
Efficiency Metrics and Bandwidth
- A Memory Perspective -
Norbert Wehn
wehn@eit.uni-kl.de
//ems.eit.uni-kl.de
ReCoSoC
July 2013, Darmstadt

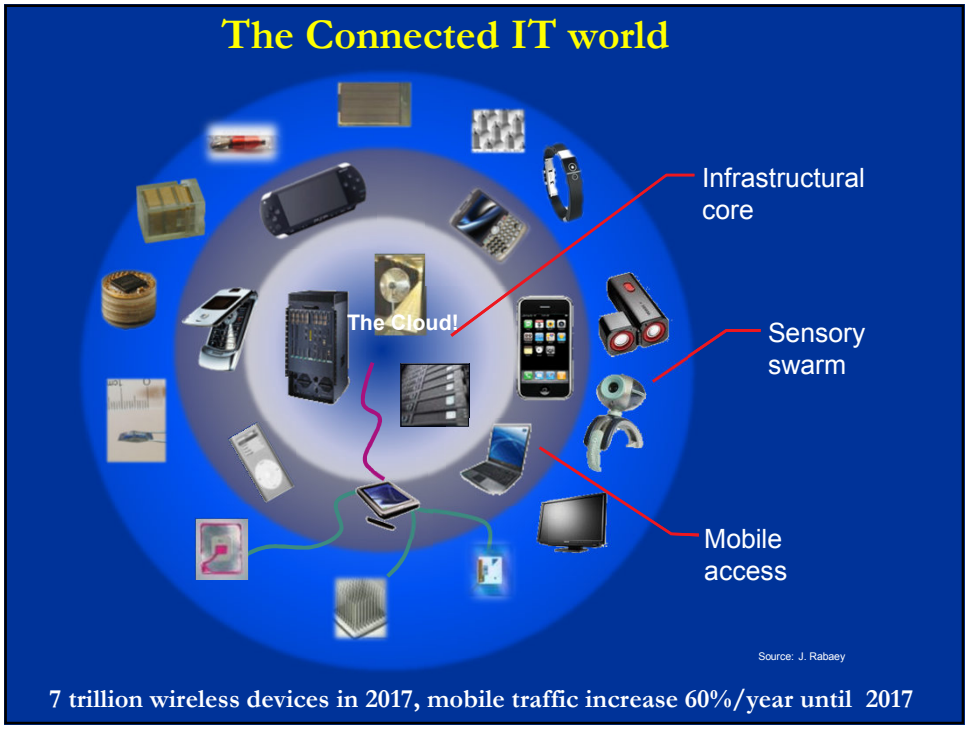
The Connected IT world
Infrastructural core
Sensory swarm
Mobile access
The Cloud!
Source: J. Rabaey
7 trillion wireless devices in 2017, mobile traffic increase 60%/year until 2017

Content

Shift from computation centric to data transfer and memory centric architectures

- Not only on large scale (IT World)
- But also on small scale (System-on-Chip)

Metrics for architecture efficiency

- Wireless baseband processing
- Impact of memories and data transfers on metrics
- Impact of application (communications) performance on metrics

3D Integration

- 3D memories and memory controllers

3.5G Digital Workload (100GOPS@1Watt)

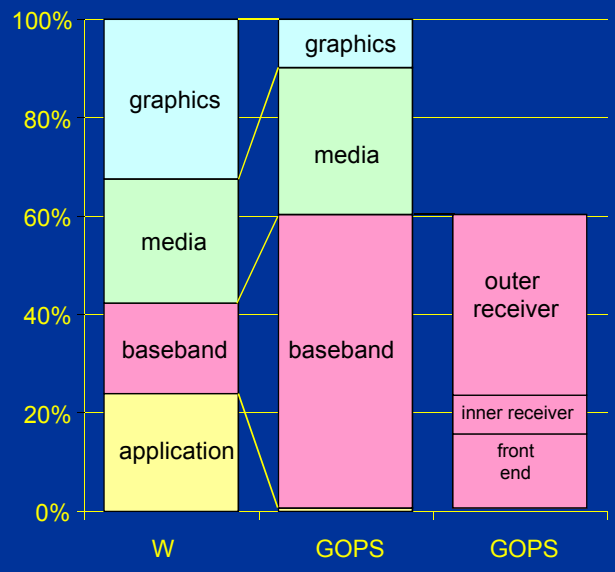

Future: 1 TOPS in 1+ Watt

Source: Kees van Berkel, MPSoC2010

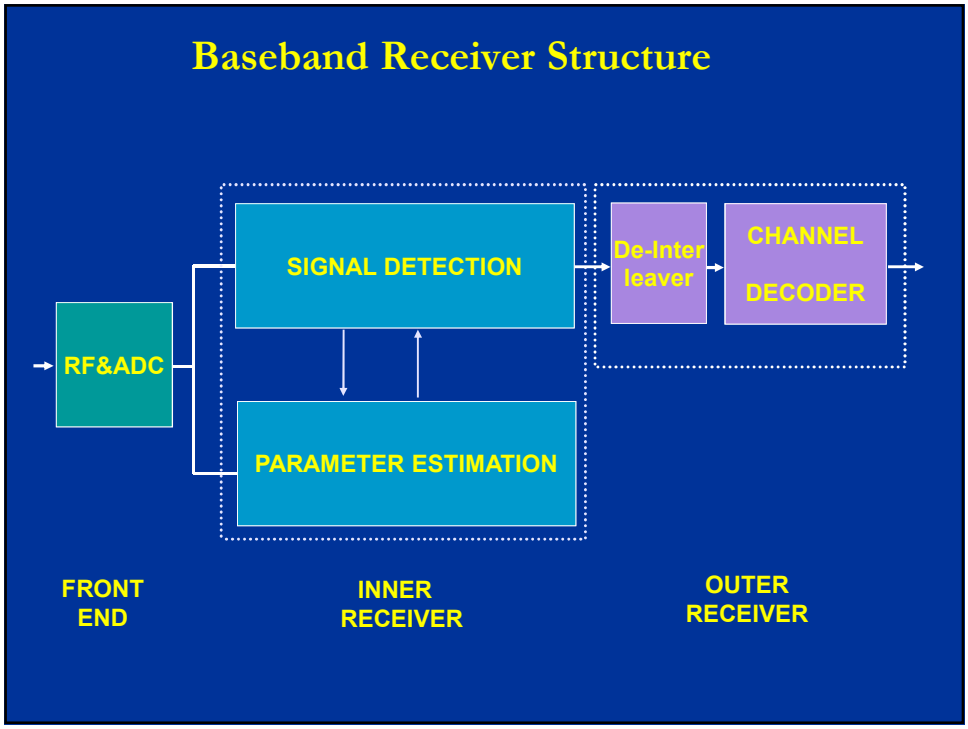
Baseband Receiver Structure
SIGNAL DETECTION
De-Inter leaver
CHANNEL DECODER
RF&ADC
PARAMETER ESTIMATION
FRONT END
INNER RECEIVER
OUTER RECEIVER

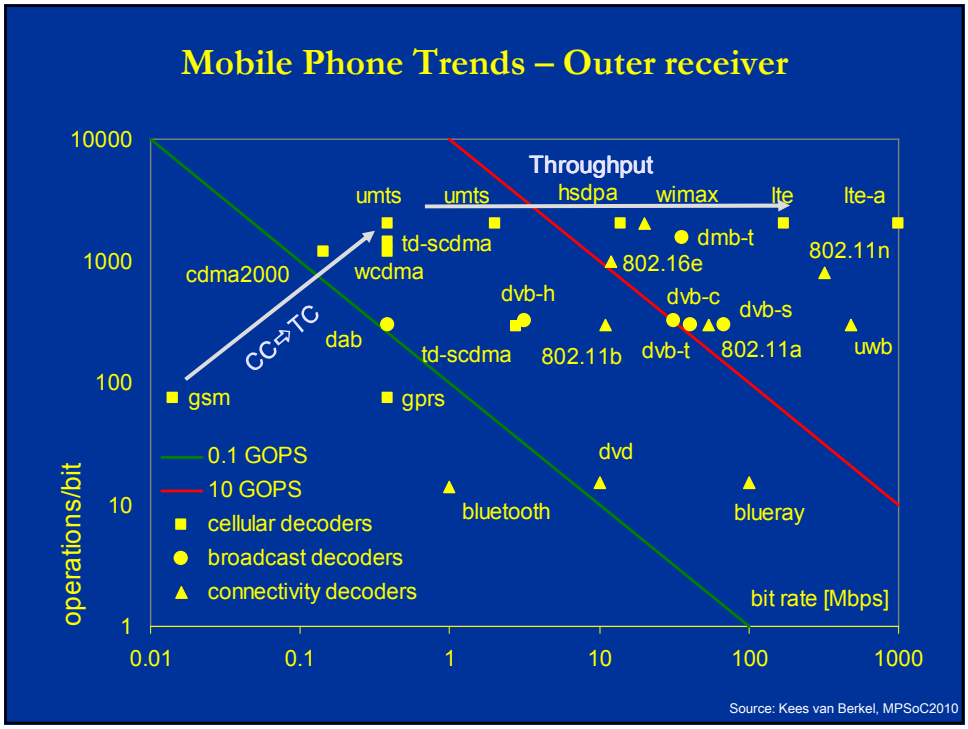
Mobile Phone Trends – Outer receiver
Throughput
umts
umts
hsdpa
wimax
lte
lte-a
td-scdma
wcdma
dmb-t
802.16e
802.11n
cdma2000
dvb-h
dvb-c
dvb-s
CC⇨TC
dab
td-scdma
802.11b
dvb-t
802.11a
uwb
gsm
gprs
0.1 GOPS
10 GOPS
cellular decoders
broadcast decoders
connectivity decoders
dvd
bluetooth
blueray
operations/bit
bit rate [Mbps]
10000
1000
100
10
1
0.01
0.1
1
10
100
1000
Source: Kees van Berkel, MPSoC2010

Recent Decoder Designs @ TU KL

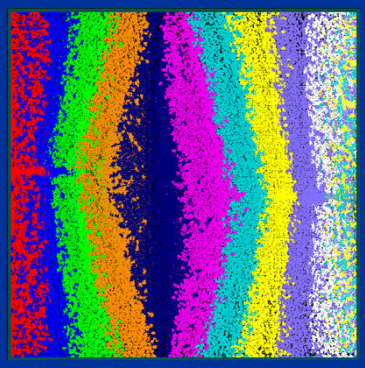

- 160.8 Gbit/s LDPC Decoder 65nm technology, 12mm^2

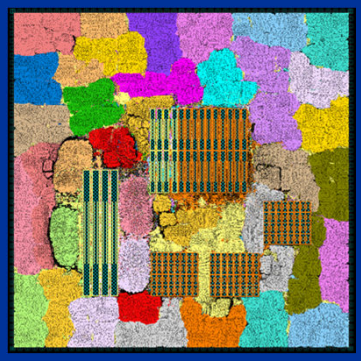

- 2.15 Gbit/s LTE TC Decoder 65nm technology, 7.7mm^2

Music Baseband SDR Chip @ 65nm

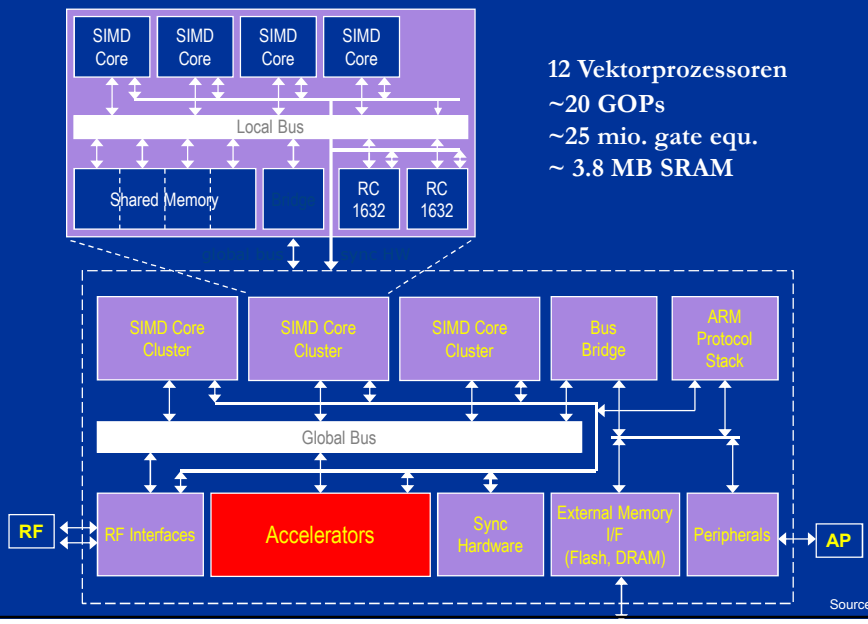

12 Vektorprozessoren
~20 GOPs
~25 mio. gate equ.
~ 3.8 MB SRAM

Source: Infineon

LETI / TU KL Magali Chip

- 477mW NoC Based Digital Baseband for MIMO 4G SDR
- 96Mtransistors, 27mm^2, 65nm technology

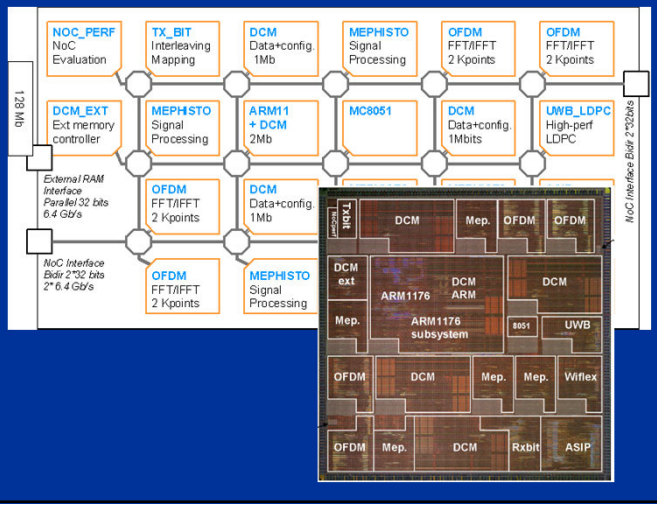

22 processing units:
5 VLIW processors
ARM11 processor
ASIP processor
HW accelerators
Distributed memory
15 asynchronous NoC router
Sophisticated power management

Source: LETI

Metric – Energy Efficiency

Example - SODA, DSP and GP Architectures

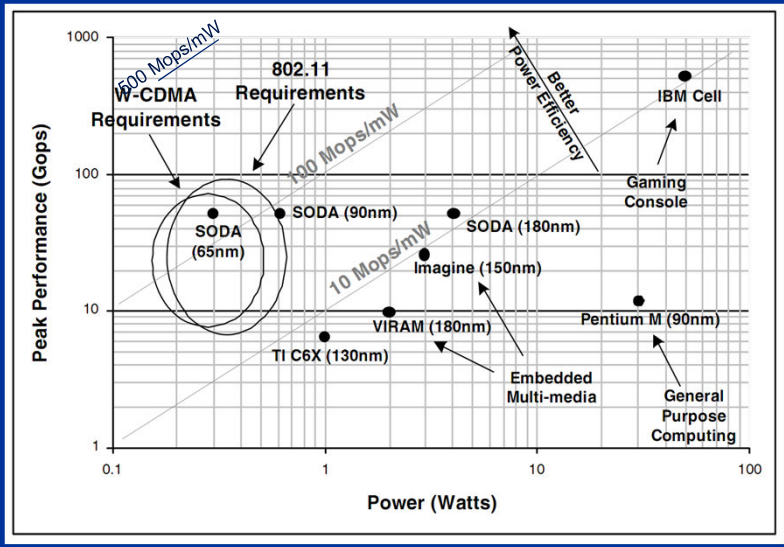

Metric Assessment - Channel Decoders

All architectures based on standard synthesis flows, 65nm technology@worst case, all data in-house available

Decoder	Flexibility	Max Block-size	Payload Throughput [Mbit/s]	Freq. [MHz]	Area [mm2]	Dynamic Power [mWatt]
ASIP (Magali)	Conv. Codes Binary TC Duo-binary TC	N=16k	40 14(6iter) 28(6iter)	385 (P&R)	0.7 (P&R)	~100
LTE Turbo (Music)	LTE turbo code	N=18k	150 (6iter)	300 (P&R)	2.1 (P&R)	~300
LDPC flex (Magali)	R=1/4 to R=9/10	N=16k	150-300 (20-10iter)	385 (P&R)	1.172 (P&R)	~389
LDPC fixed (Magali)	R=3/4	N=1.2k	480 (6iter)	435 (P&R)	0.583 (P&R)	~202
LDPC WiMedia 1.5	R=1/2-4/5	N=1.3k	640 (R=1/2,5iter) 960 (R=3/4,5iter)	265	0.51	~193
CC Decoder	64-state NSC		500	500	0.1	~37

Algorithmic Throughput Calculations [Gops]

Code	Operations per decoded information bit *normalized to ~8bit addition*		Infobit-Throughput ⇨Giga operations per second [Gops]		
			100Mbit/s	300Mbit/s	1 Gbit/s
CC: states=64	~200		**~20**	**~ 60**	**~200**
LDPC Min-Sum (x3.4 for λ-3-Min alg.)	5 iter	75/R	**~7.5/R**	**~22.5/R**	**~ 75/R**
	10 iter	150/R	**~15/R**	**~ 45/R**	**~ 150/R**
	20 iter	300/R	**~ 30/R**	**~ 90/R**	**~ 300/R**
	40 iter	600/R	**~ 60/R**	**~ 180/R**	**~ 600/R**
Turbo Max-Log	2 iter	280	**~ 28**	**~ 84**	**~ 280**
	4 iter	560	**~ 56**	**~168**	**~ 560**
	6 iter	840	**~ 84**	**~252**	**~ 840**

Area- and Energy Efficiency

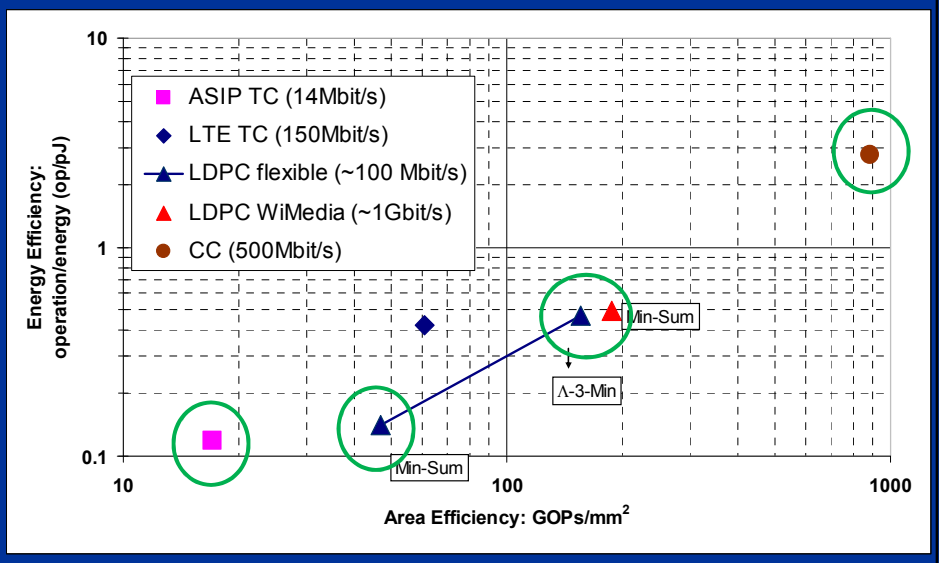

What about Memory/Data Transfers

Current metric: energy efficiency = only operations/energy

Data transfers/ accesses substantially contribute to the power consumption

Example (R=0.5)

150 Mbit/s Turbo : ~126 Gops ~40 Gaccesses

150 Mbit/s LDPC : ~90 Gops ~80 Gaccesses

Efficient data transfer is key for efficient implementation

- LTE TC: special interleaver structure to avoid access conflicts
- DVB-S2/WiMAX LDPC: special code structure to minimize access conflicts

Efficiency metrics based on operations only are not appropriate

- Power includes operations and accesses!
- Architectures are favored where operations dominate compared to accesses

Decoders in System Design Space

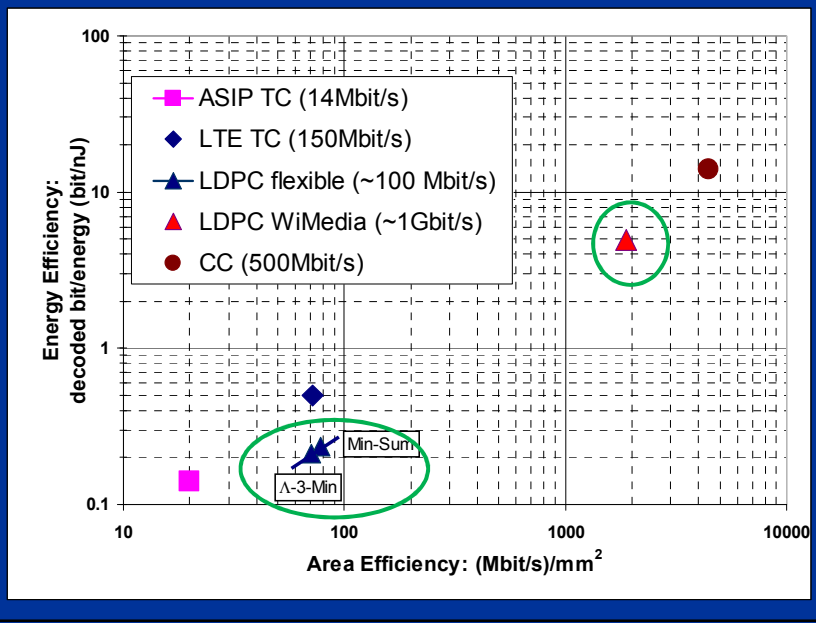

Communications Performance

Overall efficiency of a baseband receiver depends on

- Implementation performance
- Communications performance
- Flexibility

Comparison of two iterative decoders with same communications performance but different parameters (codes, code rate, iterations)

⇨ impact on implementation efficiency

Fixed Communication Performance

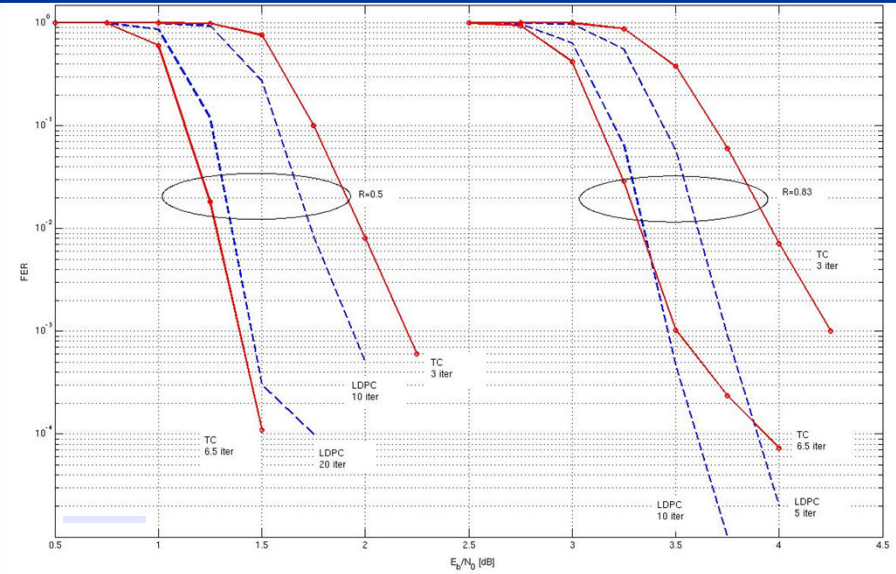

Blocksize 6145 Information bits, TC: 150Mbit/s @6.5 iterations

Implementation Efficiency

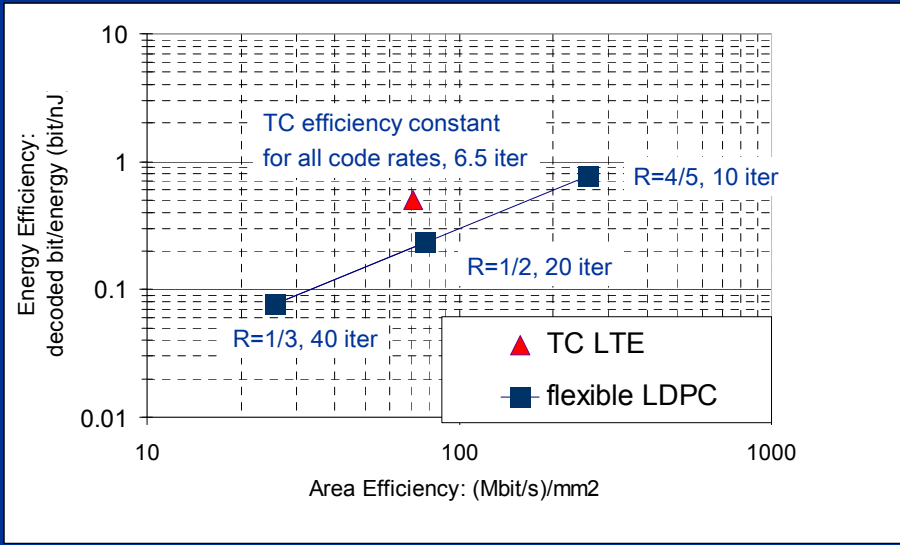

Lessons learned

- Understanding trade-offs between implementation efficiency, application performance and flexibility requirements is mandatory for efficient baseband receivers
- Operation based metrics for energy and area efficiency can be misleading
- Memory and data transfers have to be considered in metrics for design space exploration
- Implementation efficiency metrics have to be linked to application performance

Next-Generation Mobile Platform Traffic

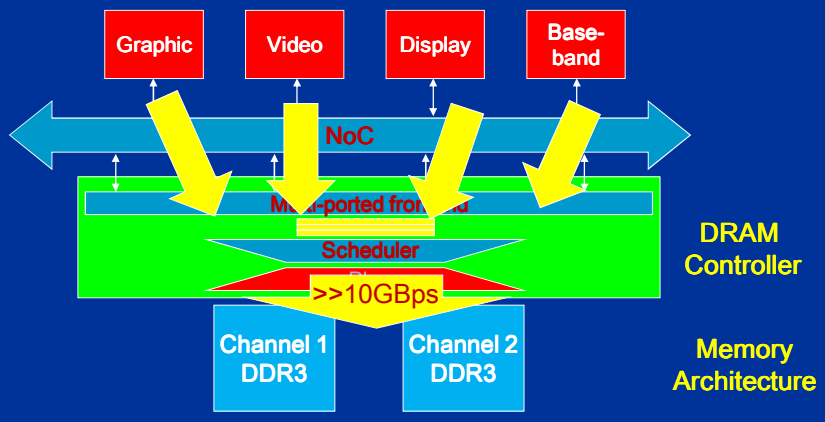

- Traditional JEDEC DRAM channels are saturating

Next Generation Teraflop Computing Platform

Year 2018
8nm Core, 10Gflop

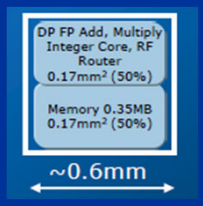

400mm² Die size
1150 Cores

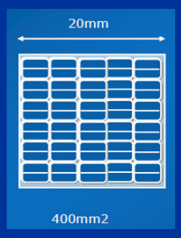

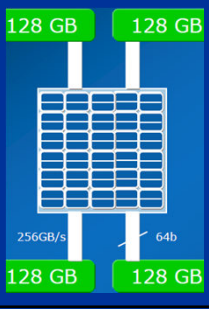

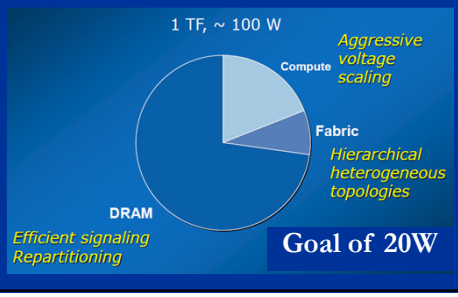

Source: Intel

Energy-Efficient Seismic Modeling

Reverse Time Migration (RTM) Method

- **Exploration area: 30km x 20km area, 10km depth**
- **Exploration ship: 10 streamer lines with 1000 receivers each, time sampling interval of 1ms and listening a total of 12 seconds each**
- **Reaching a computation time of on week**
  - >1M cores necessary with 38MW power!
- **Green Wave Computer (45nm): 2D torus NoC based, 76.000 optimized Tensilica LX2 cores**
  - 5 MW power

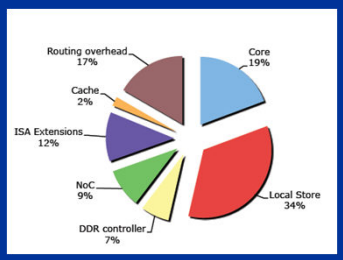

Area Breakdown

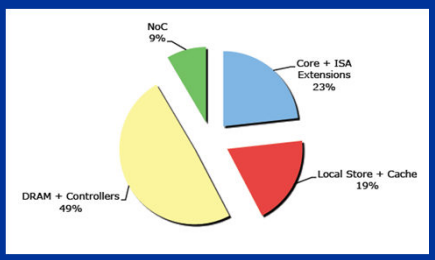

Power Breakdown

3D Integration with TSVs

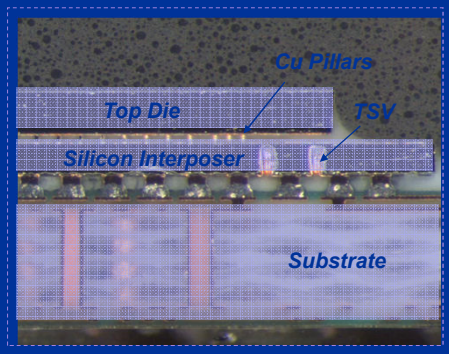

Through Silicon Vias (TSV)

- Polysilicon filled (FEOL)
  - 10.000 TSV/mm²
- Copper filled (BEOL)
  - 500 TSV/mm²

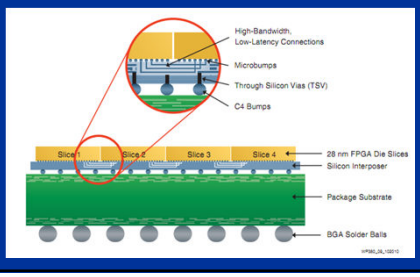

Source: XILINX

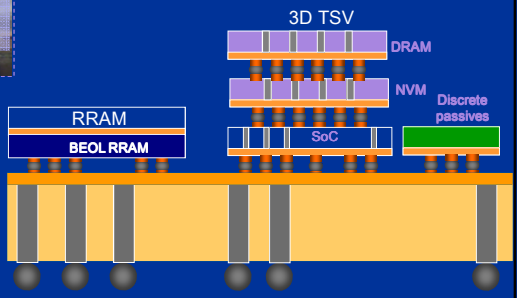

Source: LETI

Wide IO Technology (JEDEC Standard 2012)

Channel

- 4 banks with 64Mb each
- 128 bit @ 200MHz SDR
- 3.2GBps

Memory

- 4 channels
- 1Gb
- 512 bit IO
- 12.8GBps

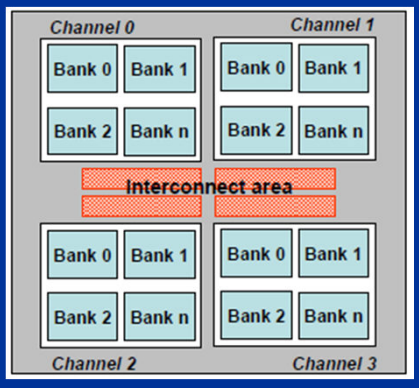

3D Magali Chip

- 65nm tech, 72mm^2, 1980TSVs for 3D NoC, 1250 TSV for wide I/O memory
- Heater, temperature sensors

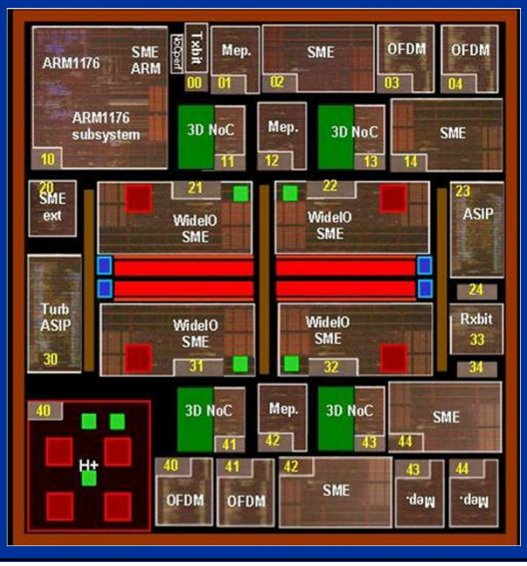

Source: LETI

Power Savings in DRAM Memory Interfaces

- Much wider I/Os possible >> 32 bits

Memory link, peak bandwidth and power consumption efficiency			Cost for 1TBps memory bandwidth	
			Number of data IO pins	Interface power consumption
Computing memory IF standard	Multi-core SoC → DDR3 → DRAM 1066 MHz I/O bus clock, 32 bits, 1.5 V, Double Data Rate	8.5 GBps 30 mW/Gbps	3800	240 W
Mobile memory IF standards	Multi-core SoC → LPDDR2 → DRAM 533 MHz I/O bus clock, 32 bits, 1.2 V, Double Data Rate	4.3 GBps 20 mW/Gbps	7700	160 W
	Multi-core SoC → Wide I/O → DRAM 200 MHz I/O bus clock, 512 bits, 1.2 V, Single Data Rate	12.8 GBps 4 mW/Gbps	41000	32 W

Source: LETI

3D-stacked DRAMs

DRAM macro cut out of a 8 bank 1Gb DDR2 SDRAM chip:

Slice it & Pack it

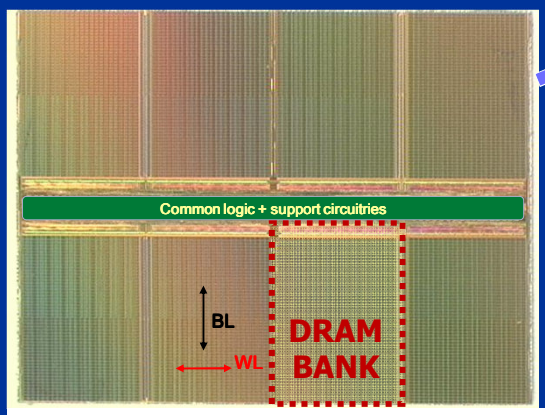

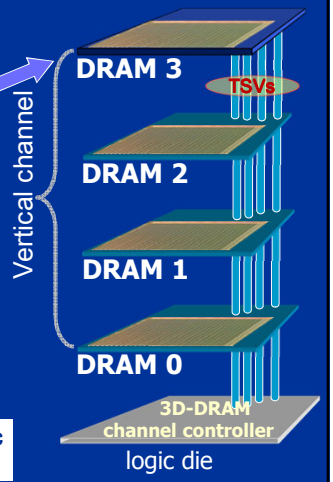

Common logic can either be placed on a separate **logic** die or on each **DRAM** layer

27

Organization and Naming

A single 3D-layer consists of 3D-DRAM banks

A bank is composed of DRAM core tiles

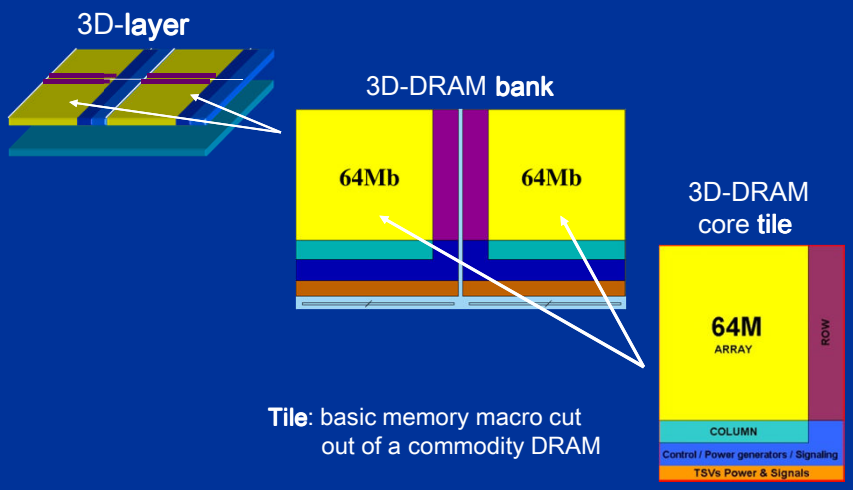

Tile: basic memory macro cut out of a commodity DRAM

Single 3D-layer Design Space

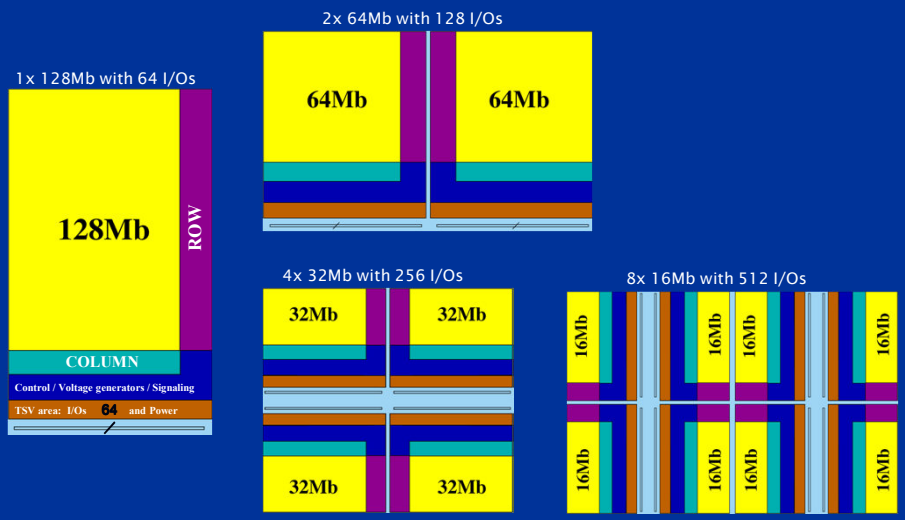

Comparison 2Gb (8 layers/8 banks)

Comparison to Micron's Hybrid Memory Cube (HMC)

- Our 3D-DRAM is more than competitive

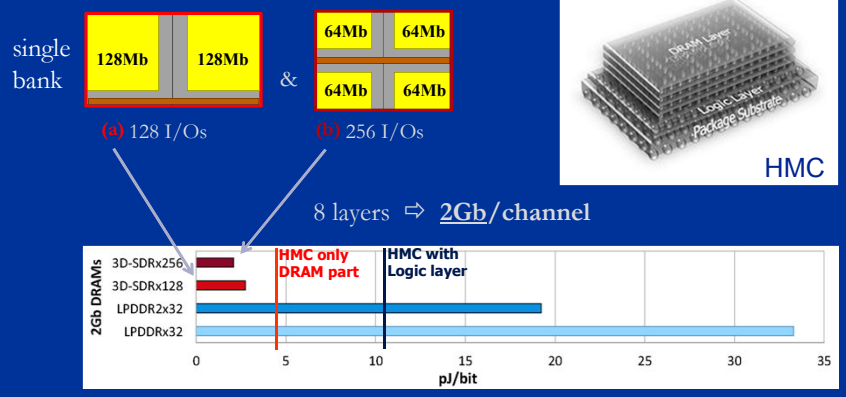

Multi-Channel 3D-DRAM Controller

Different request granularities

- But normally fixed to 128 bit (Wide IO JEDEC Standard)

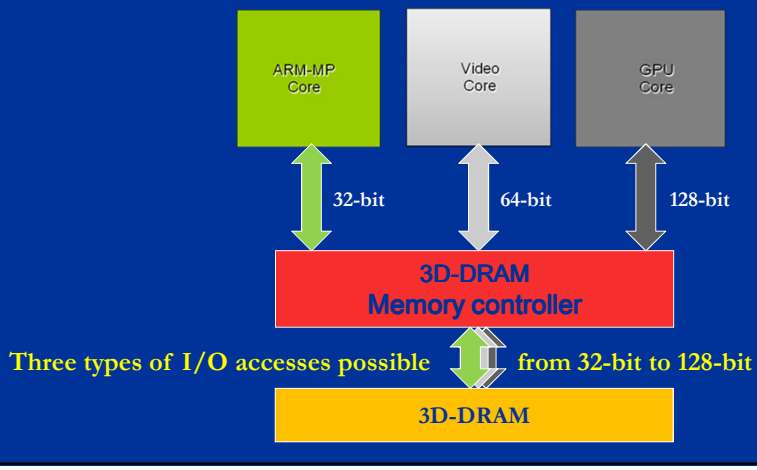

Fine-grained 3D-DRAM Access

On the fly switching of the data width: 32, 64, 128

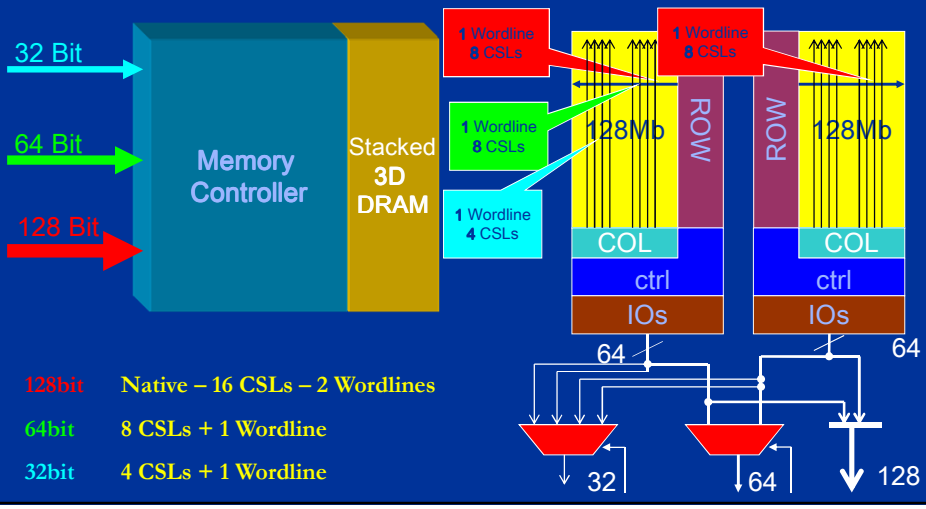

128bit Native – 16 CSLs – 2 Wordlines

64bit 8 CSLs + 1 Wordline

32bit 4 CSLs + 1 Wordline

Flexible 3D-DRAM Access

Burst Length (BL) on the fly switching

- BL can be switched on the fly (no Mode Register Set)

Combining BL and data width on the fly switching

- Additional signals needed between 3D Channel Controller and 3D-DRAM

Single 3D-DRAM access size range: 8..128 Bytes

8 Bytes (BL=2 * 32 I/Os) ➔ **128 Bytes** (BL=8 * 128 I/Os)

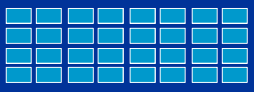

Results

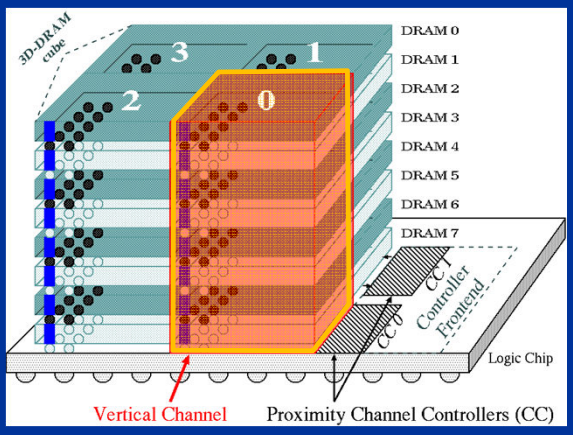

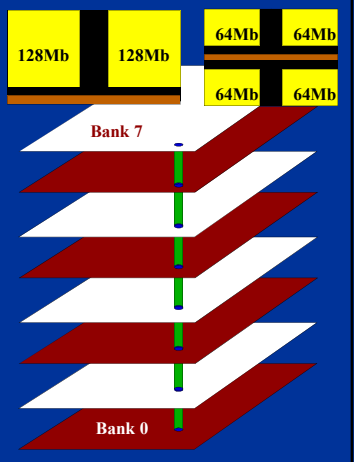

Investigated different configurations

- #Layers
- Organisation of layers
- Technologies

Investigated 3D-DRAM Configurations

3D-DRAM SINGLE CHANNEL CONFIGURATIONS

Dens. [Mb]	Architecture # lay. x [org.]	# of banks	Techn. [nm]	Cell size	A_{total} [mm^2]	Freq. [MHz]
		SDR x128				
**256	1 x [4x64Mb]	4	58	$6F^2$	16	200
512	2 x [4x64Mb]	4	58	$6F^2$	26	200
1024	8 x [2x64Mb]	8	46	$6F^2$	35	300
*2048	8 x [2x128Mb]	8	46	$6F^2$	60	167
4096	8 x [4x128Mb]	8	45	$4F^2$	97	200
		DDR x128				
256	1 x [4x64Mb]	4	58	$6F^2$	22	200
512	2 x [4x64Mb]	4	58	$6F^2$	32	200
1024	8 x [2x64Mb]	8	46	$6F^2$	44	300
*2048	8 x [4x64Mb]	8	46	$6F^2$	69	300
4096	8 x [4x128Mb]	8	45	$4F^2$	98	200

** *Density emulates the published Samsung 1Gb WIDE IO chip [15].*

Simulation Set-Up

Emulation of workload via 3 traffic generators

- Traffic A – Cache misses of a 2 core ARM ~ 100 MB/s per core
- Traffic B – DMA accesses in a SoC (Imaging) ~ 0.8 GB/s
- Traffic C – HD Video DMA accesses ~ 1.5 GB/s

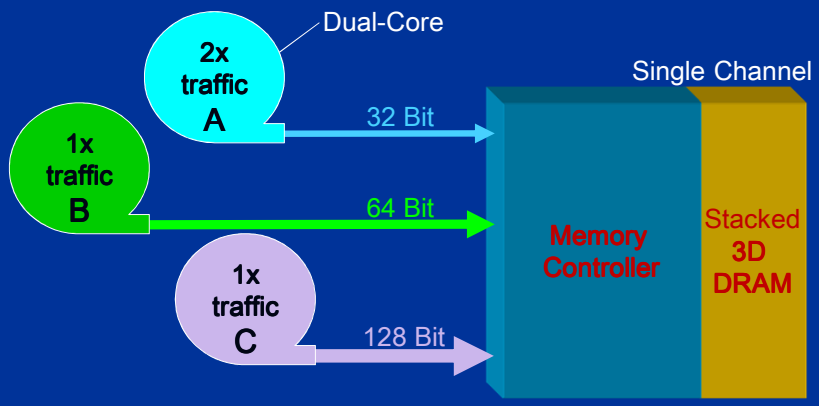

Results with BL and Data Width Switching

Factor 3x energy improvement of 2Gbit 3D compared to 2Gbit LPDDR2 (PageHitRate=50%)

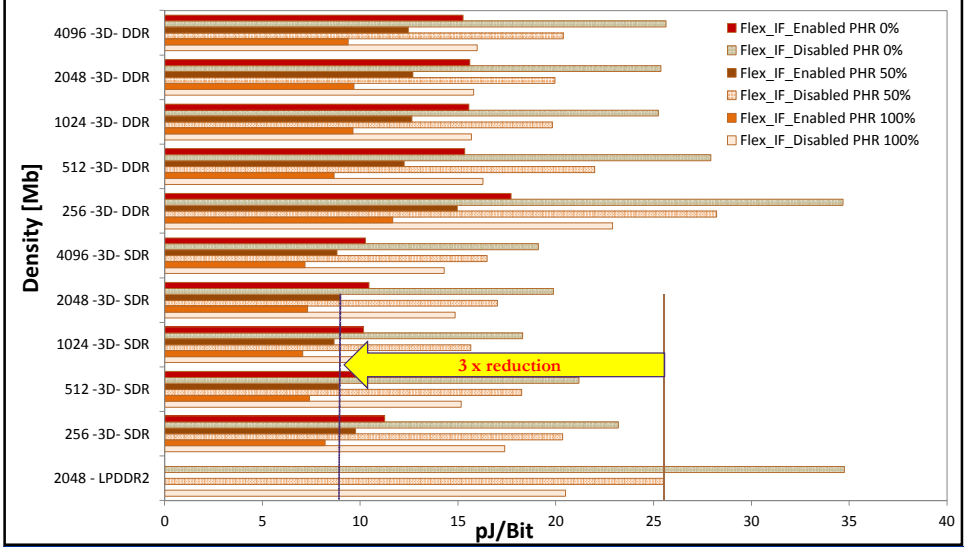

Outlook: Multichannel DRAM Controller

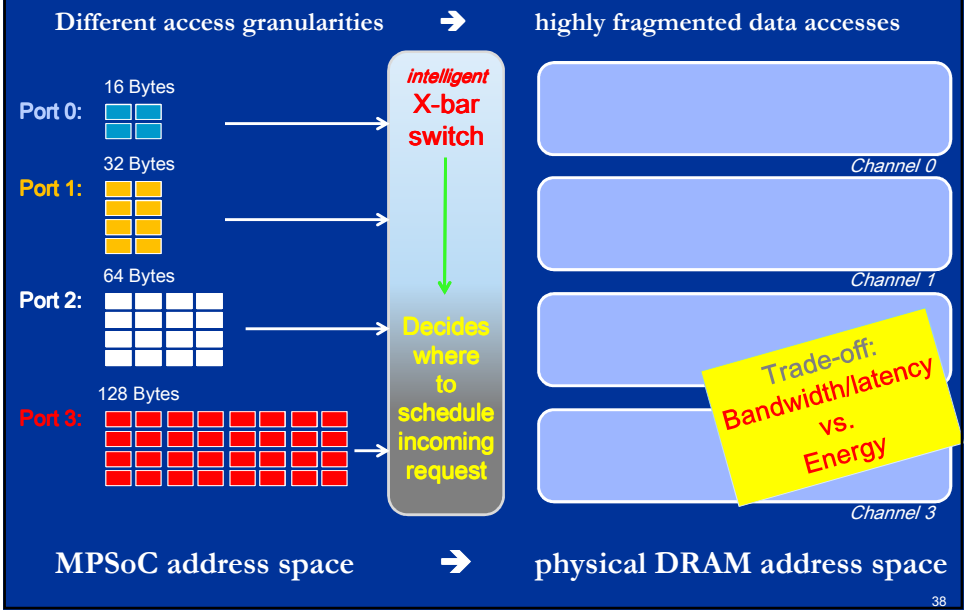

Thank you for attention!

For more information please visit

http://ems.eit.uni-kl.de

On Complexity, Energy- and Implementation Efficiency of Channel Decoders

F. Kienle, N. Wehn, H. Meyr. *IEEE Transactions on Communications*, Vol. 59, Nr. 12 , pages 3301 – 3310, December 2011, New York

Exploration and Optimization of 3-D Integrated DRAM Subsystems

C. Weis, I. Loi, L. Benini, N. Wehn. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems,* Vol. 32, Issue 4, Pages 597 - 610, April, 2013, New York