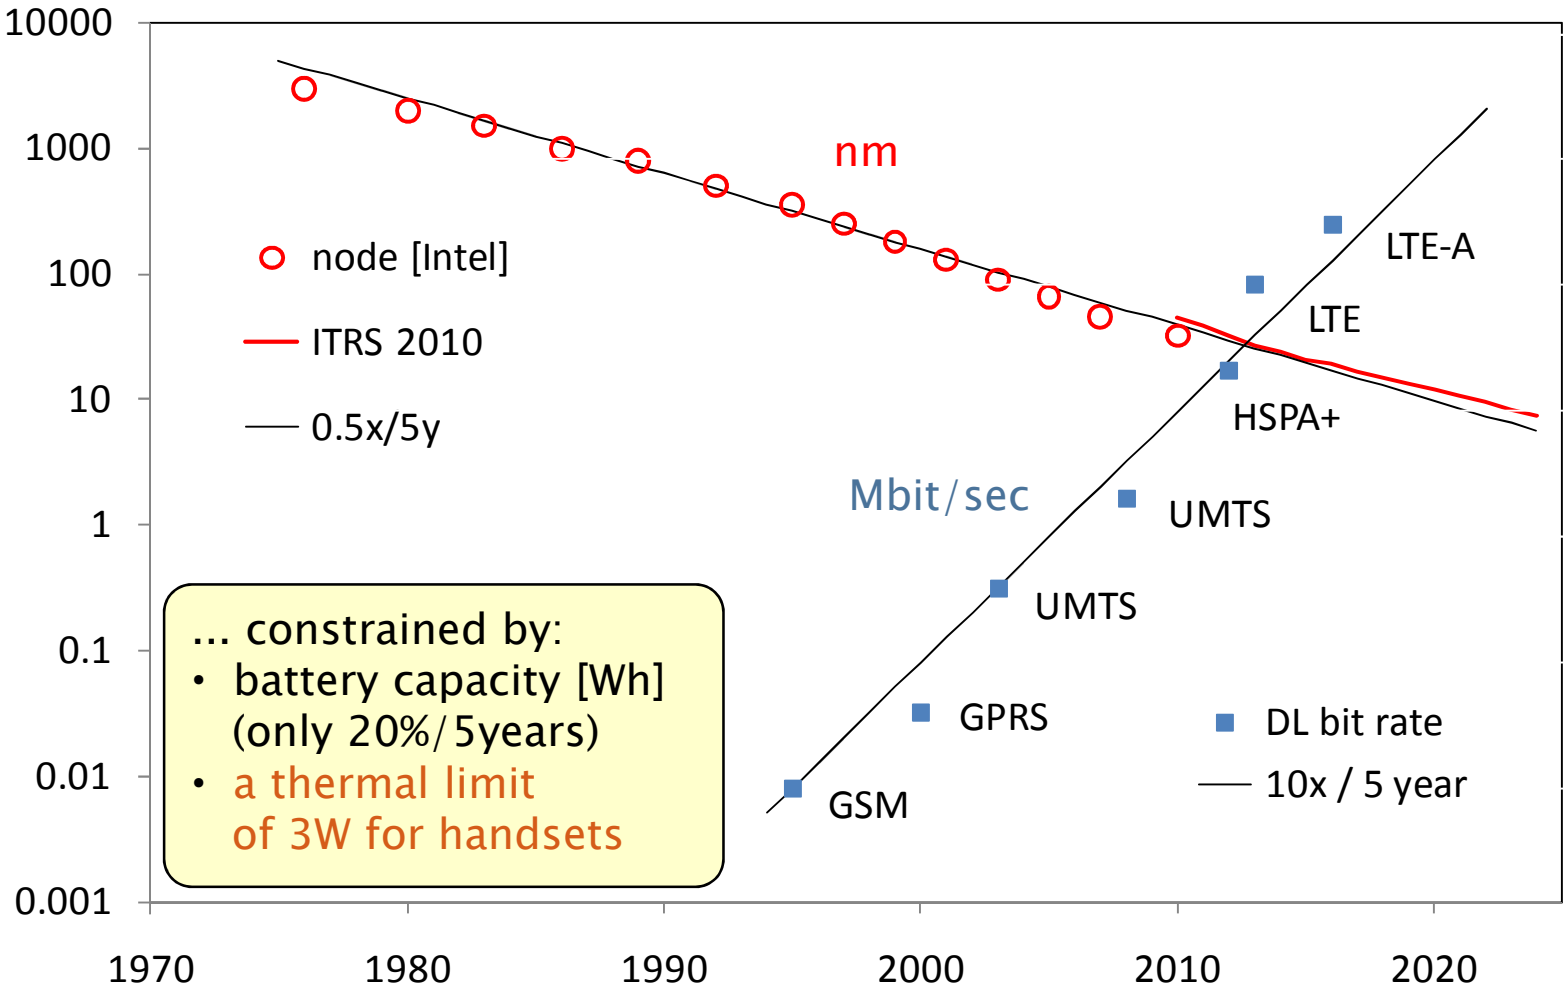

Can ITRS keep up with 3GPP ?

TU/e ST ERICSSON

Plan: "GMAC[16b]/200mJ" (= GMAC/s/200mW)

Focus on "Multiply-ACumulate" (MAC) part of baseband processing

Assume power *budget* of 200 mW for "MAC part"

Quantify:

- available GMAC/200mJ [16b] for "ITRS year"
- required MAC/b (16b-MAC/received bit) for "3GPP year"

$\Rightarrow$ required GMAC/s for high-end data rate for "3GPP year"

Robert H. Dennard [1974]:

$L \propto \alpha$, $V \propto \alpha$ ("constant field")

$\Rightarrow C \propto \alpha$, $I \propto \alpha$ (constant I/μ)

$\Rightarrow$ delay $= CV/I \propto \alpha$

$\Rightarrow$ energy $= CV^2 \propto \alpha^3$

With Dennard scaling: /5 year

CMOS	α	0.5
$\Rightarrow$ Energy	α^3	0.125
$\Rightarrow$ GMAC/J	α^{-3}	8×
3GPP bit rate		10×

... the outlook seems promising!

TU/e ST ERICSSON

Dennard scaling? Does "$V \propto L$" ?

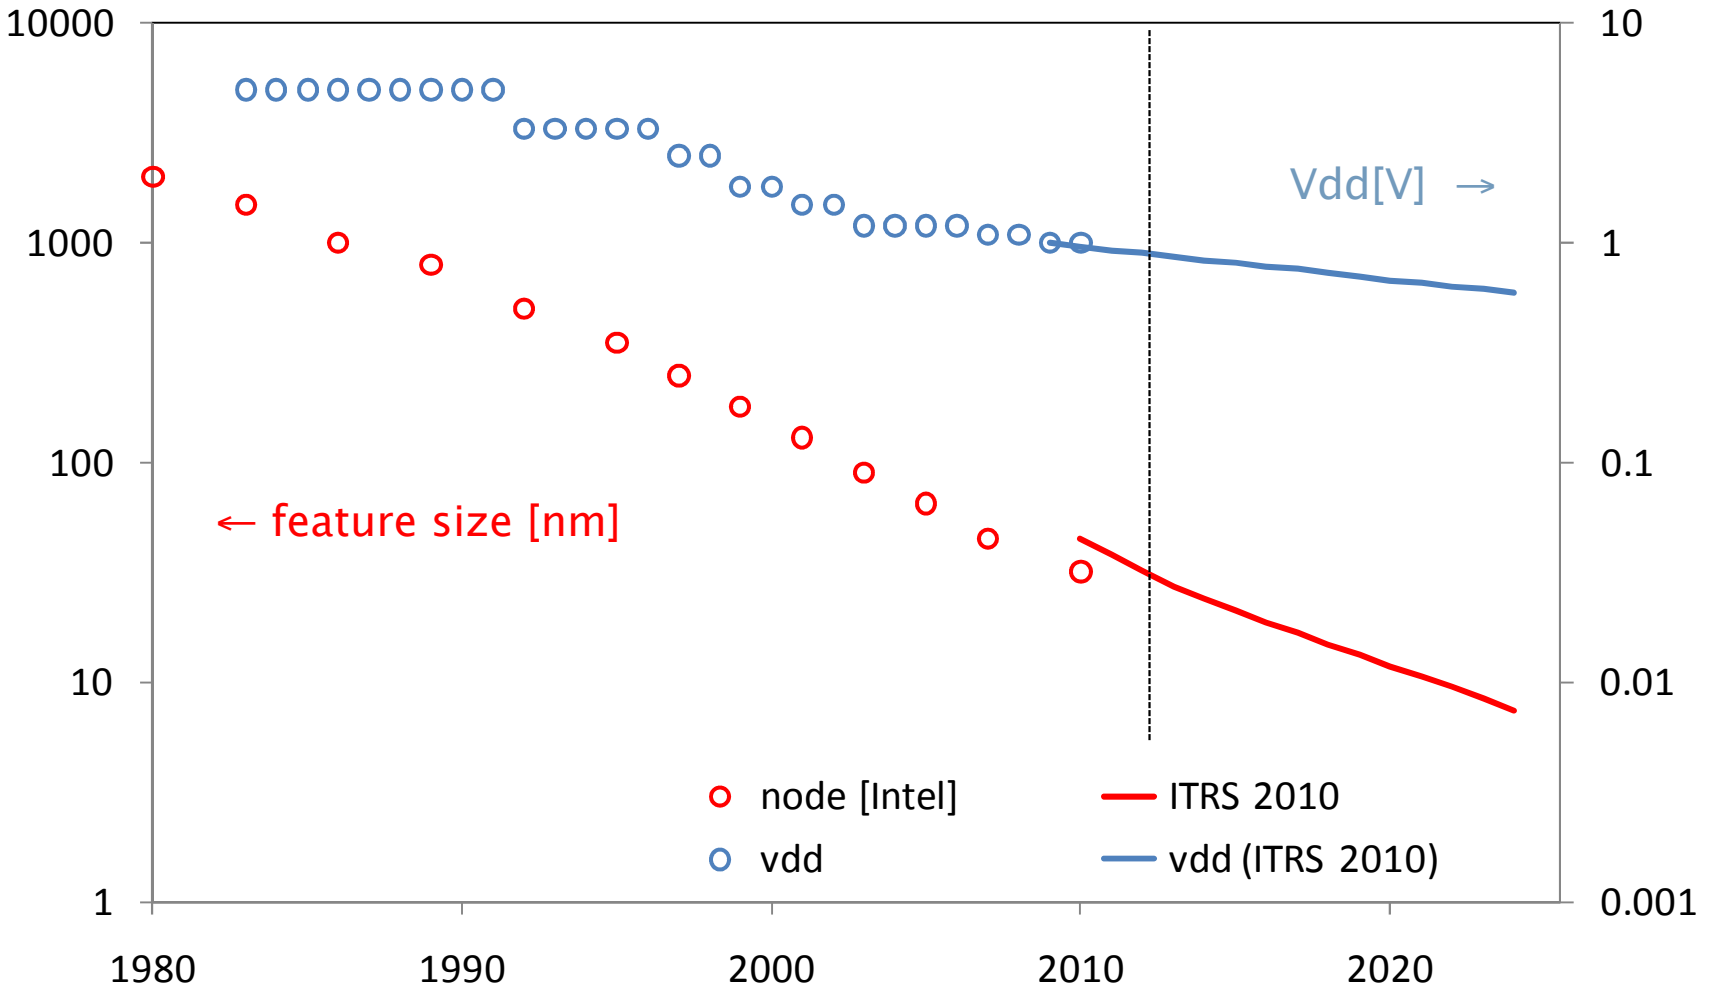

No! Feature size/voltage: we have lost an order of magnitude!
... and the gap is widening

GMAC[16b]/200mJoule

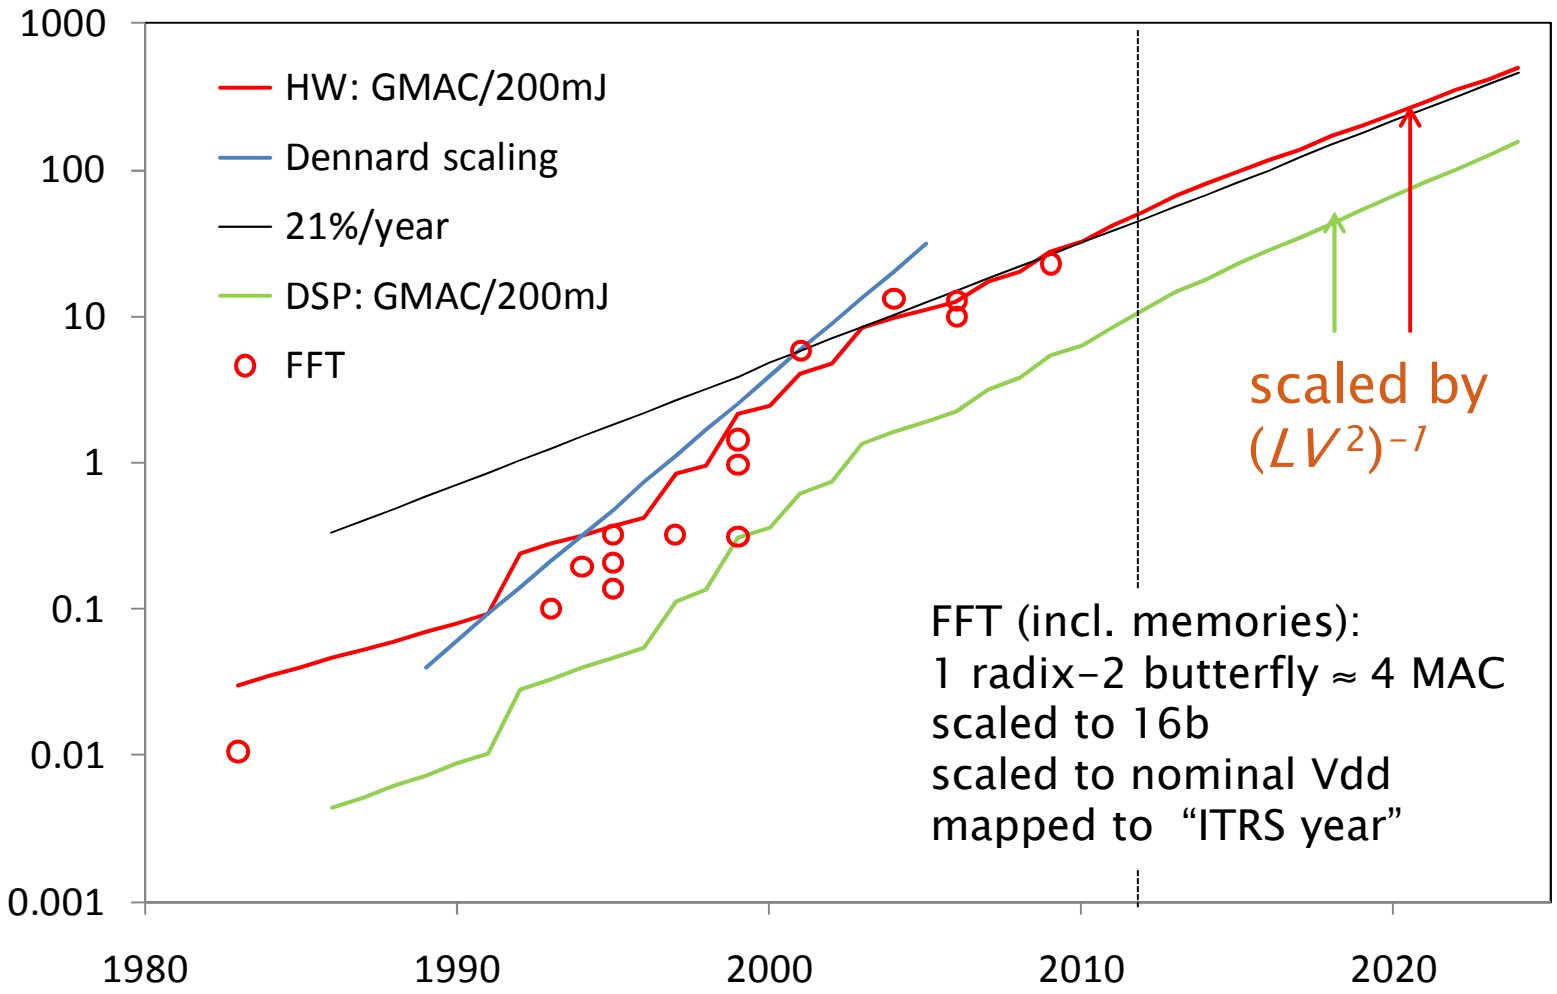

However, Dennard scaling lasted only 1 decade: 1991 – 2002

Baseband (simplified): MAC/bit

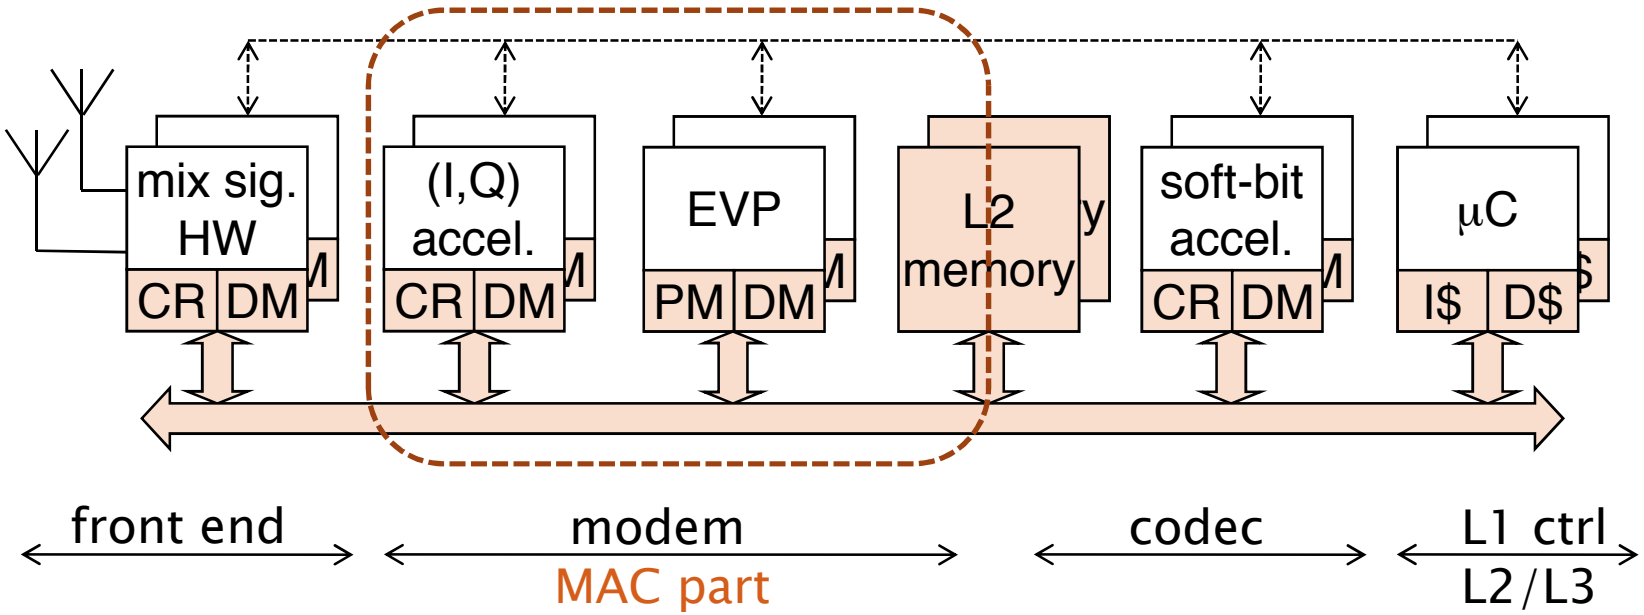

- "MAC part": mostly complex numbers, 2×16b:
  FIRs, IIRs, FFTs, correlators, $M \times V$, $M \times M$, M^{-1}, data selection, ...
- WCDMA: a rake receiver also uses "1b×16b" complex MAC (additions)
- total: 100 ↔ 300 MAC/bit (simple ↔ advanced algorithms)
- trend: towards more advanced algorithms, to mitigate interference

3GPP versus ITRS: an increasing GMAC gap

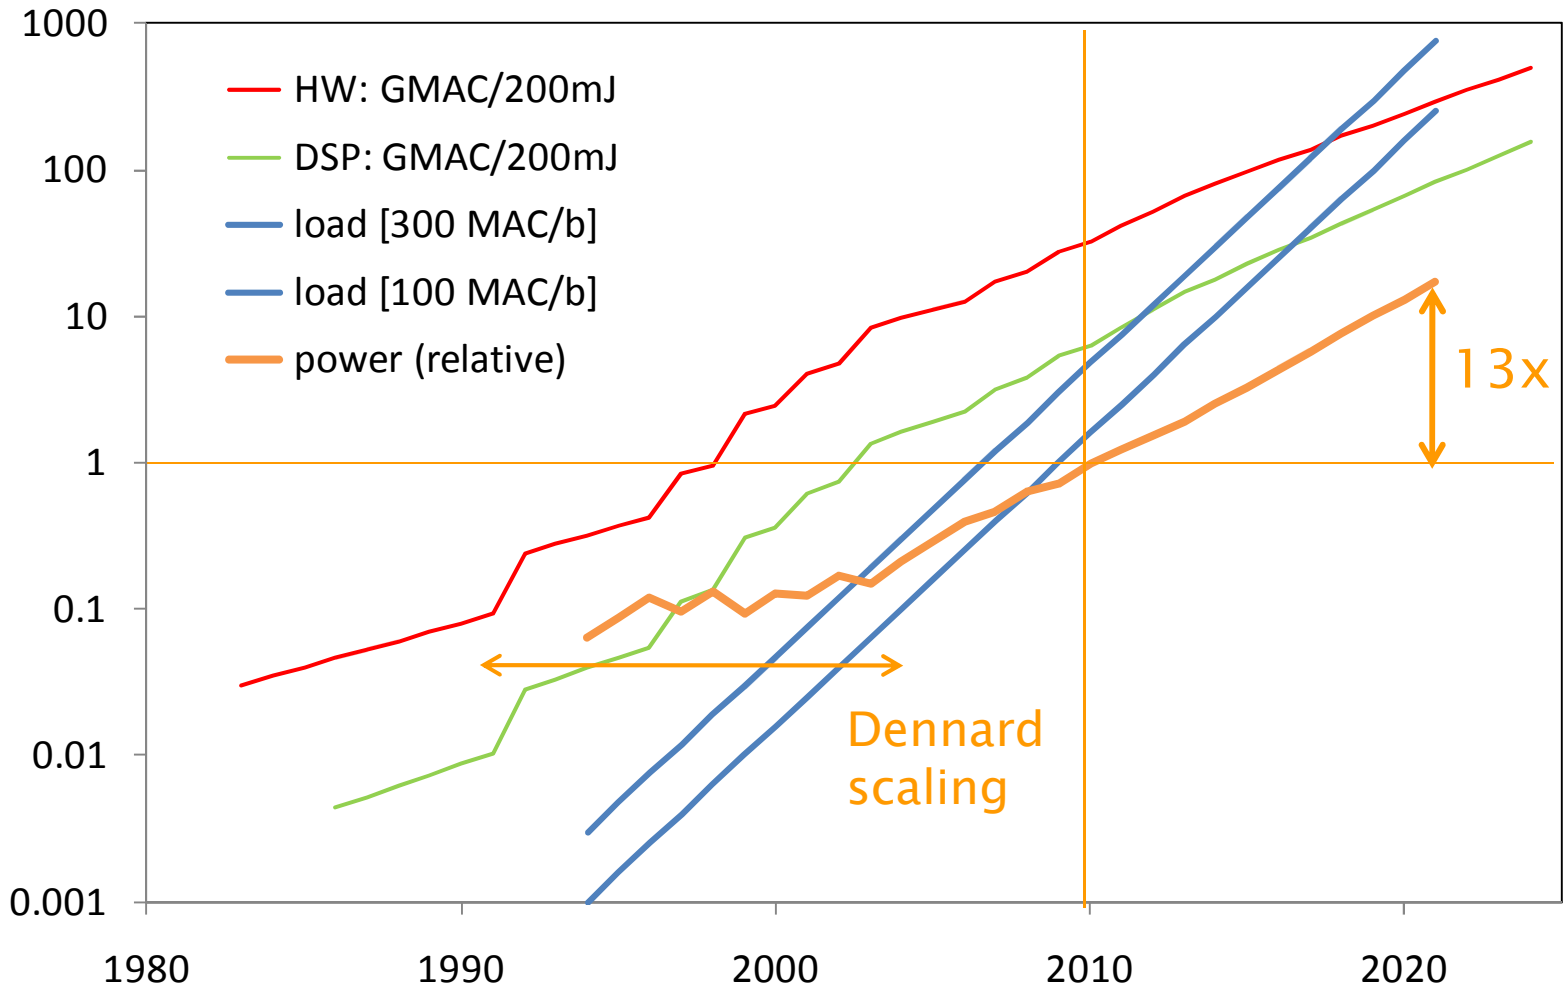

- 2010 – 2020: modem power up 13x (more with heavier algorithms!)

Closing the gap 1: increase power budget

- power of MAC part will increase 13× in the current decade:
- unlikely for smart phones, but
- some increase likely for tablets

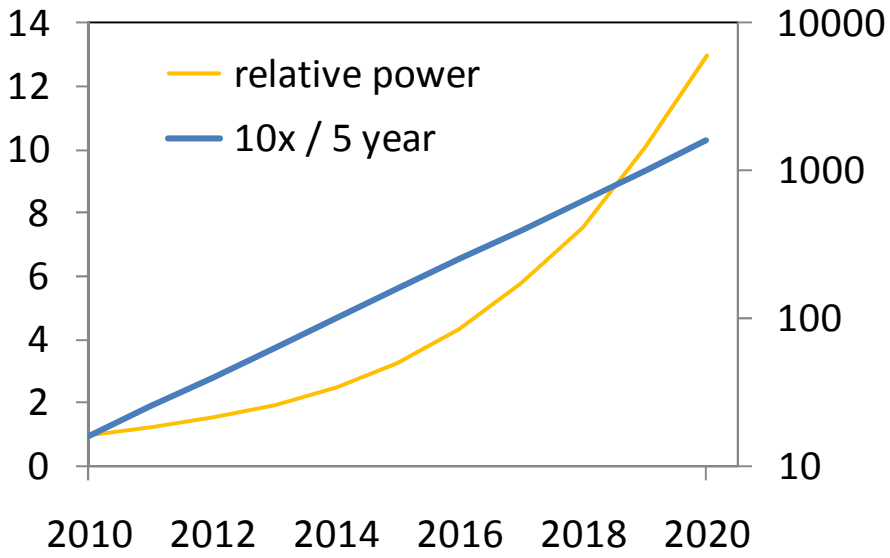

Larger surface area of tablets:

- more room for batteries
- larger area to emit heat
- more room for antennas

Closing the gap 2: peak rate for bursts only

- keep average bit rate
  @ constant power of 200mW
- allow peak rate for bursts only
  (throttling the DL stream)
- period <<
  thermal time constant handset
- period < user response time
  (content dependent)

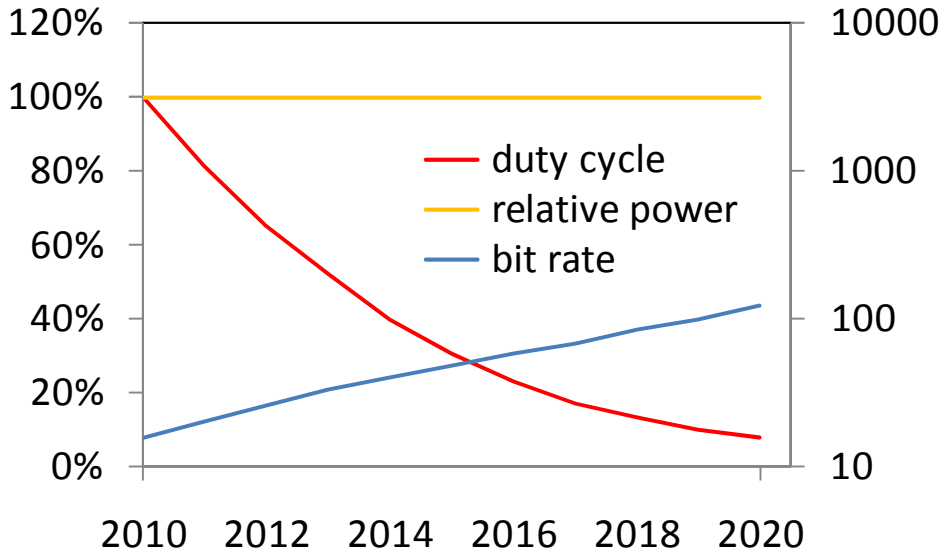

- This results in a decreasing duty cycle of the baseband processing.
- Bonus: fewer DSP resources needed.
- This requires a standardized protocol with the base stations.

Closing the gap 3: optimize algorithms

Equivalent algorithms may require fewer MACs/bit, e.g

- Fast FIR
- frequency-domain filtering
- ...

... but unlikely to provide 13×

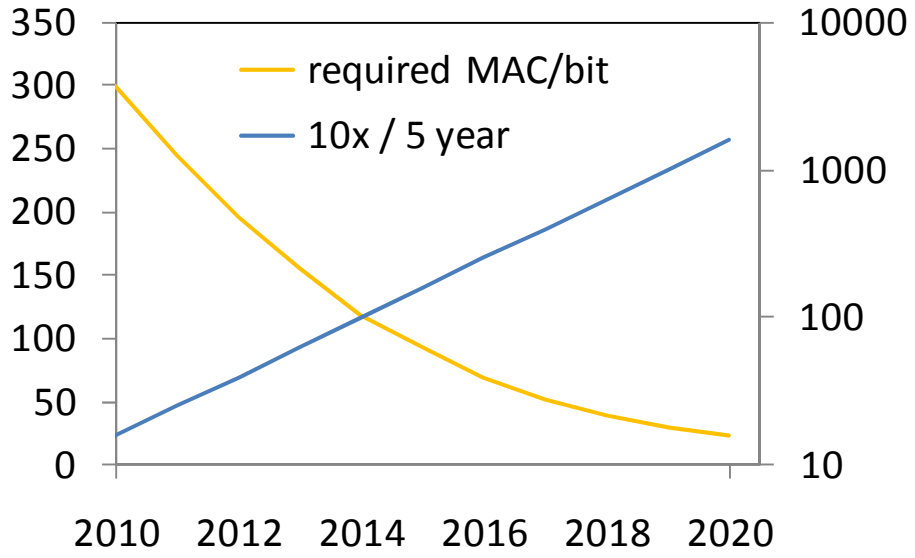

Simpler algorithms for high bit rates: (*scalable algorithms, adaptation*):

- high bitrates only feasible when channel is "clean";
- .. allowing for simpler algorithms that require fewer MAC/bit.

That is, high MAC/bit algorithms only when channel is challenging.

Closing the gap 4: more parallelism at low Vdd

1. scale Vdd with $\sqrt[-1]{load\ increase}$
   to keep dynamic power constant
2. calculate f_{clock} slow down
3. compensate lower f_{clock}
   by increase in parallelism
4. calculate area & watch in horror!

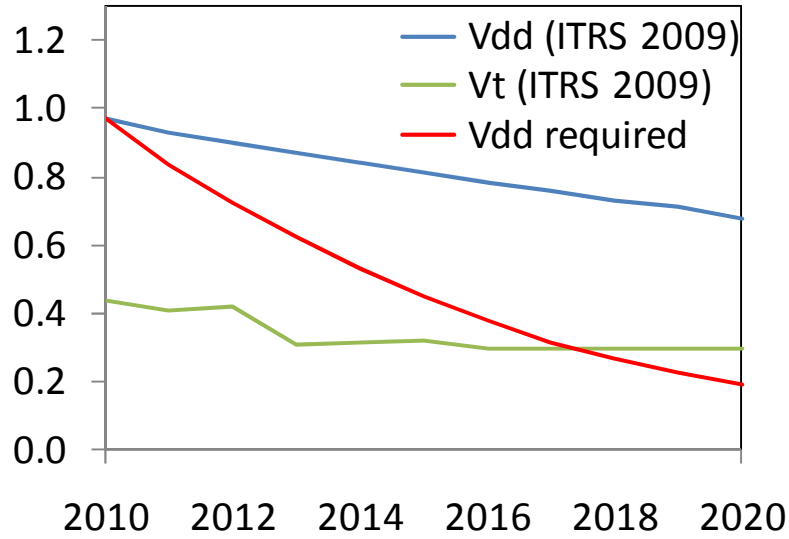

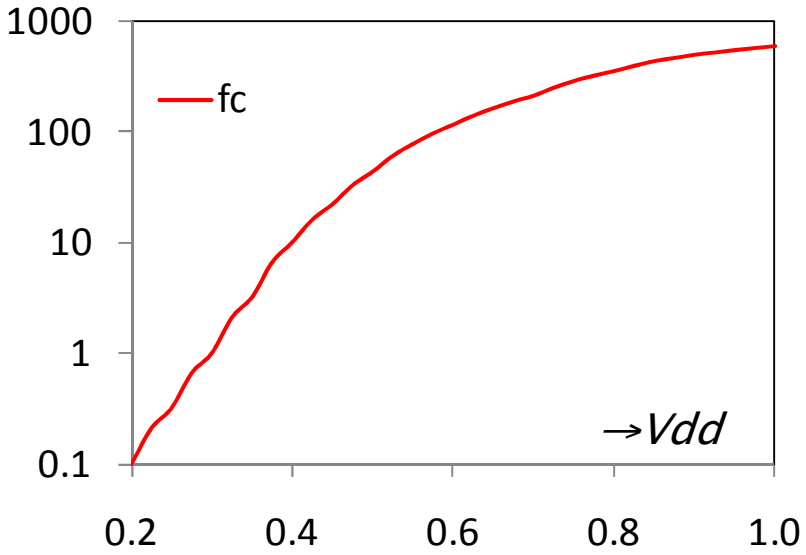

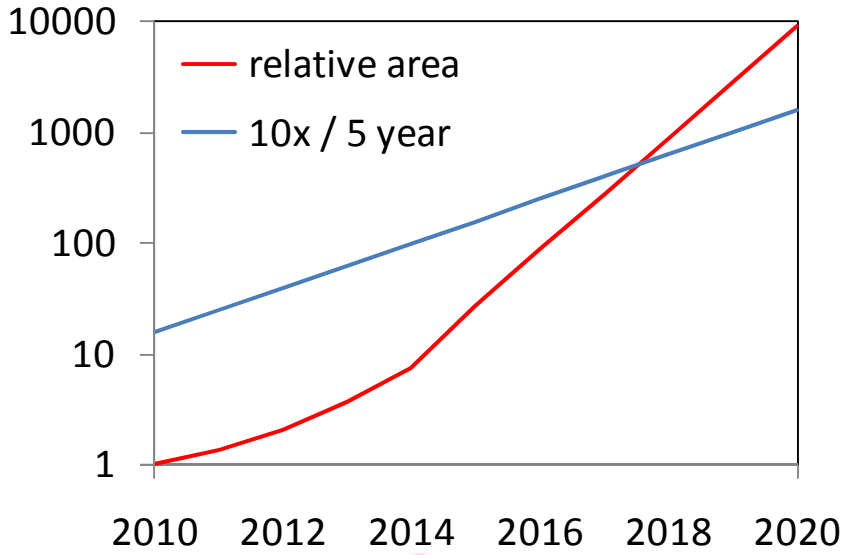

Closing the gap 5: optimize HW architecture

Conflict:

- multi-standard, multi-channel
  push: HW → SW
- to close the power gap
  push: SW → HW

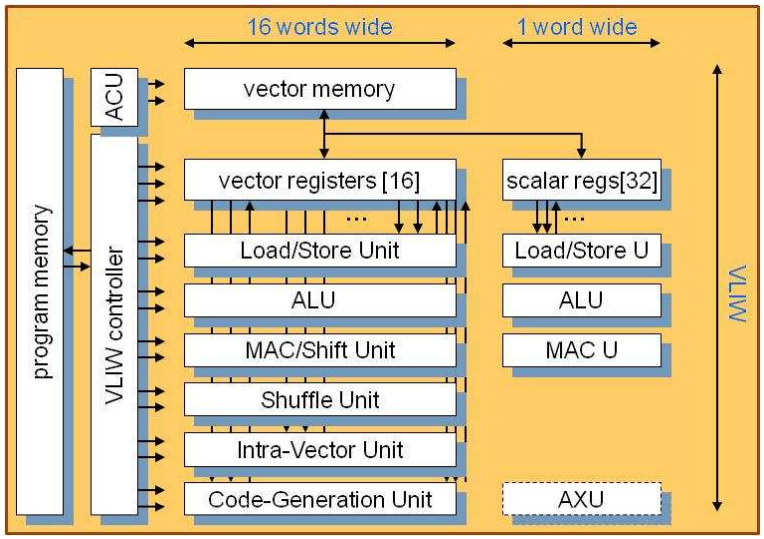

- Efficiency of DSPs [GMAC/J] is improving relative to HW
  SIMD, SIMD width, complex number support, special instructions, …
  … historically a few % per year, likely to continue.
- share of load on DSP to decrease over time?
- DSP flexibility often overkill; more tailored flexibility needed
- (how to quantify flexibility, versatility …?)

#accumulators, area [3 scenarios]

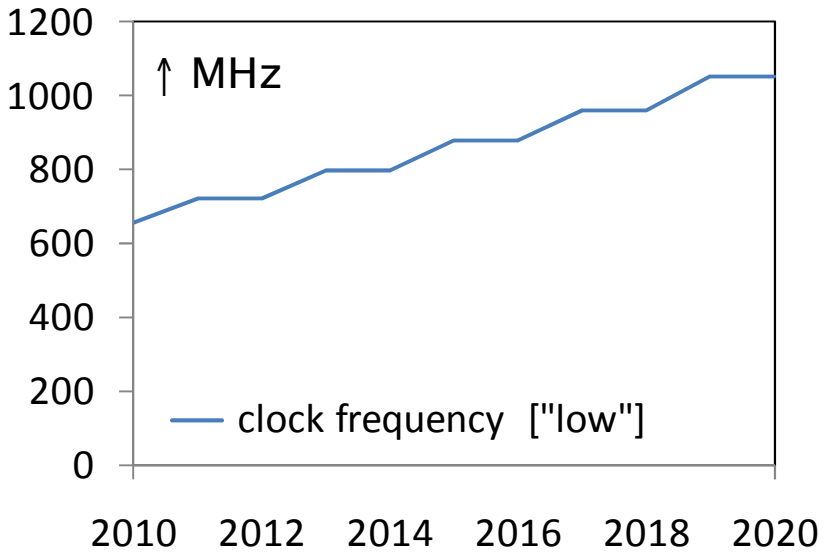

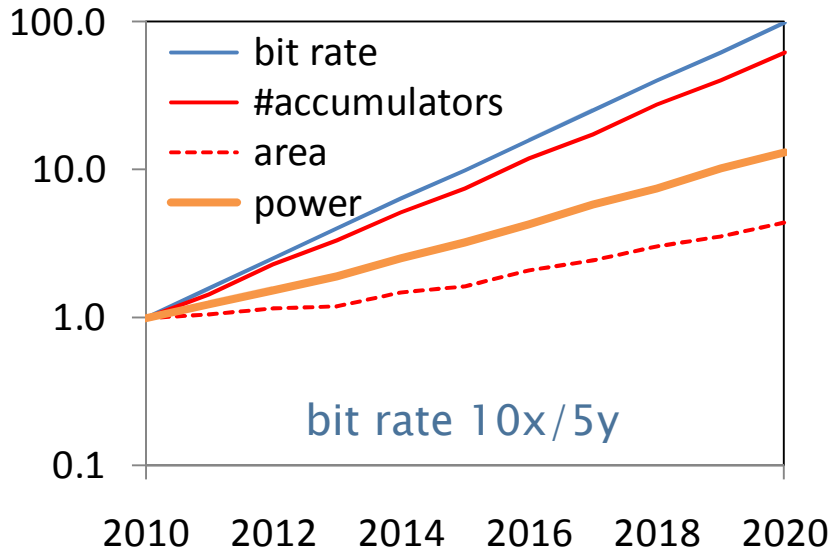

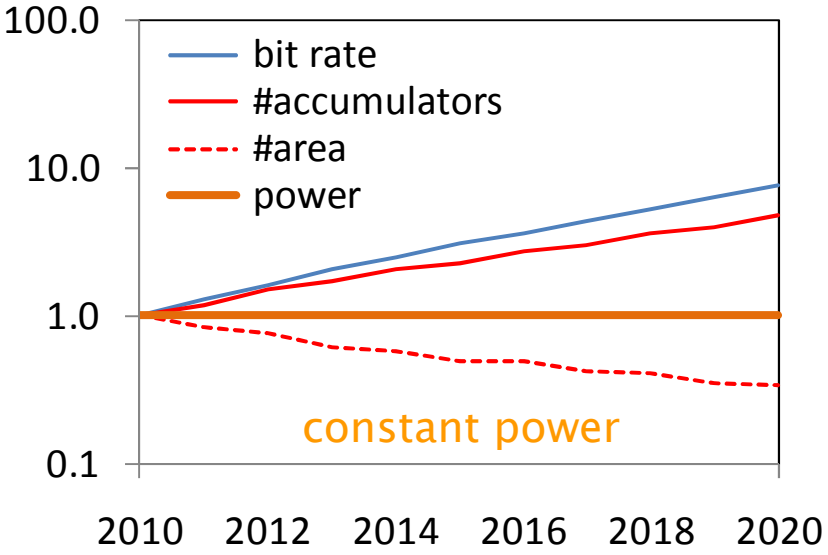

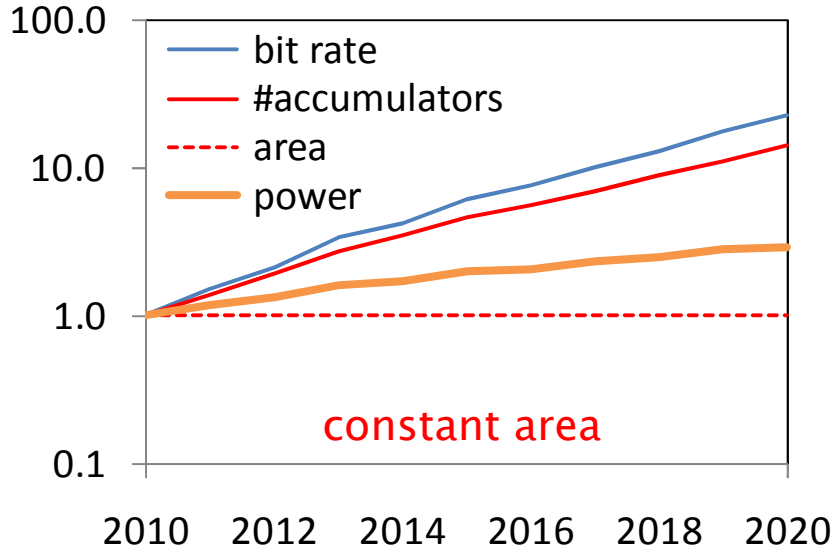

- power, without discussed measures

3GPP versus ITRS : conclusion

In the decade 2010–2020:

- the cellular peak downlink data rates will increase 100x
  (the associated baseband workload may grow even more)
- whereas CMOS feature size will decrease by only 4x
  and, as a result, MAC/200mJ will increase by only 8x

To close this >13x gap between available and required GMAC/200mJ, we need to:

1. allow for a somewhat higher power budget for modems in tablets,
2. restrict peak rates to a controlled duty cycle,
3. optimize baseband algorithms & adapt them to channel conditions,
4. operate at a (slightly) lower Vdd by using more parallelism (?),
5. optimize the hardware: by (less?) usage of more efficient vector DSPs.

Bibliography

1) 3GPP, http://en.wikipedia.org/wiki/3gpp,

2) LTE-Advanced, http://www.3gpp.org/LTE-Advanced

3) S. Sesia et al, LTE – The UMTS Long Term Evolution, from Theory to Practice, Wiley 2009
REV-090003r1 IMT-Advanced Evaluation Workshop, 2009, Beijing

4) Ghosh et al, LTE-Advanced: Next Generation Wireless Broadband Technology, IEEE Wireless Communications, June 2010

5) International Technology Roadmap for Semiconductors, http://www.itrs.net/reports.html

6) Hiroshi Iwai, "Technology Roadmap for 22nm and Beyond", 2009 *2nd Intl. Workshop on Electron Devices and Semiconductor Technology*

7) Kwangok Jeong and Andrew B. Kahng, "A Power-Constrained MPU Roadmap for the International Technology Roadmap for Semiconductors (ITRS)", *ISOCC 2009, pp. 49-52*

8) R. Dennard, et al., "Design of ion-implanted MOSFETs with very small physical dimensions," *IEEE Journal of Solid State Circuits*, vol. SC-9, no. 5, pp. 256-268, Oct. 1974.

9) Kees van Berkel, et al, Vector Processing as an Enabler for Software-Defined Radio in Handheld Devices, "EURASIP Journal on Applied Signal Processing," vol 2005, issue 16, pp. 2613-2625

10) J. Berkmann et al, On 3G LTE Terminal Implementation – Standard, Algorithms, Complexities and Challenges, IEEE Journal … 2008,

11) G. Gammie et al, A 28nm 0.6V Low-Power DSP for Mobile Applications, ISSCC 2011, p. 7.5

12) Kees van Berkel, Multi-Core for Mobile Phones, *DATE'09, April 2009, Nice (invited paper)*


LET'S
CREATE
IT
THANK YOU
ST
ERICSSON