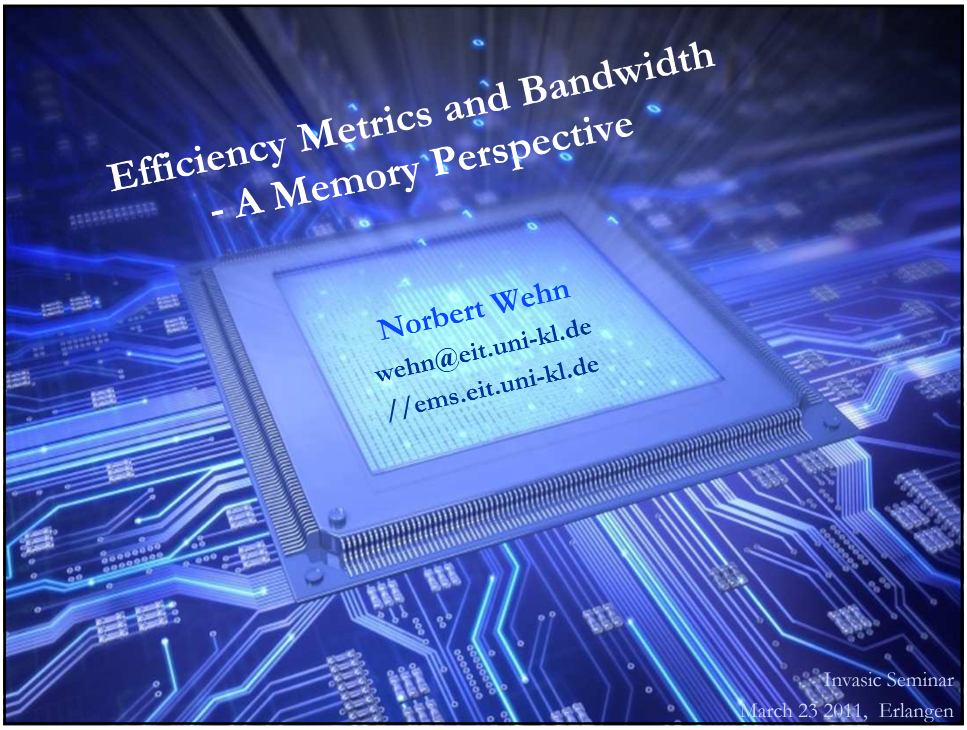

Efficiency Metrics and Bandwidth - A Memory Perspective

Norbert Wehn
wehn@eit.uni-kl.de
//ems.eit.uni-kl.de

Invasic Seminar
March 23 2011, Erlangen

Content

Computing increase and power challenge in (embedded) computing
- Heterogeneous multi-core architectures with dedicated accelerators
- New paradigm e.g. invasive computing

New Challenges
- Memory and bandwidth

Metrics for design space exploration
- Wireless baseband processing
- Impact of memories and data transfers on metrics
- Impact of application (communications) performance on metrics

3D MPSoCs
- 3D memories and memory controllers

Communication Centric World

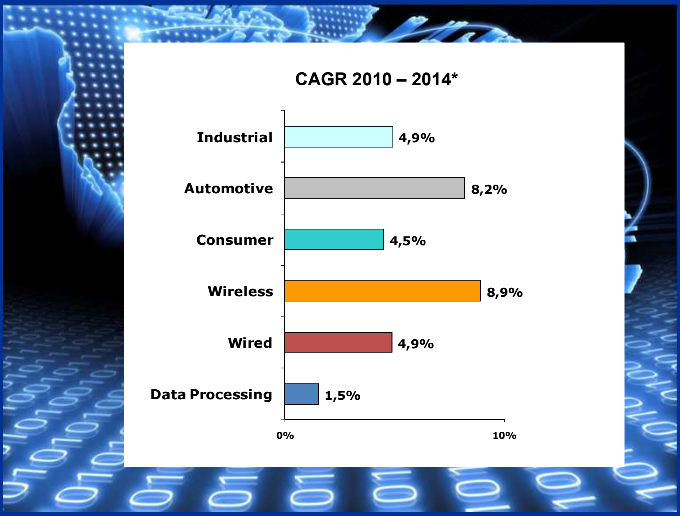

- New cellular mode is added every 3 years
- A new frequency band is added every year
- Continuous demand for higher data rates and more services

Baseband Receiver Structure

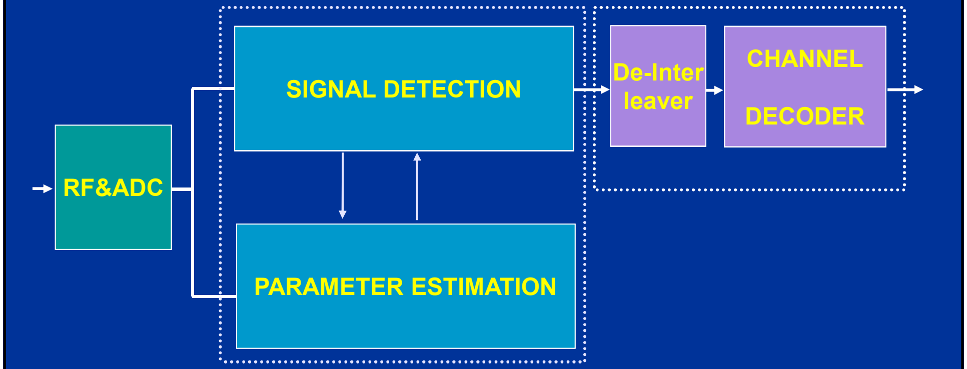

FRONT END

INNER RECEIVER

OUTER RECEIVER

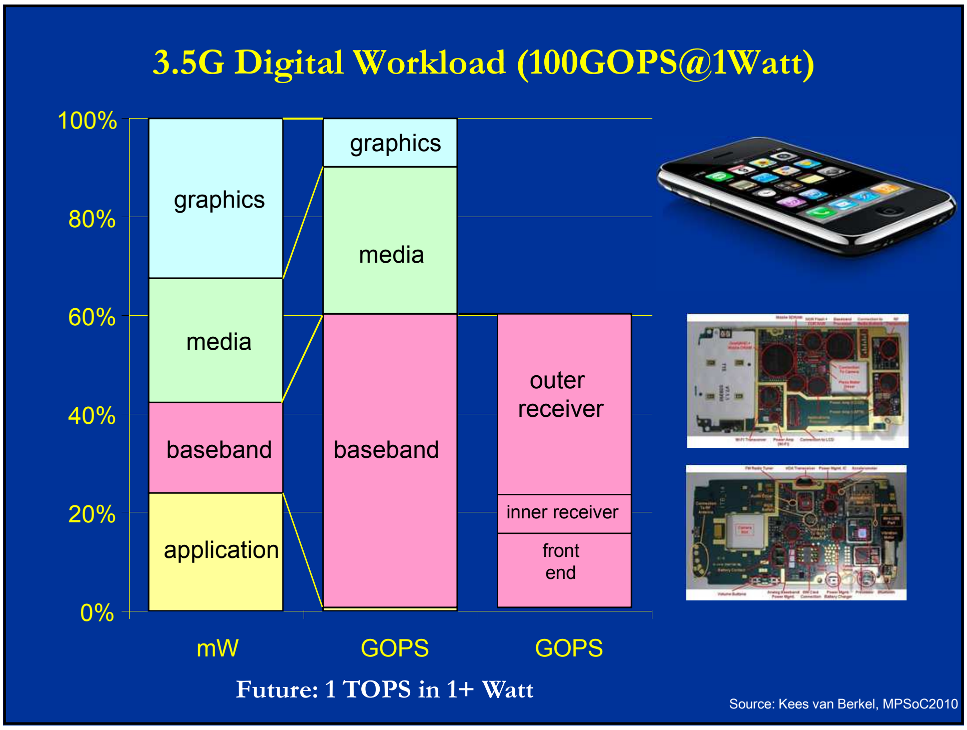

3.5G Digital Workload (100GOPS@1Watt)

Future: 1 TOPS in 1+ Watt

Source: Kees van Berkel, MPSoC2010

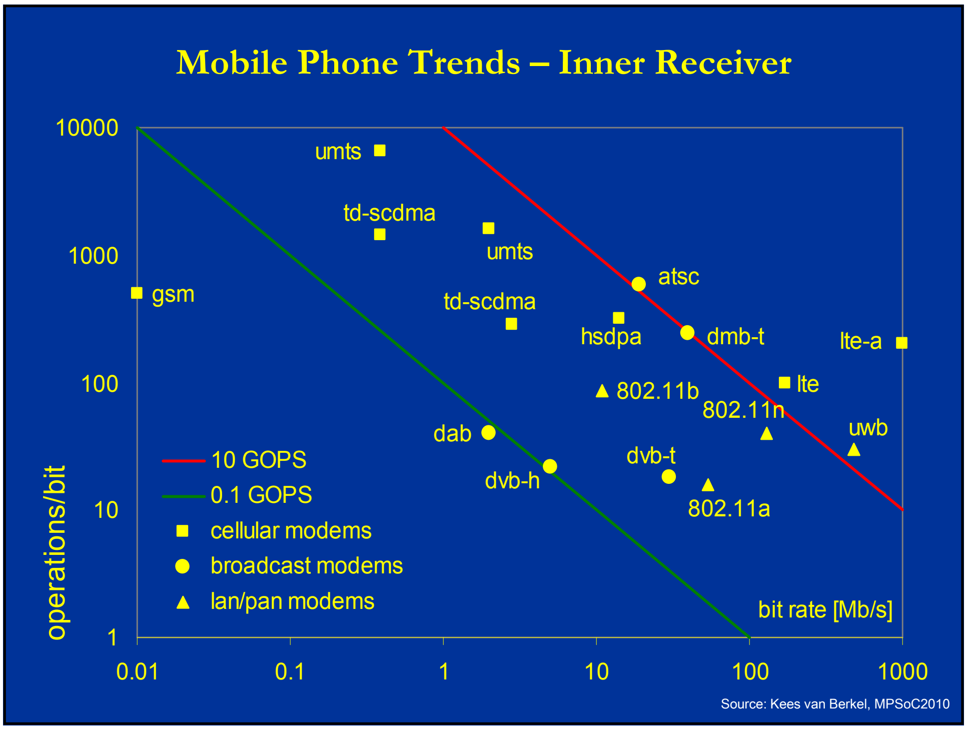

Mobile Phone Trends – Inner Receiver

Source: Kees van Berkel, MPSoC2010

Mobile Phone Trends – Outer receiver

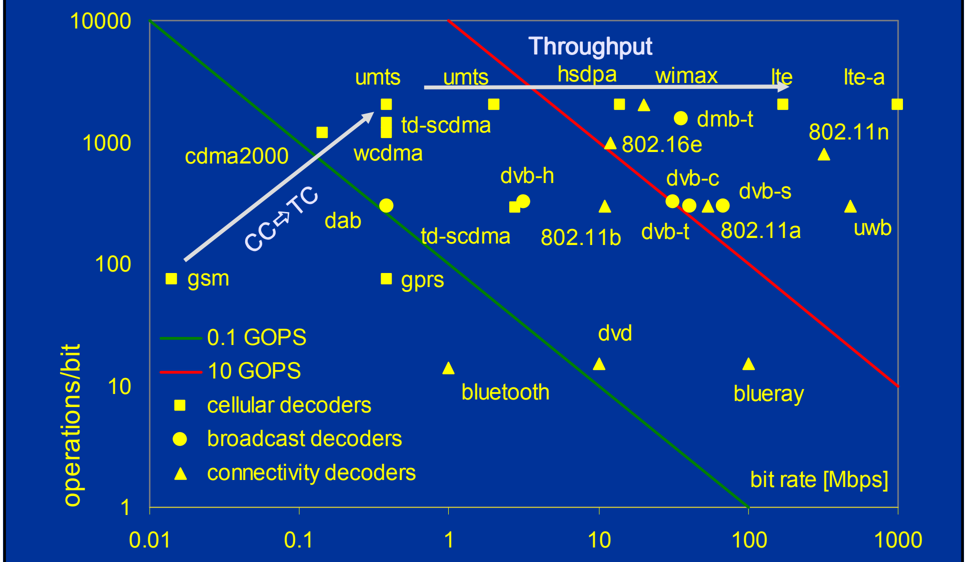

3G ⇨ 3.9G (LTE): 130 x decoding Mops/mW

Source: Kees van Berkel, MPSoC2010

Music Baseband SDR Chip @ 65nm

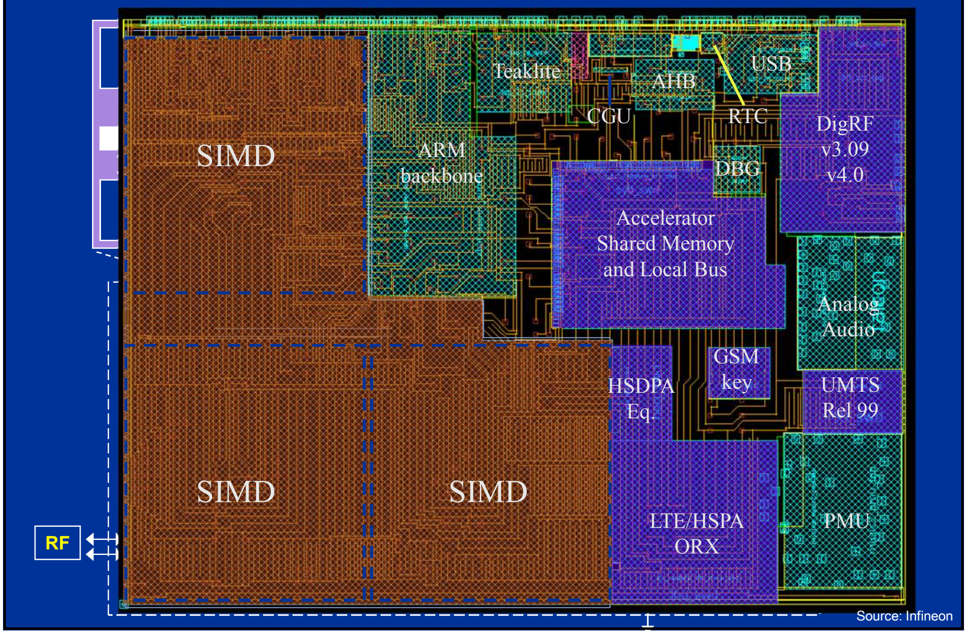

Source: Infineon

LETI / TU KL Magali Chip

- 47 4G SDR
- 96

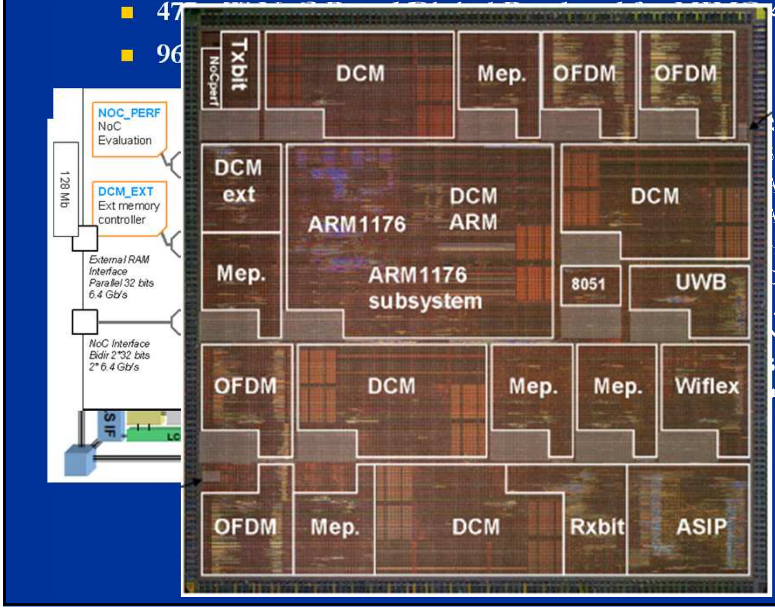

2 processing units:

VLIW processors

RM11 processor

SIP processor

IW accelerators

Distributed memory

5 asynchronous NoC router

ophisticated power nanagement

Source: LETI

Metric – Energy Efficiency

Example - SODA, DSP and GP Architectures

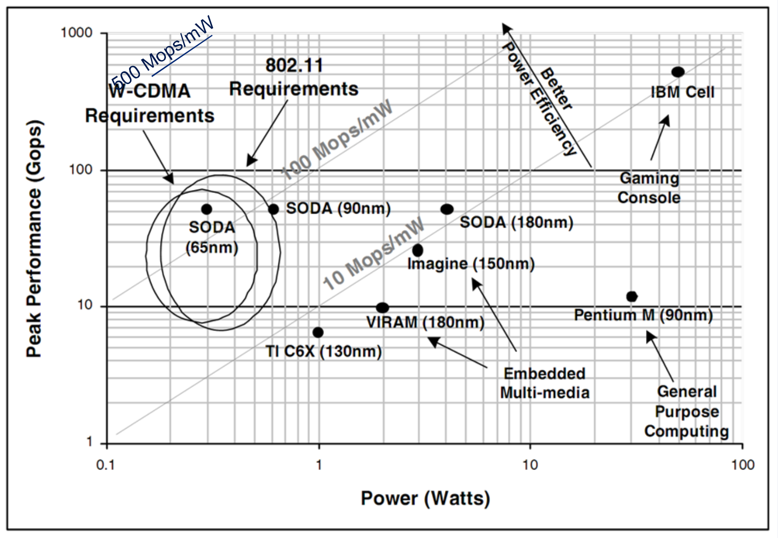

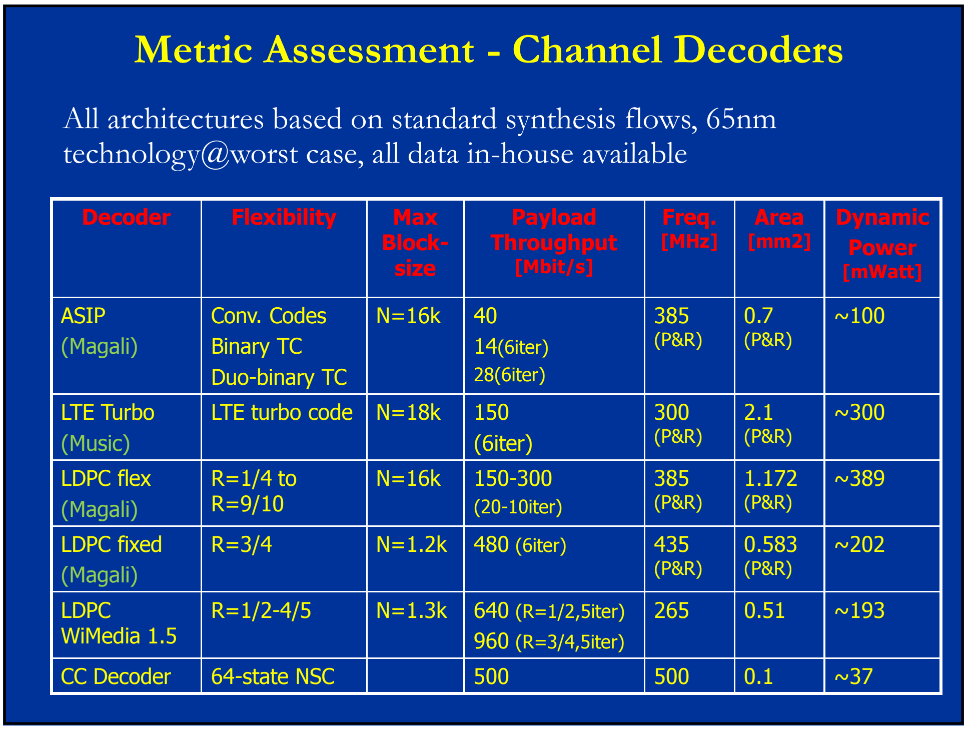

Metric Assessment - Channel Decoders

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

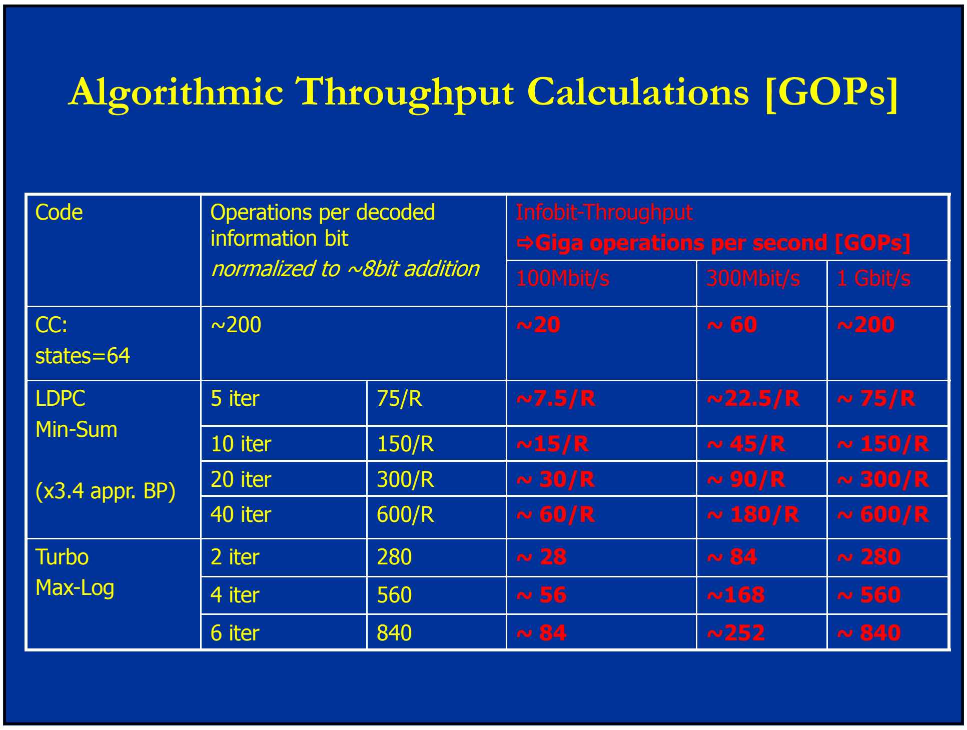

Algorithmic Throughput Calculations [GOPs]

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 appr. BP)	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

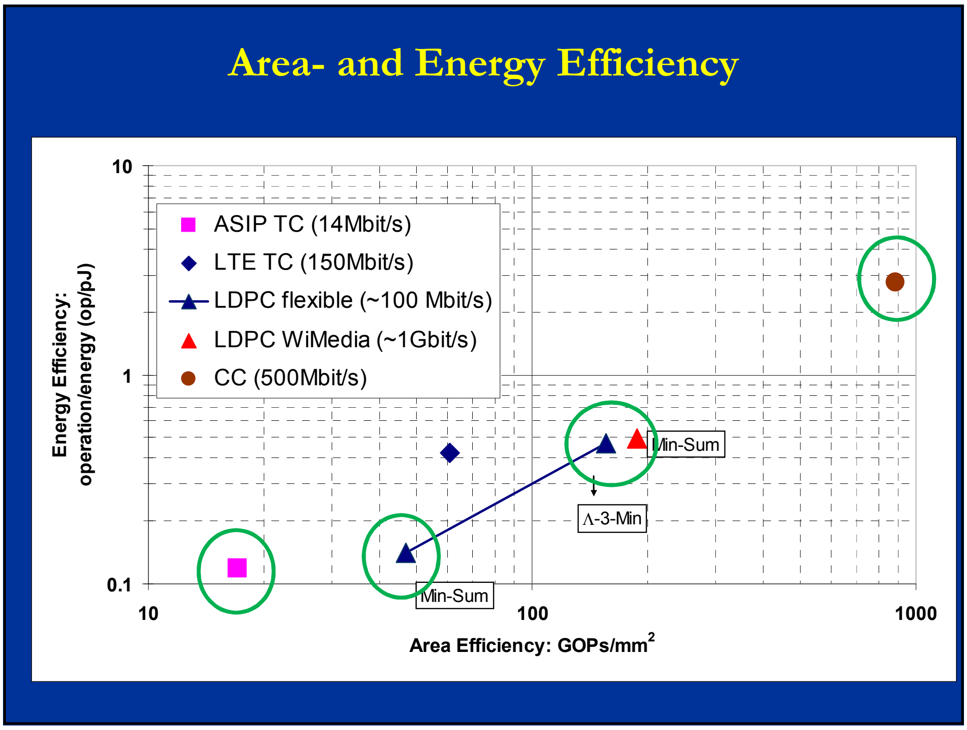

Area- and Energy Efficiency

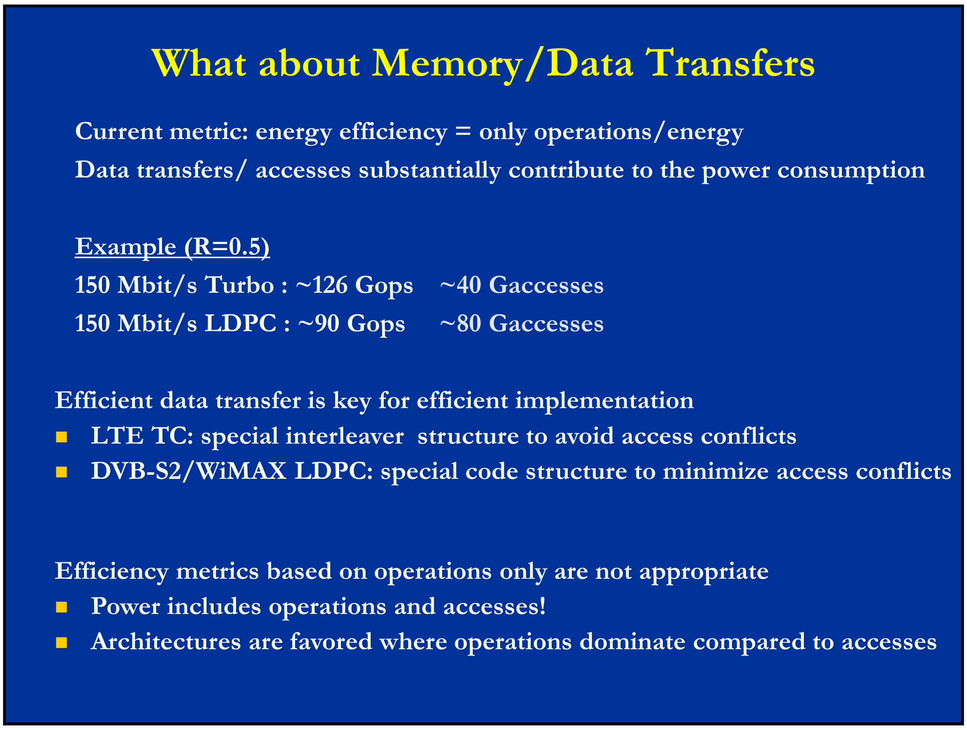

What about Memory/Data Transfers

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

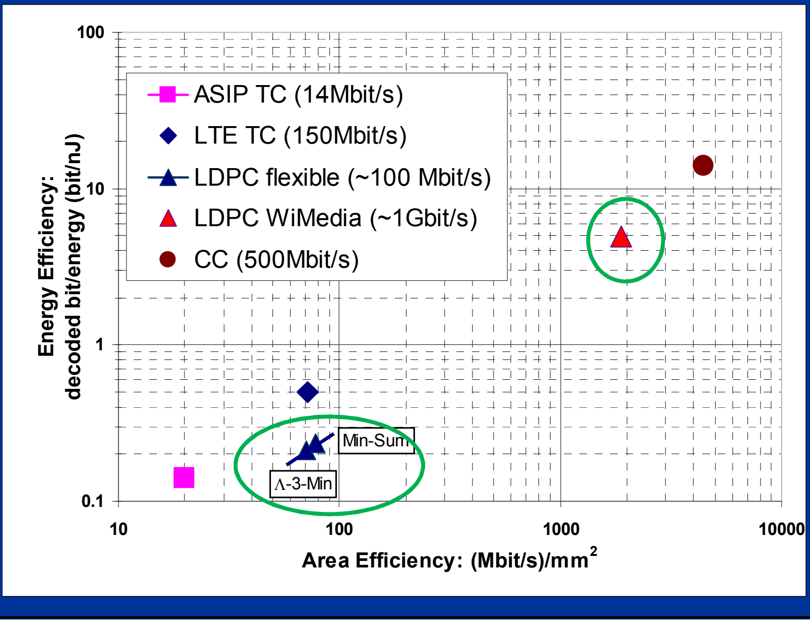

Communications Performance

Overall efficiency of a baseband receiver depends on

- Implementation performance
- Communications performance
- Flexibility

Scenario 1: Fixed Communication performance

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

⇨ impact on implementation efficiency

Scenario 2: Implementation driven

- Comparison of iterative and non-iterative decoders with varying communications performance
- 64-state convolutional code 960 Mbit/s (WiMedia 1.2) and WiMedia 1.5 LDPC decoder

⇨ impact on implementations efficiency

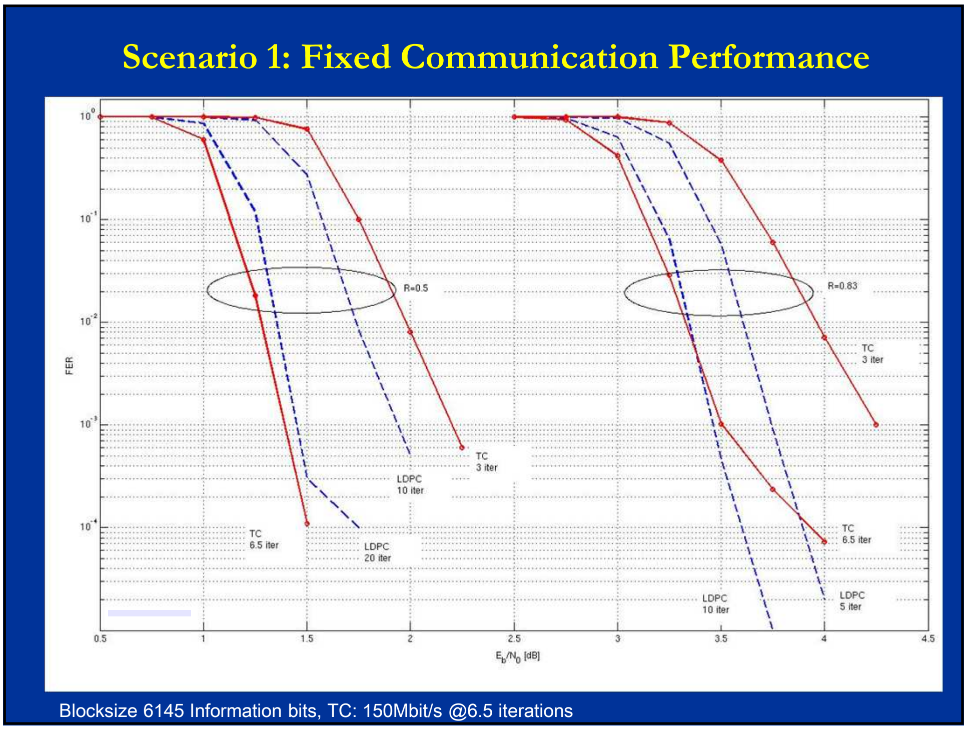

Scenario 1: Fixed Communication Performance

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

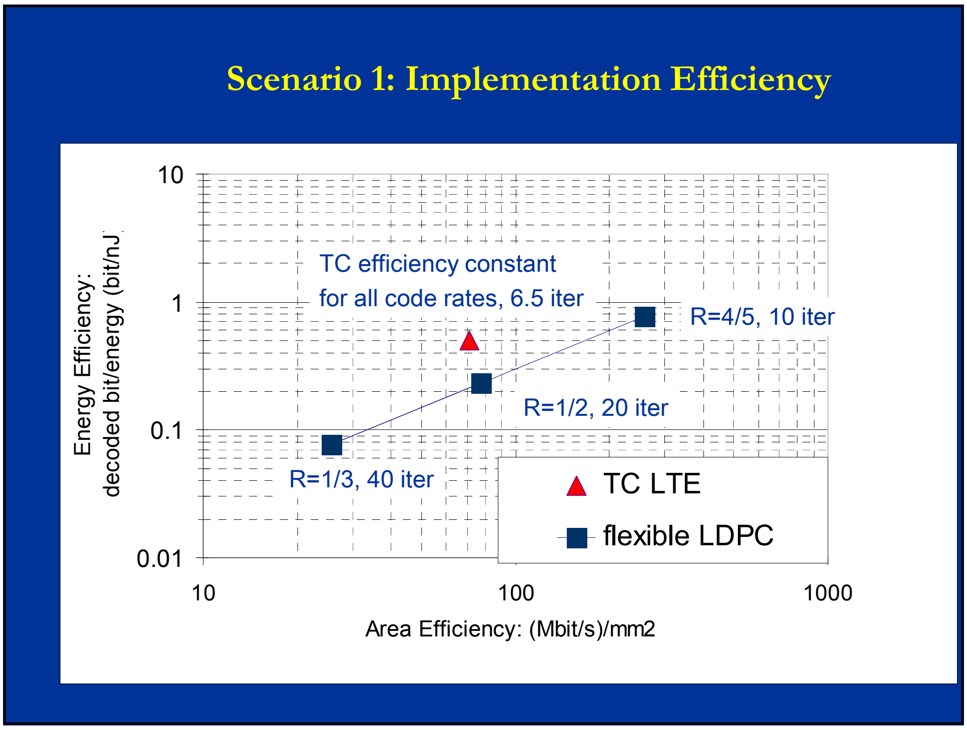

Scenario 1: Implementation Efficiency

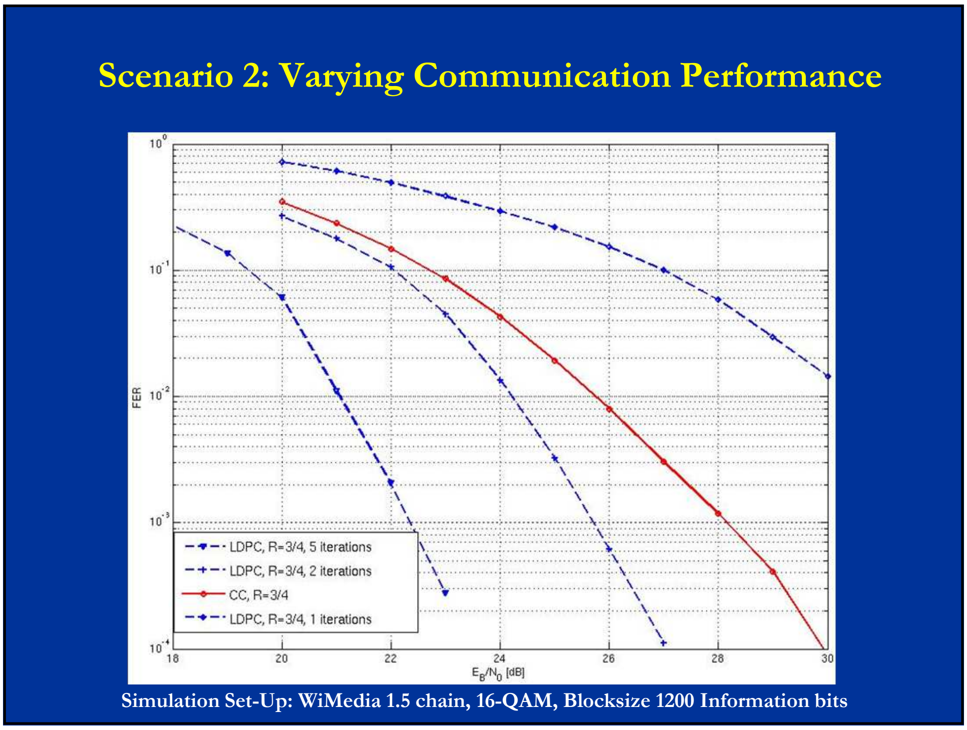

Scenario 2: Varying Communication Performance

Simulation Set-Up: WiMedia 1.5 chain, 16-QAM, Blocksize 1200 Information bits

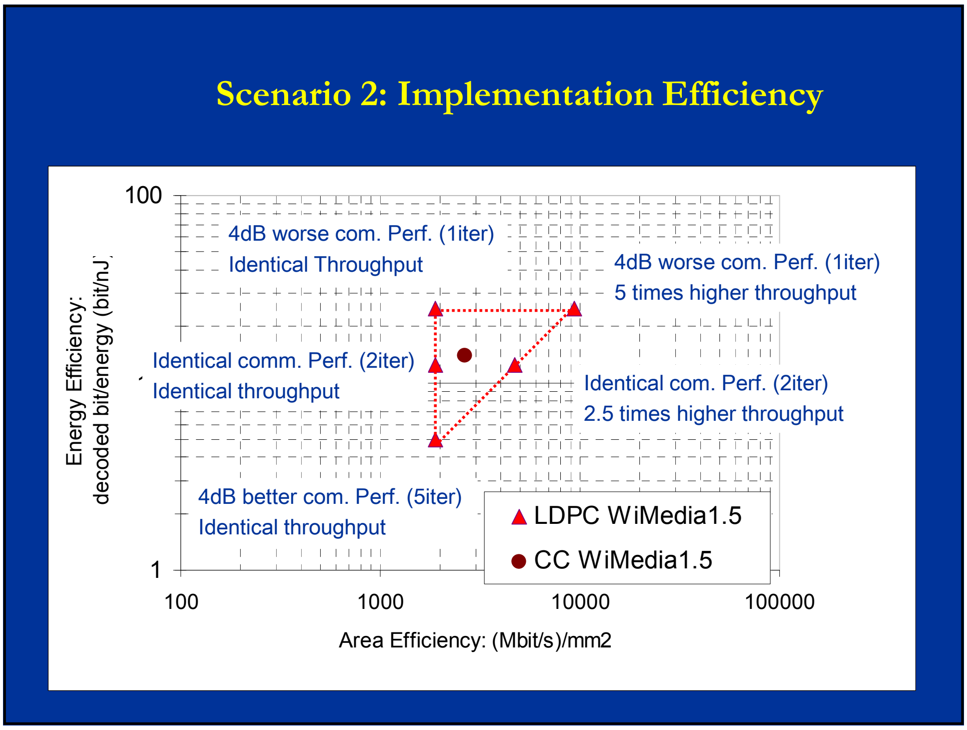

Scenario 2: Implementation Efficiency

4dB worse com. Perf. (1iter)
Identical Throughput

4dB worse com. Perf. (1iter)
5 times higher throughput

Identical comm. Perf. (2iter)
Identical throughput

Identical com. Perf. (2iter)
2.5 times higher throughput

4dB better com. Perf. (5iter)
Identical throughput

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 ⇨ trajectory

Off-chip Memory Bandwidth

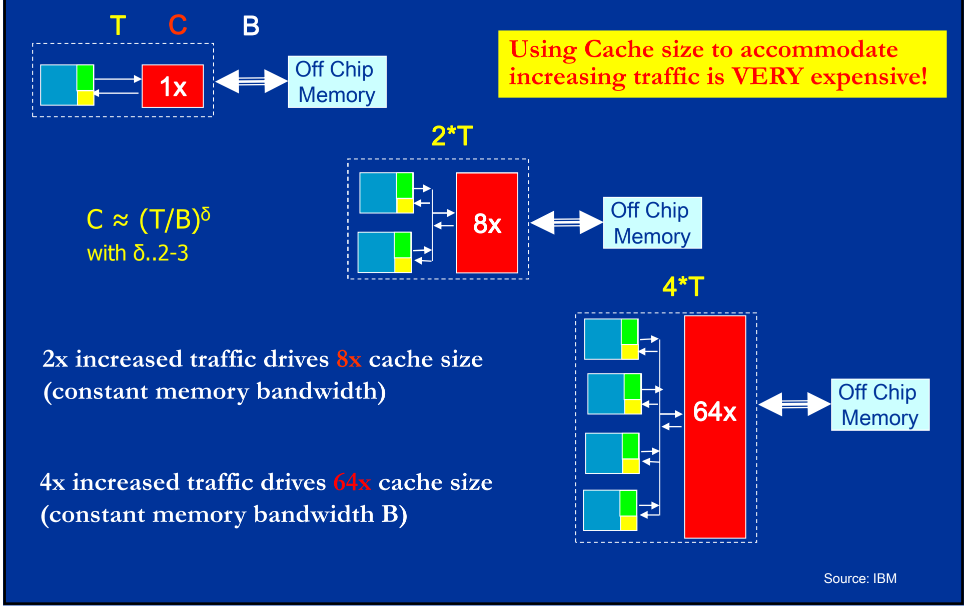

Using Cache size to accommodate increasing traffic is VERY expensive!

$C \approx (T/B)^{\delta}$
with δ..2-3

2x increased traffic drives 8x cache size (constant memory bandwidth)

4x increased traffic drives 64x cache size (constant memory bandwidth B)

Source: IBM

Next-Generation Mobile Platform Traffic

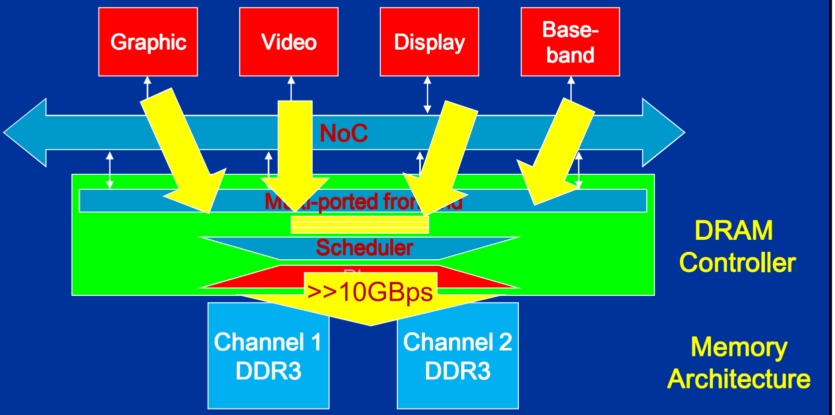

- Traditional JEDEC DRAM channels are saturating

Next Generation Teraflop Computing Platform

Year 2018
8nm Core, 10Gflop

$400mm^2$ Die size
1150 Cores

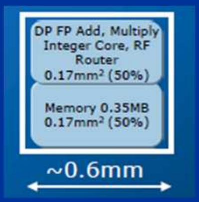

20mm
400mm2

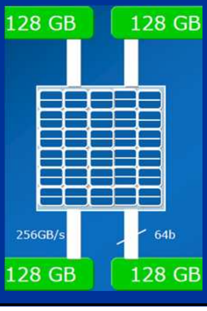

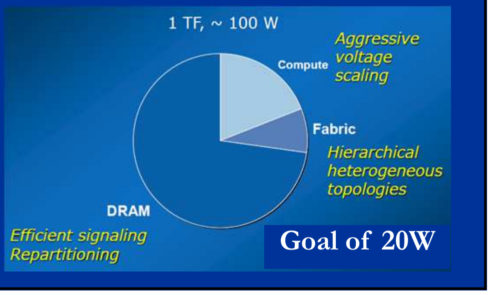

Goal of 20W

Source: Intel

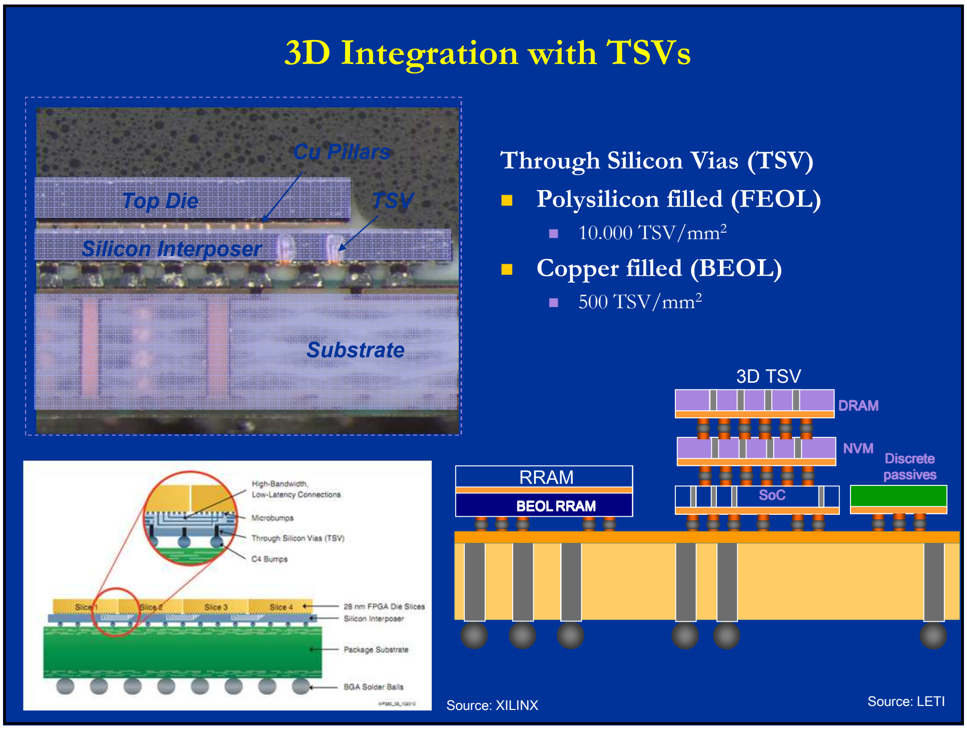

3D Integration with TSVs

Through Silicon Vias (TSV)

- Polysilicon filled (FEOL)
  - 10.000 TSV/mm^2
- Copper filled (BEOL)
  - 500 TSV/mm^2

Source: XILINX

Source: LETI

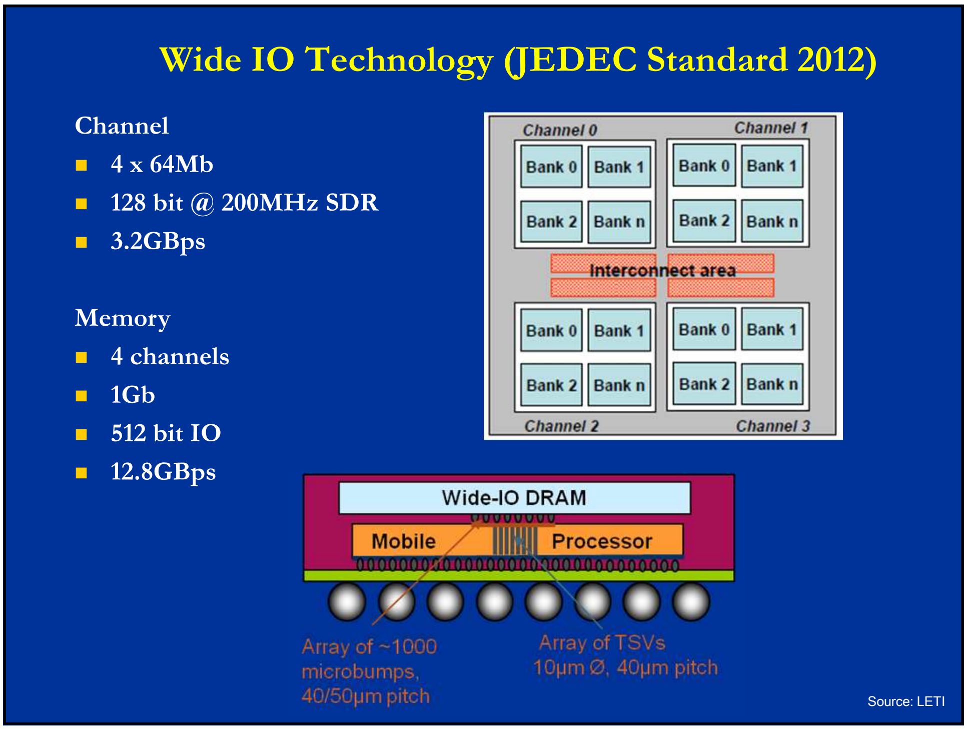

Wide IO Technology (JEDEC Standard 2012)

Channel

- 4 x 64Mb
- 128 bit @ 200MHz SDR
- 3.2GBps

Memory

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

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	8.532 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	4.264 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

Transition to 3D-DRAMs

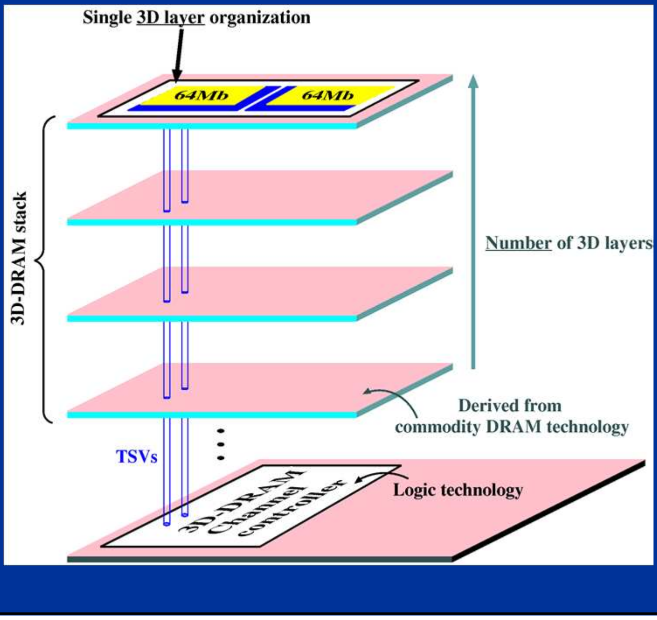
Single 3D layer organization
64Mb
64Mb
3D-DRAM stack
Number of 3D layers
Derived from commodity DRAM technology
TSVs
3D-DRAM Channel controller
Logic technology

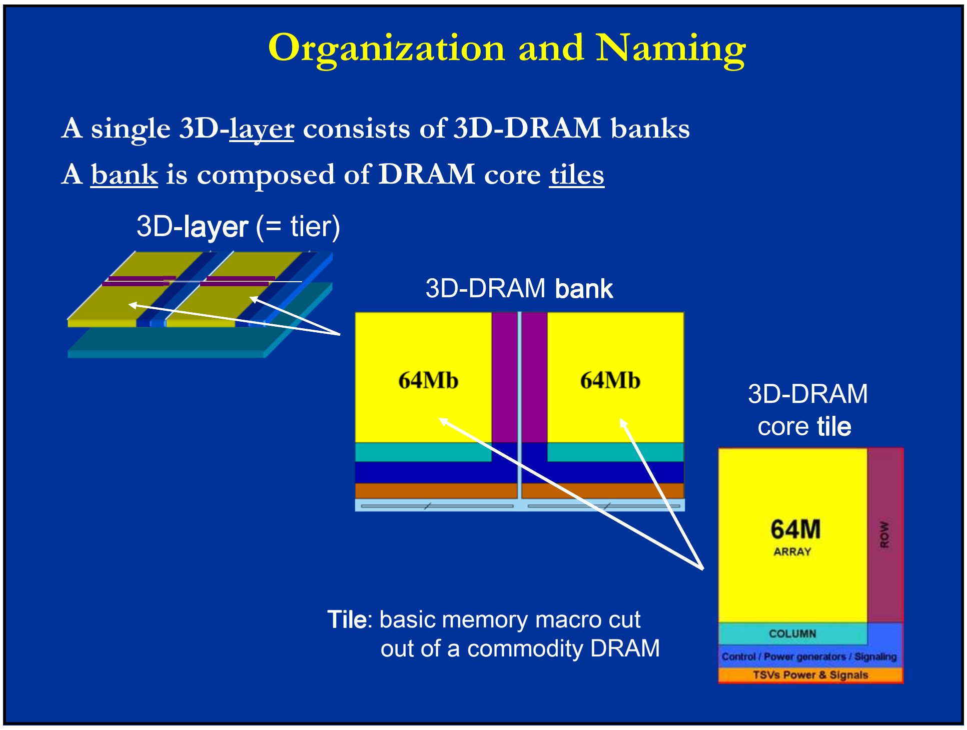

Organization and Naming

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

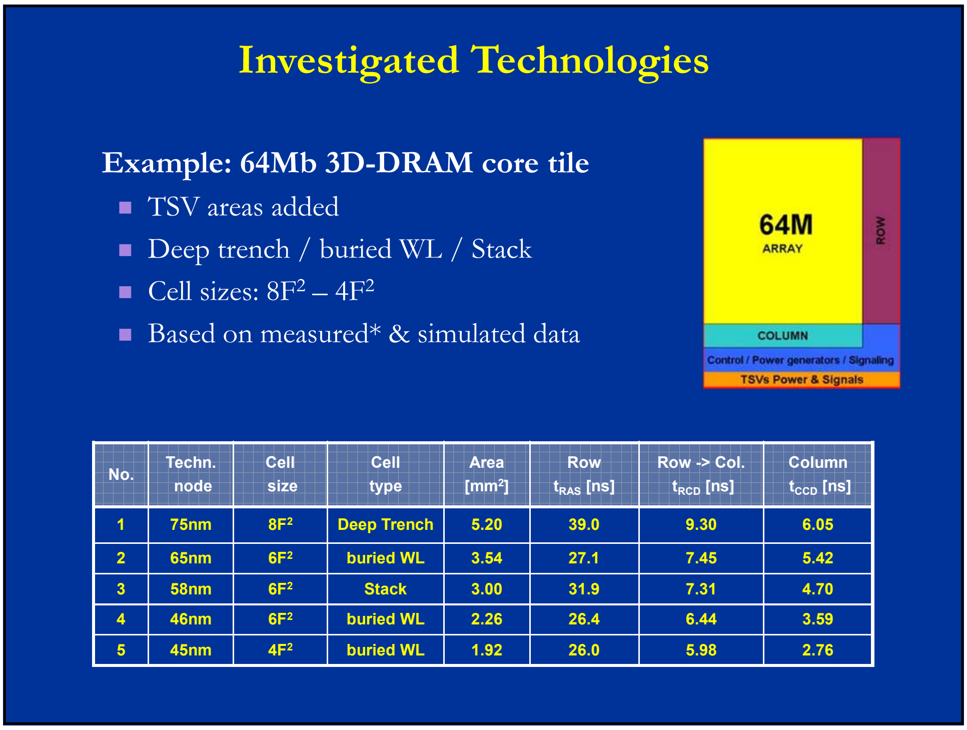

Investigated Technologies

Example: 64Mb 3D-DRAM core tile

- TSV areas added
- Deep trench / buried WL / Stack
- Cell sizes: $8F^2 – 4F^2$
- Based on measured* & simulated data

64M
ARRAY
ROW
COLUMN
Control / Power generators / Signaling
TSVs Power & Signals

No.	Techn. node	Cell size	Cell type	Area [mm²]	Row t_{RAS} [ns]	Row -> Col. t_{RCD} [ns]	Column t_{CCD} [ns]
1	75nm	$8F^2$	Deep Trench	5.20	39.0	9.30	6.05
2	65nm	$6F^2$	buried WL	3.54	27.1	7.45	5.42
3	58nm	$6F^2$	Stack	3.00	31.9	7.31	4.70
4	46nm	$6F^2$	buried WL	2.26	26.4	6.44	3.59
5	45nm	$4F^2$	buried WL	1.92	26.0	5.98	2.76

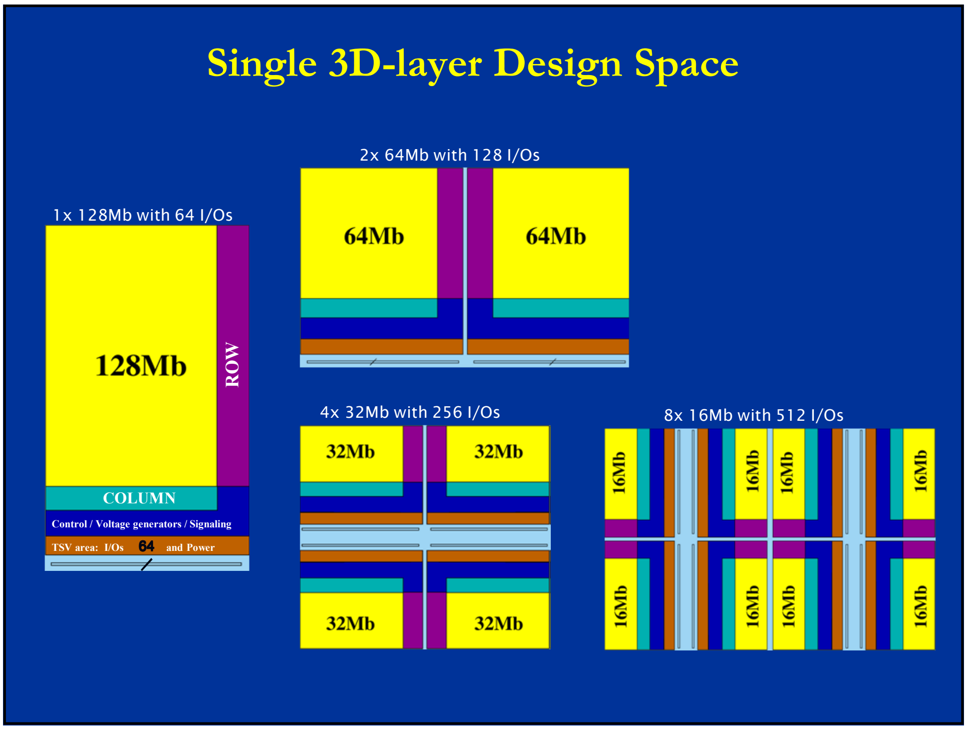

Single 3D-layer Design Space

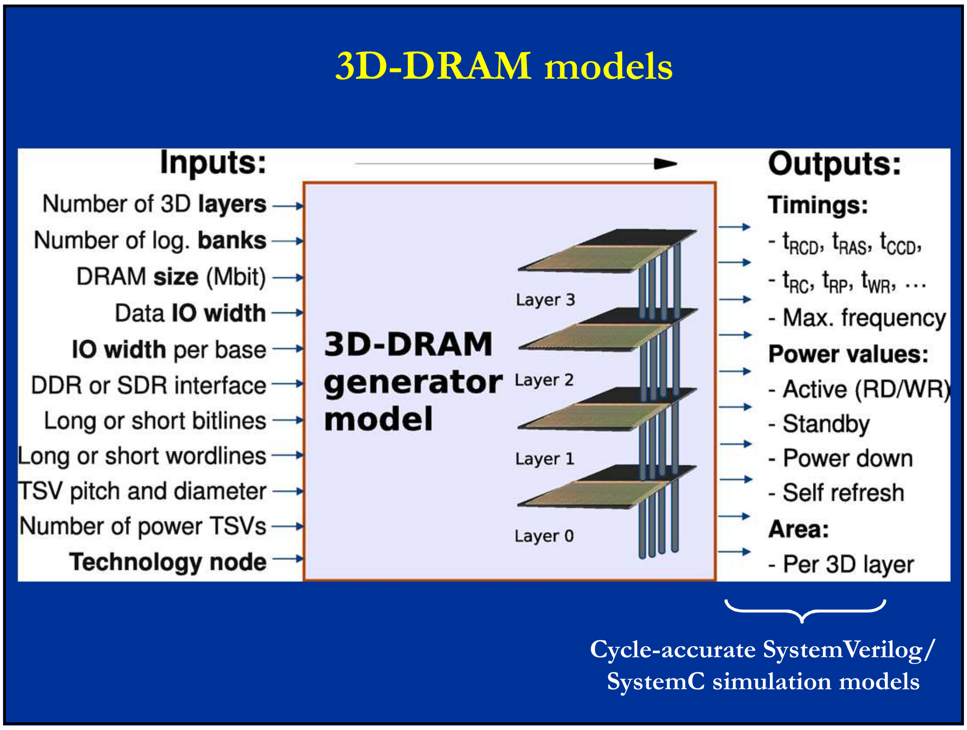

3D-DRAM models

Inputs:

- Number of 3D **layers**
- Number of log. **banks**
- DRAM **size** (Mbit)
- Data **IO width**
- **IO width** per base
- DDR or SDR interface
- Long or short bitlines
- Long or short wordlines
- TSV pitch and diameter
- Number of power TSVs
- **Technology node**

3D-DRAM generator model

Outputs:

Timings:
- t_{RCD}, t_{RAS}, t_{CCD},
- t_{RC}, t_{RP}, t_{WR}, …
- Max. frequency

Power values:
- Active (RD/WR)
- Standby
- Power down
- Self refresh

Area:
- Per 3D layer

Cycle-accurate SystemVerilog/
SystemC simulation models

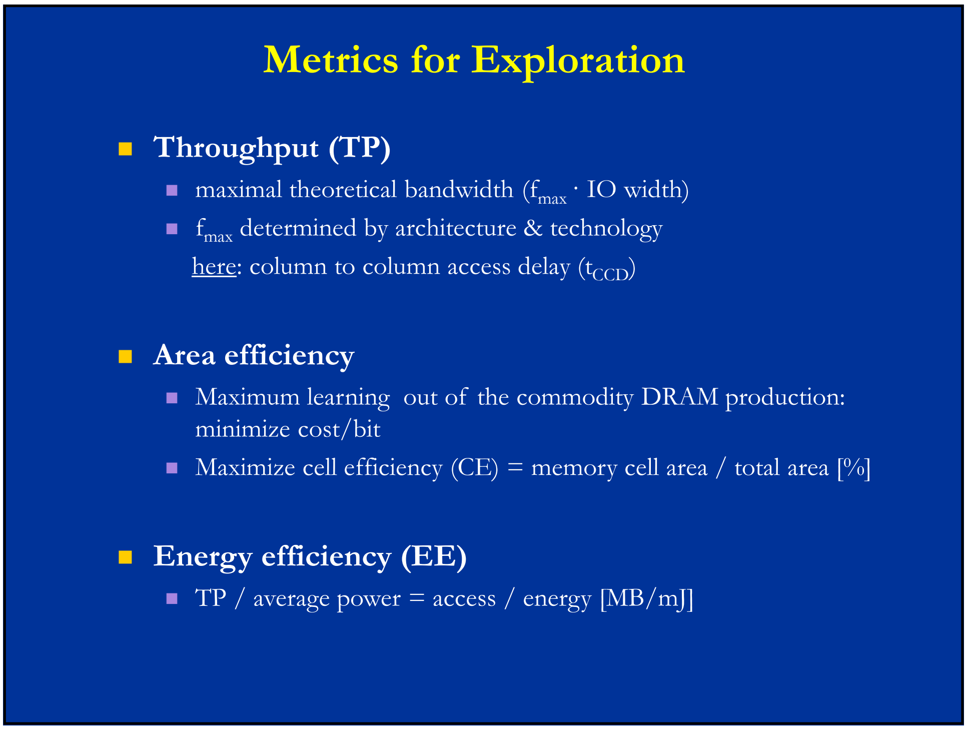

Metrics for Exploration

- **Throughput (TP)**
  - maximal theoretical bandwidth (f_{max} · IO width)
  - f_{max} determined by architecture & technology
    here: column to column access delay (t_{CCD})
- **Area efficiency**
  - Maximum learning out of the commodity DRAM production: minimize cost/bit
  - Maximize cell efficiency (CE) = memory cell area / total area [%]
- **Energy efficiency (EE)**
  - TP / average power = access / energy [MB/mJ]

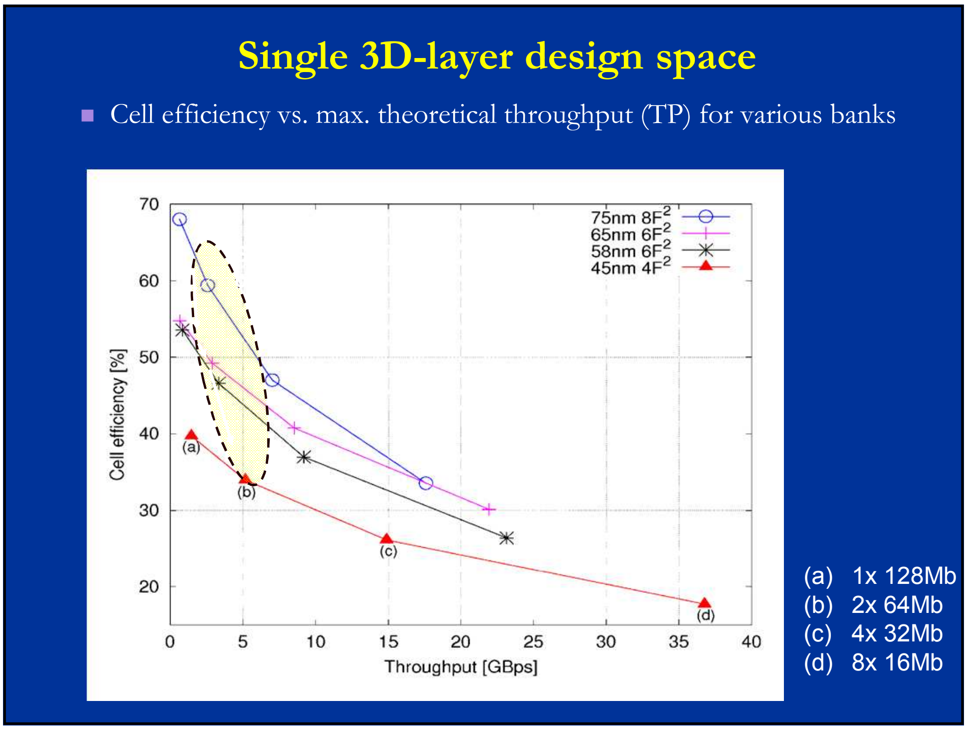

Single 3D-layer design space

- Cell efficiency vs. max. theoretical throughput (TP) for various banks

(a) 1x 128Mb
(b) 2x 64Mb
(c) 4x 32Mb
(d) 8x 16Mb

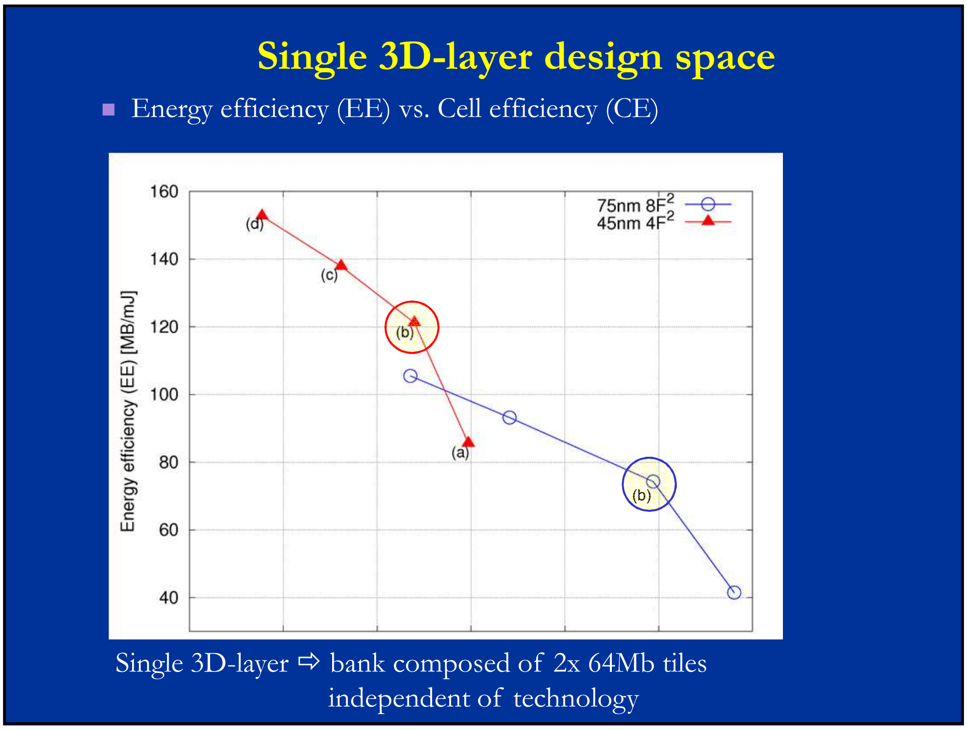

Single 3D-layer design space

- Energy efficiency (EE) vs. Cell efficiency (CE)

Single 3D-layer ⇨ bank composed of 2x 64Mb tiles independent of technology

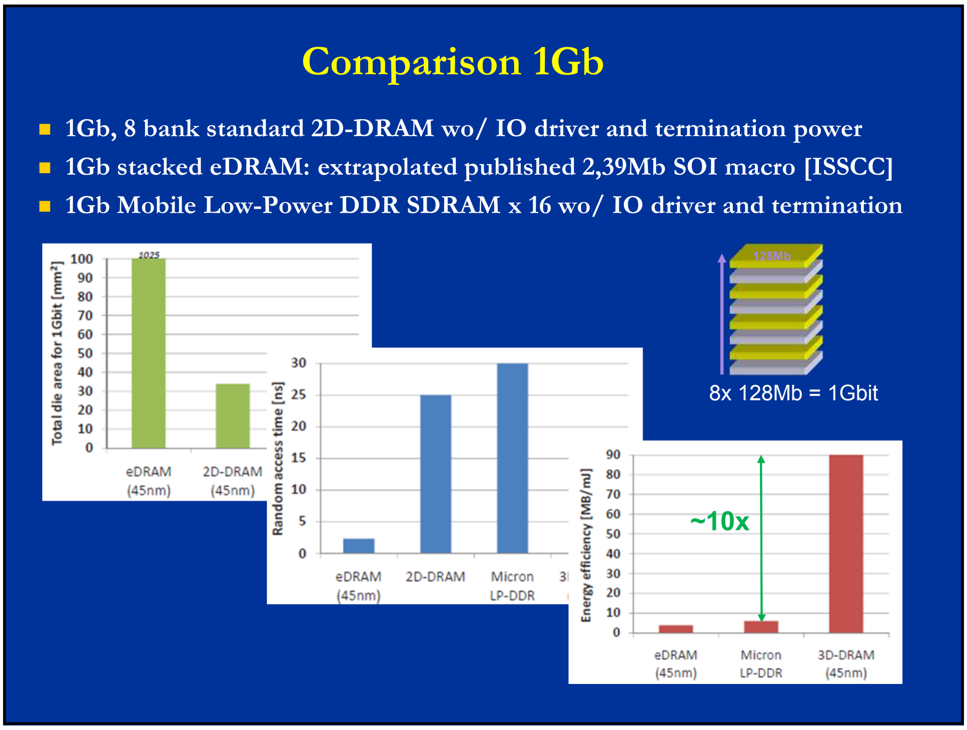

Comparison 1Gb

- 1Gb, 8 bank standard 2D-DRAM wo/ IO driver and termination power
- 1Gb stacked eDRAM: extrapolated published 2,39Mb SOI macro [ISSCC]
- 1Gb Mobile Low-Power DDR SDRAM x 16 wo/ IO driver and termination

8x 128Mb = 1Gbit

Multi-Channel 3D-DRAM Controller

Different request granularities

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

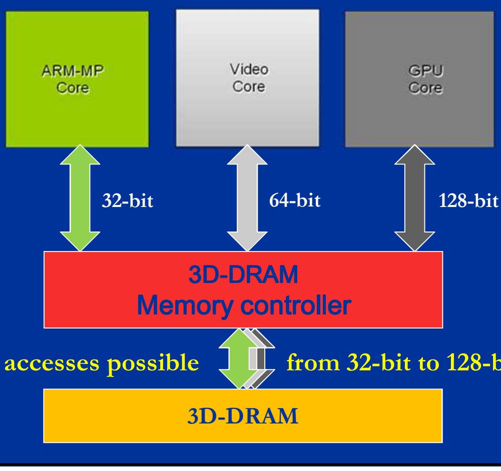

Three types of I/O accesses possible from 32-bit to 128-bit

Multi-Channel 3D-DRAM Controller

Front End:

- Synchronization with Dual Clock FIFOs
- Arbitration
- Buffering, Scheduling, Reordering

Back End

- 3D DRAM command Encoding
- Tracking of the BANK status
- Multi IO reconfiguration and data latching for 32/64/128 bit

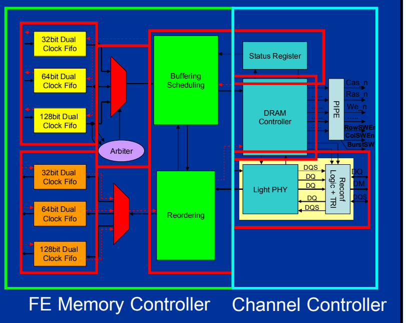

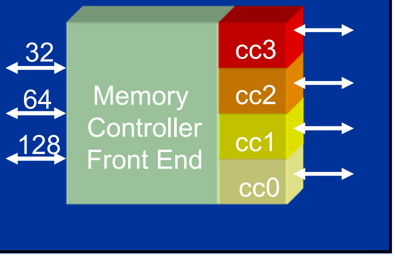

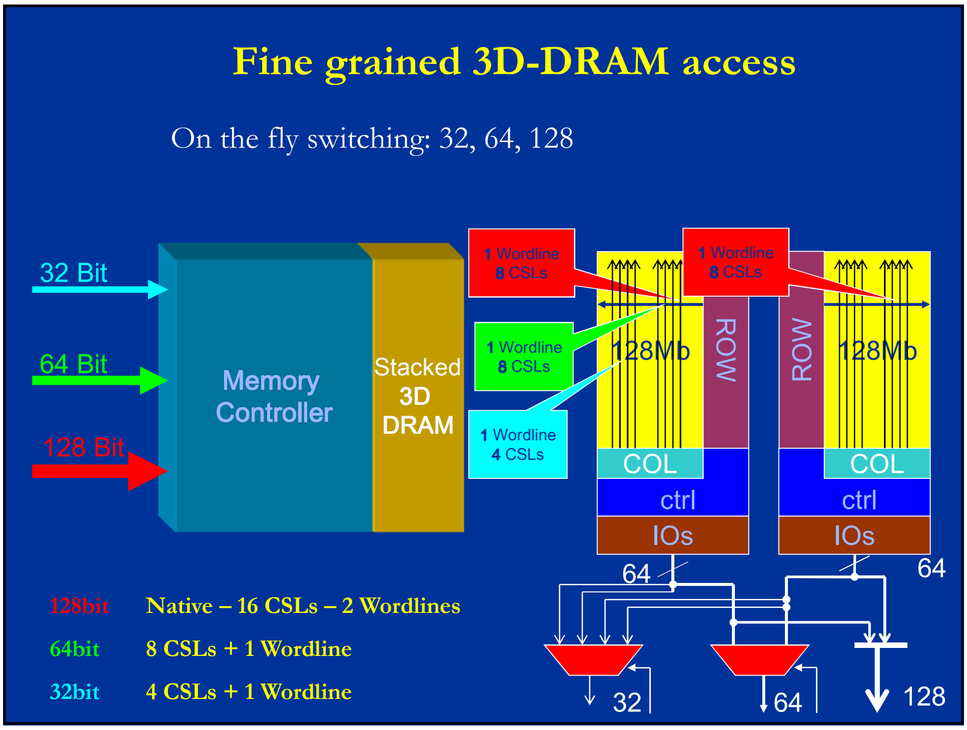

Fine grained 3D-DRAM access

On the fly switching: 32, 64, 128

128bit	Native – 16 CSLs – 2 Wordlines
64bit	8 CSLs + 1 Wordline
32bit	4 CSLs + 1 Wordline

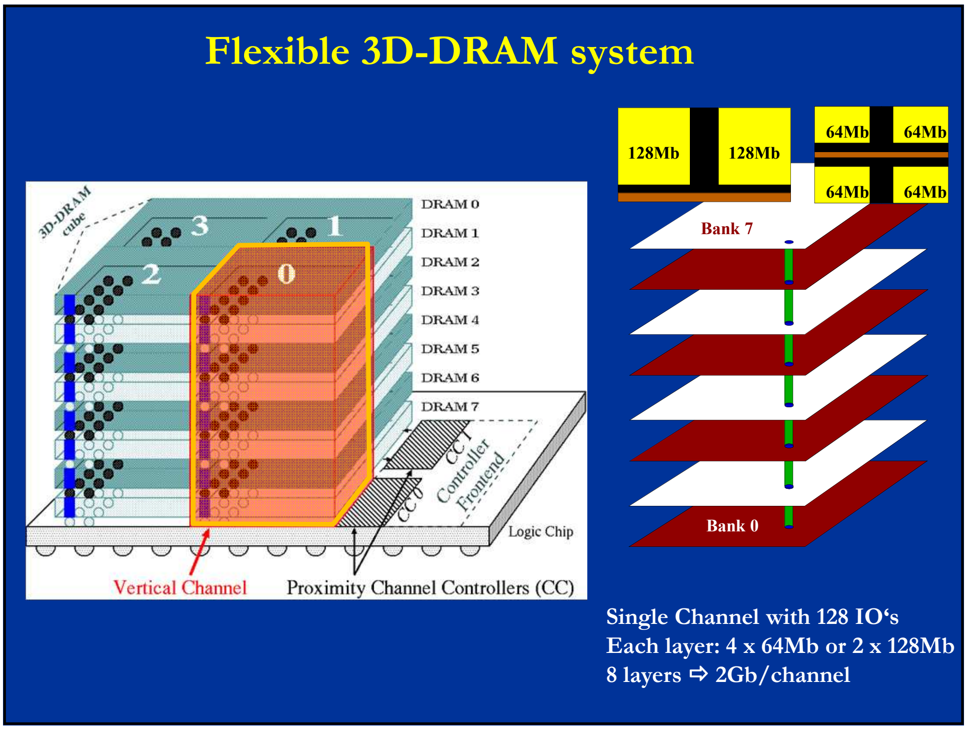

Flexible 3D-DRAM system

Single Channel with 128 IO's
Each layer: 4 x 64Mb or 2 x 128Mb
8 layers ⇨ 2Gb/channel

Investigated 3D-DRAM Configurations

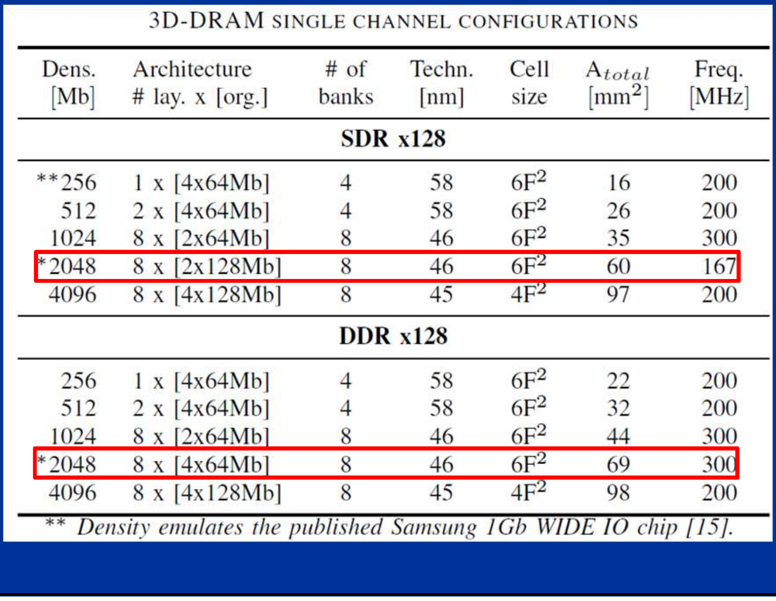

3D-DRAM SINGLE CHANNEL CONFIGURATIONS

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

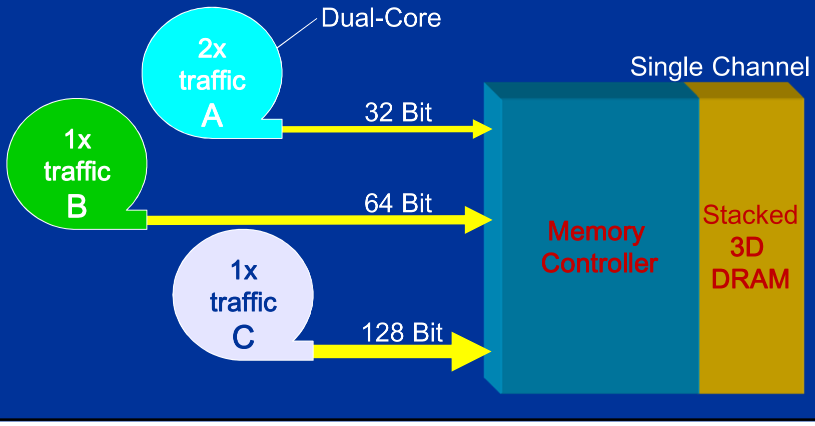

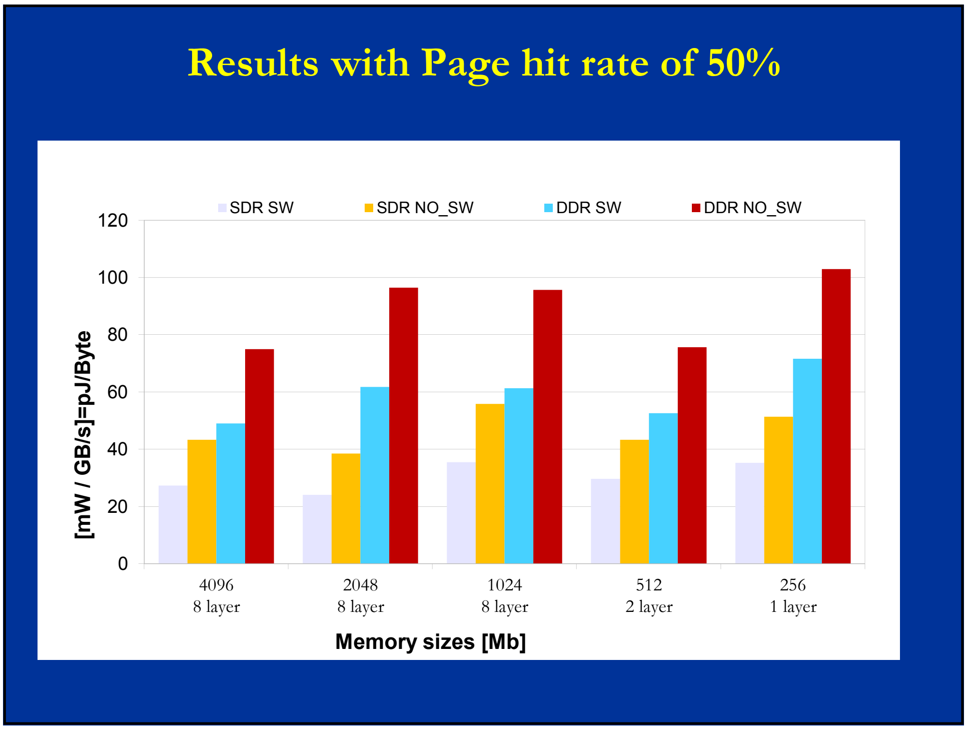

Results with Page hit rate of 50%

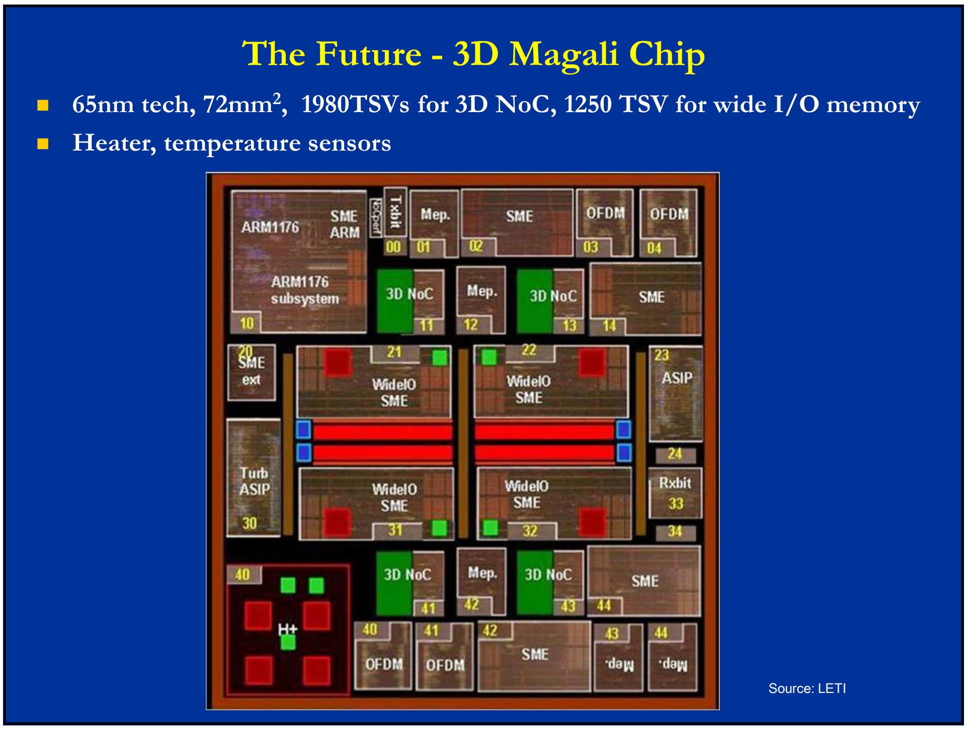

The Future - 3D Magali Chip

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

Source: LETI

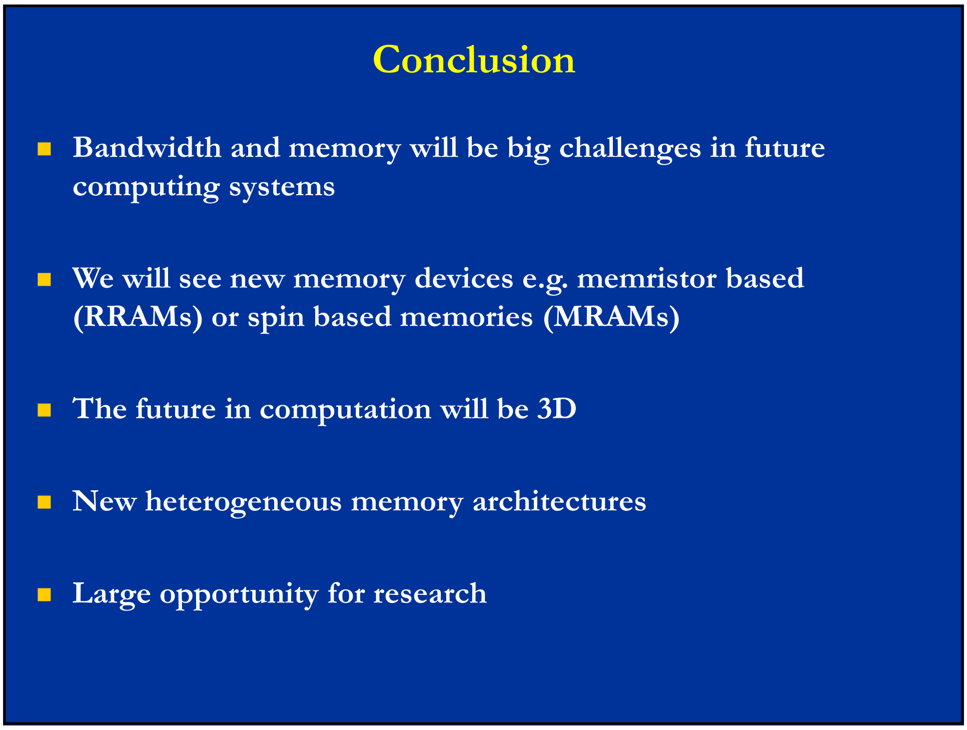

Conclusion

- Bandwidth and memory will be big challenges in future computing systems
- We will see new memory devices e.g. memristor based (RRAMs) or spin based memories (MRAMs)
- The future in computation will be 3D
- New heterogeneous memory architectures
- Large opportunity for research