DESIGNCON® 2014
JANUARY 28-31, 2014
SANTA CLARA CONVENTION CENTER

An Implementer's Guide to Low-Power and High-Performance Memory Solutions

John Eble , Rambus Inc.
Ming Li , Rambus Inc.
Wendem Beyene , Rambus Inc.

Rambus

UBM Electronics

Memory Interfaces are Important … Everywhere

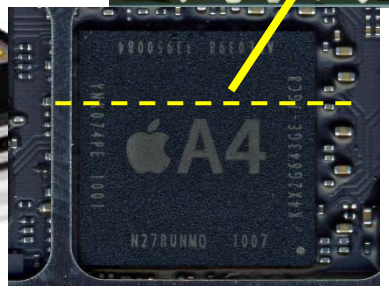

2x32 LPDDR2 ~ 6.4GB/s

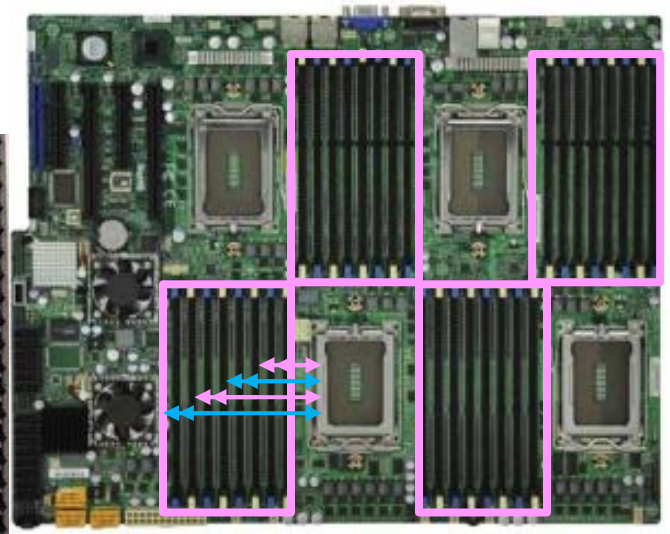

AMD & Intel Quad-Channel:
4×64 DDR3 ~ 51.2GB/s/CPU
(hundreds of DRAMs)

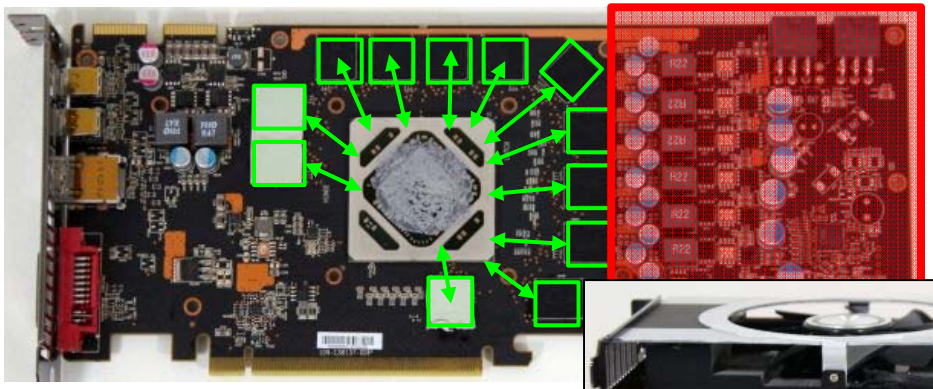

AMD Radeon R7950 Black Edition
384-bit bus, 264GB/s, 12 DRAMs
source: anandtech.com

Memory Interfaces are Important … Everywhere

Quad-Channel:
~ 51.2GB/s/CPU
of DRAMs)

AMD Radeon R7950 Black Edition
384-bit bus, 264GB/s, 12 DRAMs
source: anandtech.com

Diversity in Memory Systems

Different system considerations…

- Cost: # devices, packaging, interconnect, assembly
- Performance: Bandwidth, latency, working sets/locality
- Capacity: Min/max, configurability, upgradability
- Power: Thermal limits, energy cost, battery life

⋮

… lead to different device & system solutions

- Topologies: Point-to-Point, Multi-Rank, Buffered
- Assemblies: Module, Solder-down, SiP, PoP
- Interfaces: GDDR, DDR, LPDDR, XDR, HMC

Memory BW Requirement Is Growing Rapidly

SmartPhone Example

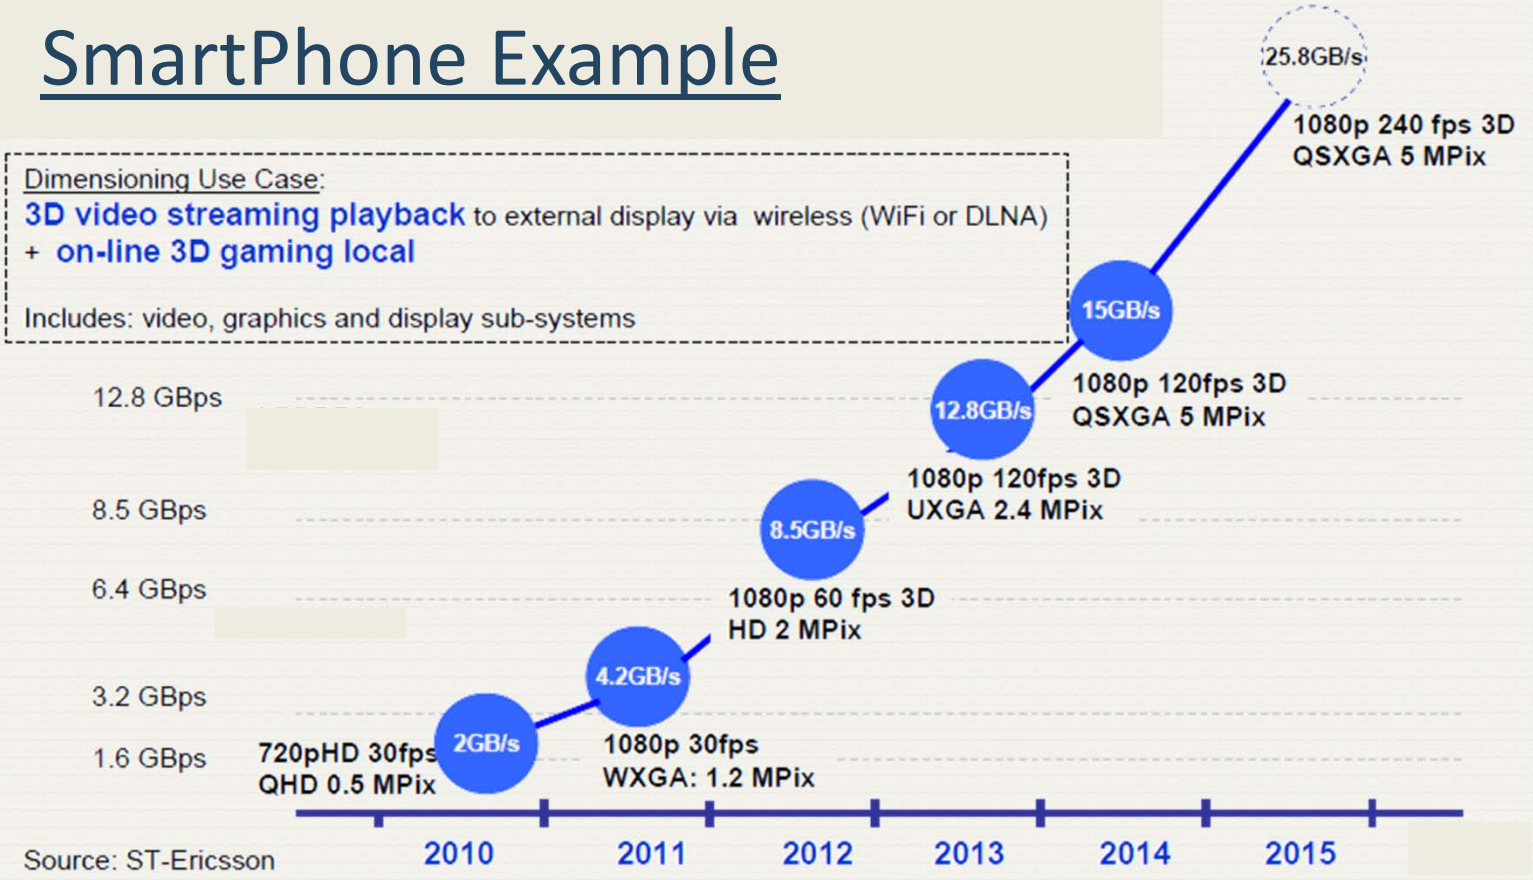

- Similar trends in CPU/GPU
  - Exacerbated by many-core trend

Memory Interface Needs to Continue Improving

- Conventional interface pin-count is not growing as fast
  - Package-on-Package (PoP) staying at 2x32 DQ bus in near future
    - Combination of PoP ball pitch, DRAM bond pads, package routing
  - High-end GPUs slowly increasing from 256→384→512 bit
  - Servers have slowly migrated to 4×64 bit channels

  →*BW/pin must increase to keep up with BW demand*

- Virtually all spaces have reached critical power limit
  - Thermal dissipation, battery life, energy cost, or combination
  - Energy/bit must be reduced to achieve higher BW
  - Complicated by BW/pin increase

3D Integration: Potentially Disruptive Technology

- Break the BW/power tradeoff with advanced packaging (e.g. Wide IO)
  - What are challenges for memory interfaces?
  - When and where will 3D integration take off?

- Compare *Solutions*, not *Technologies*
  - Solutions will vary across diversity of platforms
  - Conventionally packaged solutions are entrenched, evolving
  - Ultimately, a complex economic decision:
    - Overall cost/benefit, reuse across applications, migration path, …

Tutorial Goal

- Provide an overview of today's solution space and why the standards have evolved the way they have with an emphasis on interface ***signaling*** and ***clocking***
- Focus on ***3D packaging*** as a promising avenue to continue achieving necessary BW and Power
  - Understand the opportunities and key challenges/barriers that exist
- Discuss the role of ***signal integrity analysis*** in choosing your next solution
- Summarize implementer choices

Outline of Tutorial

- DRAM review [John]
- Conventionally packaged interfaces [John]
  - Evolution of standard interfaces
  - Signaling and clocking at peak bandwidth
  - Bandwidth variation
- 3D Packaging for Memory Applications [Ming]
  - Traditional 3D Packaging Technologies
  - New TSV Based 3D Packaging Technologies
- Signal and power Integrity [Wendem]
  - Signal and power integrity challenges
  - Memory interface analysis methodology
  - Main features of memory interfaces
  - Comparison of emerging standards-based solutions
- Conclusion and Final Comparison [John]

Commodity DRAM Architecture

Vogelsang, T. "Understanding the Energy Consumption of Dynamic Random Access Memories," MICRO, 2010

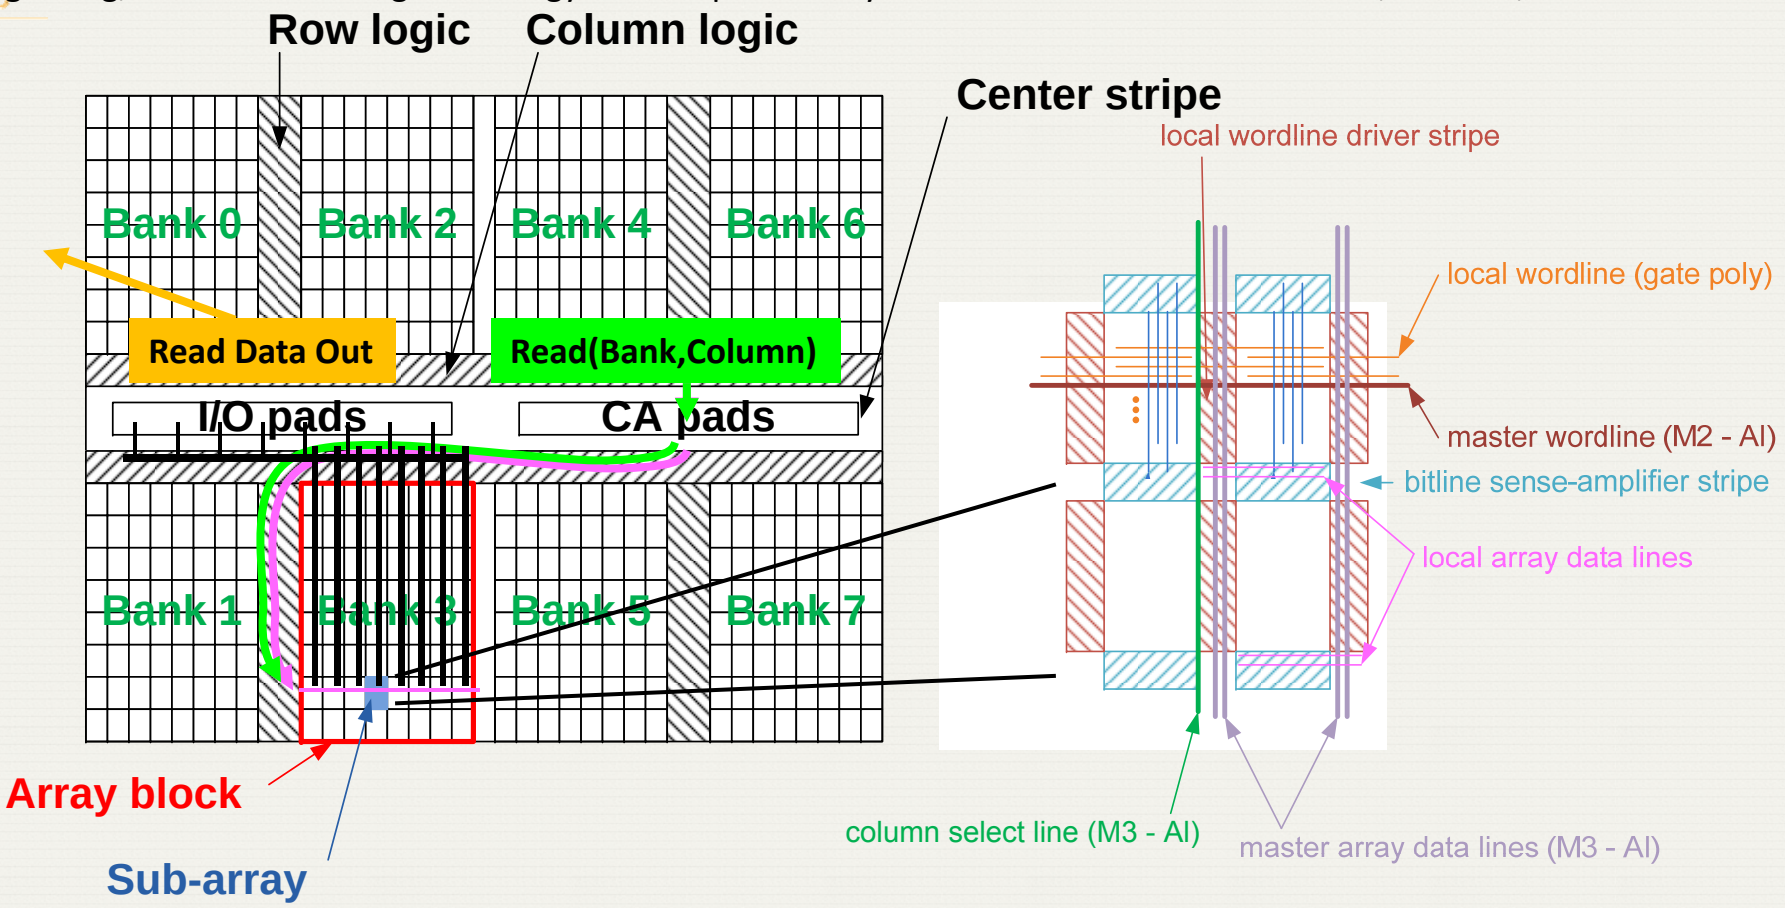

- DRAM core is RC limited
  - ~50ns row cycle, 5ns column cycle
  - $P = CV^2f$, wires not getting any shorter, V scaling slowly

Challenges in Increasing DRAM Core Bandwidth

- A few general approaches:
  - Finer array segmentation → GDDR ~40% more area than DDR
  - Interleaving bank accesses → access restrictions
  - More parallel data on each cycle → Worse *granularity*

- *Core design may impact efficiency as much as I/O*
  - Maintain balance, can't focus on one item in isolation
  - Even more tradeoffs at system design level

Challenges to Decreasing DRAM Core Power

- Cell capacitance nearly constant with generation to keep refresh times unchanged
- Optimal die size for yield and to drive $/bit down is ~fixed
  - Wire lengths also ~fixed with no minor improvements in fF/um
  - Increase in banks can help this at higher $
- Process constraints (cheap, low leakage) dictate higher voltages than logic processes, where V scaling has slowed
- Core/IO power split ~65%/35%
- LPDDR4 note:
  - Core reorganized to reduce power
  - Single x32 transitions to 2 x16

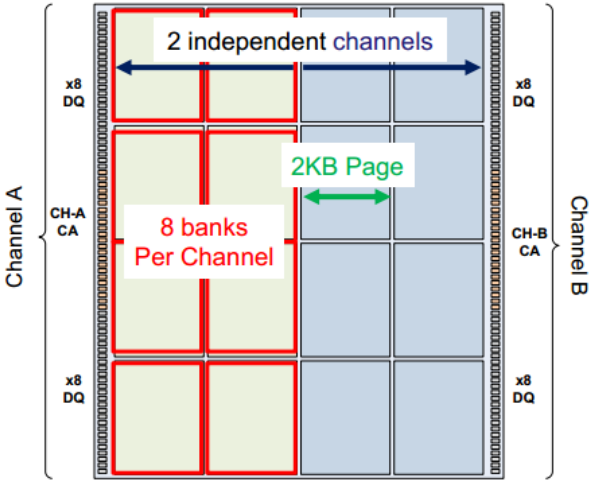

(Source: Micron, Mobile Forum 2013)

Outline of Tutorial

- DRAM review [John]
- Conventionally packaged interfaces [John]
  - Evolution of standard interfaces
  - Signaling and clocking at peak bandwidth
  - Bandwidth variation
- 3D Packaging for Memory Applications [Ming]
  - Traditional 3D Packaging Technologies
  - New TSV Based 3D Packaging Technologies
- Signal and power Integrity [Wendem]
  - Signal and power integrity challenges
  - Memory interface analysis methodology
  - Main features of memory interfaces
  - Comparison of emerging standards-based solutions
- Conclusion and Final Comparison [John]

Industry Roadmaps – BW and power efficiency

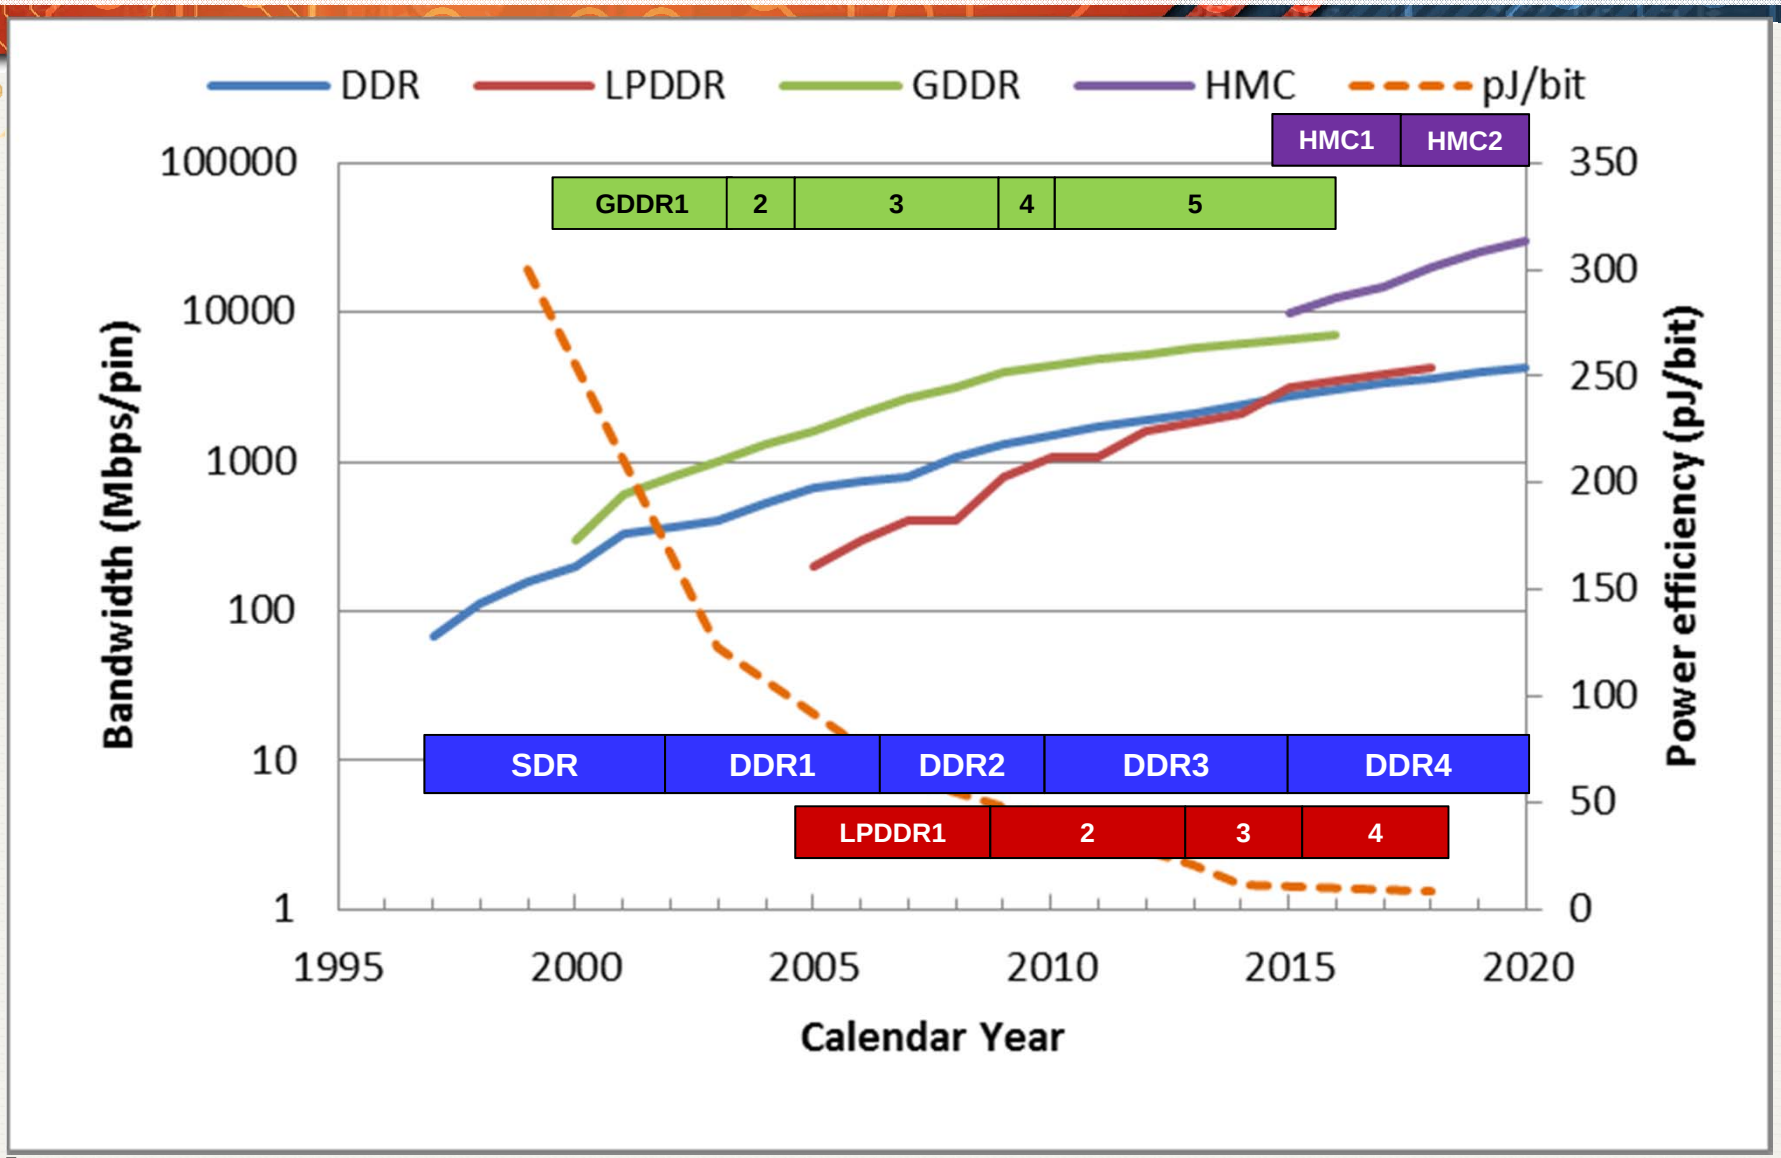

Power efficiency estimates include DRAM IO and core

Today's two mainstream memory system topologies and types

- DDR
  - High capacity (modules) or cost sensitive
  - Mainstream DRAM process
  - General and broad application space

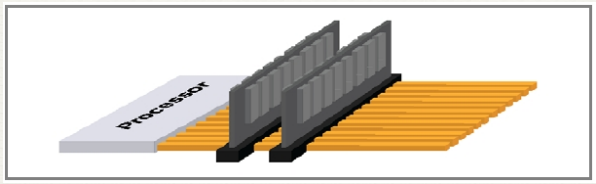

- LPDDR
  - Low capacity PoP or short reach direct attach
  - Lower leakage process
  - Battery life and stringent TDP
  - Aggressive power states, self-refresh, and transition times

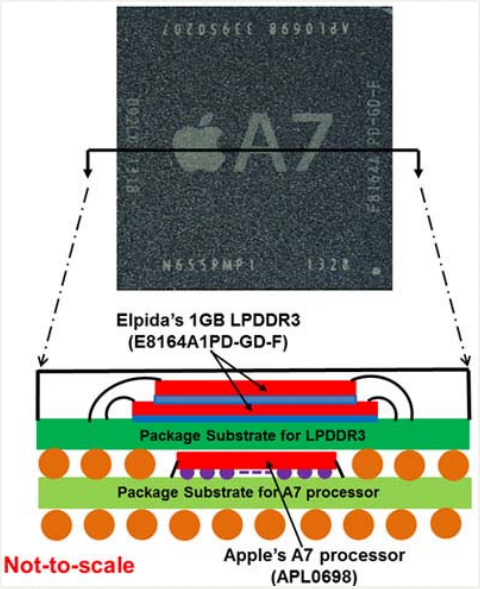

Market/topologies lead to different tradeoffs and decisions

Key architecture differences - DDR vs. LPDDR

	DDR	LPDDR
Device widths	x4, x8, x16	x32, dual x16
Ranks per channel	Up to 4	Up to 2
Channel width	Up to x72	x32
Command/Addressing	Explicit single cycle, high pin count	Encoded multi-cycle, low pin count
Training	ZQ, DQ W leveling and R/W eye, VREF, read gate, CA training	~same + periodic read gate training
Prefetch	8-bits (with bank grouping in DDR4)	8-bits
Die Floorplan	Die pads in center	Die pads on edge
Power Modes	Base	+ Deep Power Down and Clock Stop
Refresh modes	Base	+ Temp-comp. refresh and Partial Array Self Refresh
RAS features	CRC and CA parity	None

Multiple generations of memory technology

- Generations overlap and controller PHYs tend to be backwards compatible
- Memory interfaces are DC-coupled and, thus far, require agreement on the VDDQ voltage
  - Driven by technology shrinkage and power reduction requirements
  - E.q. DDR3 has three variations (base + DDR3L, DDR3U)
- Densities increase requiring more address pins
- Speeds increase requiring interface feature changes as well as changes in the “prefetch” to accommodate the fixed core speed
- DDR has gone through 5 generations over ~20 years
  - DDR4 specification released in 2012
- LPDDR will have gone through 4 generations **~8** years!
  - LPDDR4 specification expected to be publically released this year

Evolution of DDR

Variables	SDRAM	DDR1	DDR2	DDR3	DDR4
VDD/VDDQ/VDP	3.3/2/5	2.5	1.8	1.5/1.5/-	1.2/1.2/2.5
Data Rate (Mbps)	166	400	800	1600	3200
Clock	Single ended	Differential	Differential	Differential	Differential
Strobes	unsupported	Single ended	Single ended or Differential	Differential	Differential
Vref	N/A	External	External	Ex/Internal	Ex/Internal*
ODT	unsupported	unsupported	supported	supported	supported
Interface	LVTTL	SSTL	SSTL	SSTL	PODL

- Key implementer choices to consider
  - System configuration and topology
  - Data rate and power envelope required
  - VDDQ voltage
  - Cost and availability through product life cycle
  - Available CPHY IP or internal design – complexity (SSTL/PODL)

* External for CA

Evolution of LPDDR

Variables	LPDDR1	LDDR2	LDDR3	LDDR4†
VDD1, VDD2, VDDQ (V)	1.8	1.8/1.2/1.2/1.2	1.8/1.2/1.2	1.8/1.1/1.1
Data Rate (Mbps)	400	1066	2133	4266
CA	SDR	DDR	DDR	SDR
Clock	Single ended	Differential	Differential	Differential
Strobes	Single ended	Differential	Differential	Differential
Vref	External	External	External	Ex/Internal
ODT	unsupported	Unsupported	supported	supported
Interface	SSTL	HSUL	HSUL/PODL	LVSTL

- Key implementer choices to consider
  - System configuration and topology, data rate and power envelope required
    - Unterminated (HSUL) vs. Terminated (PODL/LVSTL)
  - VDDQ voltage
  - Cost and availability through product life cycle
  - Available CPHY IP or internal design - complexity

†Based on available public information

Common features for BW/power efficiency improvements

- Silicon level (CPHY, DRAM)
  - Improved processes and voltage scaling (but severe slow-down)
  - Improved termination and signaling choices
  - Low jitter clocking techniques
  - Controller PHY complexity – asymmetric system
  - Calibration and eye training (power-up and periodic)
  - Coding - DBI
  - Noise management (SSO, Vref, crosstalk, reference planes)
  - RAS features: Parity and CRC
- System Level (channel, topology)
  - Improved package for reducing inductance (L dI/dt) and reducing crosstalk
  - Board & connector or co-packaging improvements
  - Reduction of supported ranks; move to point-to-point

Outline of Tutorial

- DRAM review [John]
- Conventionally packaged interfaces [John]
  - Evolution of standard interfaces
  - Signaling and clocking at peak bandwidth
  - Bandwidth variation
- 3D Packaging for Memory Applications [Ming]
  - Traditional 3D Packaging Technologies
  - New TSV Based 3D Packaging Technologies
- Signal and power Integrity [Wendem]
  - Signal and power integrity challenges
  - Memory interface analysis methodology
  - Main features of memory interfaces
  - Comparison of emerging standards-based solutions
- Conclusion and Final Comparison [John]

Importance of *Signaling* and *Clocking* - Power

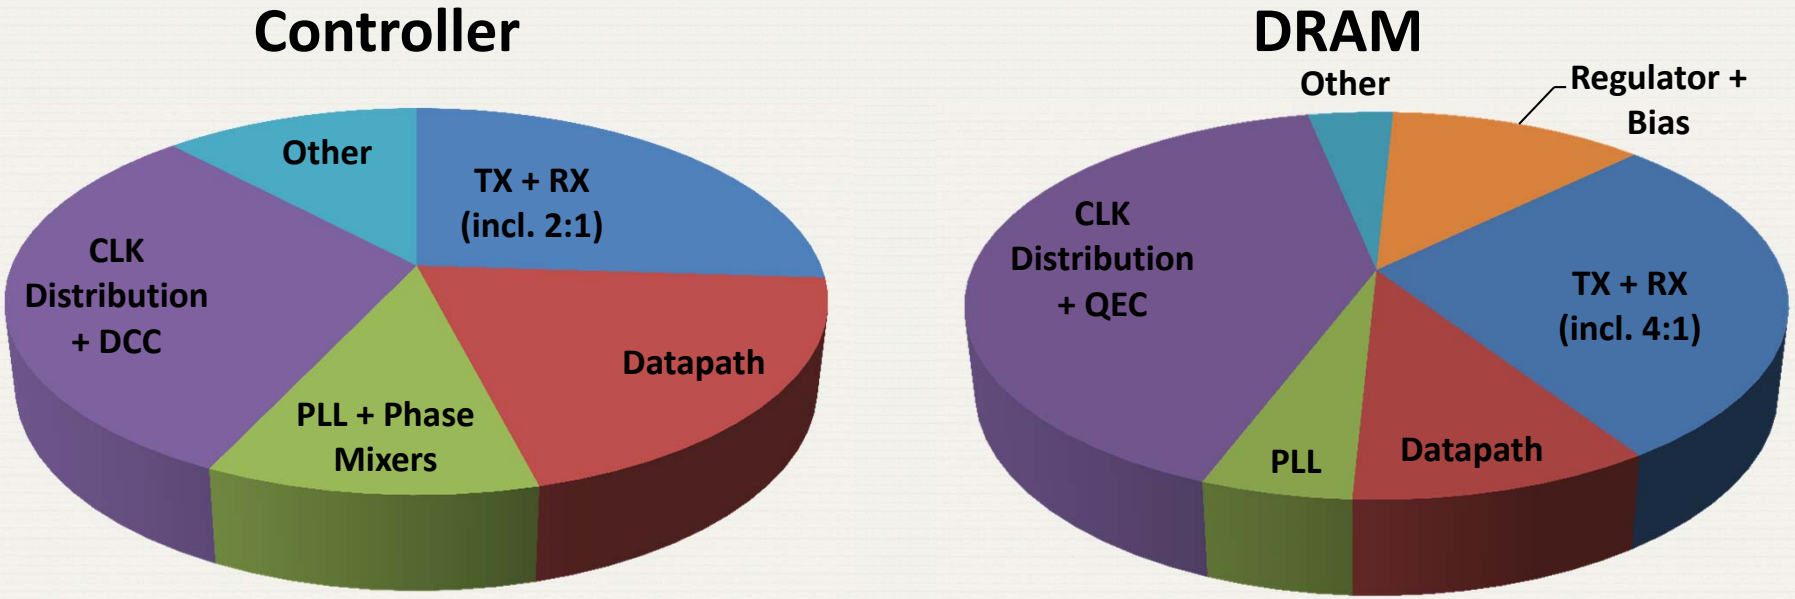

Lee, et.al. JSSC'08

- 16 Gb/s graphics memory interface prototype
  - **CTRL:** 65nm G+ ASIC process **DRAM**: "emulated" 40nm DRAM process
  - Write: 5-tap TX-FIR with CML type driver
  - Read: Continuous time linear equalizer

High-Speed Signaling

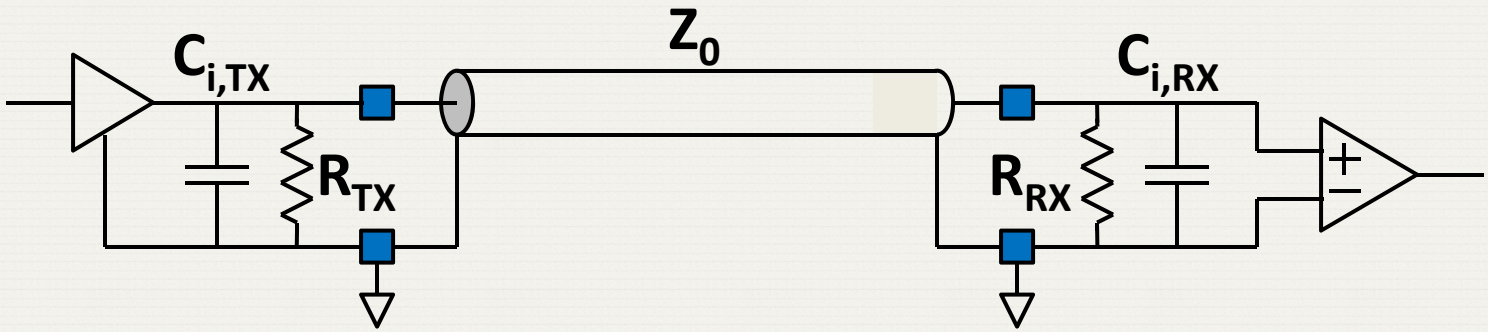

- Must treat channels as transmission lines
  - Typically terminated to reduce reflections
  - Impedance discontinuities lead to reflections (C_i, vias, connectors, trace/termination impedance tolerances, …)
  - Dielectric loss and skin-effect lead to pulse dispersion
  - Crosstalk from neighboring pins and traces
- Most memory channels are bi-directional
- Channel equalization approaches:
  - Linear equalizer, transmit FIR, decision feedback equalizer (DFE)
  - Well known in SerDes, just starting to reach high-volume DRAM
    - ***K. Koo, et al, ISSCC 2012, Paper 2.2 (Hynix DDR4)*** shows TX pre-emphasis

Differential vs. Single Ended Signaling

Series TX Termination Examples:

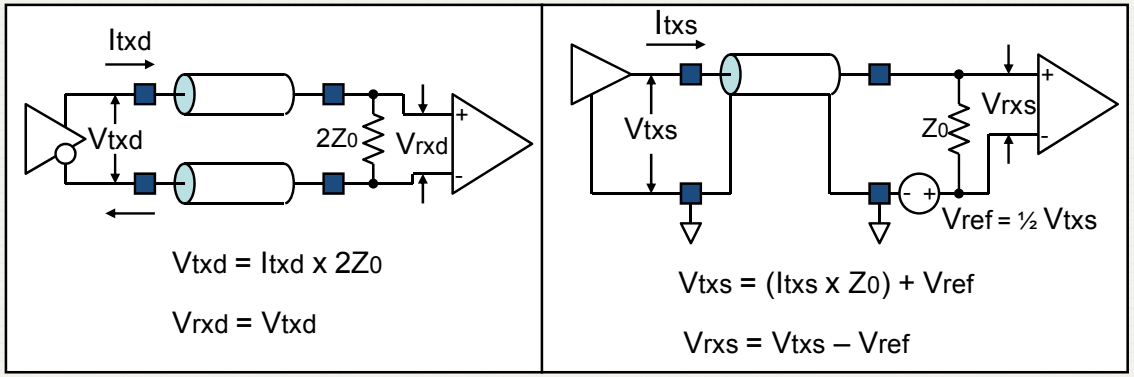

Issue	Diff	SE
Signaling current: – Differential can achieve 2× the swing with the same line current	✓	
Robustness (Performance): – SE requires reference voltage, more sensitive to noise & crosstalk	✓	
Pin efficiency: – Differential requires twice as many signal pins, or twice the baud rate – Caveat: SE typically needs more VDD/VSS pins for supply bounce		✓

- **Most SerDes, Intel/AMD parallel busses are differential**
- **Most high volume DRAMs are single ended**

Reducing Driver & Termination Power

- Lower V_{DDQ} ($P \propto V_{DDQ}^2$) or swing ($P \propto V_{swing}$)
  - Transistor headroom limitations, esp. in DRAM process
  - ***Sensitive RX is key to reducing swing***
    - Proportional errors such as ISI & crosstalk scale down
    - Fixed errors such as voltage offsets can be calibrated
    - Thermal noise is the ultimate limiter (order of 1mV,rms)

- Optimize signaling configuration
  - SSTL (DDR2/3) → PODL (DDR4, GDDR3/4/5): lower term. current
  - HSUL in LPDDR removes RX termination entirely
    - CMOS-like signaling, but TX termination still critical
    - Suitable for clean, short (wrt UI) channels and moderate data rates

- Code data (e.g., DBI-AC/DC)
  - Add one bit to a byte and code away from high-power states

Example: Low-swing driver with TX-Regulator

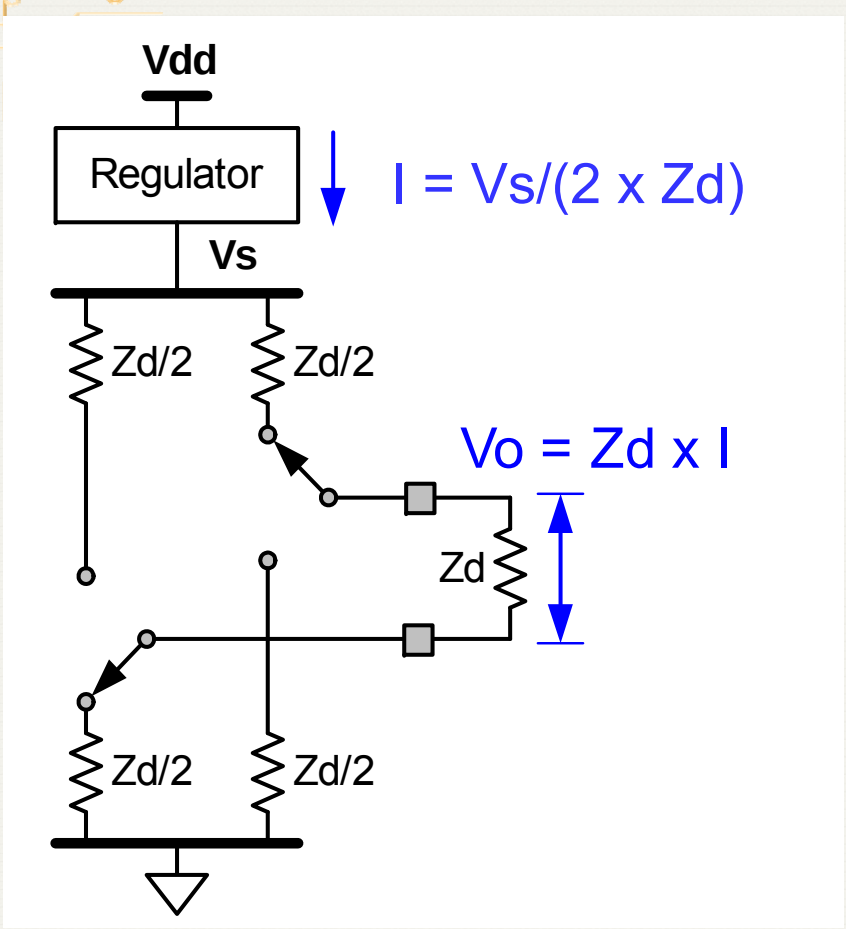

- Low-swing (±100mV,differential)
  - Sensitive RX (offset calibrated)
- Near ground common-mode
  - Can use all NMOS devices
  - Allows simple ESD structure

  → Low Ci
- Internal regulator
  - Allows series termination (a.k.a. “voltage mode” signaling)
  - Doesn’t require separate V_{DDQ}

Poulton, et al. JSSC 2007

Comparison of Signaling Standards

HSUL

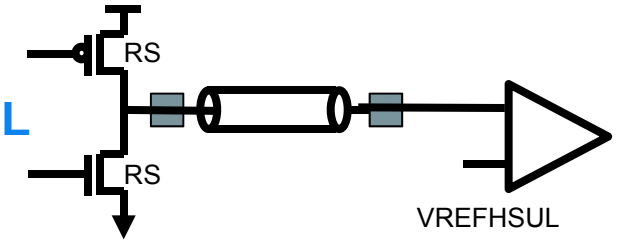

- Large voltage swings
- Power reduction in most cases
- Limited channels/data rates
- Lower Rx Ci

SSTL

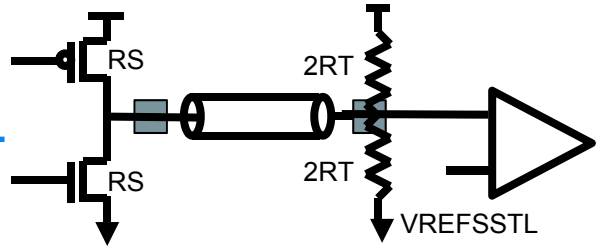

- VREF at VDDQ/2
- Much improved SI over HSUL allowing expansion of channels/data rates

PODL

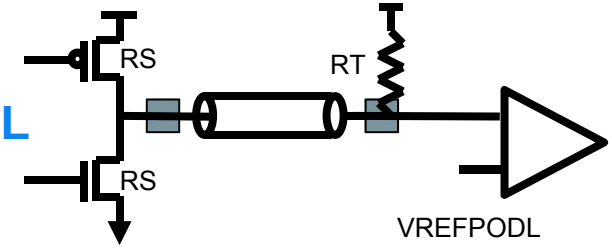

- Lower termination/DC power - high logic level has I=0
- DBI can be used
- VREF AC coupled to a single supply

LVSTL

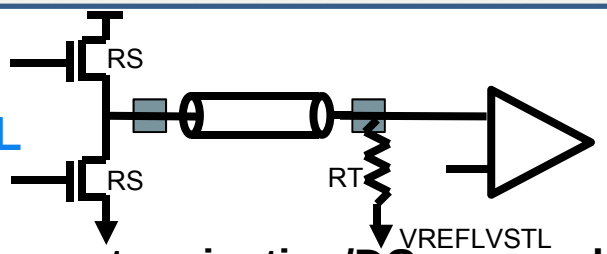

- Lower termination/DC power - low logic level has I=0
- DBI can be used
- VREF AC coupled to lowest Z supply
- VDDQ agreement may not be req'd
- N over N driver -> Ci reduction and allows aggressive VDDQ scaling

Signaling Standards – Swing, VREF, and Current

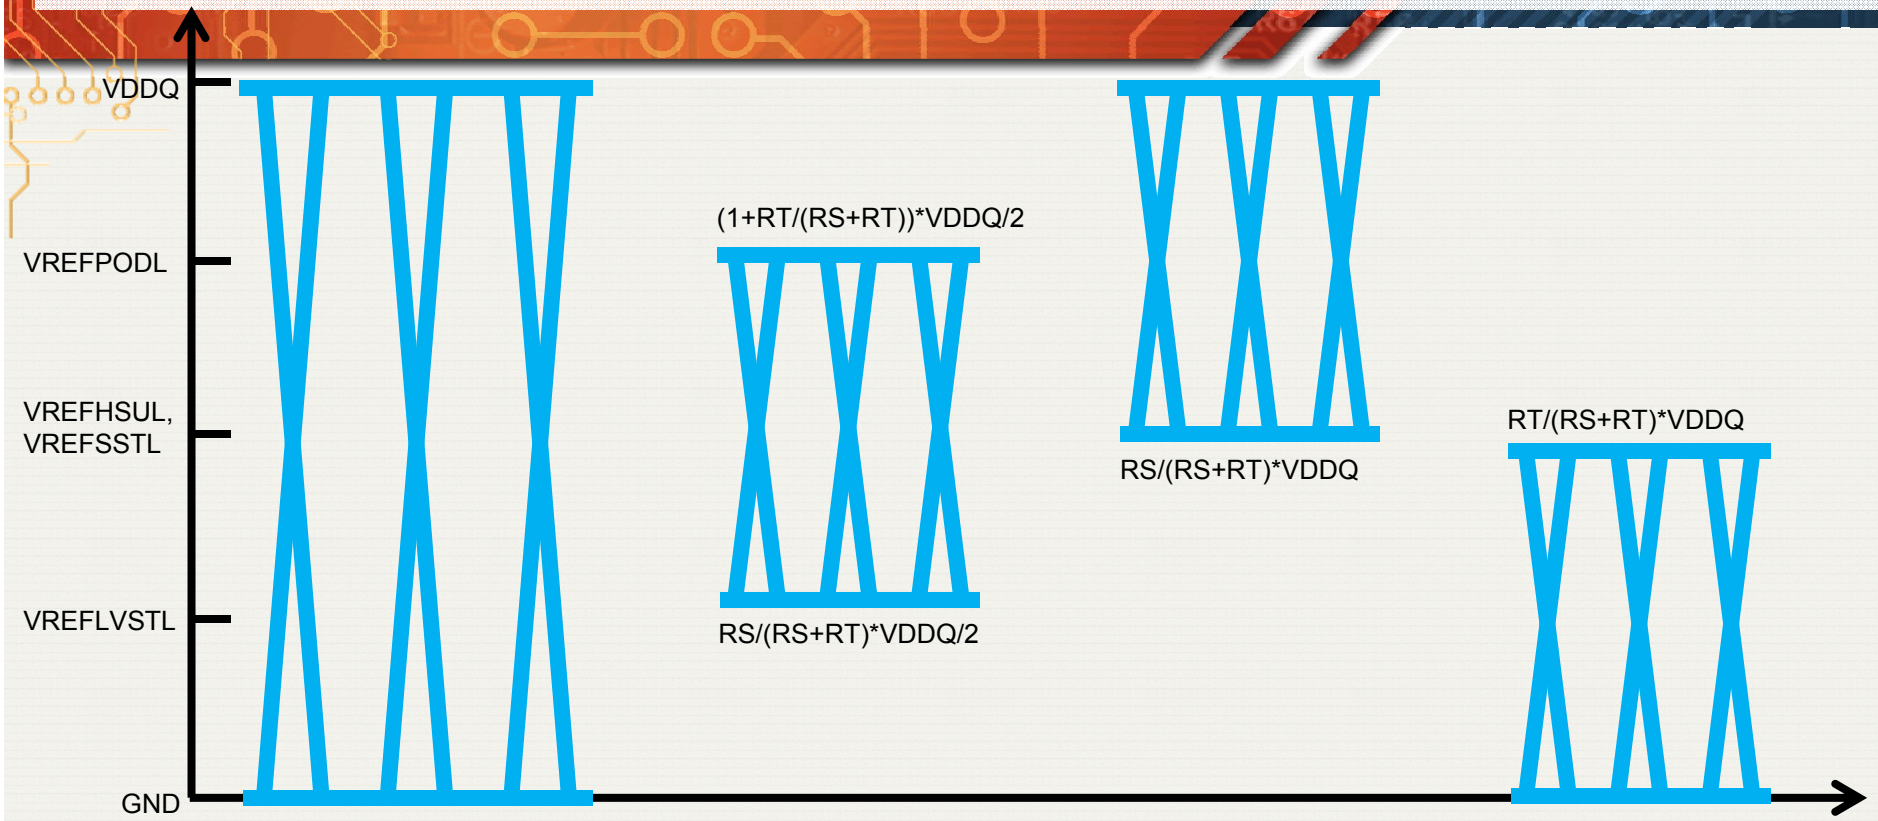

HSUL:
Swing = VDDQ
IH = 0
IL = 0
(Only dynamic switching power)
VREF = VDDQ/2

SSTL:
Swing = 2RT/(RS+RT)*VDDQ/2
IH = VDDQ / 2(RS+RT)
IL = VDDQ / 2(RS+RT)
VREF = VDDQ/2

PODL:
Swing = RT/(RS+RT)*VDDQ
IH = 0
IL = VDDQ / (RS+RT)
VREF = ((2RS+RT)/(RS+RT))* VDDQ/2

LVSTL:
Swing = RT/(RS+RT)*VDDQ
IH = VDDQ / (RT+RS)
IL = 0
VREF = RT/(RS+RT)*VDDQ/2

HSUL & PODL I/O Power Comparison

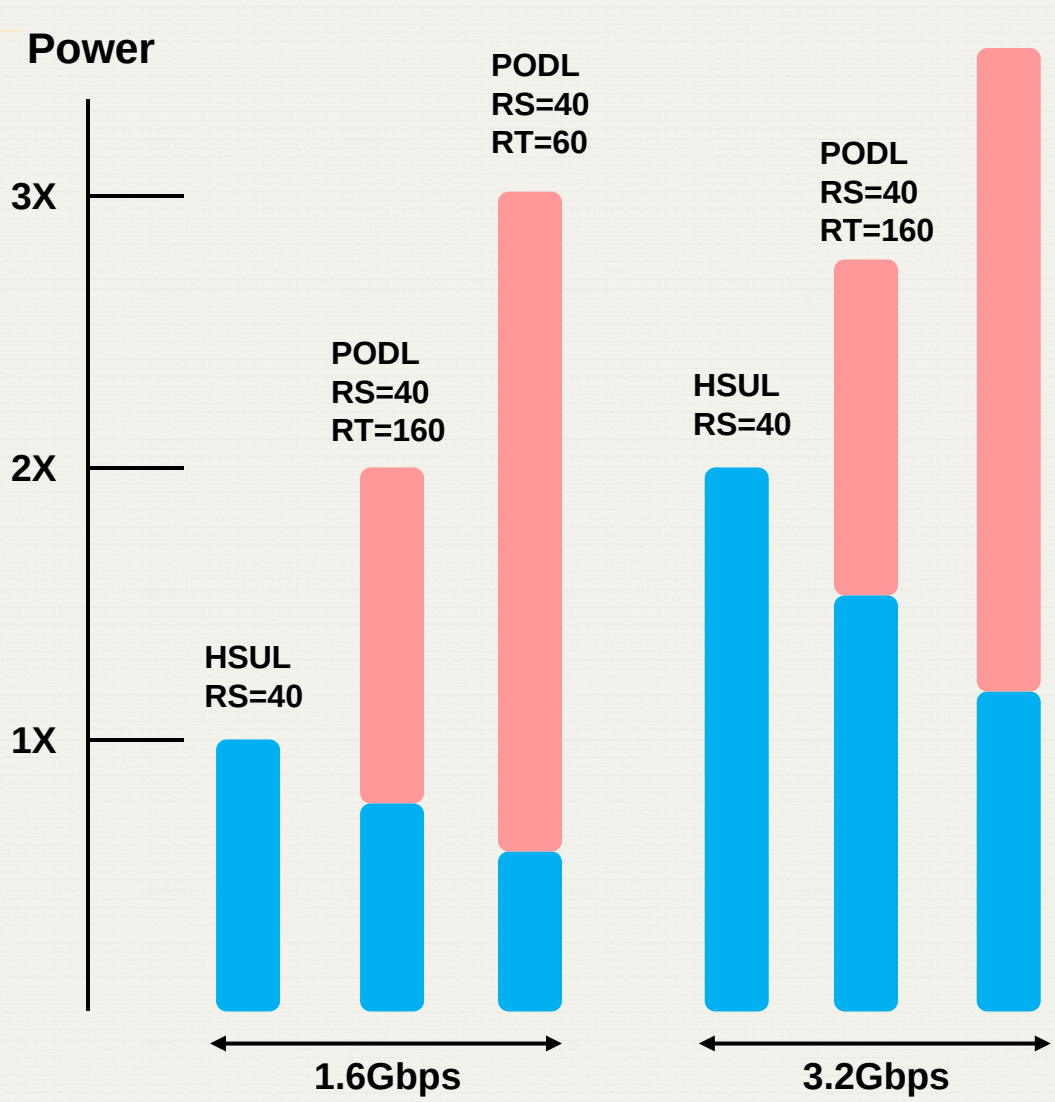

Notes :

- Weaker PODL termination reduces DC power, but increases switching power
- At 2X data-rate, switching power doubles, but DC power remains constant
- Assumption : 4.7pF total capacitance, 1.2V VDD
- DBI AC can reduce switching power
- DBI DC can reduce termination power

Termination Values (RS and RT)

- Memory system channel impedances
  - Tend to be near 50ohm standard adopted by test equipment and high-speed serial links
  - Package impedances tend to be lower than 50ohms
  - Lower impedances lead to higher bandwidth (Z_0C_i product)
- Matched termination values
  - Best signal integrity
  - Higher power than the alternative
- Under-terminated values
  - Higher voltage swings (low RS, high RT)
  - Lower power (high RT)
  - Increased reflections and more management of SI

System Level Clocking

- System timing synchronous to CK generated by CPHY
  - Used directly for CA bus which is uni-directional
  - Source of internal DRAM clock (DCLK)
- Read/writes use a bi-directional, differential strobe (RDQS/WDQS)
  - Improved source jitter tracking
- Both CK and strobes run at half the DQ transfer rate (DDR)
- Many critical timing paths to consider that could limit performance
- More timing domains introduced as BW increases
  - DDR data clock (1.6GHz for 3.2Gbps) higher than most ASIC internal clocks

System Level Clocking (cont.)

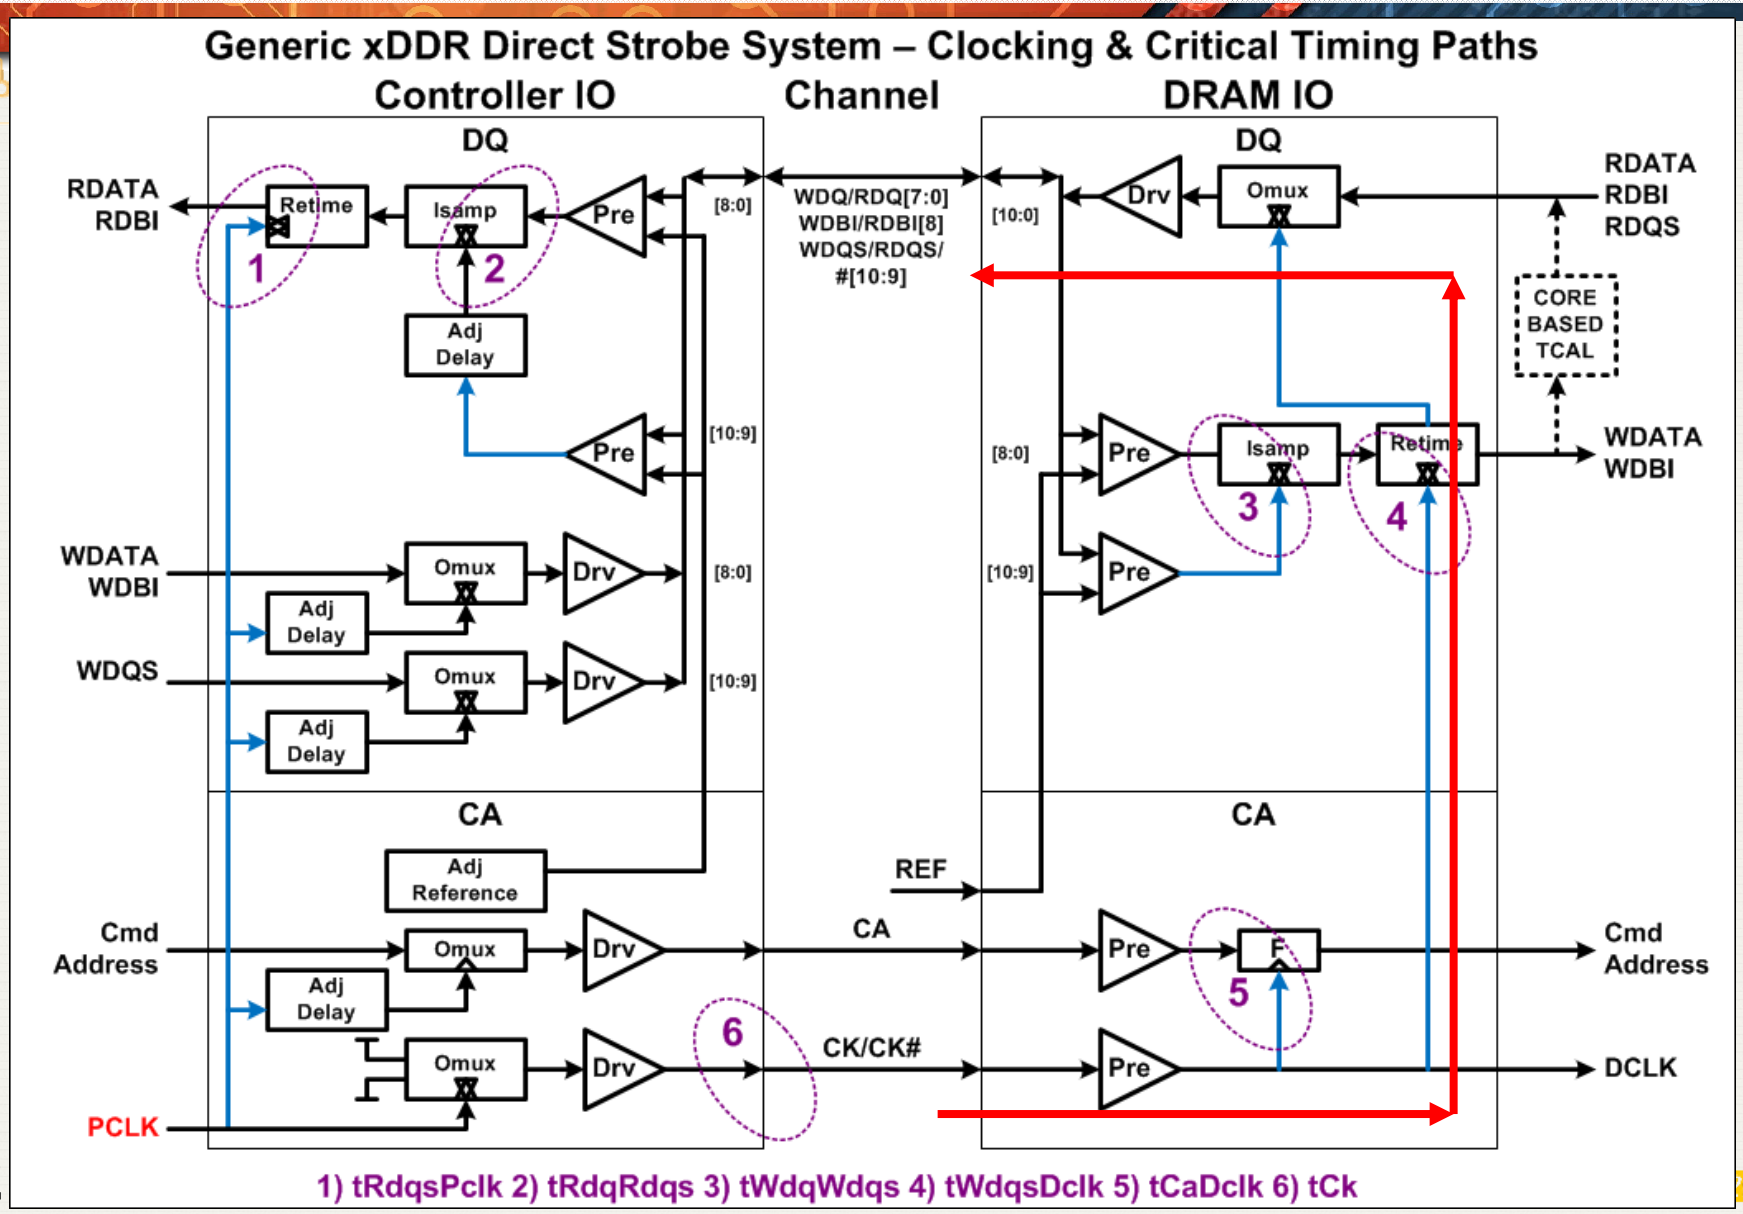

RDQS to CK uncertainty

- DDR: uses a DLL on DRAM to phase align the Read strobe, RDQS, and data, RDQ, with incoming CK
- LPDDR: foregoes the DLL to reduce power and DRAM wakeup time
  - Increased variability of read data return with respect to fixed read latency (tDQSCK)
  - Requires tracking and gating of the strobe signal by the CPHY to accommodate increased drift
  - Net: increased controller complexity and another system level timing budget item

Why is a DLL used in DRAM?

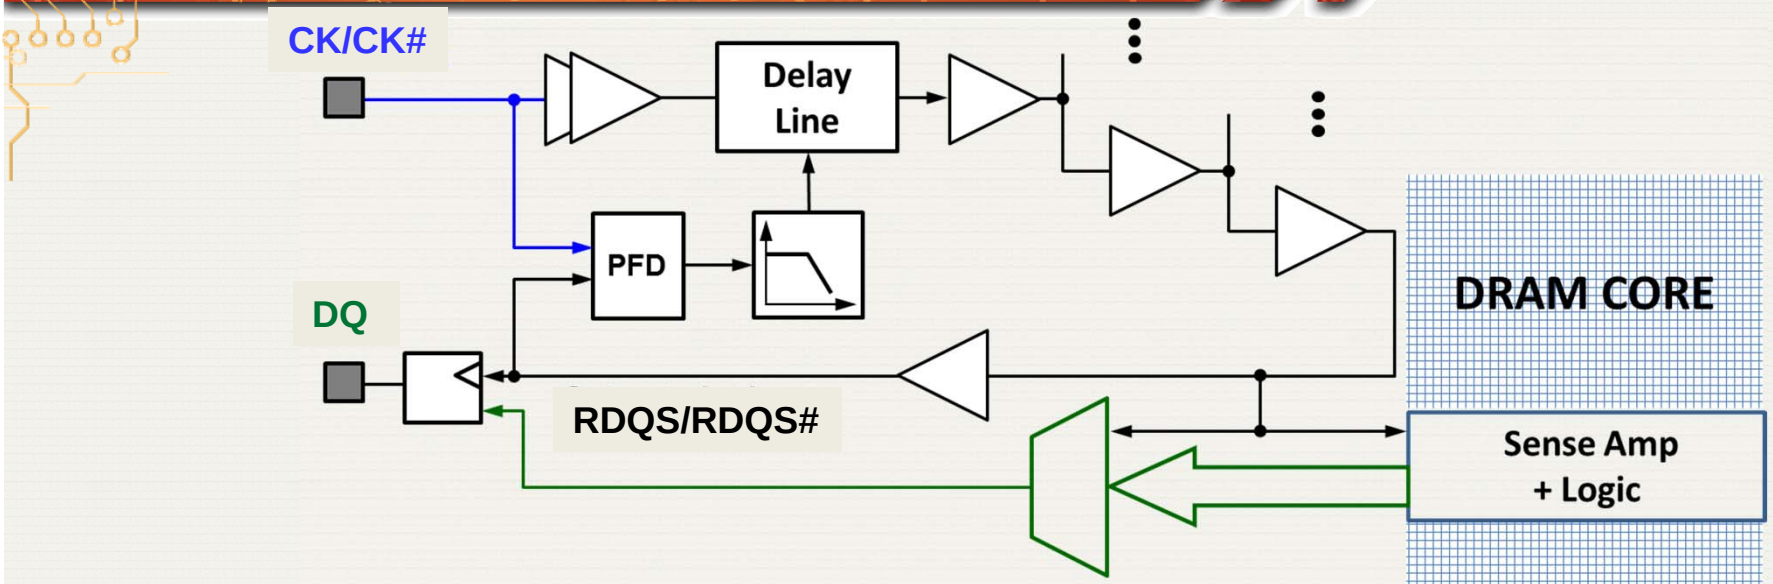

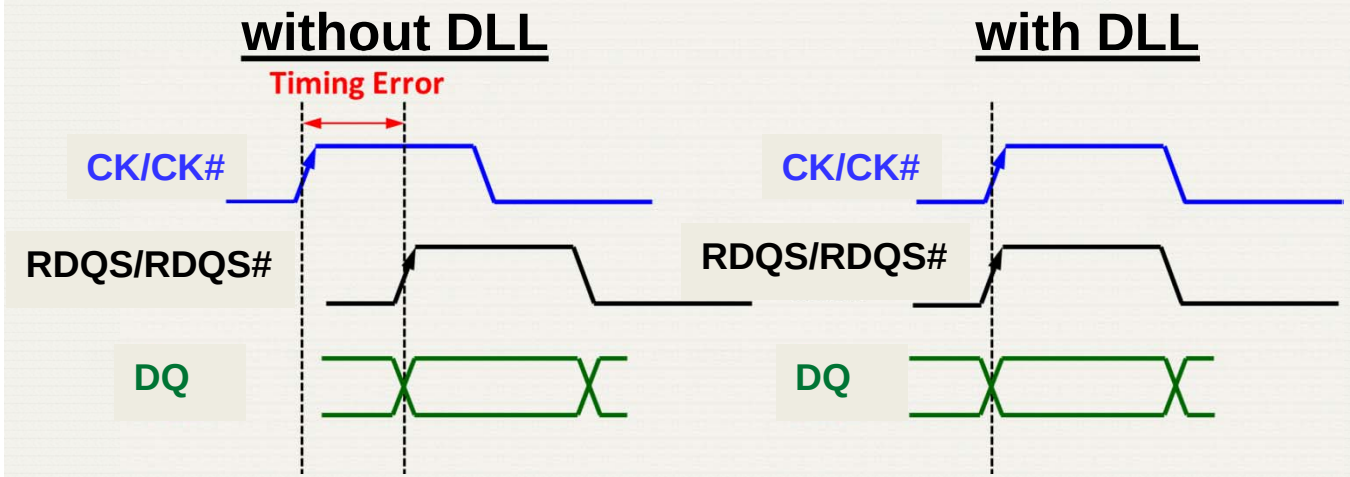

Compensates static and dynamic timing error between CK and WDQS

Internal Clock Generation and Distribution

- Every data bit needs to be clocked, often more than once
  - Amortize clock gen. across links, but still need to distribute
  - Clock has 100% activity factor
- Double-edge clocking
  - Efficient: clock frequency = ½ data rate
  - Practical: easier to trim than multi-phase clocks (just DCC)
- CMOS or CML Clocking?
  - CMOS: often preferred at moderate rates
    - Simple, works with CMOS gates and sense-amplifiers ☺
    - Supply-induced jitter (PSIJ) is bad, *not improving with process scaling* ☹
  - CML: << supply sensitivity ☺ → common at very high speeds
    - Active power not as bad as its reputation (similar to CMOS) 😐
    - Frequency scalability, clock gating are more challenging ☹
    - CML-compatible TX/RX front-end (or convert to CMOS) ☹
  - Configurable for flexibility:

K. Sohn et al, ISSCC 2012 Paper 2.1 (Samsung DDR4 DLL)

Standards Comparison Summary – Signaling & Clocking

	Signaling	VDDQ	RS	RT†	Clocking	Additional Remarks
DDR3	SSTL	1.5, 1.35, 1.25	34, 40, 48	20, 30, 40, 60, 120	DLL; SDR CA; DQ + CA eye training	
DDR4	PODL	1.2, 1.05	34, 40, 48	34, 40, 48, 60, 80, 120, 240	DLL; SDR CA; DQ + CA eye training	
LPDDR2	HSUL	1.2	NA	NA	No DLL; DDR CA; DQ eye training	
LPDDR3	HSUL/ PODL	1.2	NA/ 34, 40, 48	NA/ 120, 240	No DLL; DDR CA; DQ + CA eye training	Termination required for speeds > ~1600Gbps
LPDDR4	LVSTL	1.05-1.2?	?	?	No DLL; SDR CA; DQ + CA eye training	Based on public statements

†Multi-rank systems tend to split RT termination

Outline of Tutorial

- DRAM review [John]
- Conventionally packaged interfaces [John]
  - Evolution of standard interfaces
  - Signaling and clocking at peak bandwidth
  - Bandwidth variation
- 3D Packaging for Memory Applications [Ming]
  - Traditional 3D Packaging Technologies
  - New TSV Based 3D Packaging Technologies
- Signal and power Integrity [Wendem]
  - Signal and power integrity challenges
  - Memory interface analysis methodology
  - Main features of memory interfaces
  - Comparison of emerging standards-based solutions
- Conclusion and Final Comparison [John]

"Energy Proportional Computing"

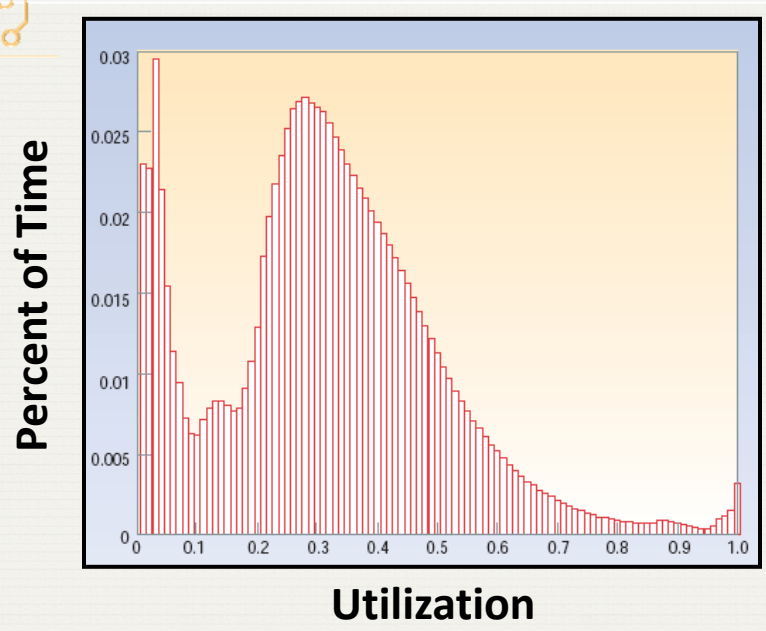

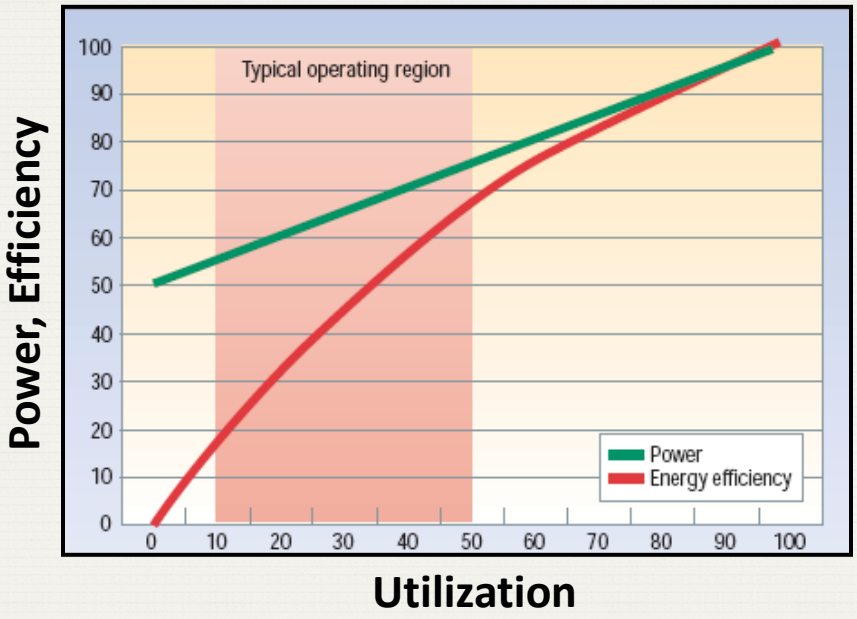

- Average efficiency often worse than peak efficiency
  - And more important!
- Utilization *varies with time*
  - Need to consider dynamics

Luiz Barroso & Urs Hölzle, IEEE Computer, December 2007

Efficiently Dealing with BW Dynamics

- May trade peak efficiency for average efficiency
  - Avoid efficient but fixed constant power techniques

Memory Interface Properties:

- Net bandwidth varies by orders of magnitude
- Time scales also vary by orders of magnitude:
  - Application launch / state change (~100 millisecond)
  - User response (~10 millisecond)
  - CPU context switch (~microsecond)
  - Random cache miss (~10 nanosecond)
- What is best approach to variable bandwidth?

Dynamic Bandwidth & Power Scaling

- Variable Bit-Rate (DVFS)†
  - Better than proportional power scaling
    - But slow response for VDD scaling, reconfiguration
    - Frequency scaling alone hurts efficiency
  - Complicated design & test, limited range (e.g., 3-5×)
  - Alt: Scaling # of active links has some similar challenges

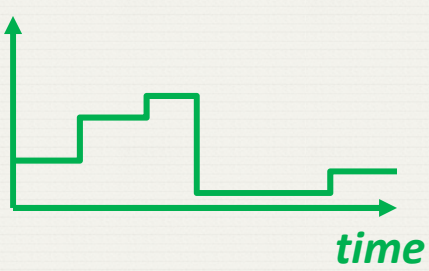

- **Burst Transactions at Peak Rate**
  - Suited to very fast, wide dynamic shifts
  - Optimize link at one data rate
  - Common in current mobile memory systems
  - *Key metric:* wakeup time vs. burst length

† ***G. Balamurugan, et al., JSSCC, April 2008*** is a good example

Dynamic BW Memory Systems

- Support for burst-mode (clock stop mode) present in LPDDR
- All memory systems implicitly support frequency scaling (both half-rates and non-integer frequency hopping)
  - Only big difference is the transition time supported by the DRAM and the controller
- So far, no DVFS
  - Scaling of IO voltages on DRAM would be significant change

Burst Mode Challenges in High Speed Links

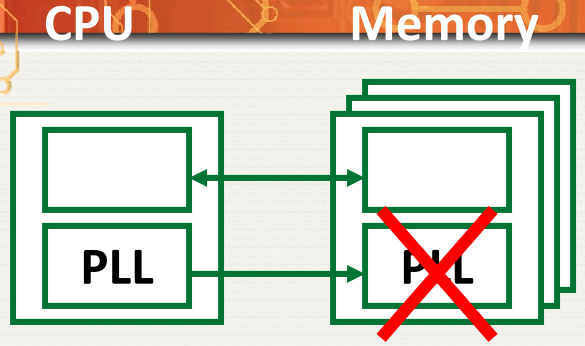

PLL/DLL:

- → Typically slow startup
- → Eliminate, at least on DRAM

Resonant Clocking:

- → Can't synchronously gate resonant node
- → Complex downstream clock gating

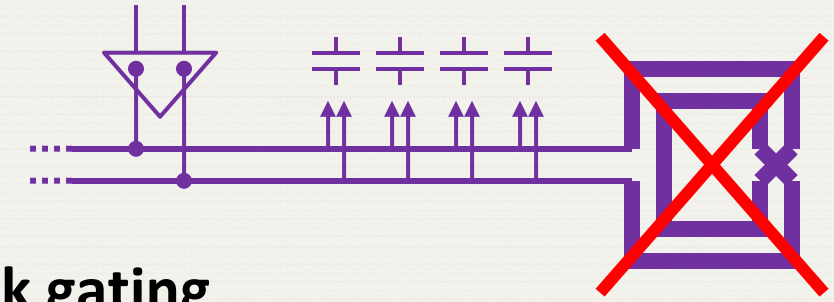

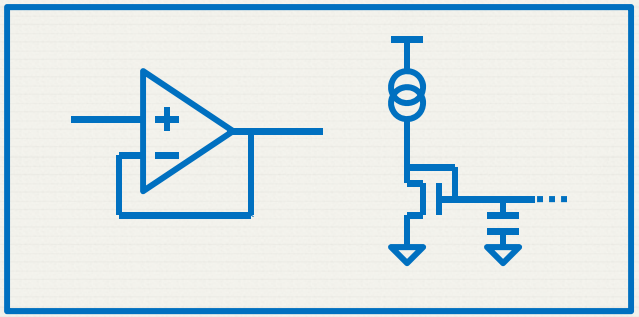

Analog Circuits:

- → Minimize Use
- → Fast startup when needed (TX regulator, CML, etc.)

Example: 4.3 GB/s Mobile Memory Interface w/ Power Efficient BW Scaling

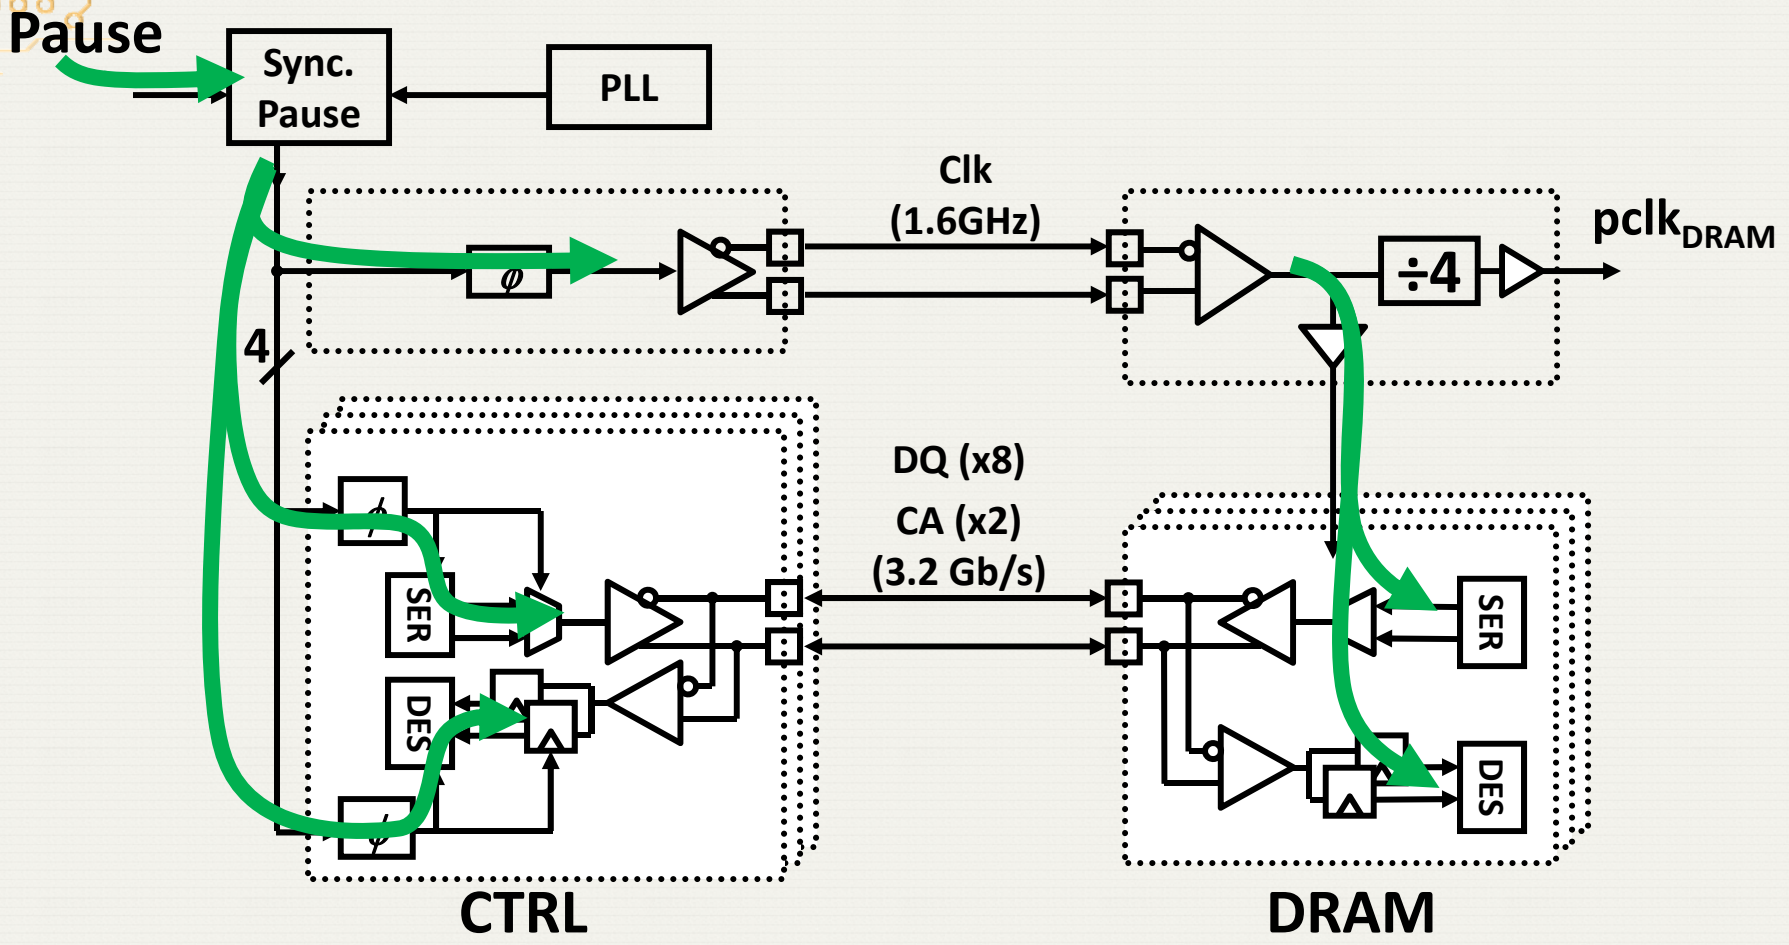

R. Palmer, et al, VLSI 2009

Conventionally Packaged Interface Summary

- Bandwidth and efficiency must continue to improve
  - Average or effective power more relevant than peak power
- Active research area for conventionally packaged interfaces:
  - Increasing use of equalization & crosstalk cancellation
  - Low-swing signaling with RX offset calibration
  - Remove expensive local delay matching buffers
  - Faster wakeup times with zero idle power
  - DVFS for slower dynamics and “drowsy” modes

Outline of Tutorial

- DRAM review [John]
- Conventionally packaged interfaces [John]
  - Evolution of standard interfaces
  - Signaling and clocking at peak bandwidth
  - Bandwidth variation
- 3D Packaging for Memory Applications [Ming]
  - Traditional 3D Packaging Technologies
  - New TSV Based 3D Packaging Technologies
- Signal and power Integrity [Wendem]
  - Signal and power integrity challenges
  - Memory interface analysis methodology
  - Main features of memory interfaces
  - Comparison of emerging standards-based solutions
- Conclusion and Final Comparison [John]

Traditional 3D Packaging for Memory Application – System-in-Package (SiP)

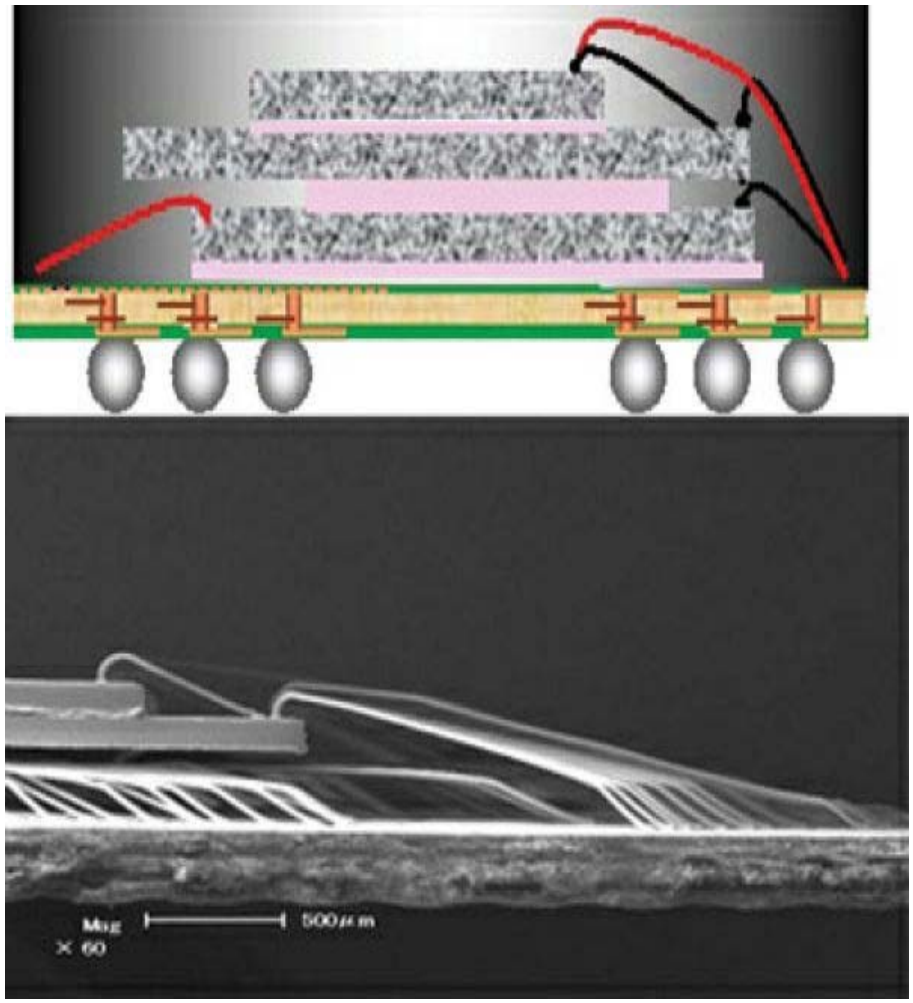

(Source: Internet)

- System-in-Package (SiP)
  - Die stacking using wire bond (WB) interconnects
  - Same memory die stacking for memory capacity (NAND or DRAM)
  - Logic/memory die stacking for memory bandwidth (BW)
  - Low cost solution
  - Limited BW scalability due to limited WB interconnect density
  - Know-Good-Die (KGD) concern

Traditional 3D Packaging for Memory Application – Package-on-Package (PoP)

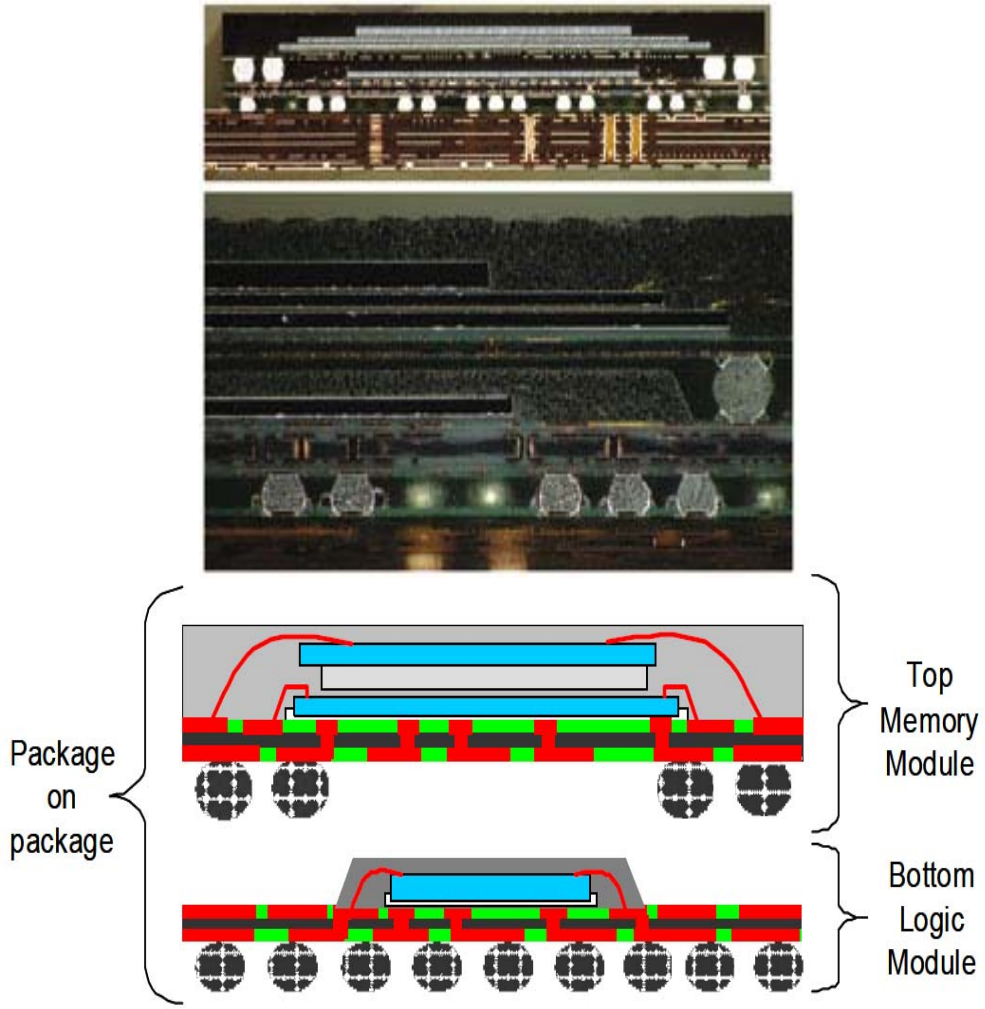

(Source: Internet)

- Package-on-Package (PoP)
  - Top memory package
  - Bottom logic package
  - Logic and memory packages separated for logistic control and KGD
  - JEDEC defined standard memory interface between top and bottom
    - Recently more non-standard approaches in Apple and Samsung products
  - Limited BW scalability due to limited interface memory bus width
    - Scaing with memory technology scaling, LPDDR2 -> LPDDR3 -> LPDDR4
  - Mainly for mobile application
    - Dominating in high end smart phones

Memory PKG Trends for Smartphone Application

PoP for middle to high-end, MCP (SiP) for low to middle-end

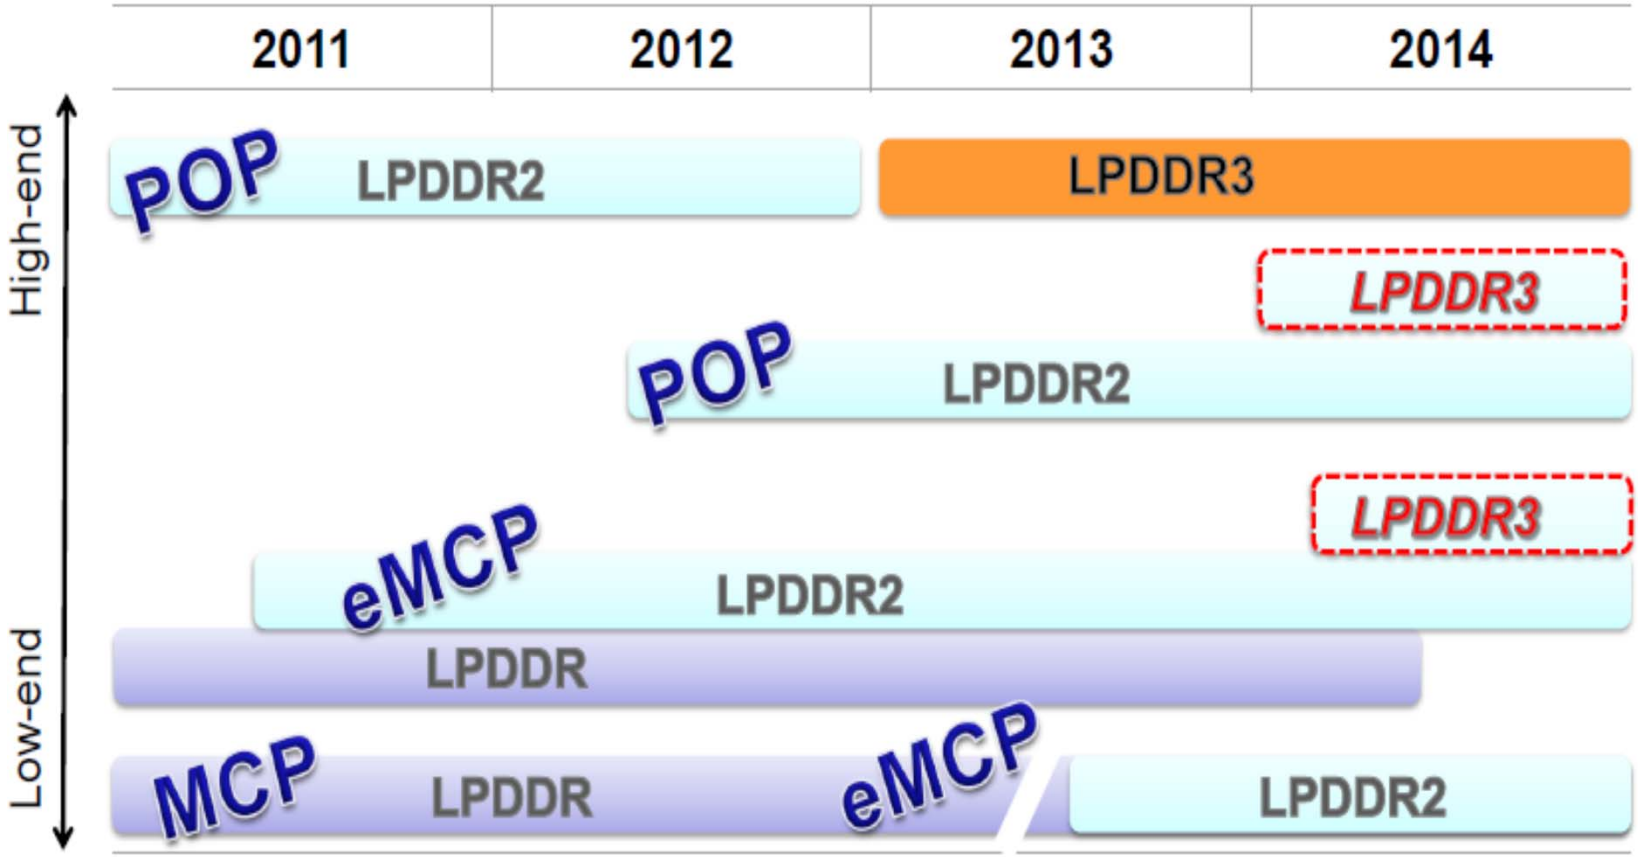

MCP = SLC NAND and Mobile DRAM
eMCP = eMMC and Mobile DRAM
POP (Package on Package) = Mobile DRAM

(Minho Kim, Sk hynix, SemiCon West 2013)

PoP Basic Parameters

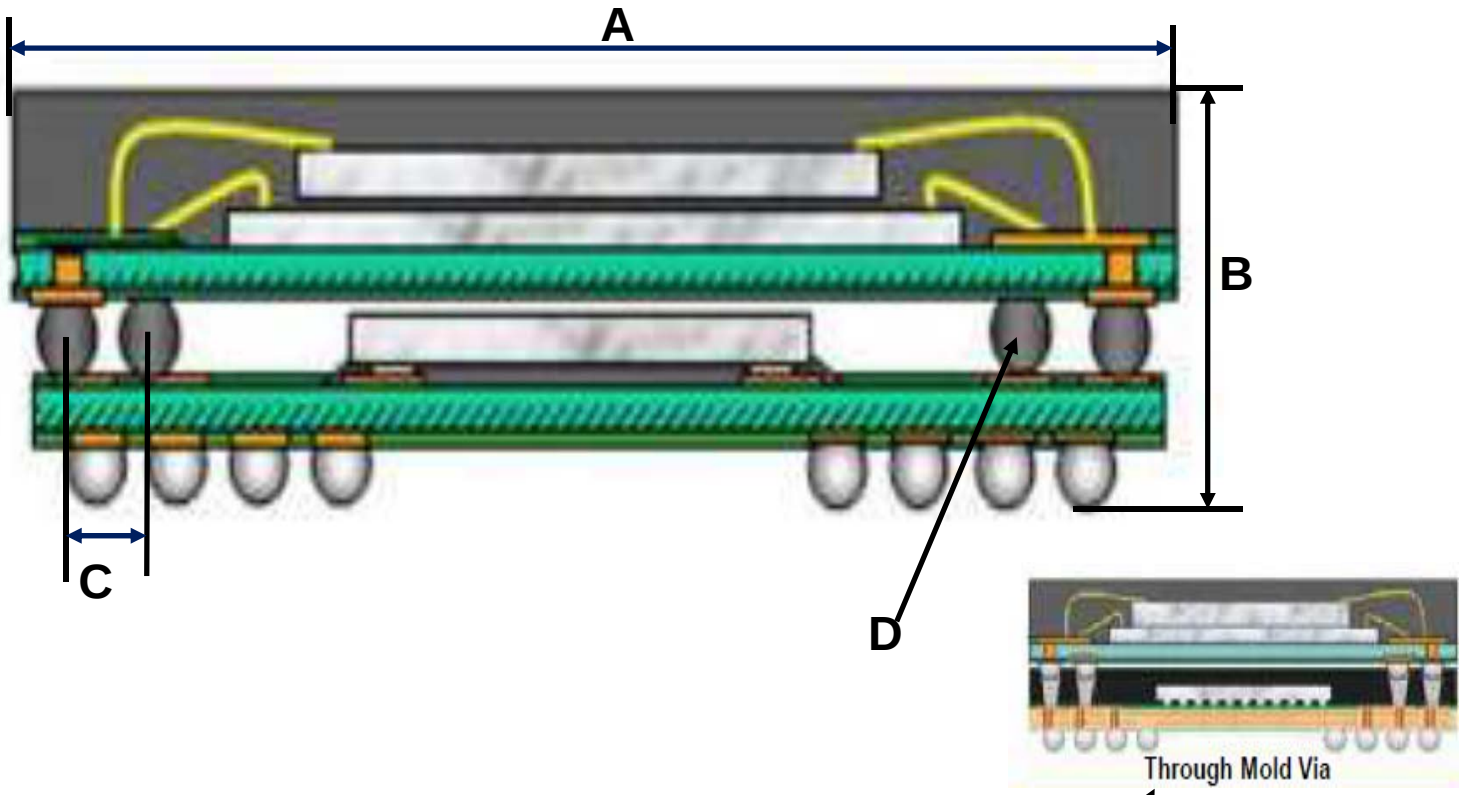

A – PoP Pkg size, JEDEC defined 12 or 14 mm
B – PoP total height, trends from 1.4 -> 1.2 -> 1.0 mm
C – Interface BGA ball pitch, JEDEC defined 0.5 or 0.4 mm, TMV for smaller pitch
D – Interface BGA ball count, most import parameter for BW scaling

JEDEC Defined PoP Ball Count vs. Memory BW

Pkg Size (mm^2)	Interface BGA Ball Pitch (mm)	Interface Ball Count	Memory BW for 2Ch x32 (GB/s)
12 x 12	0.4	216	6.4 for LPDDR2
14 x 14	0.4	256	12.8 for LPDDR3
14 x 14	0.5	240*	6.4 for LPDDR2
15 x15	0.5	216	12.8 for LPDDR3

* Three interface BGA rows, 3rd row depopulated

JEDEC 12 mm/0.4 mm, 216 Ball PoP for 2CH x32 LPDDR2 Ballout

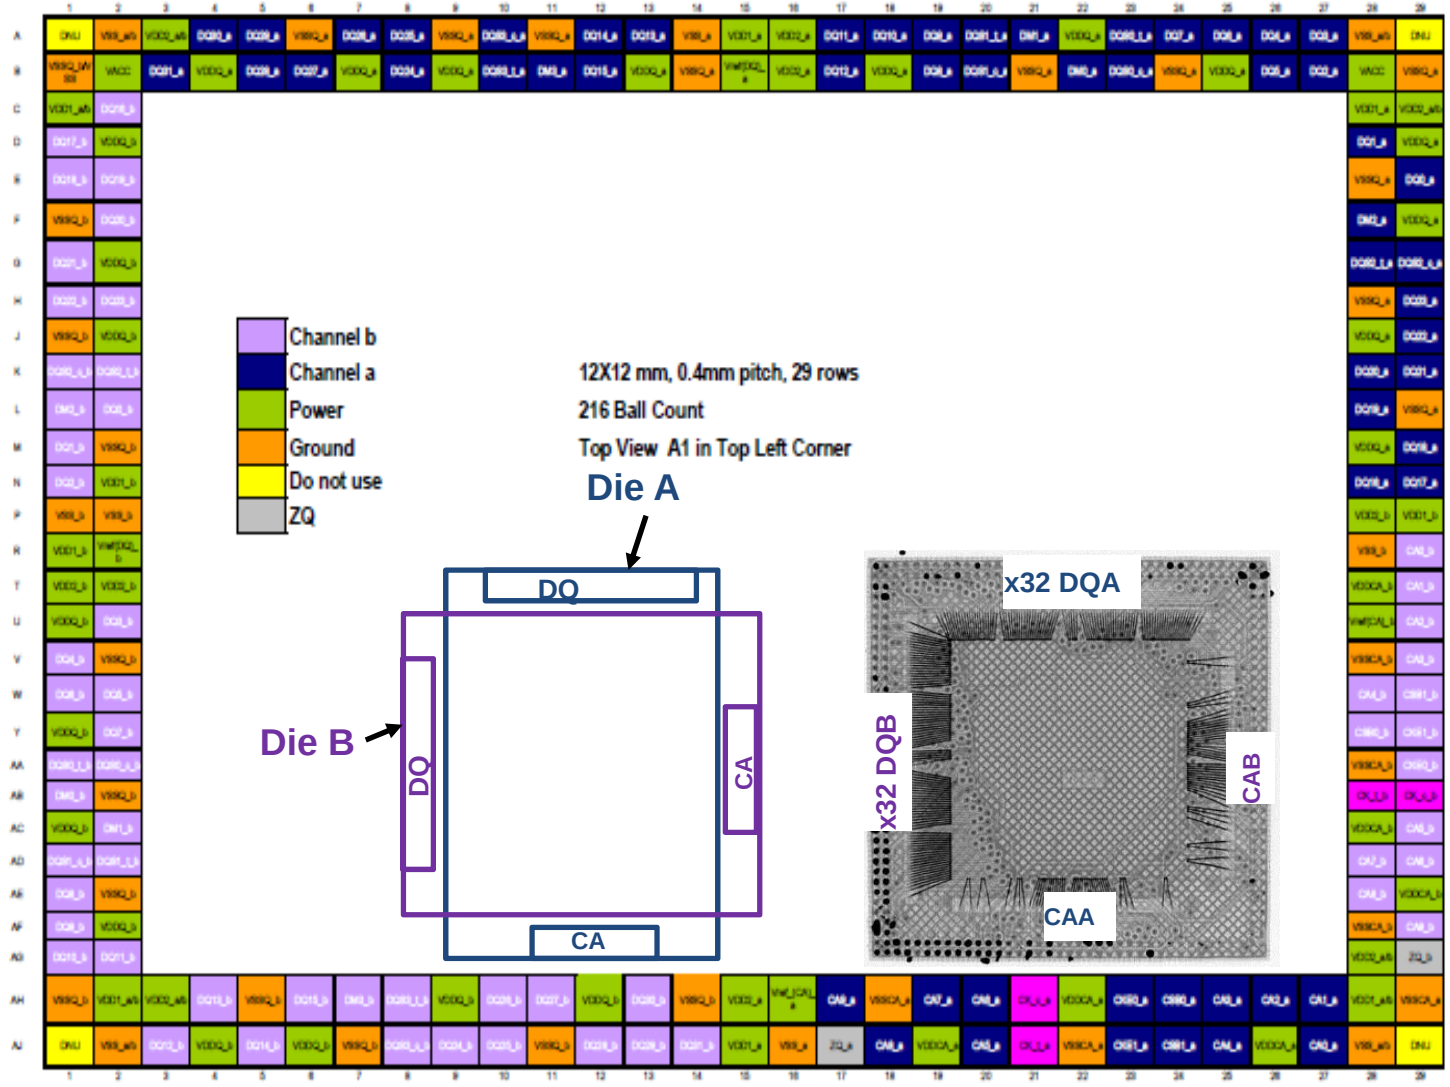

PoP New Trends: Non-JEDEC Formfactor

iPhone 5S (A7 SoC) – 1GB LPDDR3
14x15.5 mm, 456-Ball, 0.35mm pitch

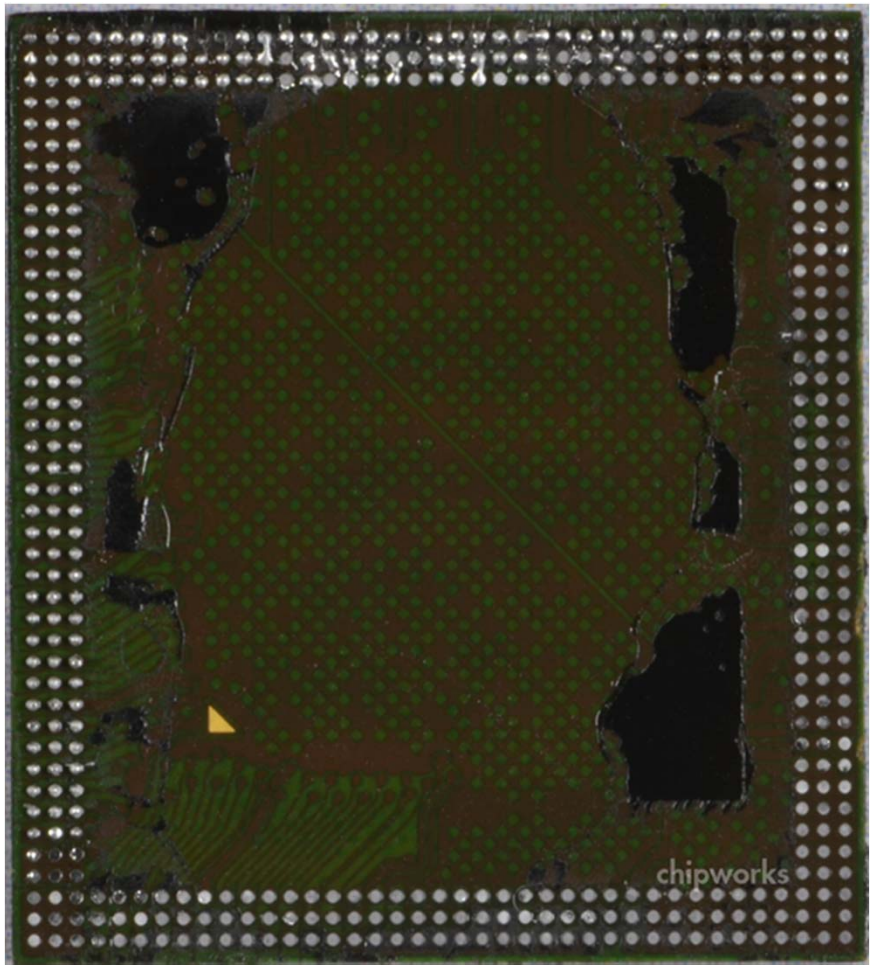

BW = 12.8 - 17 GB/s

iPhone 5 (A6 SoC)– 1GB LPDDR2
14x15.5 mm, 272-Ball, 0.4mm pitch

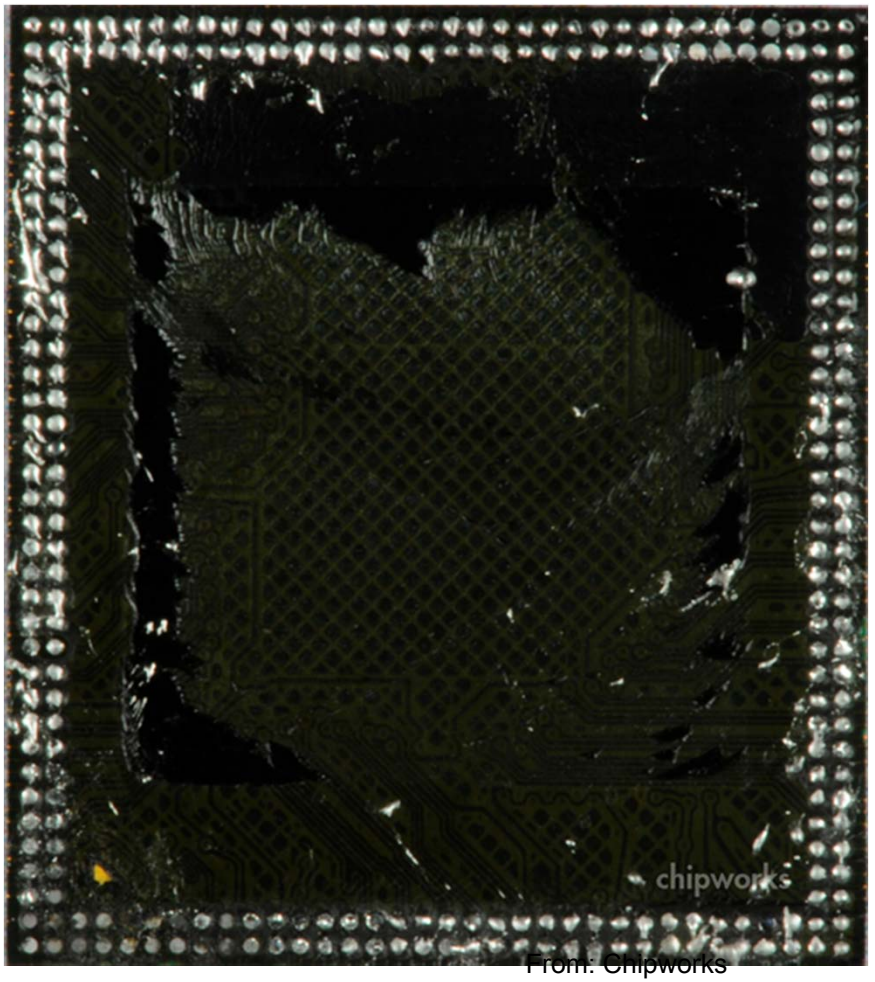

From: Chipworks

BW = 6.4 -8.5 GB/s

Graphics is driving memory BW requirement in mobile systems

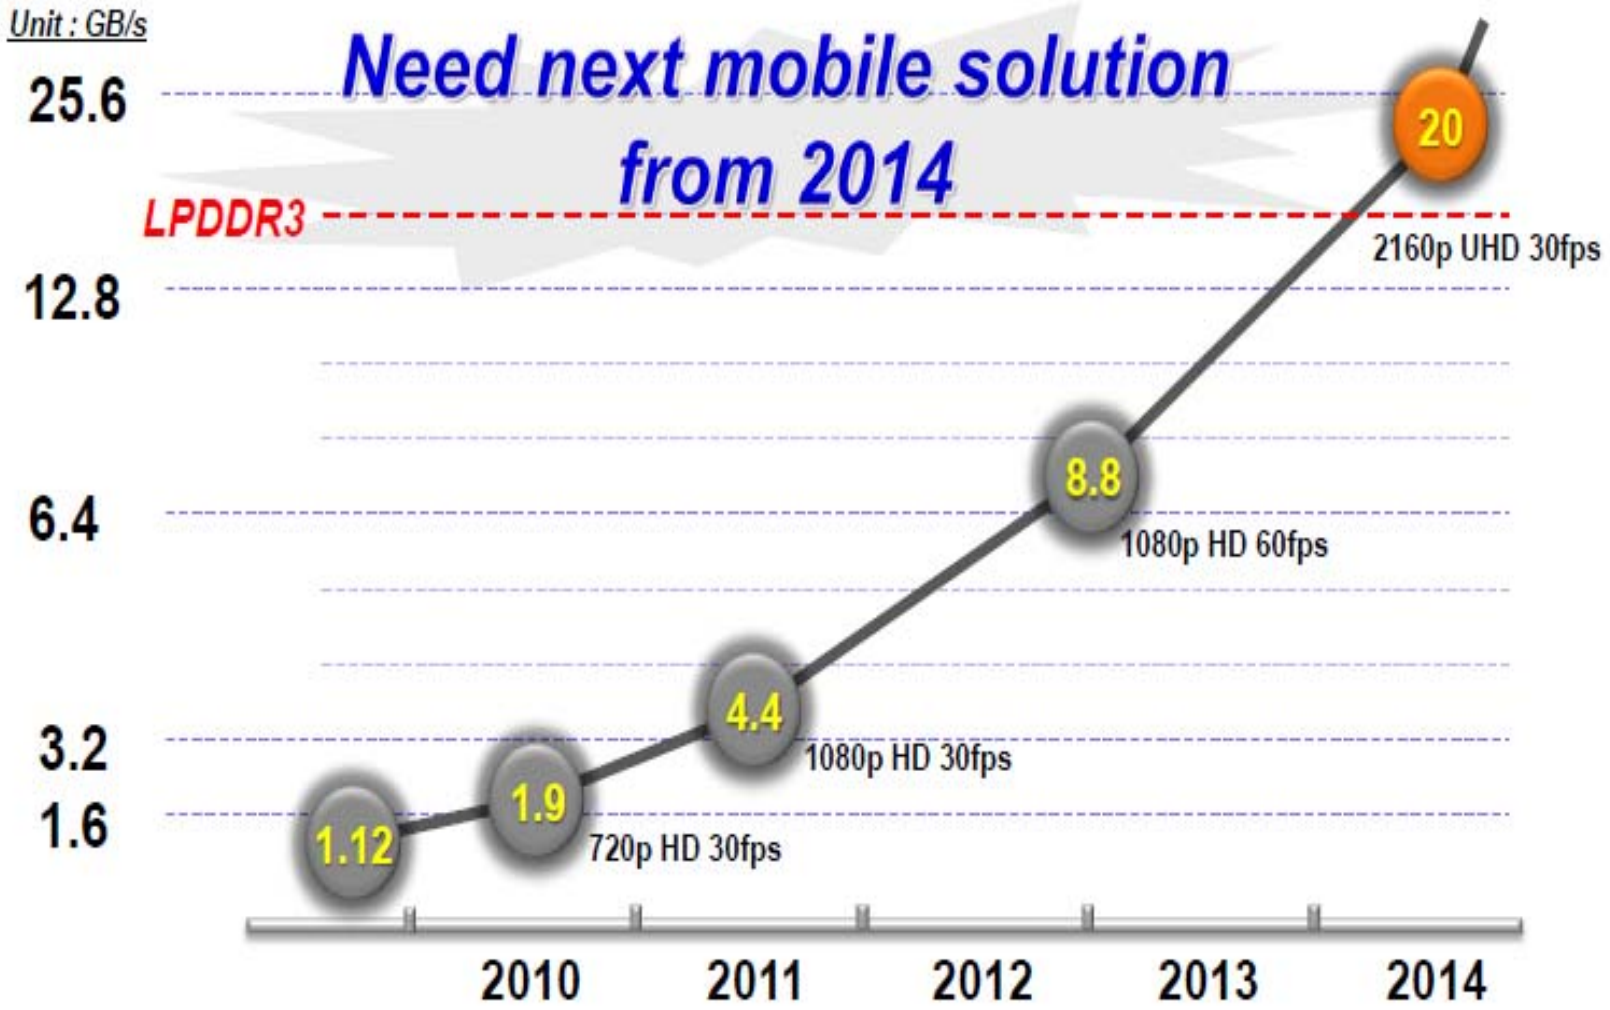

(Minho Kim, SK hynix)

Next Generation Mobile DRAM

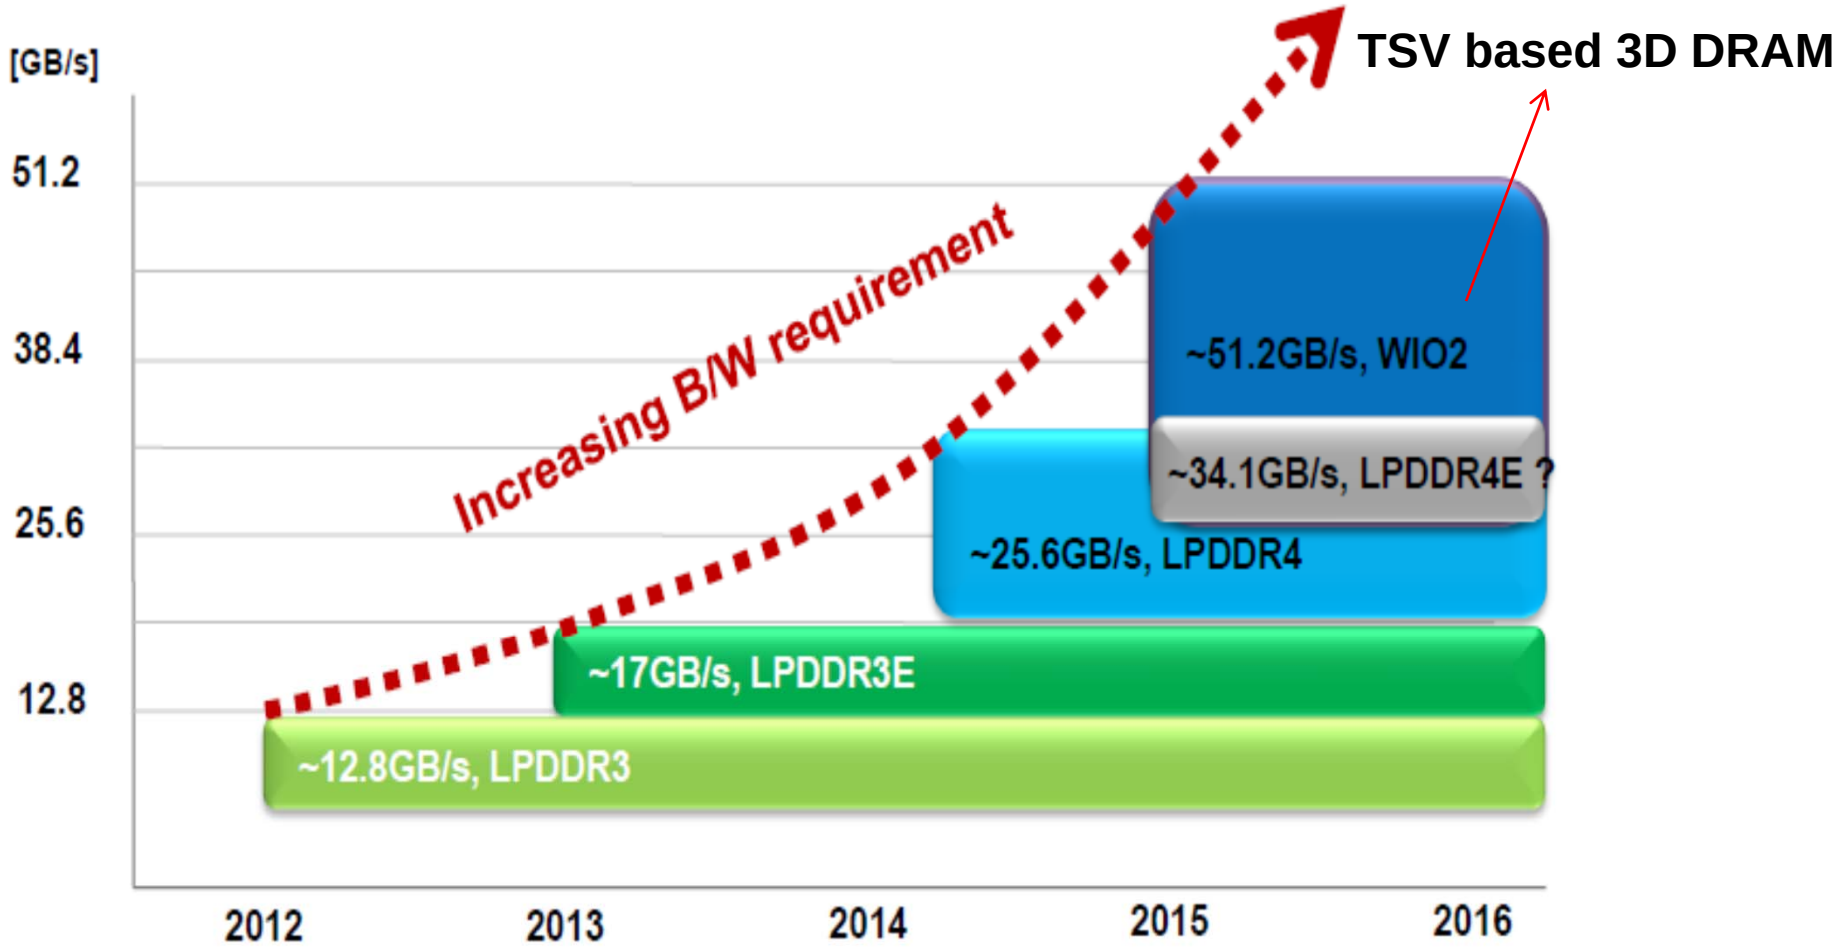

(Minho Kim, SK hynix)

TSV Based 3D Packaging Technology for Memory Application – Wide IO DRAM

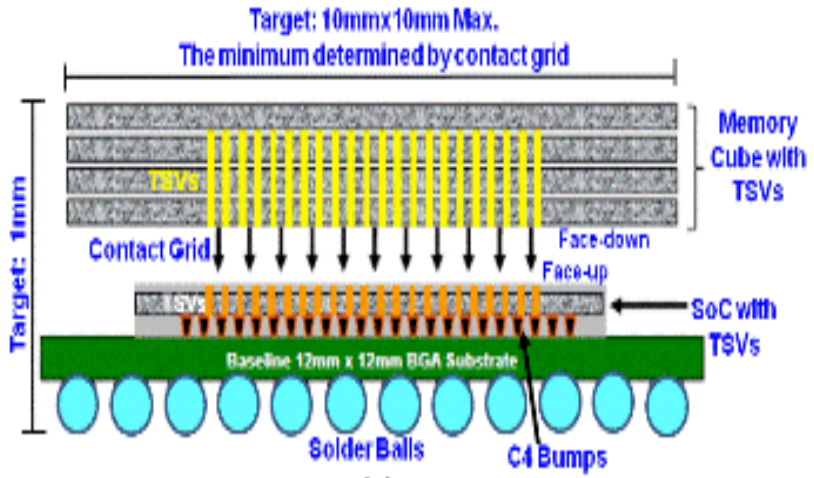

(c)

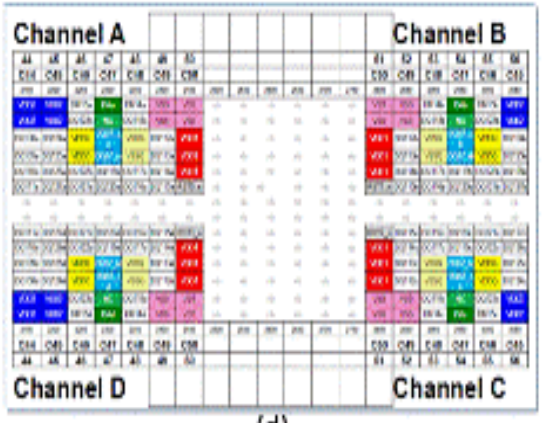

(d)

(Source: Internet)

	WideIO	LPDDR3
Power	1.2 volts	1.2 volts
Speed	200 Mbps	1600 Mbps
Data width	512 bits	32 or 64 bits
Data rate	single	double
Maximum channel bandwidth	12.8 GB/s	@32 bits -- 6.4 GB/s

- **Low power for mobile application**
  - **Much better performance/power, power efficiency than LPDDRx**
- **Small formfactor**
  - **Thin profile and smaller x/y footprint**
- **Better thermal performance than LPDDRx PoP**
- **High cost associated with TSV**
- **Complicate business model**
  - **Yield loss**
  - **Scrap management**

Wide IO2 vs. LPDDR3/4

	LPDDR3 & LPDDR3E	LPDDR4	Wide IO2
Die Organization	1ch X 8 banks X 32 IO 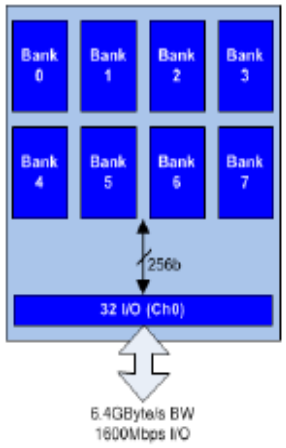	2ch X 8banks X16 IO 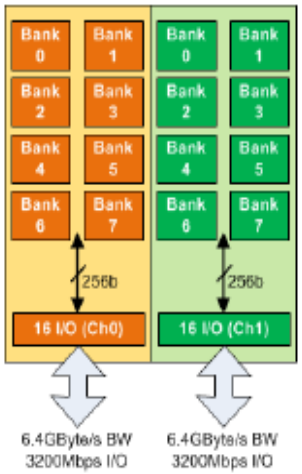	4ch X 8banks X 64 IO 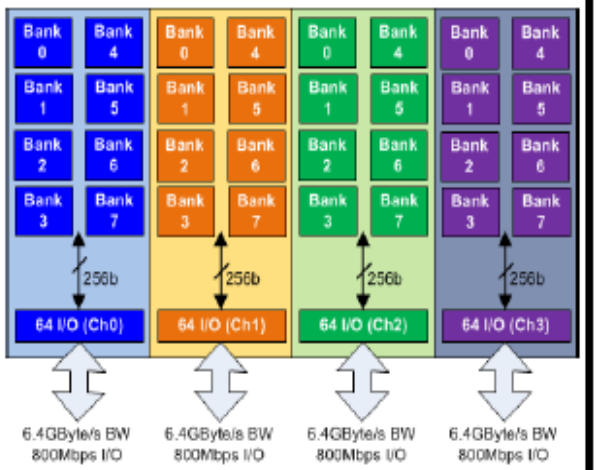
Channel #	1	2	4 & 8
Bank #	8	8 per channel (16 per die)	32 per die
Density	4Gb – 32Gb	4Gb – 32Gb	8Gb – 32Gb
Page Size	4KByte	2KByte	4KByte (4ch die), 2KB (8ch die)
Max BW per die	6.4GB/s, 8.5GB/s (overclocking)	12.8GB/s, 17GB/s (overclocking)	25.6GB/s & 51.2GB/s 34GB/s & 68GB/s(overclocking)
Max IO Speed	2133Mbps	4266Mbps	1066Mbps
Signal Pin #	62 per die	66 per die	~430 per die (4ch die), ~850 per die(8ch die)
Package	POP, MCP	POP, MCP	KGD,

(Hung Vuong, Qualcomm, JEDEC Mobile Forum 2013)

Wide IO vs. PoP

WIO2 takes revolutionary path, showing significant difference than LPDDRx

LPDDRx

- POP Structure -

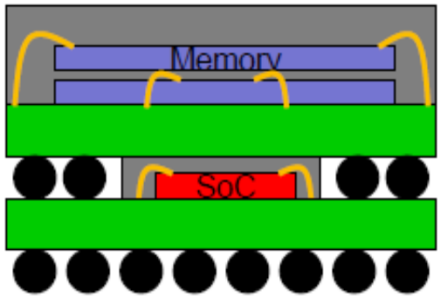

WIO2

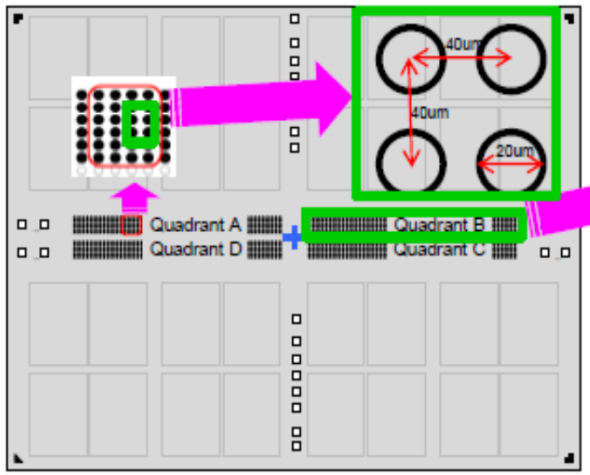

- WIO TSV Technology -

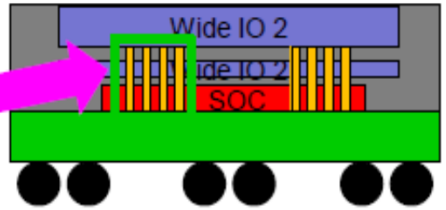

-17-

(Minho Kim, Sk hynix, SemiCon West 2013)

SK hynix DESIGNCON 2014

TSV Based 3D Packaging Technology for Memory Application – Hybrid Memory Cube (HMC)

Fast process logic and advanced DRAM design in one optimized package

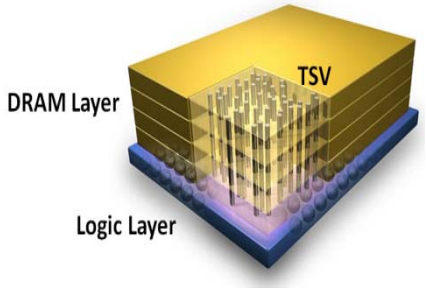

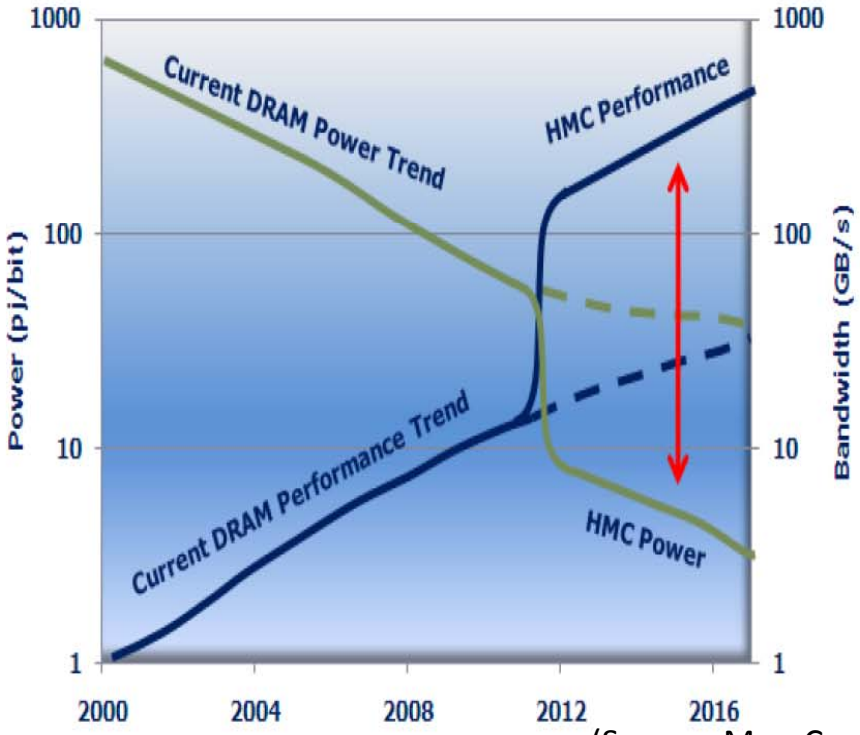

- Power Efficient
- Smaller Footprint
- Increased Bandwidth
- Reduced Latency

(Source: MemCon Memory Conference 2012)

HMC Memory Interface

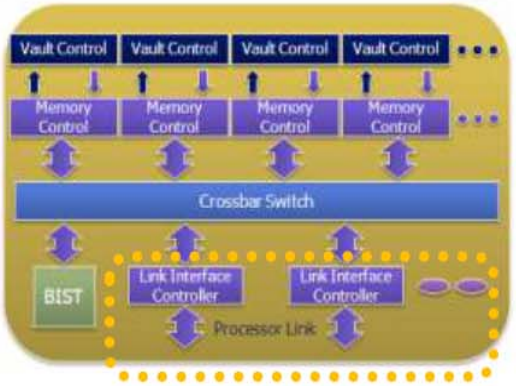

HMC (Gen2) has Four Links

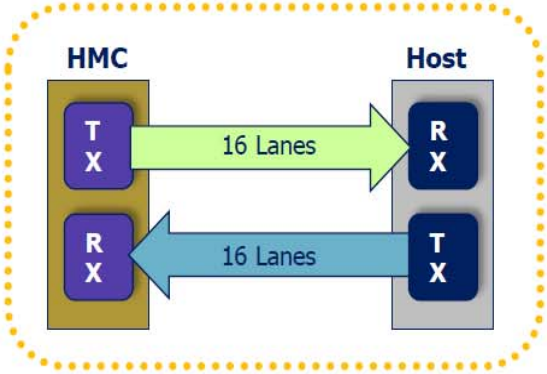

Each Link has 32 differential Lanes

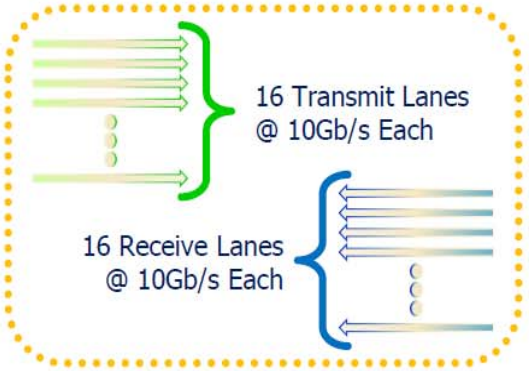

(*) HMC-SR supports up to 15Gbps SERDES options

Designed with Off-the-Shelf, High Speed SerDes Interface

Link Interface Examples

4 Link Example (160GB/s)
- 10Gb/s per lane
- 32 lanes per link (320Gb/s = 40GB/s)
  - 16 TX and 16 RX
- 4 Links (40GB/s x 4) = **160GB/s**

(*) Aggregate DRAM peak bandwidth remains 160GB/s

- Up to 15X the bandwidth of a DDR3 module
- 70% less energy usage per bit than existing technologies
- Occupying nearly 90% less space than today's RDIMMs

(Source: Mike Black, Micron, EDPS 2013)

HMC Applications

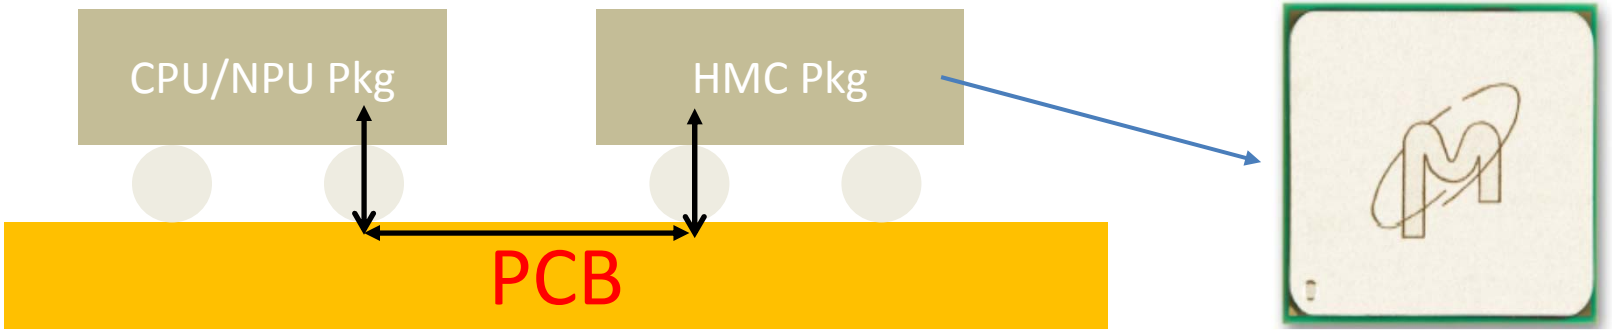

- HMC SR-15 (Gen2) for HPC/Server (CPU) & Network (NPU) Applications
  - Short reach channel with 12 dB of insertion loss
  - 160 GB/s Memory BW
  - 10, 12.5, 15 Gb/s SERDES interface
  - 2/4GB memory density
  - 31 x 31 mm 4-link package
  - 16.5 x 19 mm 2-link package

TSV Based 3D Packaging Technology for Memory Application - High Bandwidth Memory (HBM)

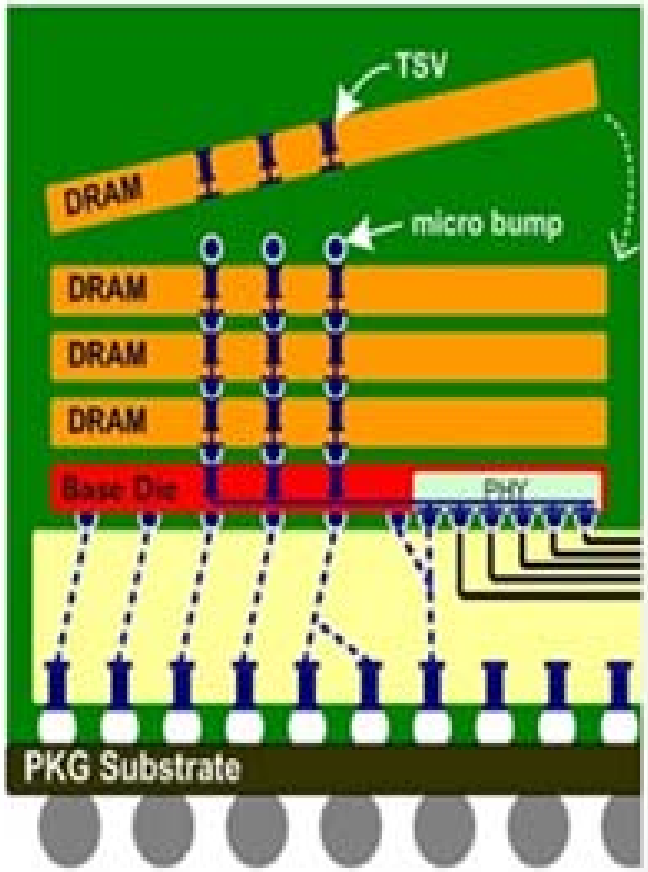

ITEM	TARGET
Burst Length	2, 4
Stack Density	1GByte per stack (2Gbit per slice)
Channel / Slice	2
Banks / Channel	8
IO / Channel	128
Prefetch / Channel	32B (128x2bit)
Channels / Stack	8
Total TSV Data IO Width	1024
Clock Speed	500MHz
Peak Read BW / Stack	128 GB/s
Page Size	2KB
Data Parity	1 bit / 32 bit
DRAM Core Voltage	1.2V
Logic Buffer IO Voltage	1.2V

(Source: SK/Hynix, Jan 2013)

HBM Applications

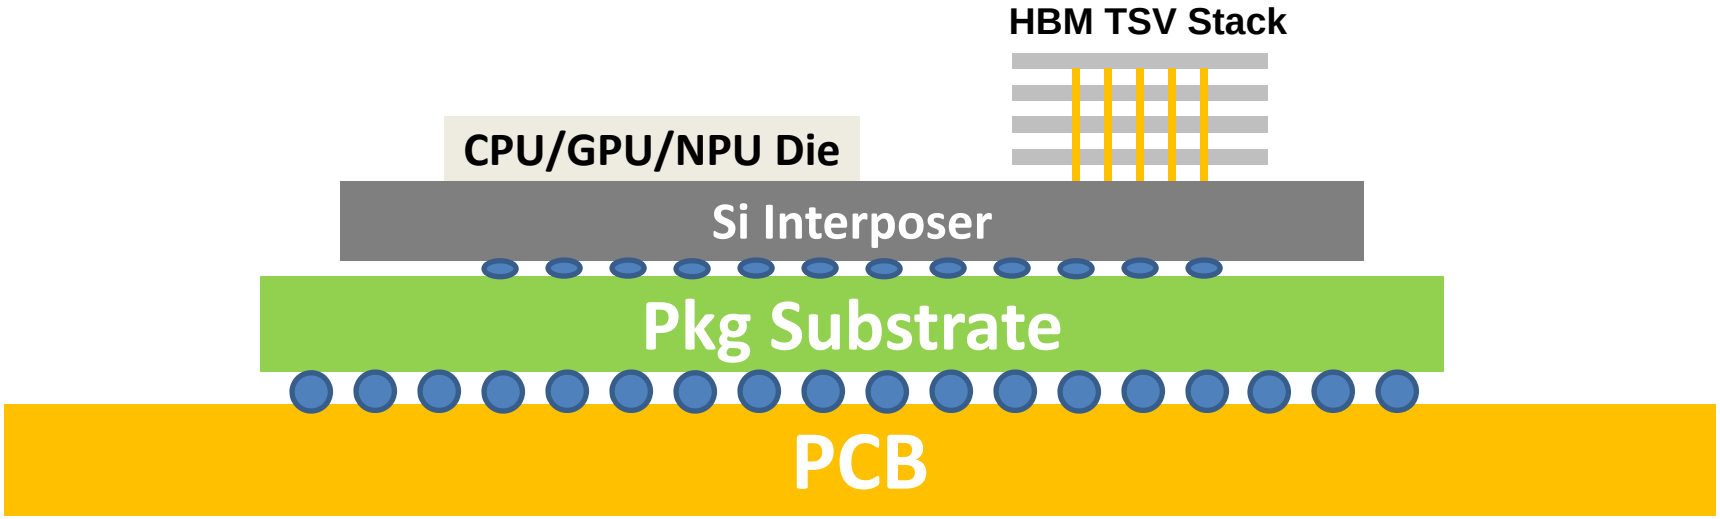

- CPU for HPC applications, GPU for graphic application, and NPU for network application
- 1024 bit bus width, 1 Gb/s signaling (Gen1), DDR memory interface, total 128 GB/s memory BW
- On a Si interposer, need probing on micro-bump

HBM vs. HMC

	HBM	HMC
Max Signaling (Gb/s)	2 (Gen 2)	20 - 30 (Gen 3)
Bus Width (bit)	1024	4 Links (16 lane per link, bi-directional)
Max BW (GB/s)	256	320 - 480
PHY	Memory PHY only	Memory/Controller PHYs in serial links
Format	In a Si Interposer	Stand alone as a complete package, No KGD issue
Standardization	JEDEC	HMC Consortium
Applications	Graphic, HPC, Network	Network, HPC, Server

HBM vs. Wide IO2

ITEM		Mobile WIO2	HBM (High B/W Memory)
		DRAM 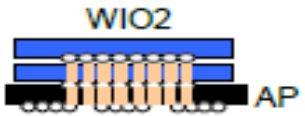	Base die + DRAM 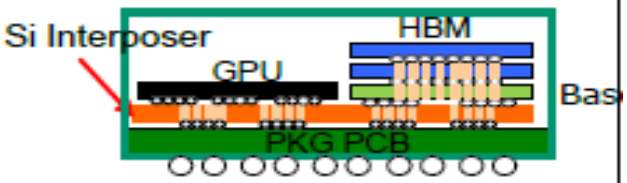
Bottom die		N/A	Buffering & Signal re-routing
BW (GB/s)		25.6~51.2	128~256
Pin	Speed	0.4~0.8 Gbps	1~2 Gbps
	# I/O	512	1,024
#Bump	Logic	1~2K	6K~8K
	DRAM	1~2K	~3K
Cube (GB)		1 / 2	1 / 2 / 4
# TSV stack		1 / 2 / 4	1 / 2 / 4
DRAM density		8Gb	8Gb
Application	GFX card	○	○
	ULT	○	-
	HPC	-	○
	Server	-	○(Cache)
	Mobile	○	-

(Bob Brennan, Samsung, Memcon 2013)

Markets for TSV Based 3D Packaging Technologies

Market	SmartPhone	Tablet	Networking	Graphics
Processor	Apps Processor	Apps Processor	Networking Proc	Graphic Proc
Power	1-2W	1-5W	20W +	20W +
Memory Type	Wide I/Ox	Wide I/Ox or LPDDRx	HMC / HBM	HBM
Memory Size	2 → 4 GB	4 → 8 → 16 GB	4 → 8 → 16 GB	4 → 16 GB
Interface	Wide I/O2	Wide I/O or DDR	SerDes / Parallel	Parallel
I/O	1000	1000 or 500	<500 / >1000	1600 +
Min Bump Pitch	40x40µm rows	40x50µm rows or 80µm rows	1mm / 96x55µm array	96x55µm array, staggered rows
Packaging	3D	2.5D medium L/S density or 3D with heat management	Off chip memory or 2.5D high density interposers	2.5D high density interposers

(Source: Internet)

Acknowledgments

- I would like to acknowledge the Rambus SSE engineers, especially Will Ng and Dave Secker, for contributions and many helpful discussions over the years

Outline of Tutorial

- DRAM review [John]
- Conventionally packaged interfaces [John]
  - Evolution of standard interfaces
  - Signaling and clocking at peak bandwidth
  - Bandwidth variation
- 3D Packaging for Memory Applications [Ming]
  - Traditional 3D Packaging Technologies
  - New TSV Based 3D Packaging Technologies
- Signal and power Integrity [Wendem]
  - Signal and power integrity challenges
  - Memory interface analysis methodology
  - Main features of memory interfaces
  - Comparison of emerging standards-based solutions
- Conclusion and Final Comparison [John]

Challenges of Memory Systems

- Memory data rates have doubled every four years for main memory and every two years for mobile memory at a similar or lower power envelope
- The increased performance and reduced power present very difficult SI/PI challenges
  - Inter-symbol interference (ISI) and crosstalk
  - Impedance discontinuities (stubs, connector, via, etc.)
  - Power supply noise (simultaneous switching output noise)
  - Power supply noise induced jitter
  - Pin-to-pin skew, clock and data jitter
  - Poor transistor performance
- In-depth SI and PI analyses are essential in the design of these power-efficient memory interfaces

Traditional SI/PI Analysis

- Electrical specification of memory devices consists of a set of measured electrical parameters
- The parameters are related to timing jitter, input/output voltage levels and reflect worse case performance of the complete system
  - With linear or mean square sum of voltage and timing budget analysis
  - Time domain based simulations under worse case conditions
- At the higher data rates, the traditional direct frequency and time-domain methods are no longer sufficient to design and optimize
  - Decouples voltage noise and timing jitter and leads to over design

Example of : Timing Budget

Components	WRITE (ps)	READ (ps)
Memory Controller	140	90
DRAM	140	215
Channel	300	275
Others	45	45
Total	625	625

Voltage Budget

Components	WRITE (mV)	READ (mV)
TX Swing	1000	850
Receiver Vin	300	150
Channel	700	700

Each parameter can have deterministic and random components

System-Based Approach

- BER-based method that considers the statistical distributions of both voltage and timing parameters is critical to optimize the link performance
  - to consider the interaction
  - to make trade-off between active circuitries and passive channel
  - to consider other system constraints
  - to predict the effect of the components variations on the performance of the complete system
  - and optimize the system in cost and power
  - to verify that the system meets the targeted low BER.

Forwarded Clocking of Memory Links

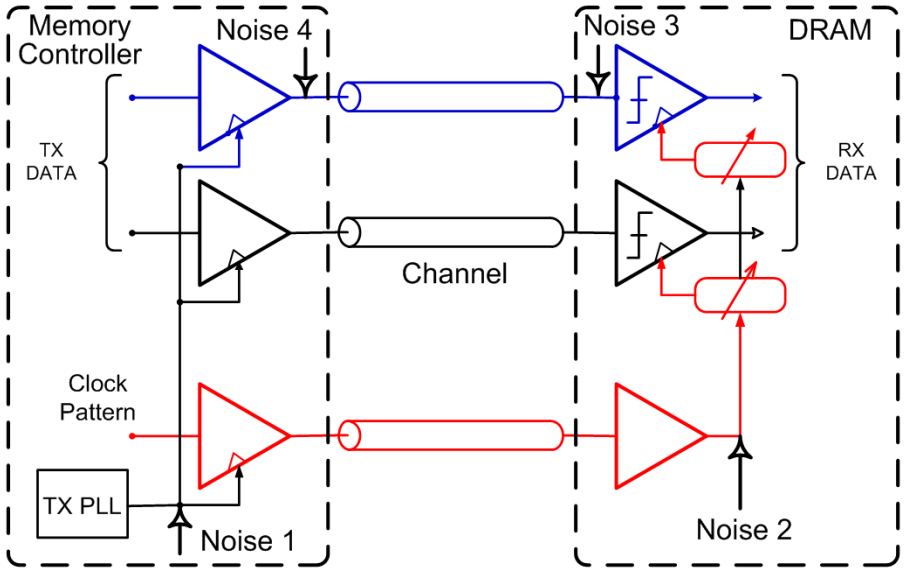

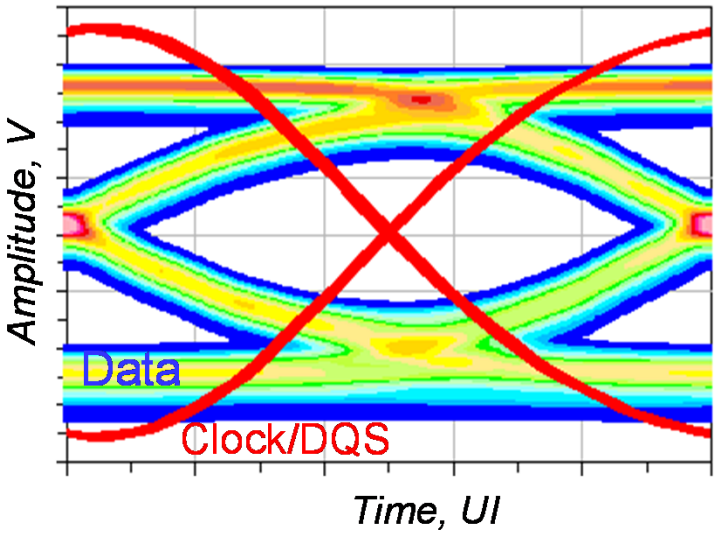

- Majority of memory interfaces are source synchronous links
  - The clock and data start from the same device and travel the same path
  - Often the clock is distributed across a byte or two
- The clock skew is adjusted to the center of the data at the receiver
- Noise 1: Common to DQ and DQS, correlated jitter tracked
  - Uncorrelated jitter is not tracked : delay difference and noise frequency
  - Noise 2, Noise 3 and Noise 4 : Jitter may not be tracked
- Power supply induced jitter is the main source of timing uncertainty

Impedance of Power Distribution Network

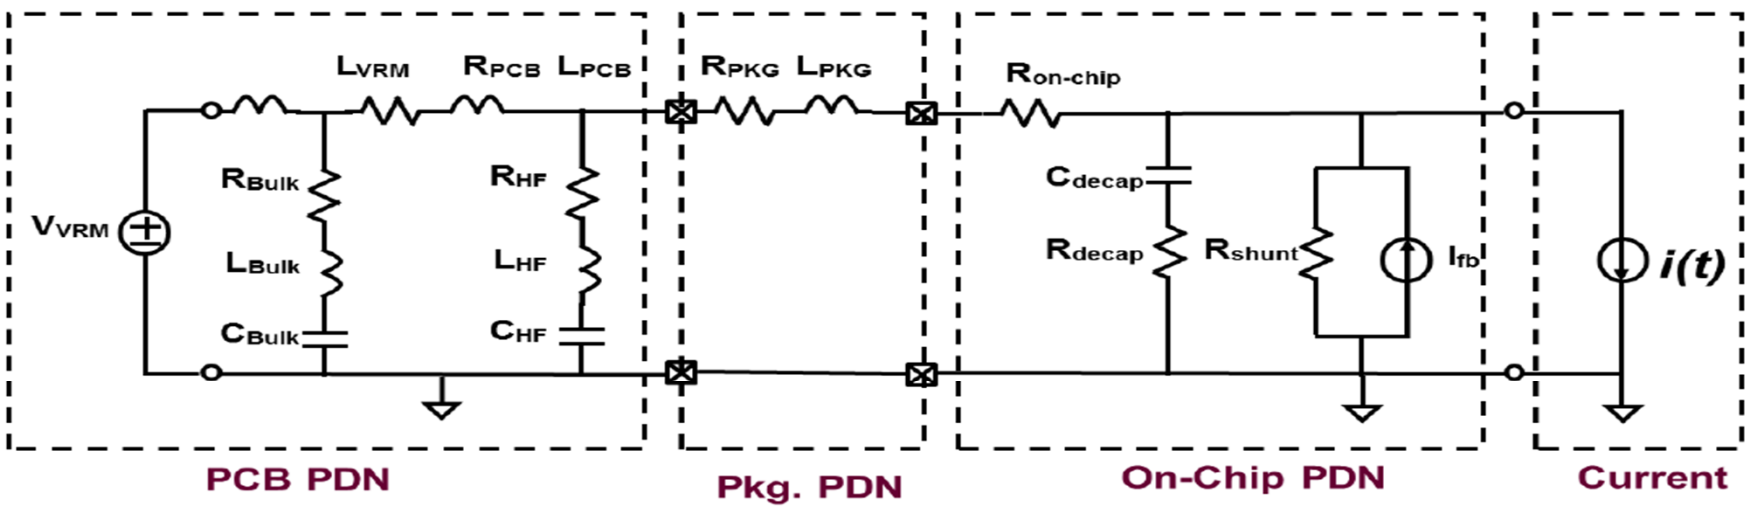

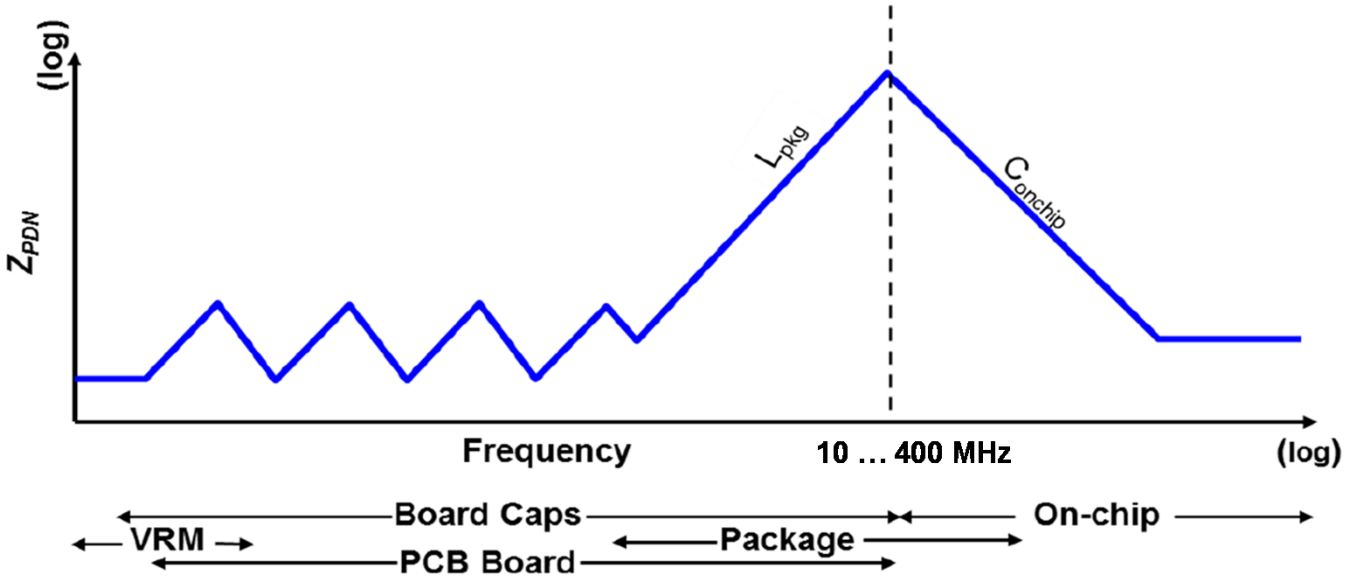

N impedance peaks
1Hz - 400 MHz
aused by variations
nts drawn from the

Supply Rails and PSIJ of Clock Path

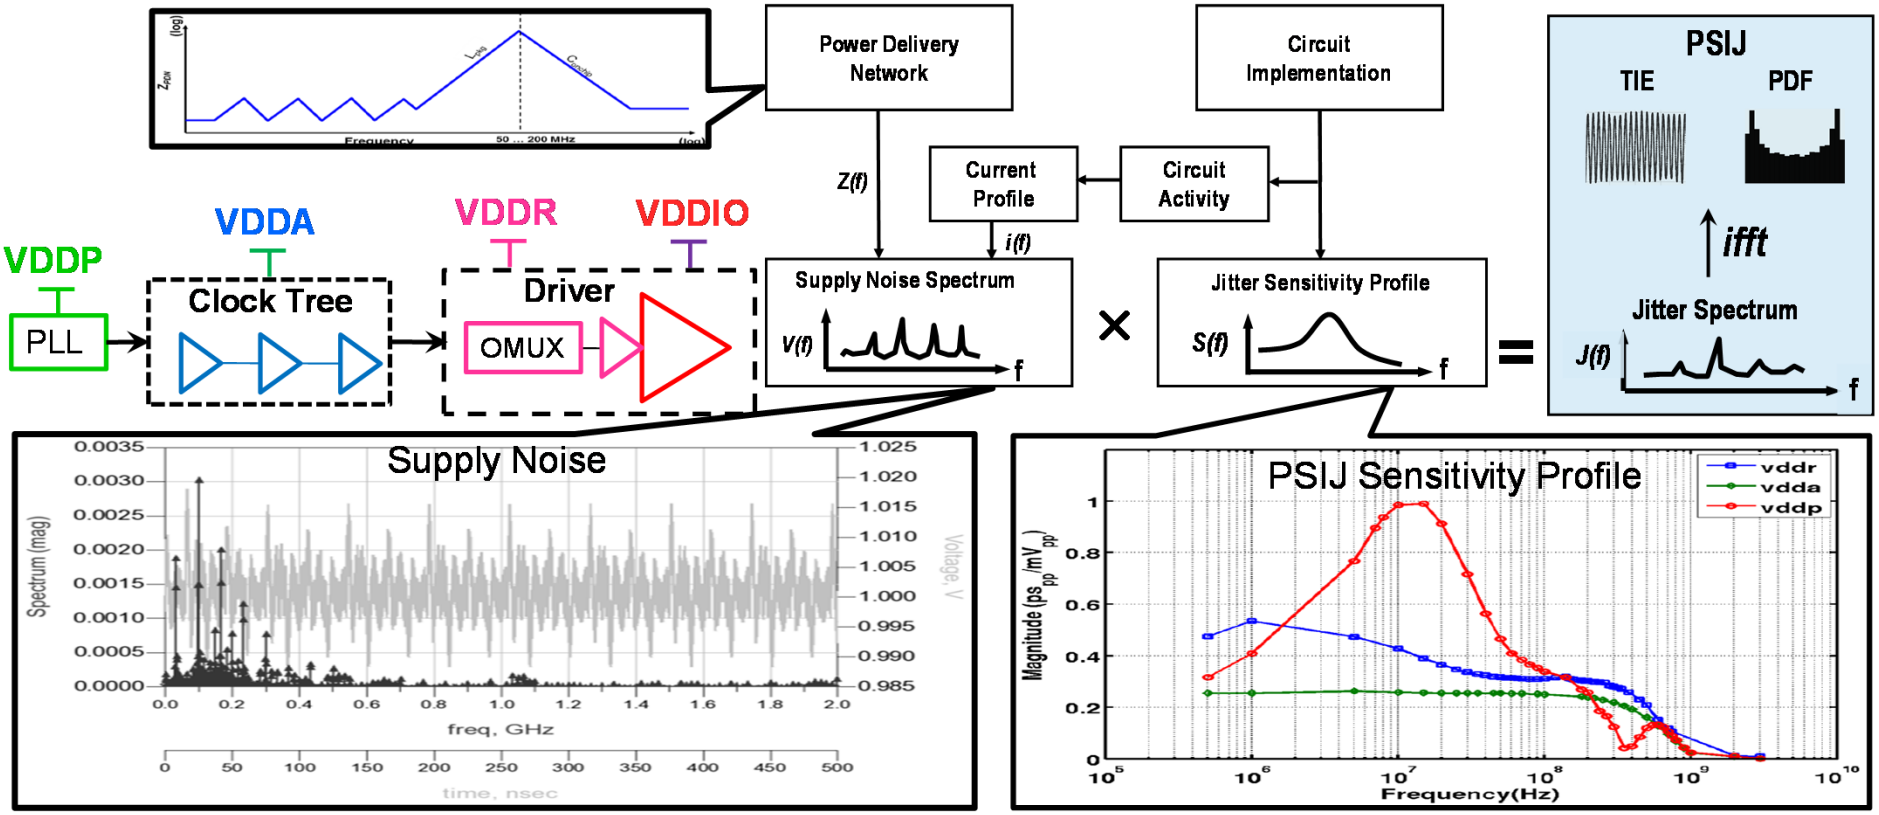

- Low-impedance power distribution network
- On-chip decoupling capacitance
- Current profile
- jitter sensitivity

Power Supply System Architecture

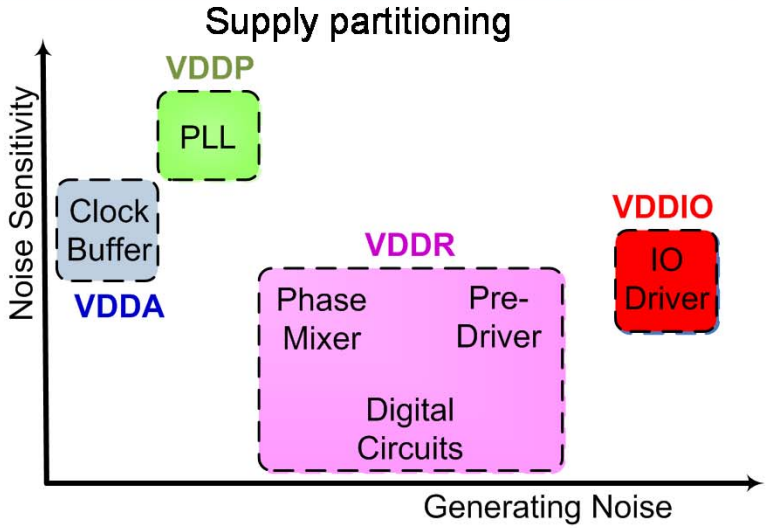

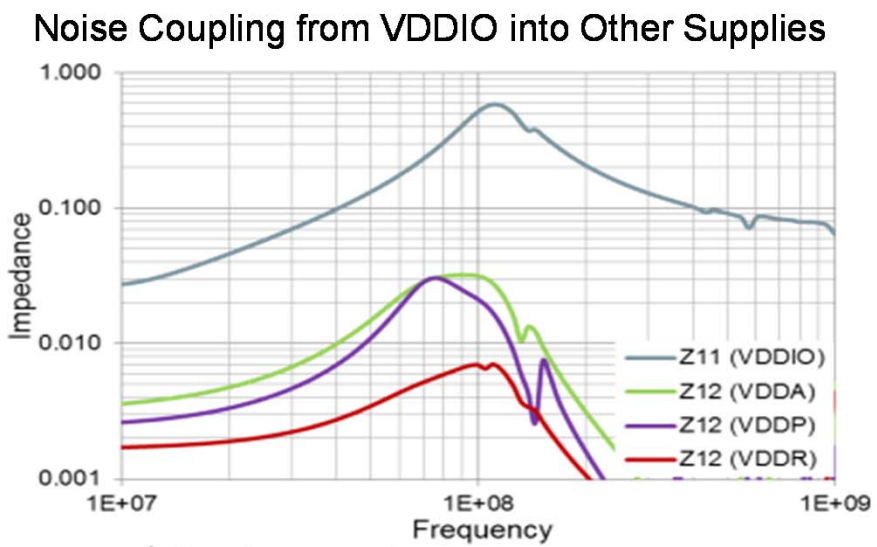

- Supply Partitioning to control noise and its impacts
  - Identify circuits generating (large) supply noise
  - Identify the circuit exposed to the noise
  - Identify circuits very sensitive to supply noise
  - Separate sensitive circuits from large noise sources
  - Evaluate the margin impact of the noise
- Minimize supply-to-supply coupling
  - Inductive coupling between bond wires and traces
  - Shared ground bond wires and traces

Power State Transition

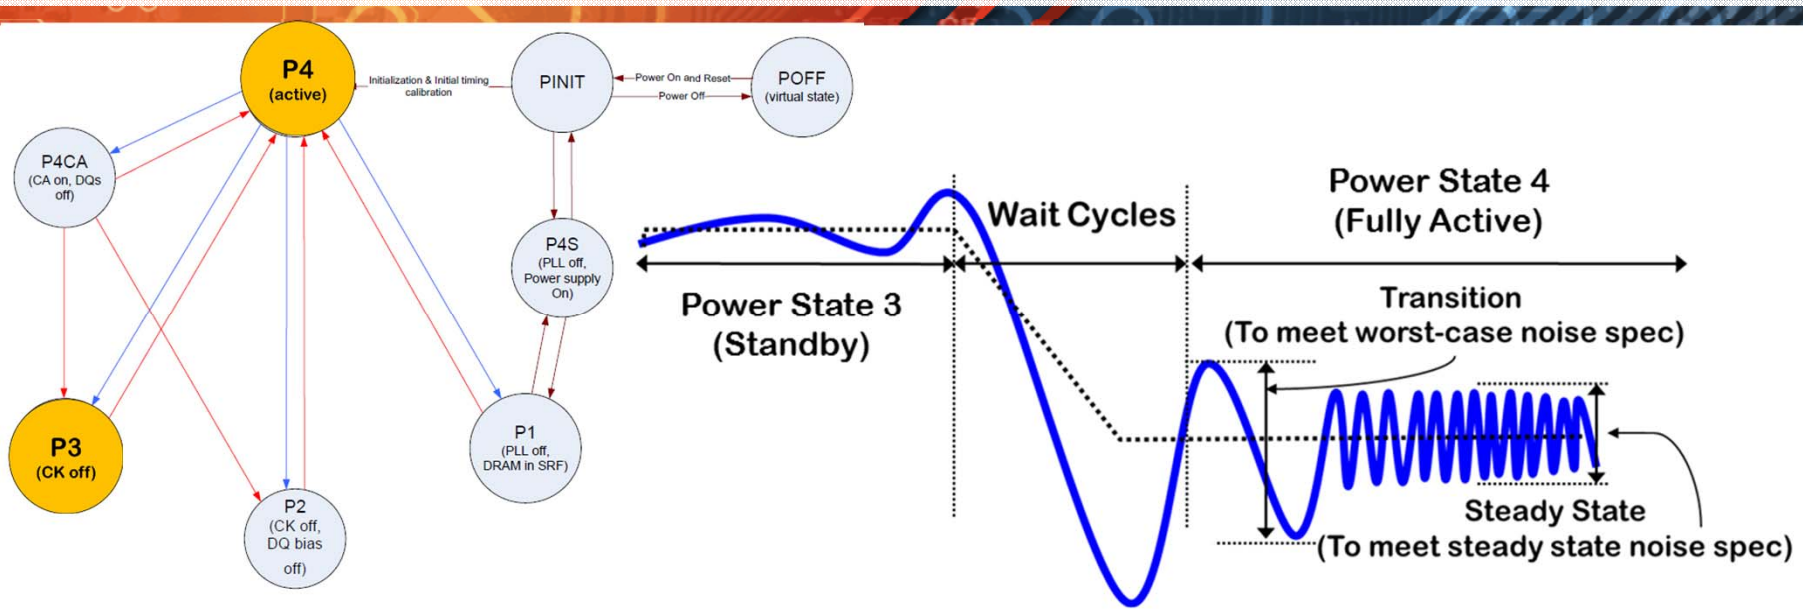

- Worst supply noise occurs during state transition, not in active states
- Example : Advanced power management of mobile memory system
  - Several power states
  - Quick exit and entry
  - Critical transition: P3→P4
- Power transition can generate large noise transients on the power supply rails
- The supply noise transient results in large power-supply induced jitters

Complete Link Analysis

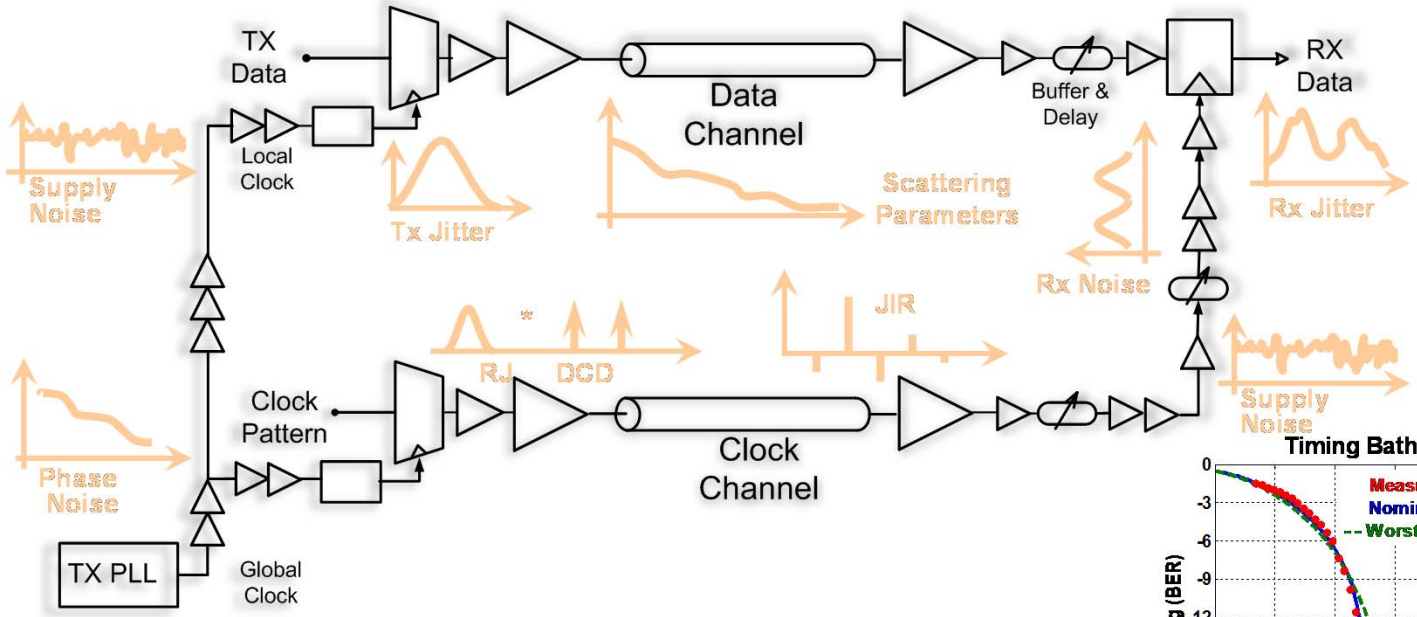

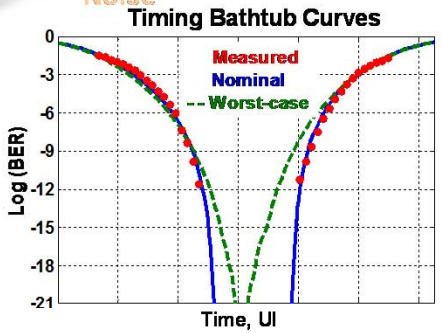

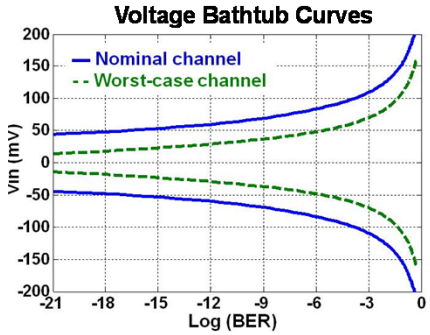

- Channel ISI, reflection, crosstalk
- Power supply noise and induced jitter,
- Deterministic and random device jitter and noise
- Transceiver bandwidth limitations
- Noise of transmitter and receivers and clocks
- Jitter correlations and amplifications

Emerging Memory Interfaces

- Four emerging memory interfaces
  - DDR3/L/U (1.6 Gbps) and DDR4 (3.2Gbps)
    - Desktop and workstations
  - LPDDR3 (2.133 Gbps) and LPDDR4 (4.266 Gbps)
    - Mobile and Tablet
  - GDDR5 (7 Gbps) / HMC (10, 12.5, 15 Gbps)
    - Graphics and main memory
  - WideIO2 (800 Mbps), HMB
    - Mobile systems
- These interfaces represent different approaches to address the bandwidth, power, and area challenges
- The channel and the system environment are different
  - Channel length, Tx swing, DRAM packaging

Evolution of DDR

Variables	SDRAM	DDR1	DDR2	DDR3	DDR4
VDD/VDDQ/VDP	3.3/2/5	2.5	1.8	1.5/1.5/-	1.2/1.2/2.5
Data Rate (Mbps)	166	400	800	2.133	3200
Clock	Single ended	Differential	Differential	Differential	Differential
Strobes	unsupported	Single ended	Single ended or Differential	Differential	Differential
Vref	N/A	External	External	Ex/Internal	Ex/Internal*
ODT	unsupported	unsupported	supported	supported	supported
Interface	LVTTL	SSTL	SSTL	SSTL	PODL

- Latest features
  - Shrinking supply
  - Point-to-point with matching termination
  - PODL with DBI
  - Internal Vref generation

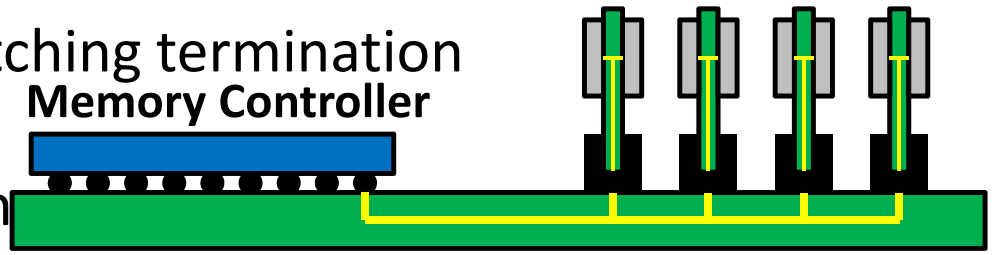

* External for CA

Main Memory Channel

- No major improvement in the passive channel
  - FR-4 board, connector, device Ci
- Signal integrity issues at higher data rates limit the number of DIMM
  - Direct attached channel
  - FBDIMM (Fully buffered DIMM) and RLDIMM (Reduced load DIMM)
- Clock and Command/Address topologies and termination changed
  - *Minimize stub length* : Tree to fly-by topology
  - On board to module termination

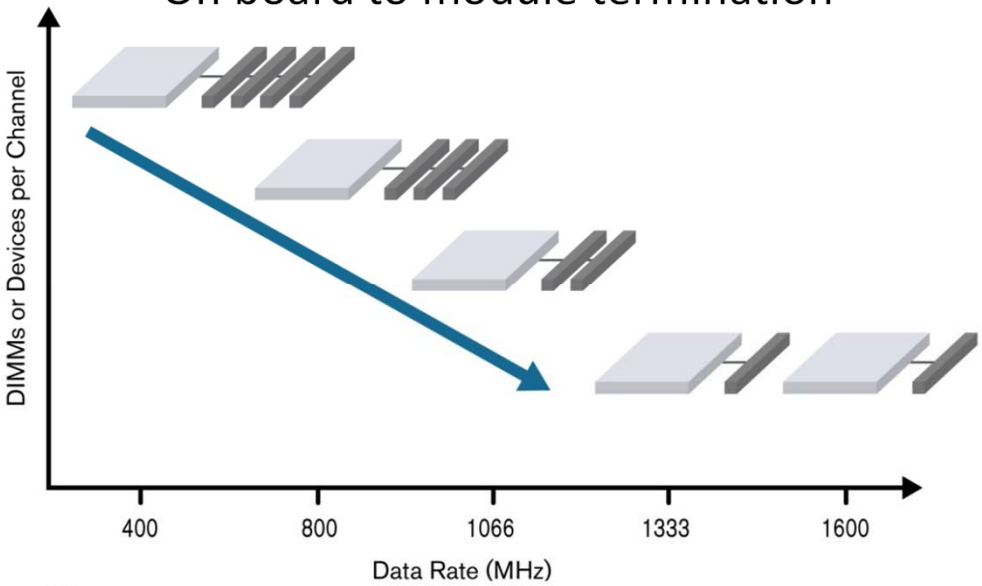

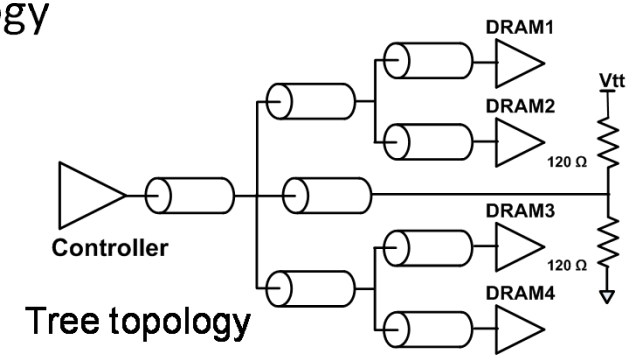

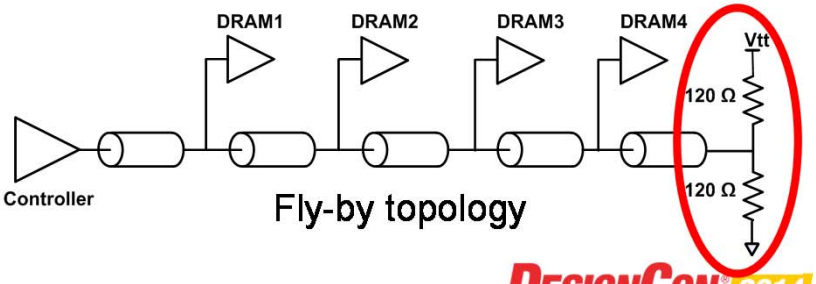

Ref: Intel® Xeon® Processor 500 Series (Nehalem) DDR3 DIMM System-level Validation Results

PODL Signaling and Supply Noise Tracking

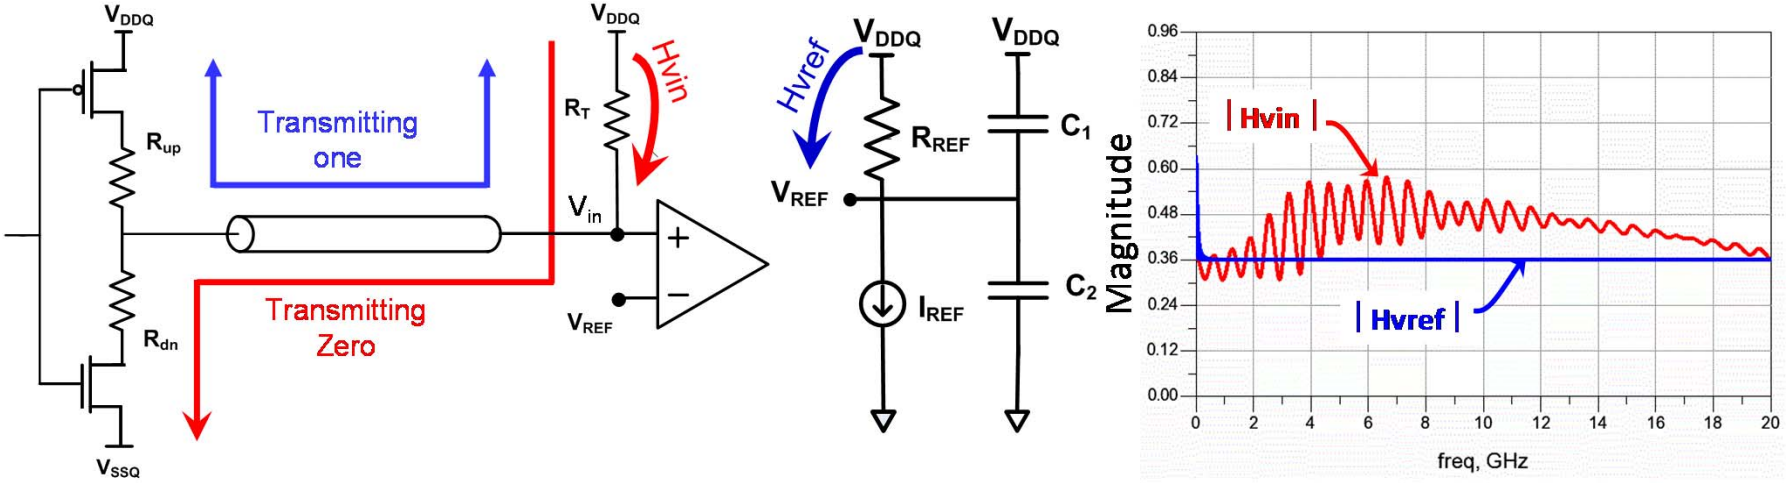

- Data Bus Inversion : one DBI pin per byte
  - DC: Limit the number of DQ lines per byte driving a Low to 4
  - AC: Number of bits switching within a byte not to exceed 4
- One reference shared between multiple pins
  - Coupling to supply network different from data
  - Couple supply noise to VREF over the frequencies where PDN peaks
- Most of the AC noise from the transmitter
  - Not tracked by the VREF at the receiver

Block Diagram of the DDR Interface

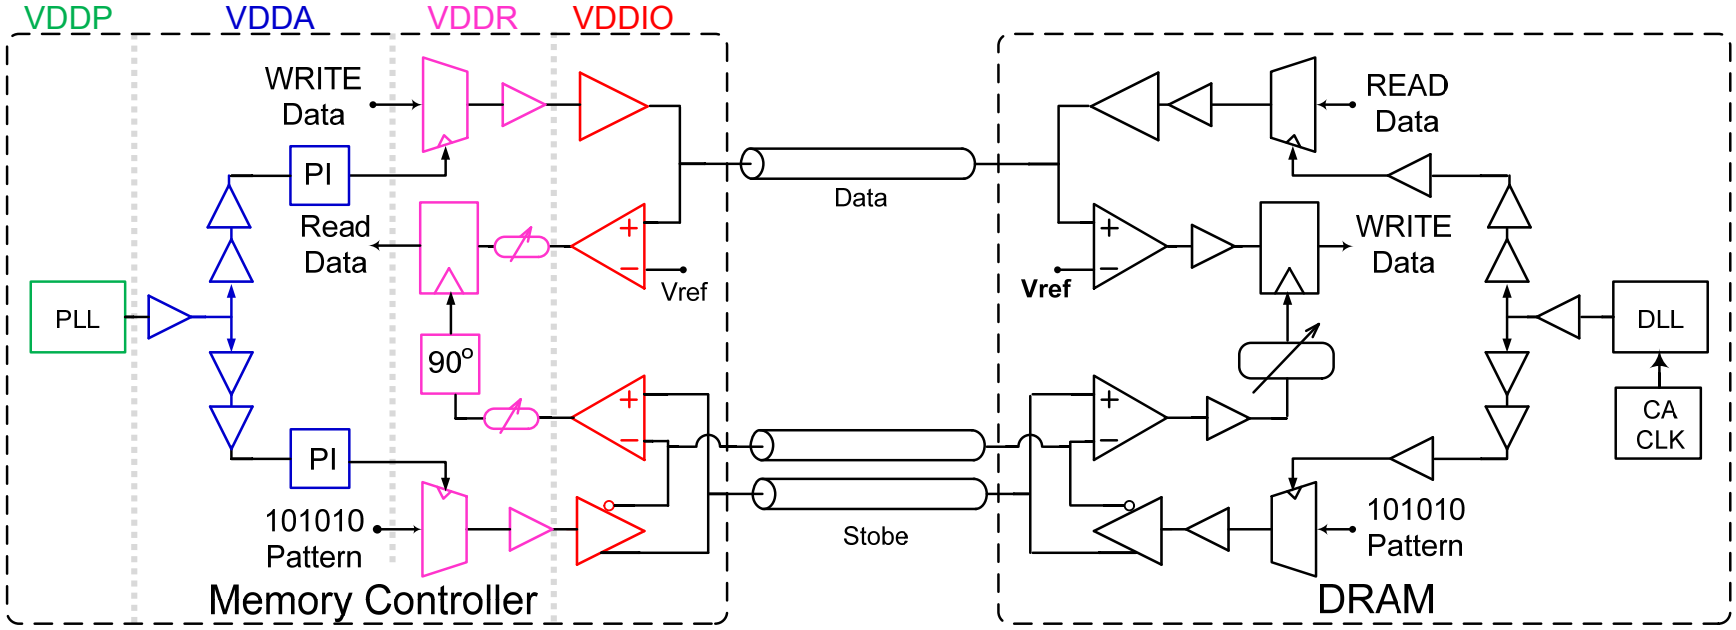

- Supply Partitioning to control noise and its impacts
  - The output drivers and front end of the receivers are on separate supply : VDDIO or VDDQ
  - The OMUX and other pre-driver stages are on VDDR
  - The clock distribution and PLL are on separate supplies
  - The partitioning are optimize using detailed on-chip and off-chip power integrity analysis
- Minimize supply-to-supply coupling
  - The supplies are separated at chip, package and even at board levels.

Evolution of LPDDR

Variables	LPDDR1	LPDDR2	LPDDR3	LPDDR4
VDD1, VDD2, VDDQ (V)	1.8	1.8/1.2/1.2/1.2	1.8/1.2/1.2	1.8/1.1/1.1
Data Rate (Mbps)	400	1066	2133	4266
CA	SDR	DDR	DDR	SDR
Clock	Single ended	Differential	Differential	Differential
Strobes	Single ended	Differential	Differential	Differential
Vref	External	External	External	Ex/Internal
ODT	unsupported	Unsupported	supported	supported
Interface	SSTL	HSUL	HSUL/PODL	LVSTL

- Latest features
  - Supply voltage shrinking
  - Internal reference voltage generation
  - Far-end termination
  - Reduced swing

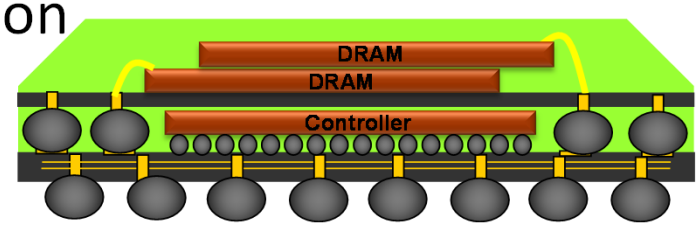

LPDDR3/LPDDR4: I-V Characteristics

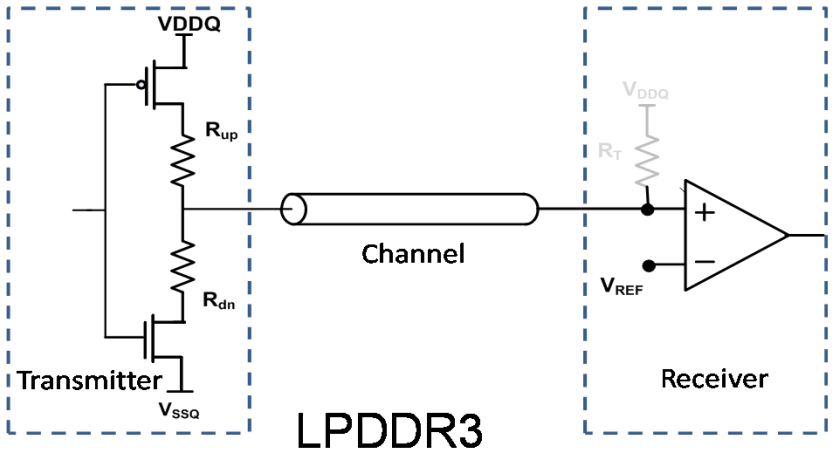

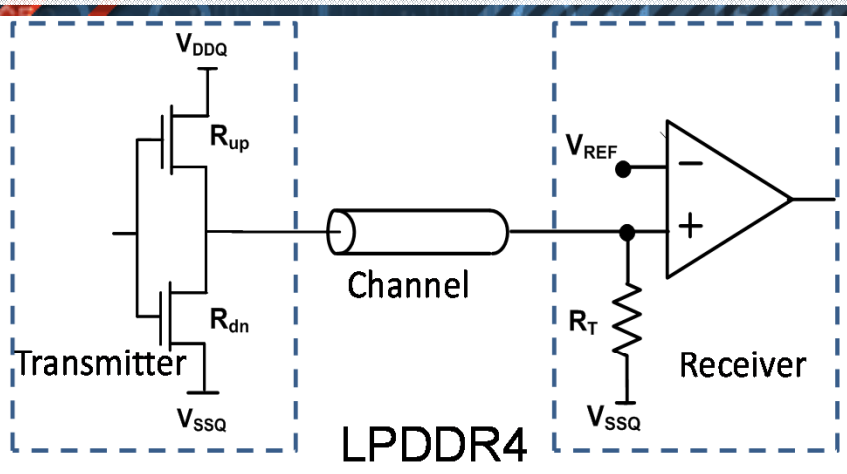

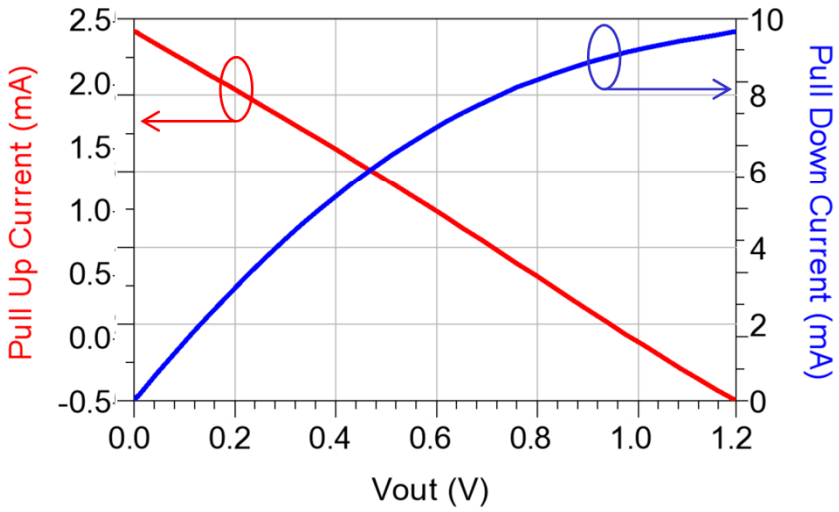

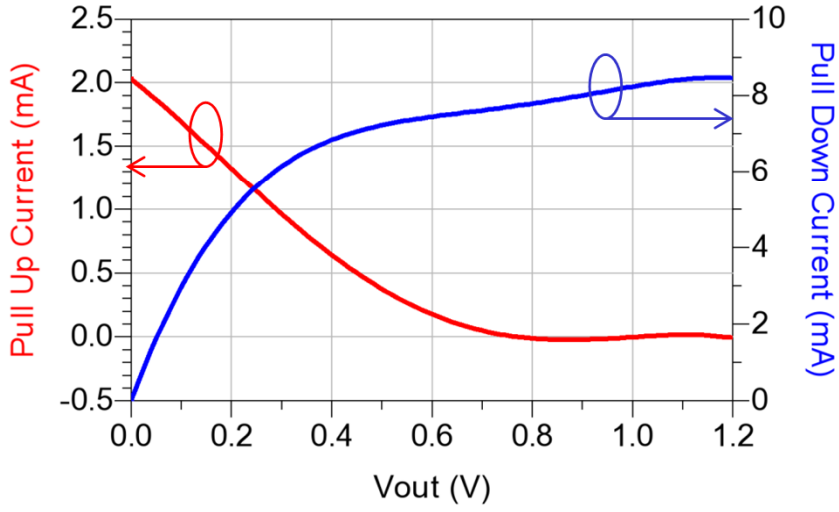

- Transmitter nonlinearity can affect the system performance

PoP Memory System

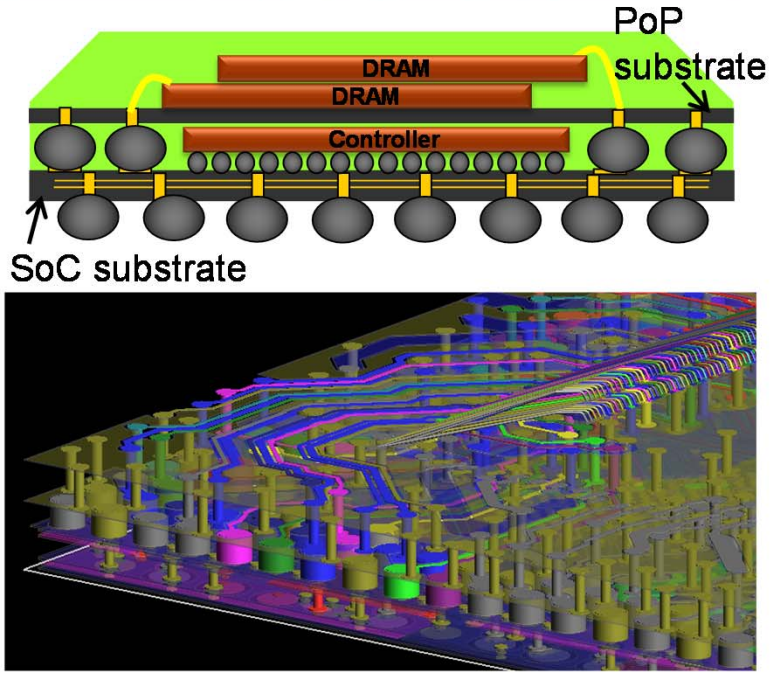

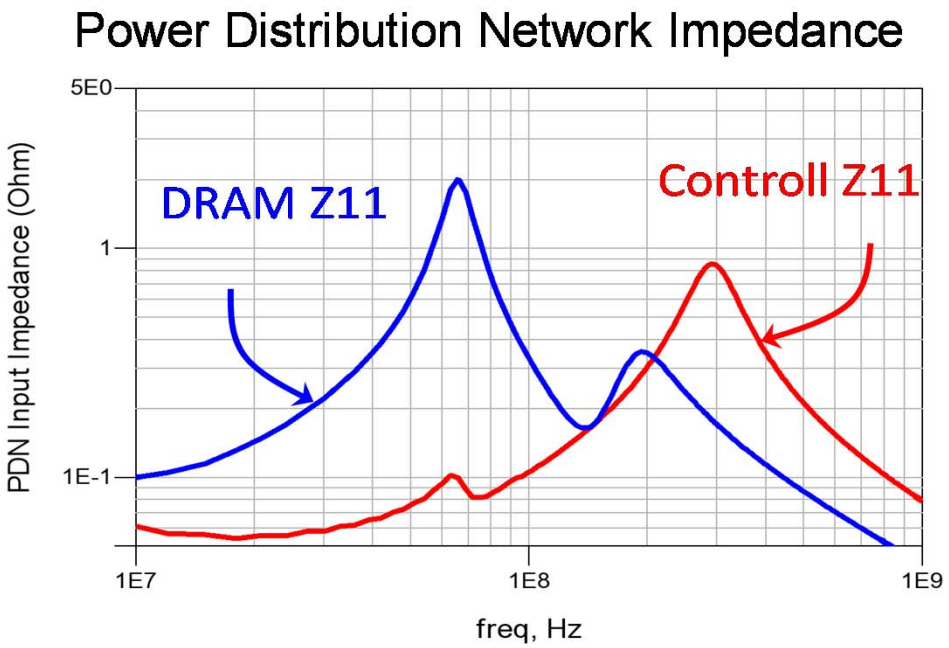

- PoP structure widely used in mobile devices
- Supply noise and package crosstalk are the major challenge to run at high data rates
  - Improve package design and wirebond length
  - Minimize PSIJ : Improve circuit sensitivity to supply noise

LPDDR3 / LPDDR4: Eye Diagrams

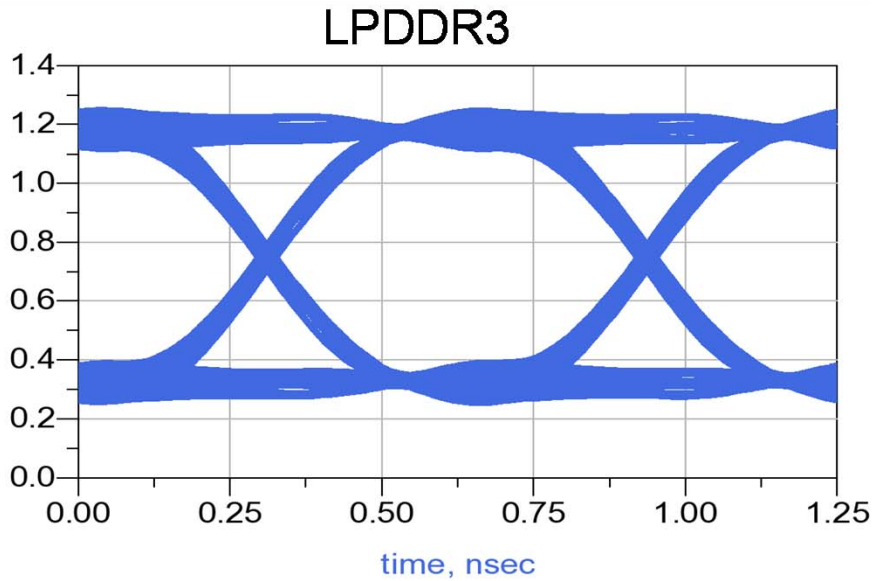

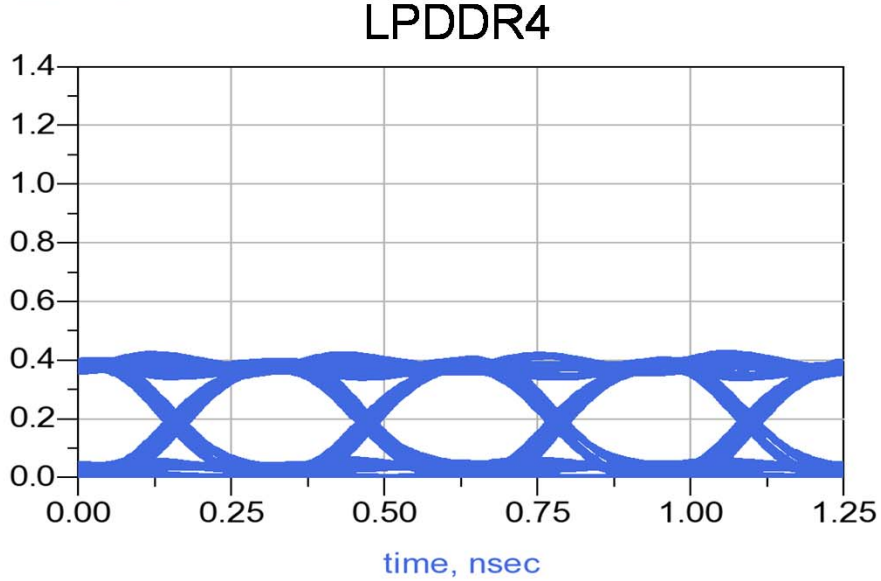

- LPDDR3: use of weak termination (240Ω→120 Ω)
- LPDDR4: Smaller Ci improves margin

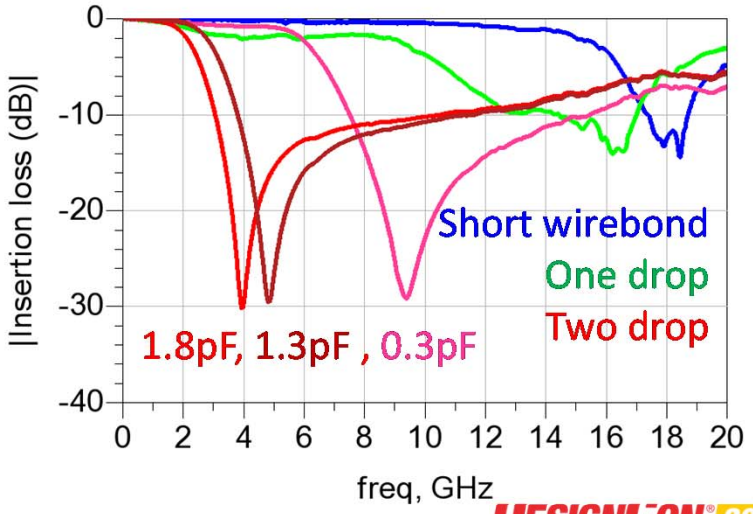

Evolution of GDDR and HMC

Variables	GDDR1	GDDR2	GDDR3	GDDR4	GDDR5	HMC
VDD (V)	2.5	1.8	1.5	1.5	1.35	1.2
Data Rate (Mbps)	0.3-0.9	0.8-1.0	0.7-2.6	2.0-3.0	3.6 - 7	10-15
DQ	Single ended	Single ended	Single ended	Single ended	Single ended	Differential
Clock	Single ended	Differential	Differential	Differential	Differential	Differential
Strobes	Single ended, DQS	Differential Bidirectional, DQS, DQSB	WDQS, RDQS, Uni-directional	WDQS, RDQS, Uni-directional	WDQS, RDQS, Uni-directional	Differential
Interface	SSTL	SSTL	PODL	PODL	PODL	SerDes

- GDDR new features
  - Data bus inversion (DBI)
  - Address bus inversion (ABI)
- HMC : Hybrid Memory Cube
  - SerDes : Tx and Rx equalization

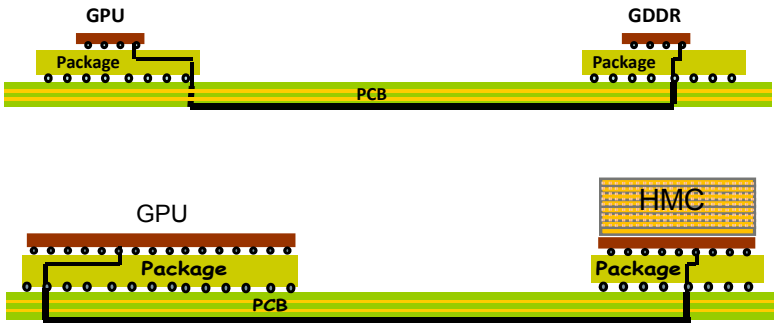

HMC Memory System

- Interface is SERDES
- Data rate : 10 to 15 Gbps
- Channel is short
  - Point to point and differential
  - Uni-directional and double-terminated
- EQ architecture : TX EQ, RX CTLE and DFE
- AC coupling : Optional

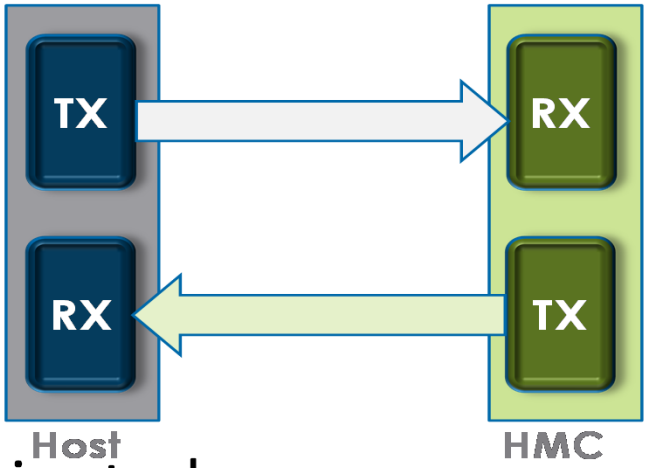

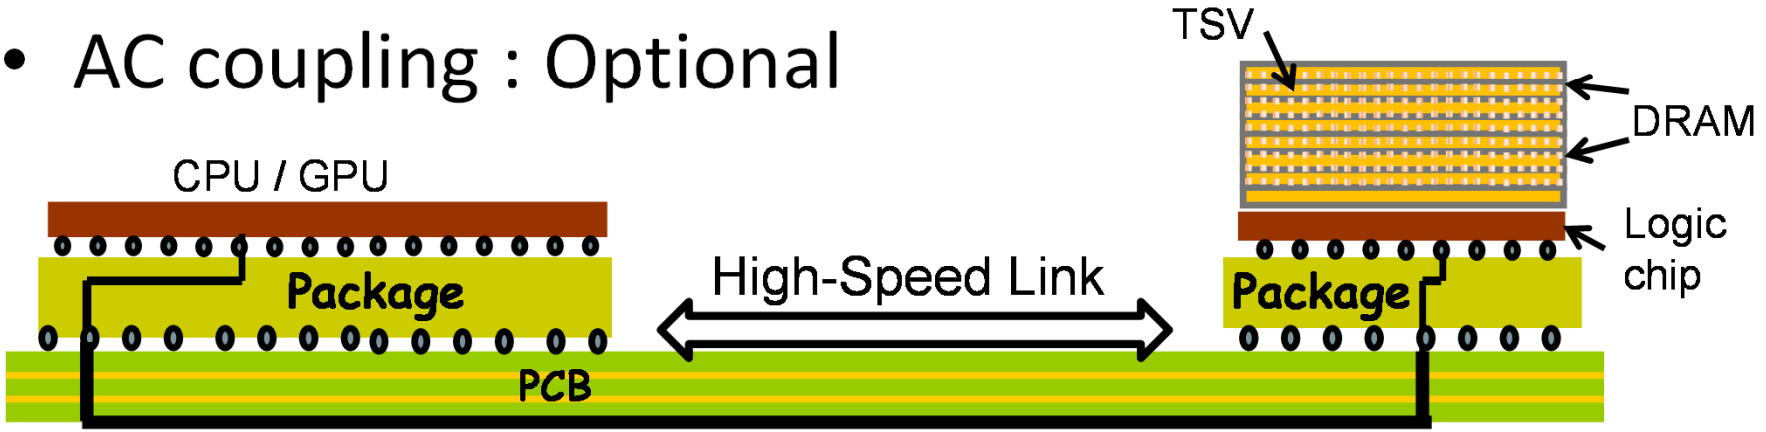

Channel Responses w/ and w/o CTLE

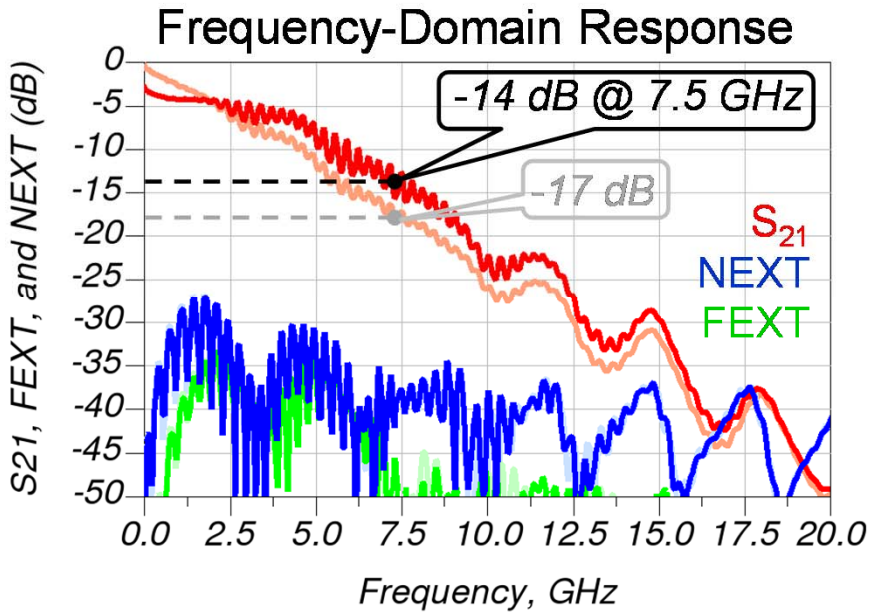

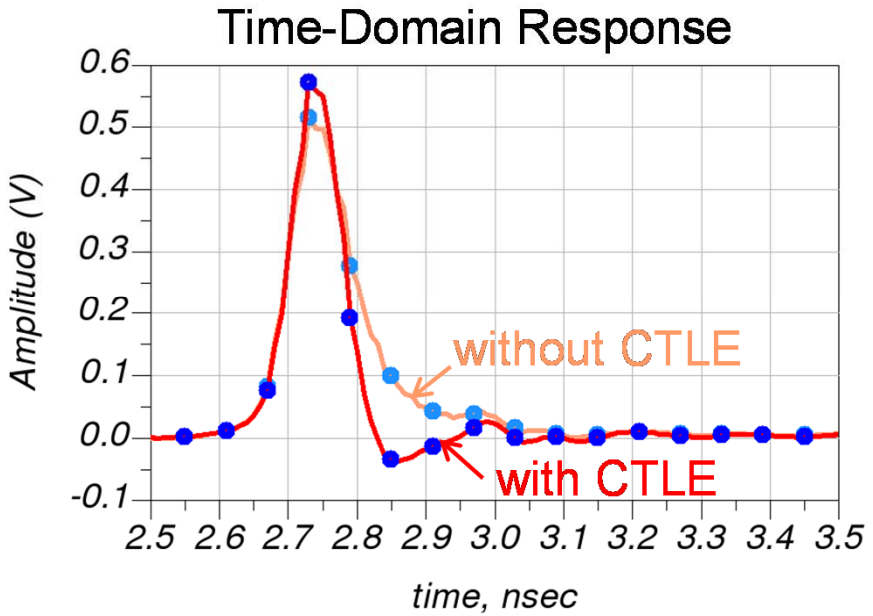

- Channel length and Ci increase attenuation
- Channel +CTLE response shows -14 dB of attenuation
- Stripline routing can further improve the crosstalk
- Significant pre (≈20%) and post cursor (>50%) ISI's

Channel Response and Eye Diagram

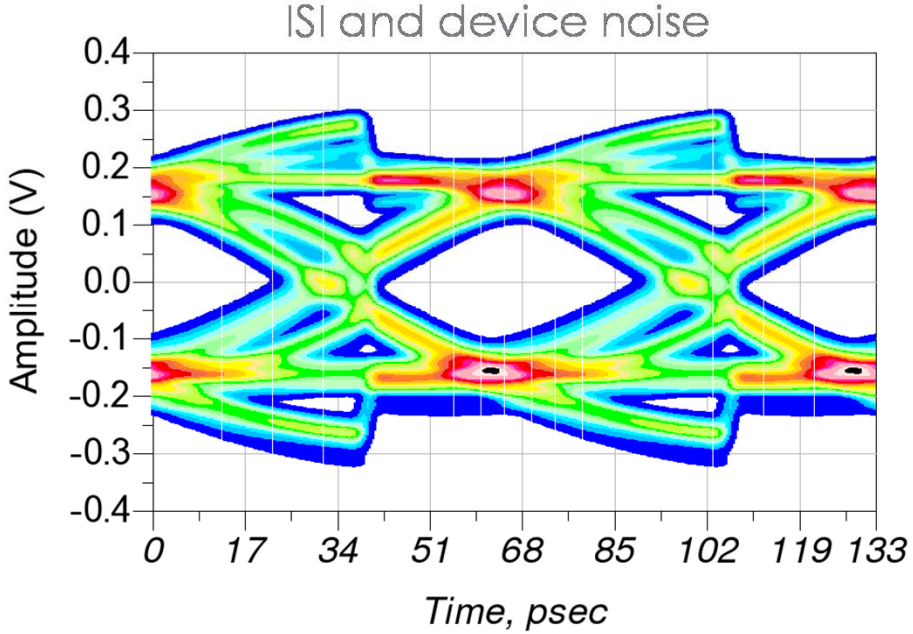

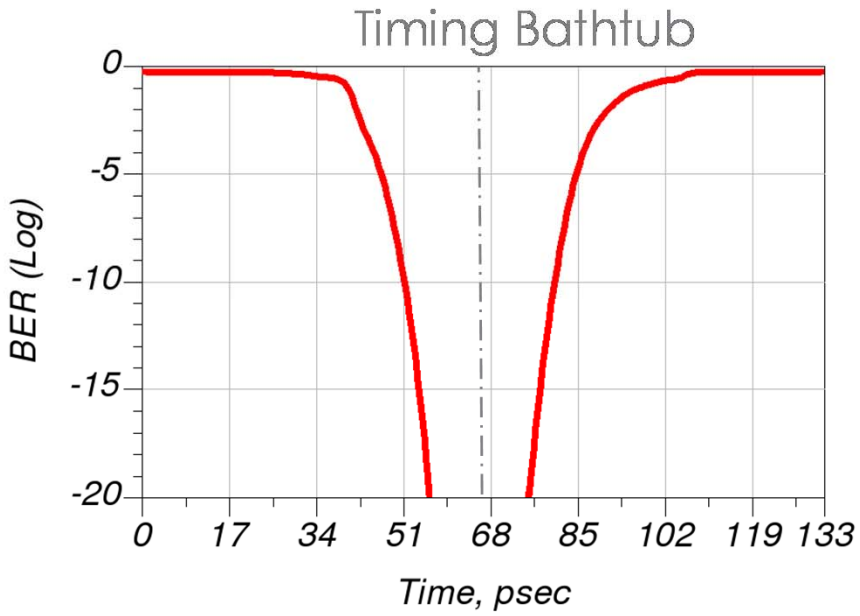

- DFE Equalization opens eye significantly
- Major contributors to margin loss
  - Channel (ISI +crosstalk) is major : *26 ps*
  - Device noise and jitter: *21 ps*
  - Timing Margin : *20 ps* at low BER

Conclusions : Signal and Power Integrity

- Increased data rate and reduced voltage supply
  - Tighter timing and voltage margin
- Topology and channel selections
  - Voltage margin improvement using simple equalization techniques
  - Timing improvement with better tracking between data and strobe
- To optimize the overall design and ensure robust link operation,
  - critical to take into account all major system impairments early
  - off-chip components, on-chip circuitry system architecture
- The power-state transition creates
  - large noise transients on the power supply rails
  - results in large power-supply induced jitter
- Cost sensitivity
  - Shortened time window for sign-off
  - Debugging failure is more difficult

Acknowledgments

- I would like to acknowledge the Rambus SI/PI engineers for their contributions and many helpful discussions over the years

Outline of Tutorial

- DRAM review [John]
- Conventionally packaged interfaces [John]
  - Evolution of standard interfaces
  - Signaling and clocking at peak bandwidth
  - Bandwidth variation
- 3D Packaging for Memory Applications [Ming]
  - Traditional 3D Packaging Technologies
  - New TSV Based 3D Packaging Technologies
- Signal and power Integrity [Wendem]
  - Signal and power integrity challenges
  - Memory interface analysis methodology
  - Main features of memory interfaces
  - Comparison of emerging standards-based solutions
- Conclusion and Final Comparison [John]

Qualitative Tradeoffs Across Interface Solutions

	DDRn		LPDDRn			3D
	DDR3	DDR4	LPDDR2	LPDDR3	LPDDR4	WideIO2
Cost	+++ Mainstream	- Not avail	++ Mainstream	+ Emerging	- Not avail	-- Not avail + Cost challenges
Pwr Efficiency	-	+ DBI, V scaling	++ Burst mode, HSUL	++ Burst mode, PODL	+++ Burt mode, LVSTL	++++ Low Ci, cmos
BW	++	+++	-	++	+++	++++
Controller complexity	++ Transaction gaps	+ New features	++ Power state transition	+ More training	+ Mult channels	+++ System variability reduced
SI/PI	+ Mulit-rank, long channel	++ DBI, more pt. to pt.	++ Short	+++ ODT	+++ ODT	++ SI improved but PI challenge?
Thermal	++++ Topology	+++ Higher BW	++ POP, mobile	+ Higher BW	+ Higher BW	-- Heat removal
Form Factor	+	+	++	++	++	+++

Even “Ideal” Interface has Finite Value to System

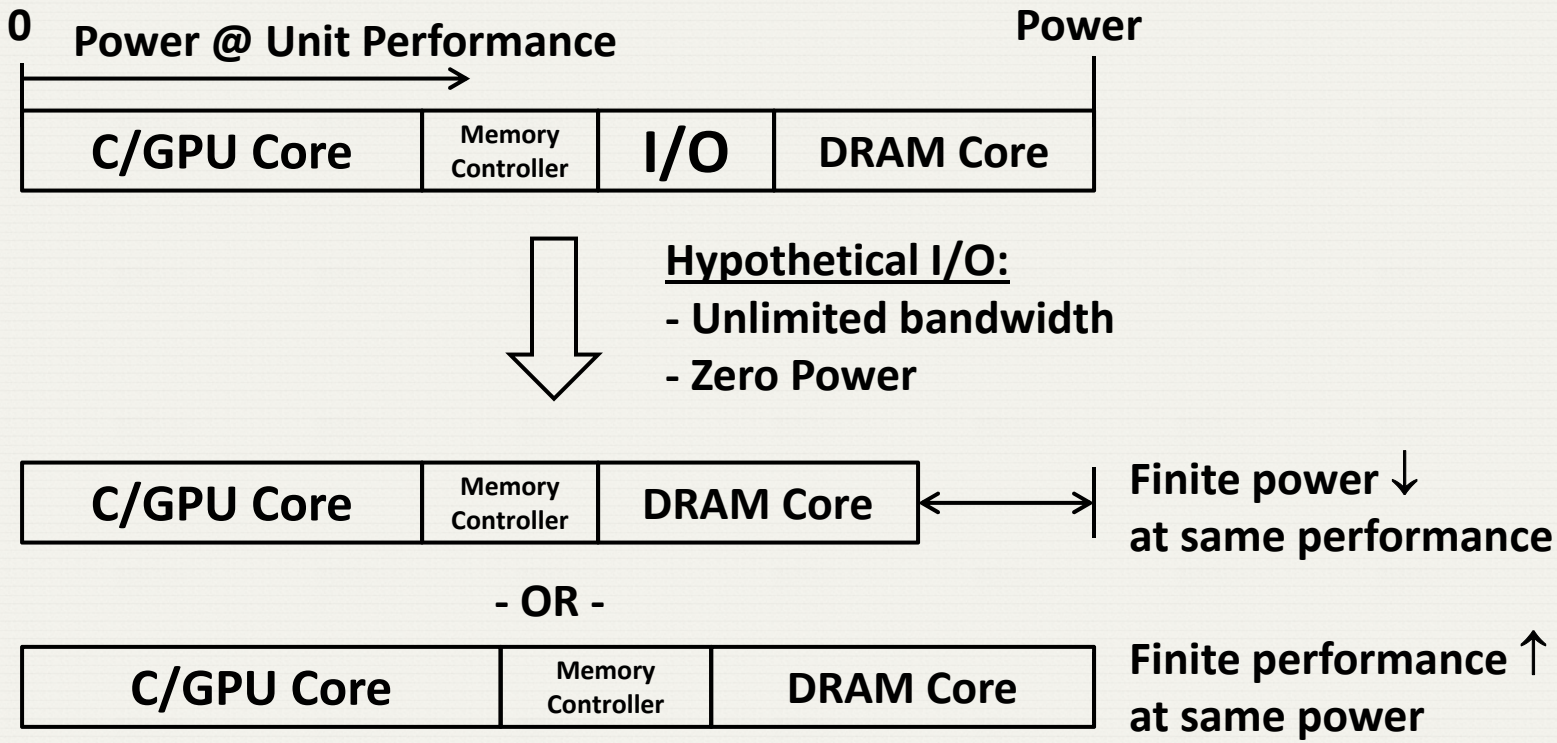

Value depends on relative power/performance of different components

DRAM Economics & Technology Transitions

- Low margins dictate high volumes
  - Any DRAM type must have significant % of market to be viable
- Multi-sourcing amortizes investment, risk
  - OEM requires multiple DRAM vendors
  - DRAM vendor needs devices to ship into multiple applications
- Price discrimination
  - DRAM devices are speed binned to monetize high & low end
  - Same GPU can ship with GDDR or DDR in high/low end systems
- This all makes abrupt technology transitions difficult
  - Even evolutionary transitions face these challenges
  - Historically, controllers are bimodal across generations

Summary

- Conventional interfaces must (will) continue evolving
  - Lower active power, lower effect power under workload
  - Sometimes needed in coordination with 3D integrated devices
- 3D integration offers disruptive interface benefits
  - Dramatic I/O efficiency improvement
  - New capabilities such as server capacity aggregation
- 3D faces uphill economic battle for market adoption
  - Value to system must justify cost (vs. alternatives)
  - Will start as niche in high-value applications:
    Server memory, HPC, later in high-end mobile
  - Slow trickle down to more mainstream applications