

DISCOBOLUS DESIGNS

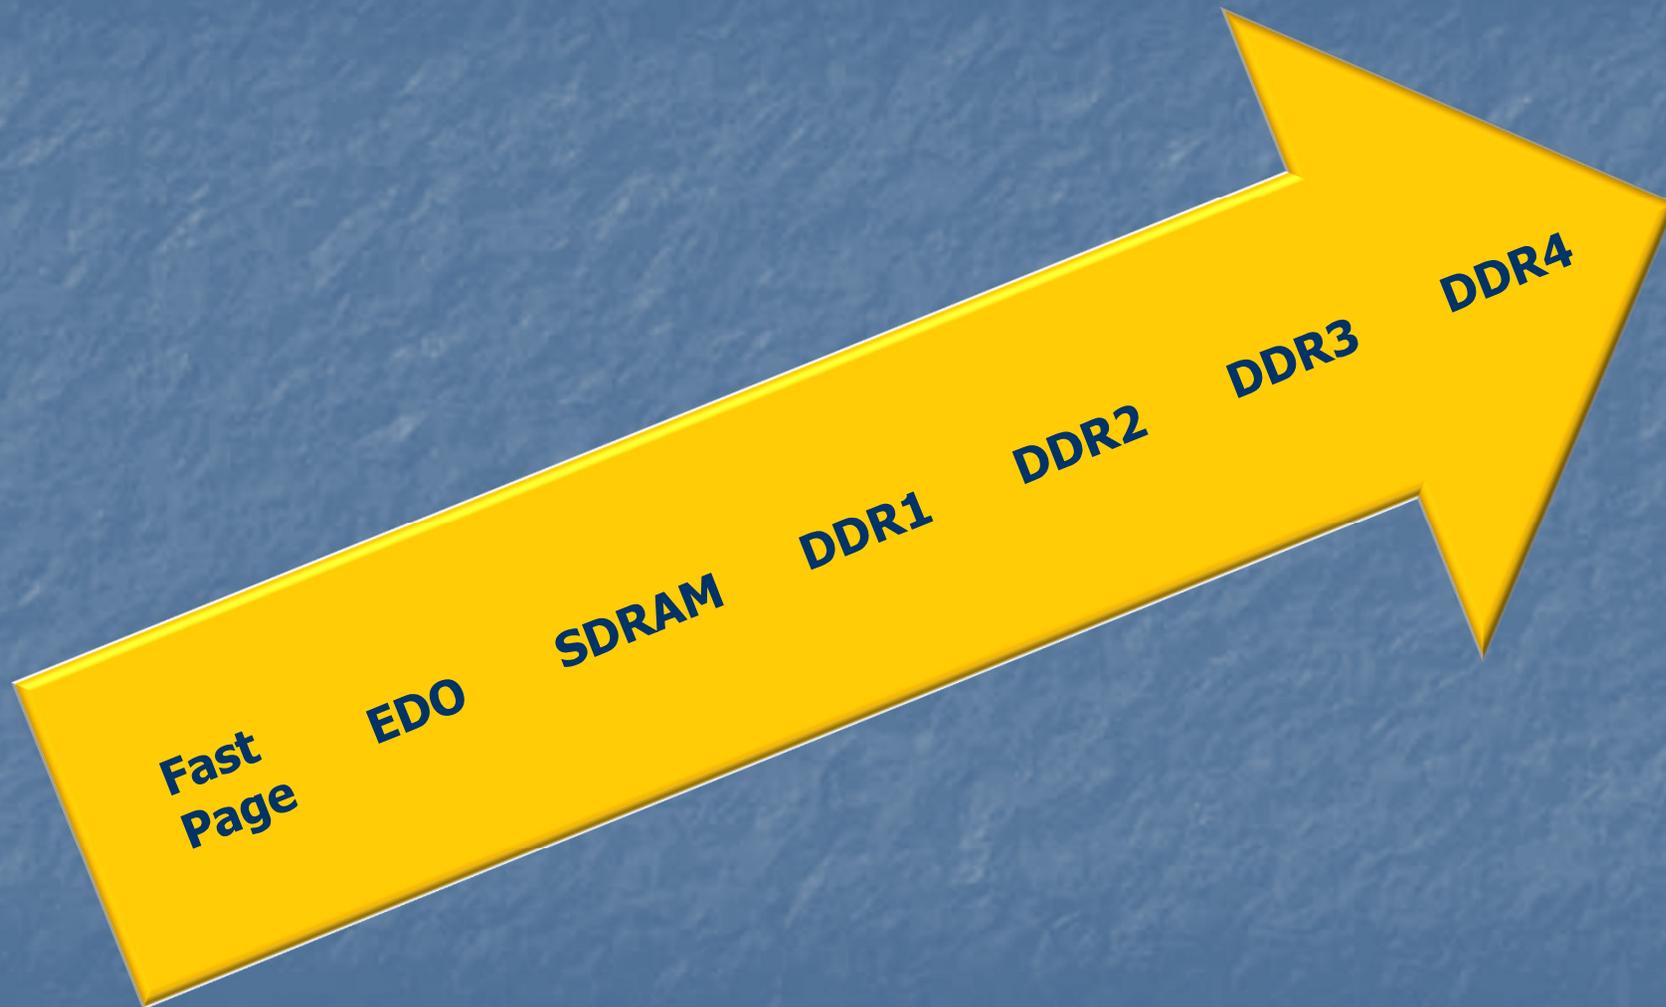
DRAM Market Forces of Fragmentation & Consolidation

Bill Gervasi
April 30, 2013





Mainstream DRAM Evolution



1990s

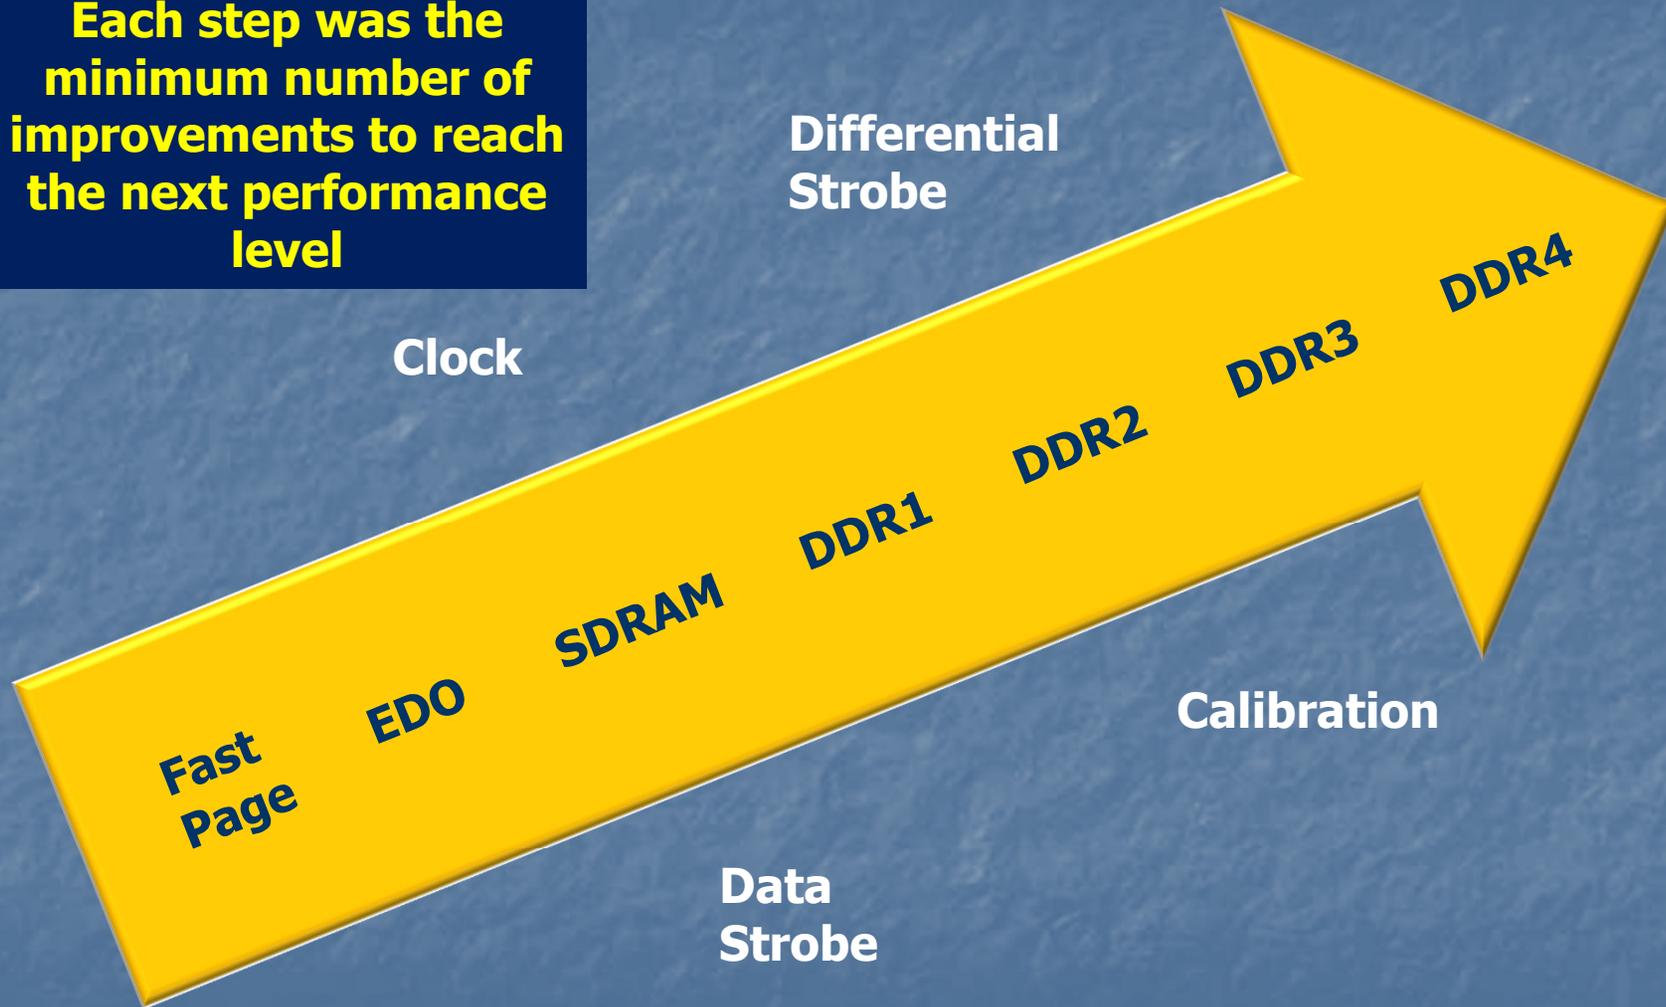
2000s

2010s



Evolution Versus Revolution

Each step was the minimum number of improvements to reach the next performance level



Timing simplification



Why Does Evolution Win?

Risk avoidance

Bridge controller between generations

Timing of generation transitions

Infrastructure changes are costly

PCB materials & methods

Sockets

Preserves knowledge base



Mainstream Market Transitions



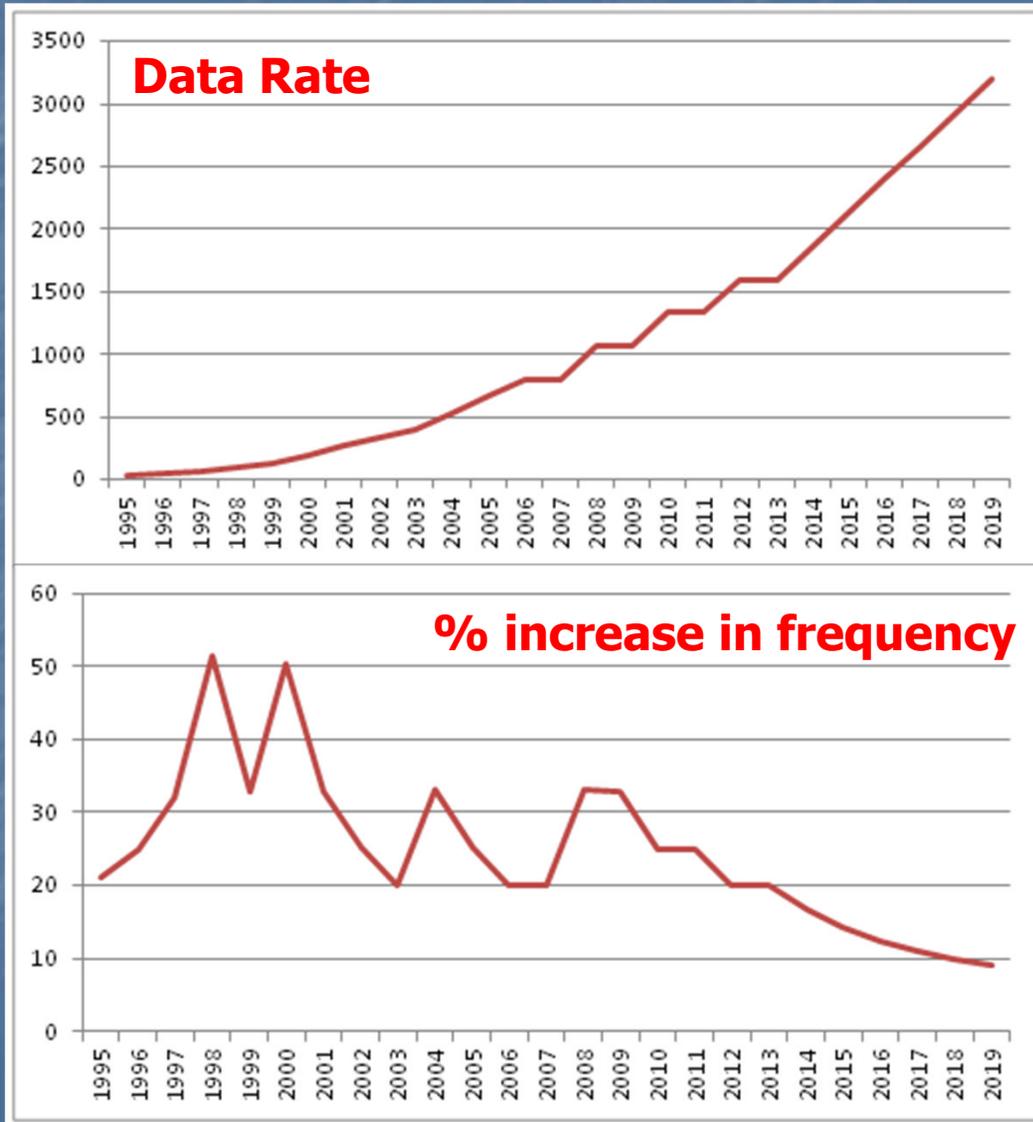
2013	2014	2015	2016	2017	2018	2019
------	------	------	------	------	------	------

1600	1866	2133	2400	2667	2933	3200	Mbps
------	------	------	------	------	------	------	------

4Gb		8Gb				16Gb
-----	--	-----	--	--	--	------



Frequency & %Frequency



Data rates continue to increase fairly linearly

However, that change becomes less of an impact on performance

Generation crossover traditionally does not occur until frequency doubles

133 → 333 → 666 → 1333

Implies that DDR4-2666 will be the crossover speed



What Has Changed?

DRAM generations used to be 3 years apart
Now they are 5-6 years apart

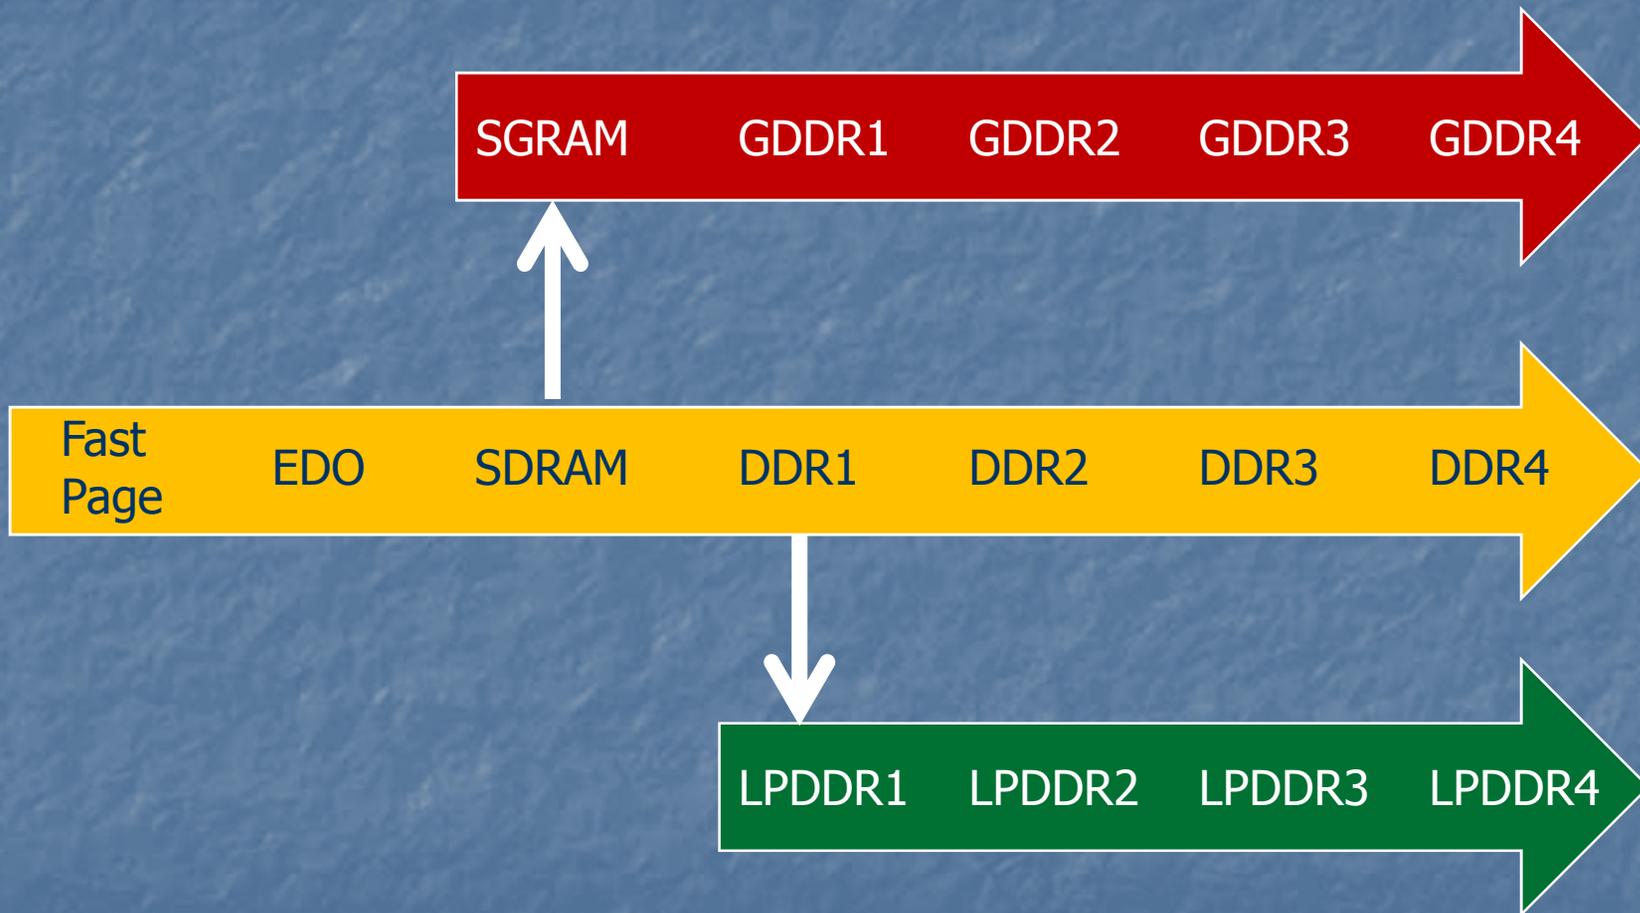
DRAM densities used to 2x every 18 months
Now they double every 3 years

DRAM speeds used to increase 33% per year
Now they are increasing 16% per year

Power consumption more important than bandwidth



Some Successful Derivatives



Why Did These Succeed?

Market size large enough

GDDR in graphics memory market

LPDDR in mobile market

Leveraged mainstream DRAM design

Fairly simple changes in core, I/O

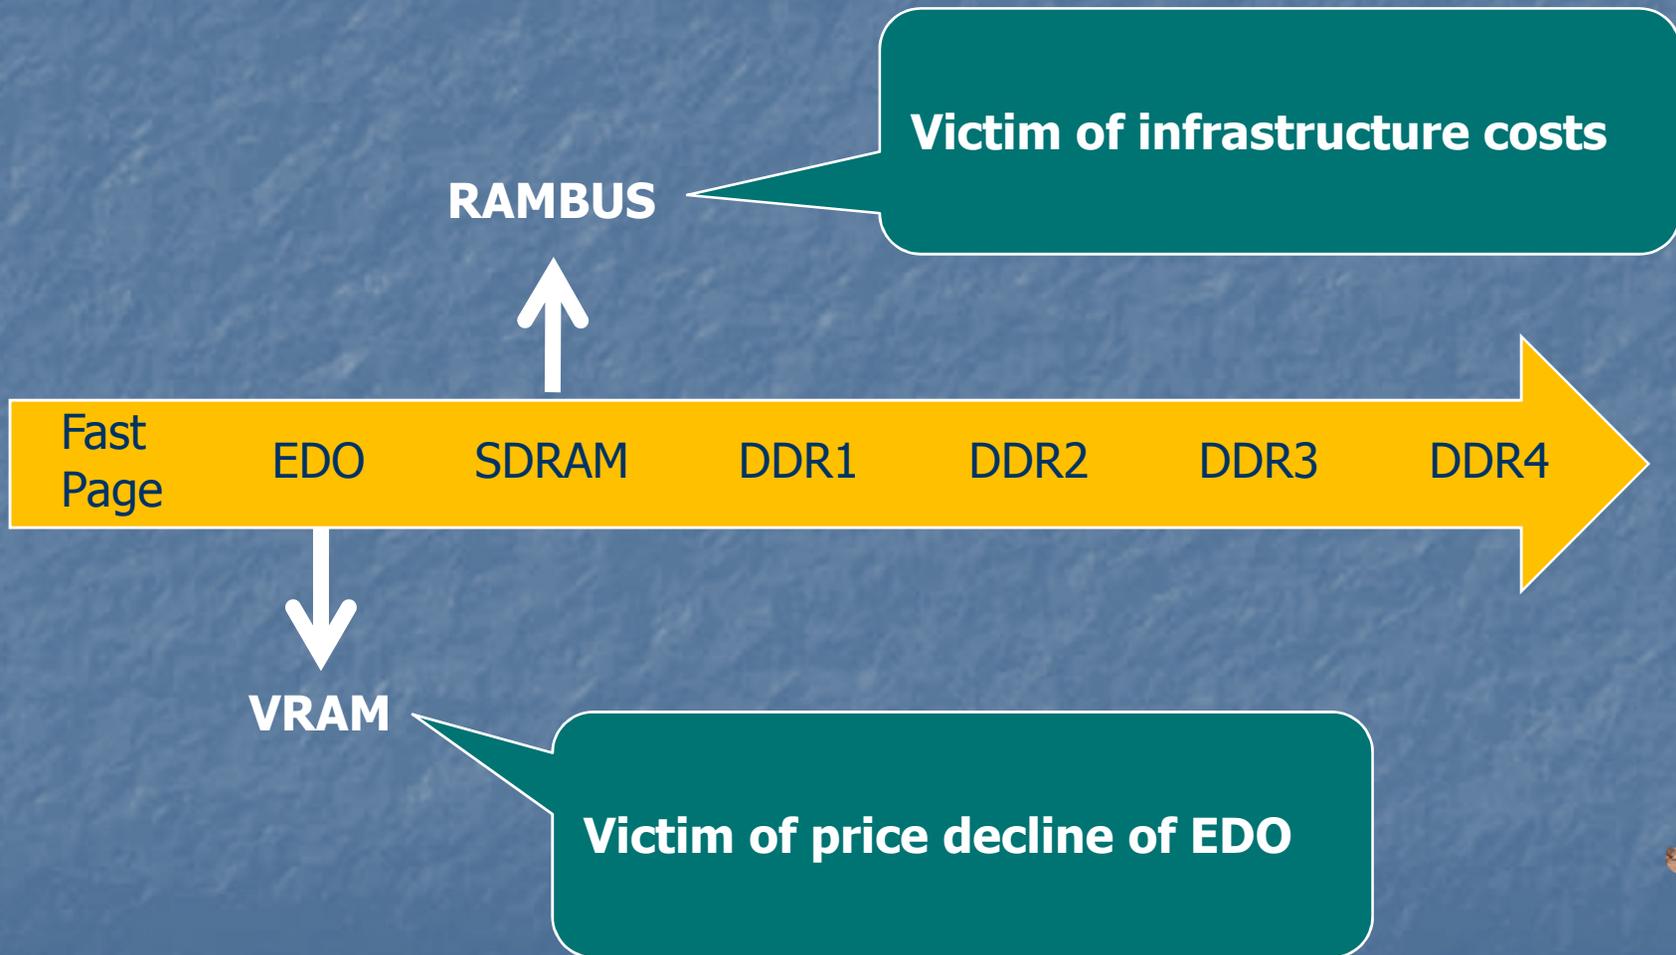
Used existing infrastructure

Standard PCB

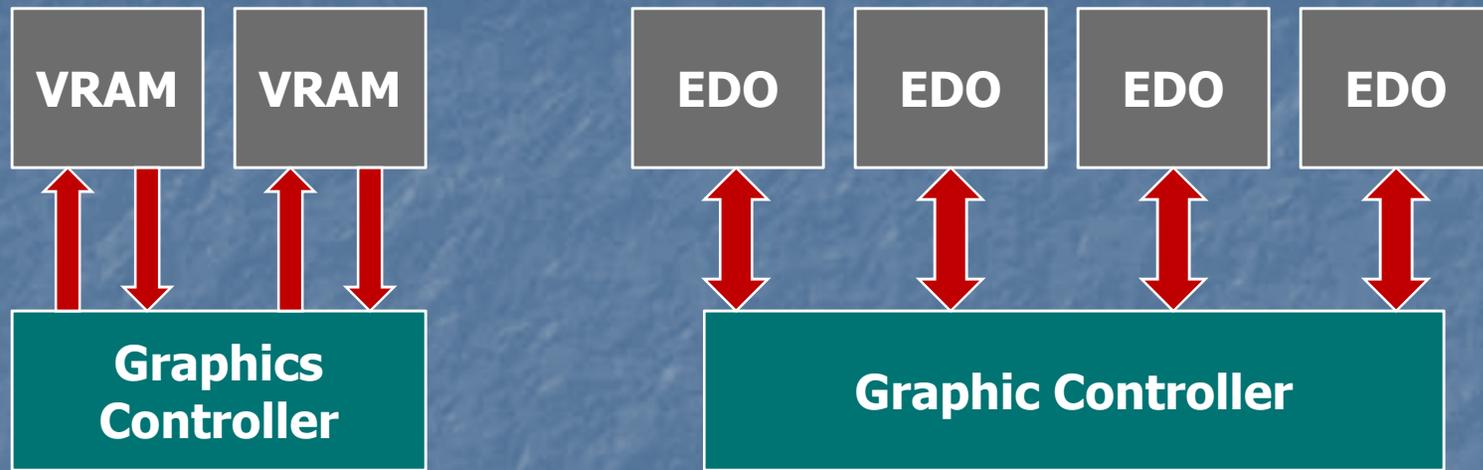
Device packaging type



Some Not So Successful Derivatives



Why Did VRAM Fail?



Not useful as a main memory

Could not generate market volume of EDO

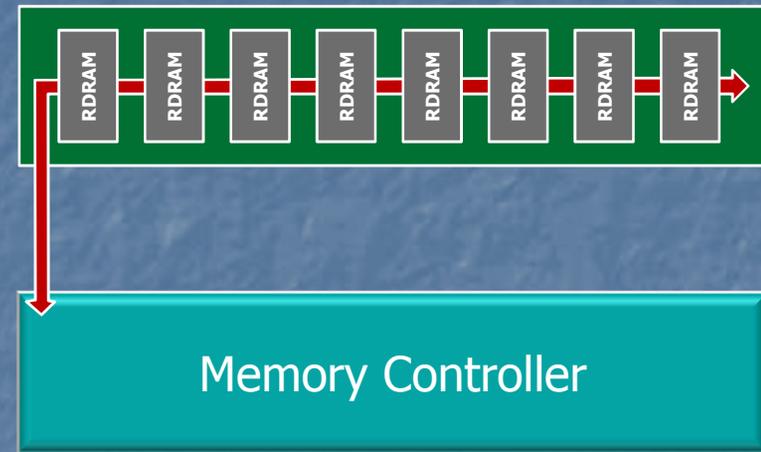
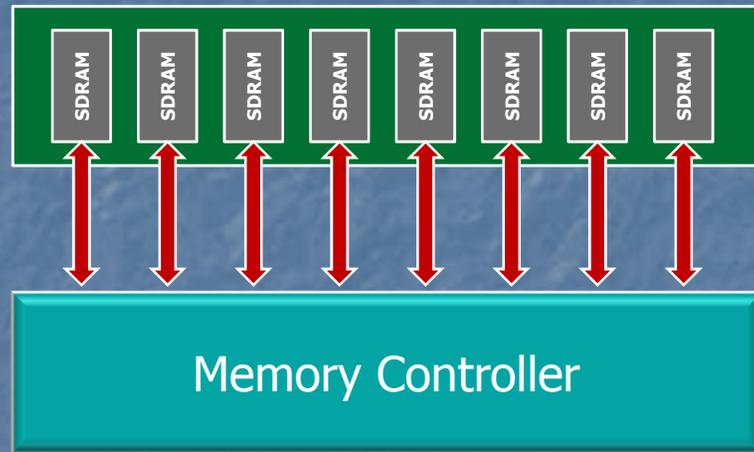
\$EDO fell to a fraction of \$VRAM

Designers worked around performance limits

Wasted memory? Who cared?



Why Did Rambus Fail?



Impossible to have bridge controller

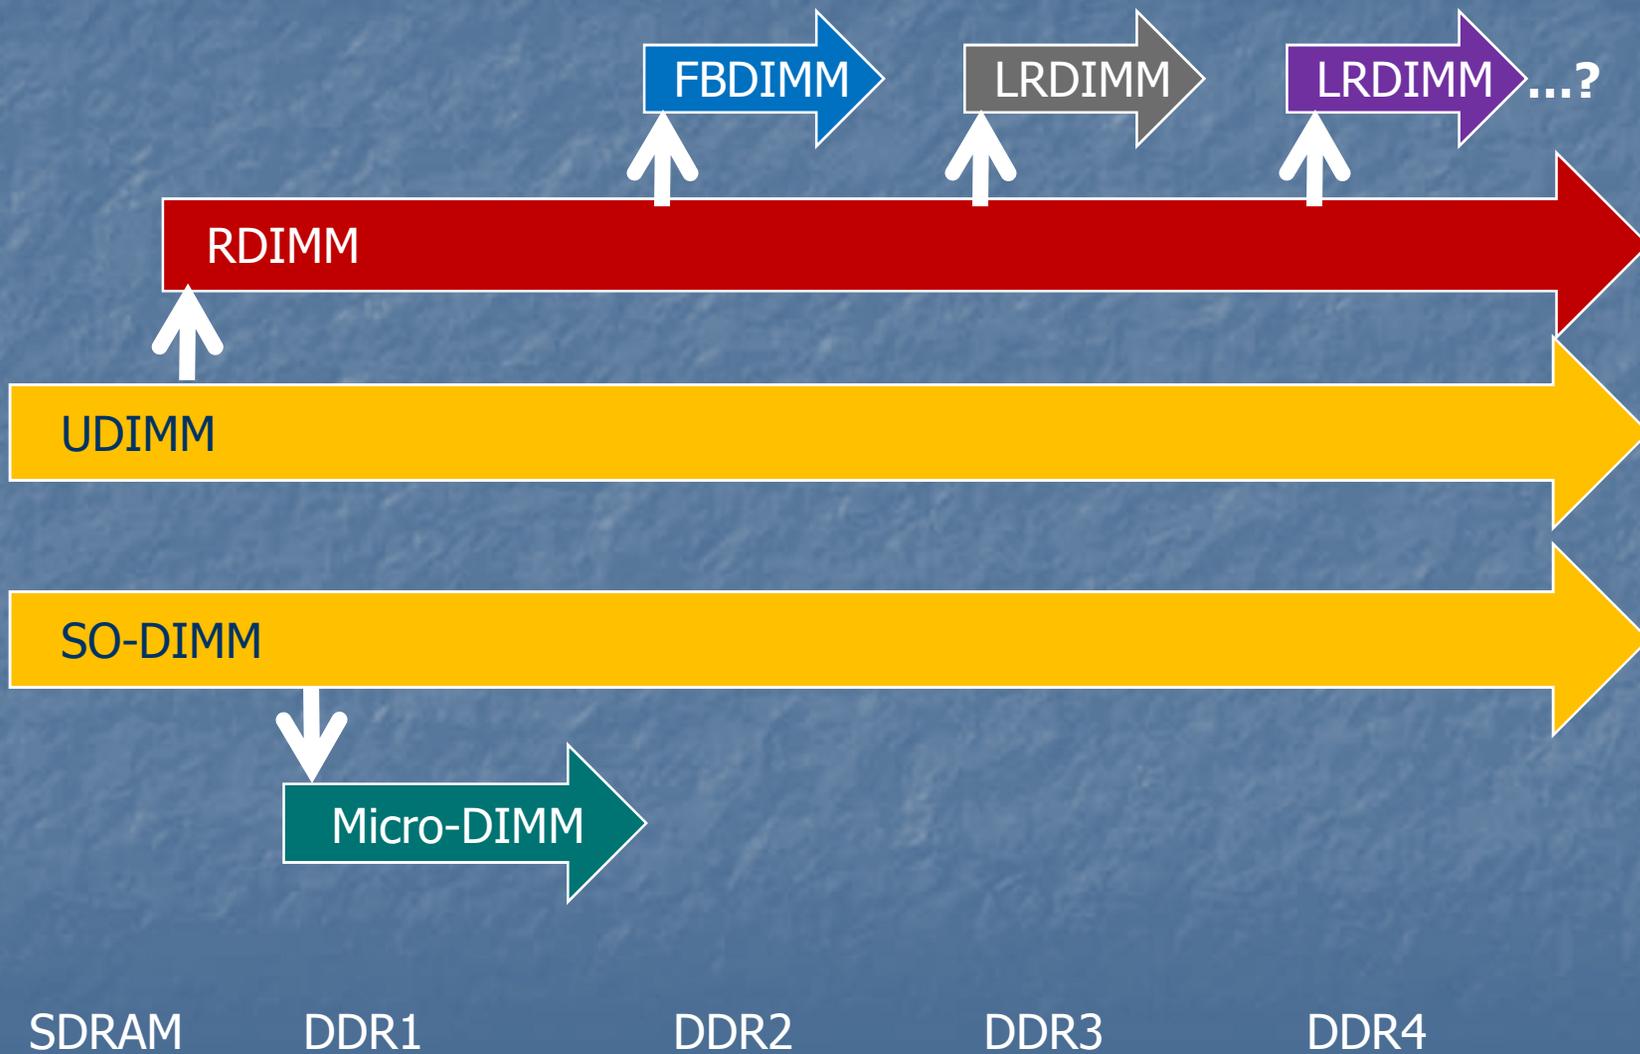
Required 28Ω @ 10% motherboard routing
(Mainstream mobo was 60Ω @ 15%)

BGA package in a TSOP world

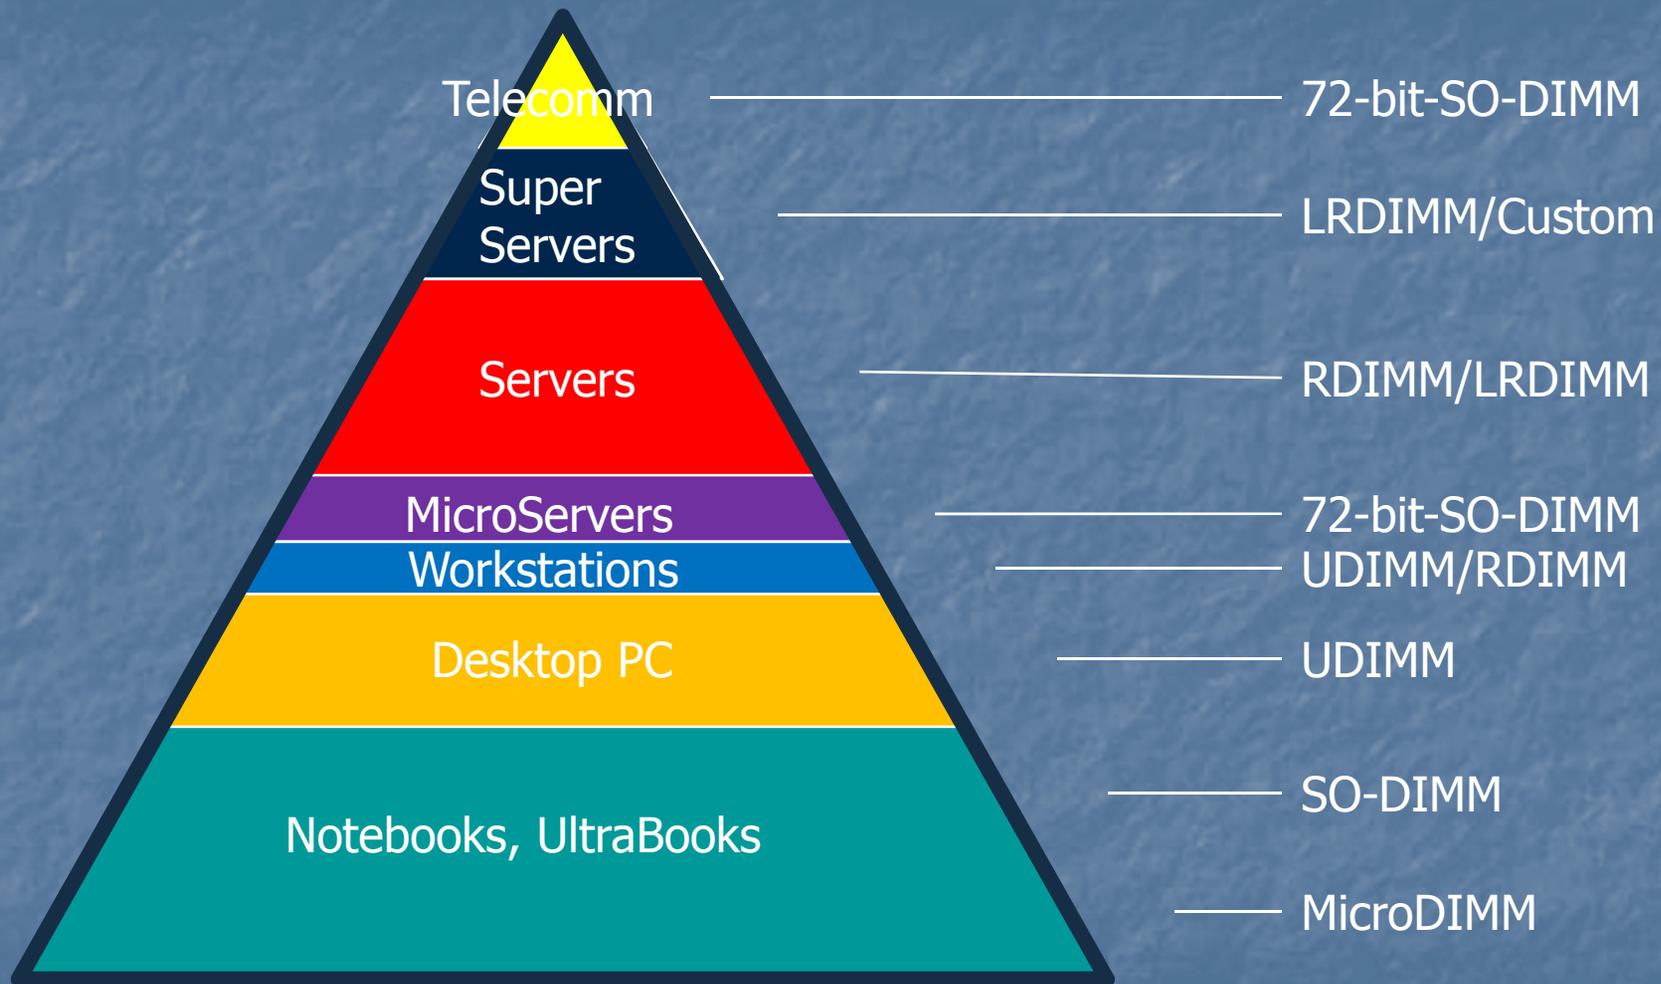
SERDES: ran hot & longer latency



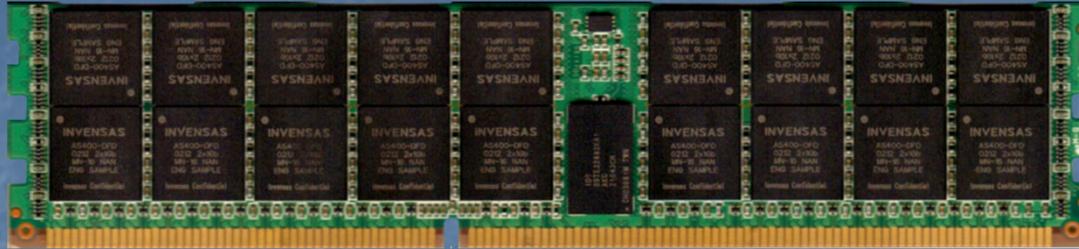
Memory Module Path



Memory Module Market



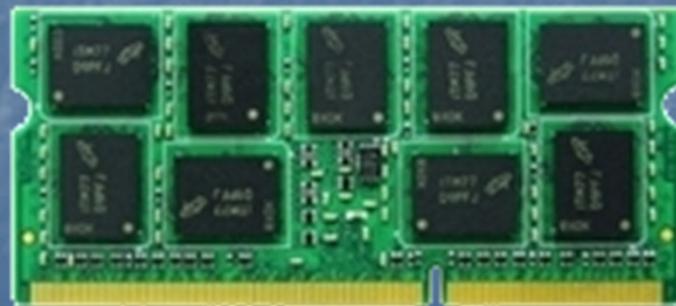
Market Opportunities



50% = LP DIMM for desktop, servers



50% = SO-DIMM for notebook market



Emerging% = 72b-SO-DIMM for Micro Server market



Successes & Failures

UDIMM & SO-DIMM served their markets
Desktop & Notebook

RDIMM = drop-in compatible with UDIMM

Micro-DIMM went out with subnote market

DDR2 FB-DIMM = Rambus for RDIMMs

Incompatible interface

SERDES: ran hot & longer latency

DDR3 LRDIMM architected poorly



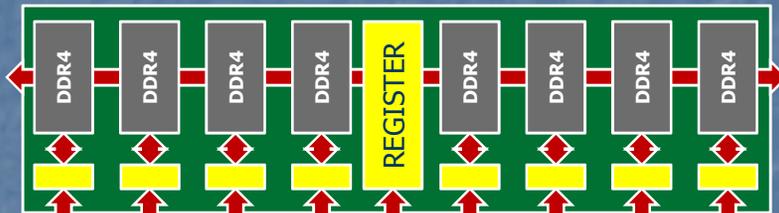
Hopefully DDR4 LRDIMM is Smarter



DDR3 Memory Controller

DDR3 LRDIMM

All data & address flowed into a single central chip
Redriven to all DDR3 SDRAMs
Long stubs hurt multi-DPC



DDR4 Memory Controller

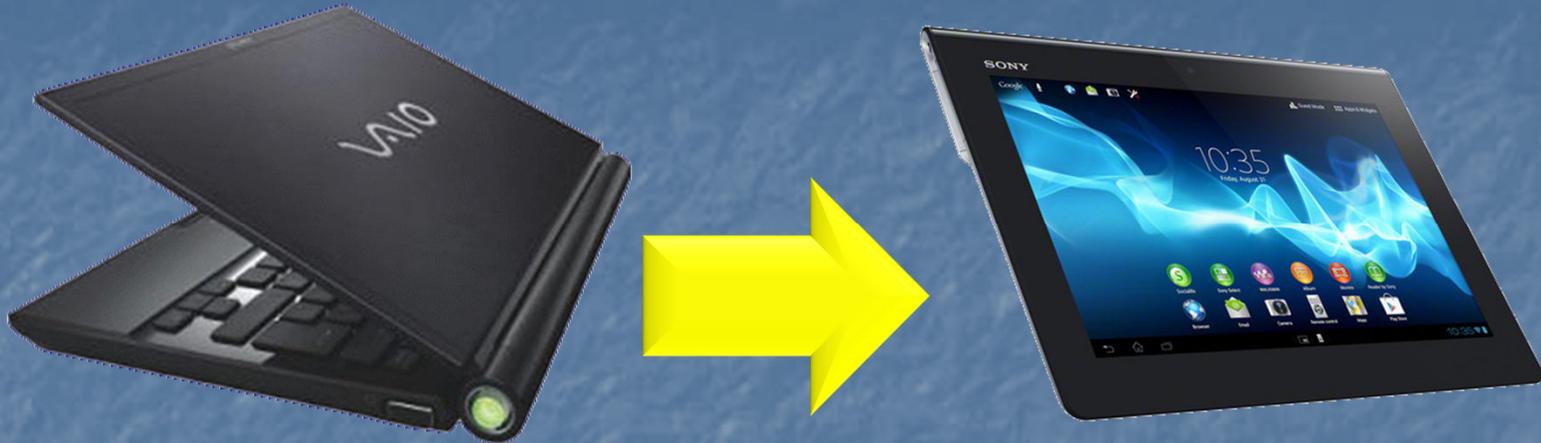
DDR4 LRDIMM

Data connected to data buffers
Address connected to register and redriven
Short stubs aid multi-DPC

(DPC = DIMMs per channel)



Market Shifts



What impact does the shift from notebook to tablets have?

For now, solder down memory instead of sockets



Do the Tablet and Notebook Merge?



Searching the web without a keyboard sucks...

...So what's the difference between a notebook and tablet...

...if you add a keyboard?



Cloud as a Mainframe?



Trend is towards cloud storage

Trend is toward cloud apps

Result:

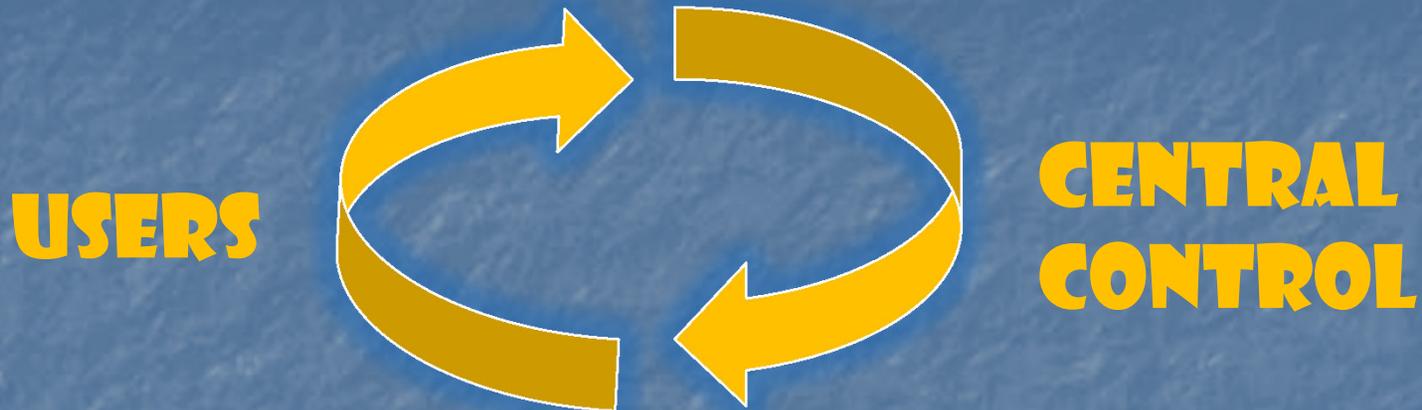
Less storage in client devices

Major increase in cloud storage

Major increase in cloud CPU cycles



Is This Another Cycle?



Cloud failures could force a recycle

MySpace

Friendster

...Facebook?

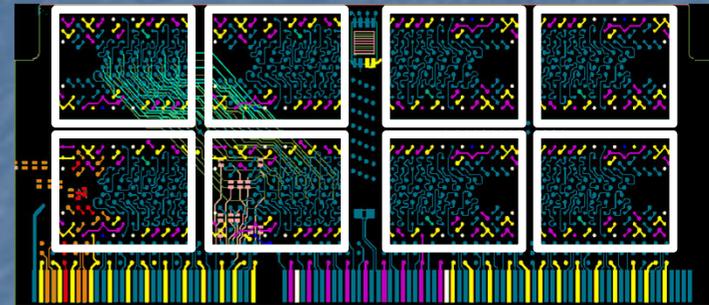
Loss of personal data could flip the cycle



Solder Down Versus Micro-DIMM



~ 17.5 x 22.5 mm (0.7 x 0.9 ")



~ 45 x 20 mm (1.8 x 0.8 ")

But rationale for socketed memory doesn't go away...

Insulation from DRAM price fluctuations

Build-to-order configurability

Likely to see Micro-DIMMs make a comeback



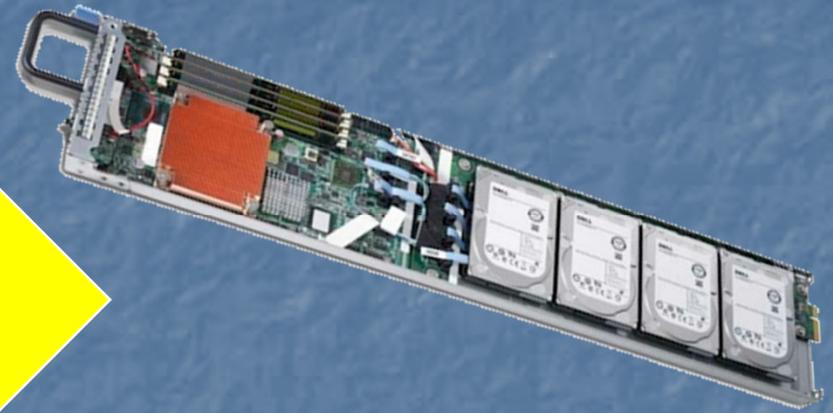
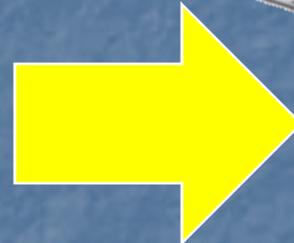
Servers to Microservers?



RDIMM



~ 133 x 30 mm (5.25 x 1.2 ")



72b-SO-DIMM



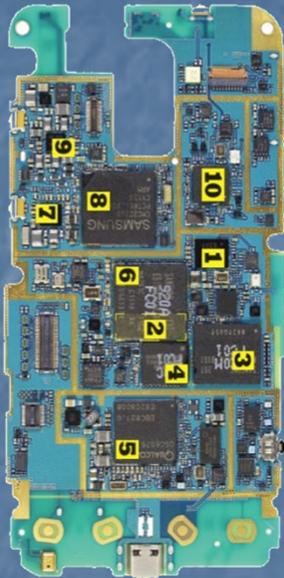
~ 68 x 30 mm (2.7 x 1.2 ")

Shrinking size with increased density requirement

Imaginative packaging will be needed



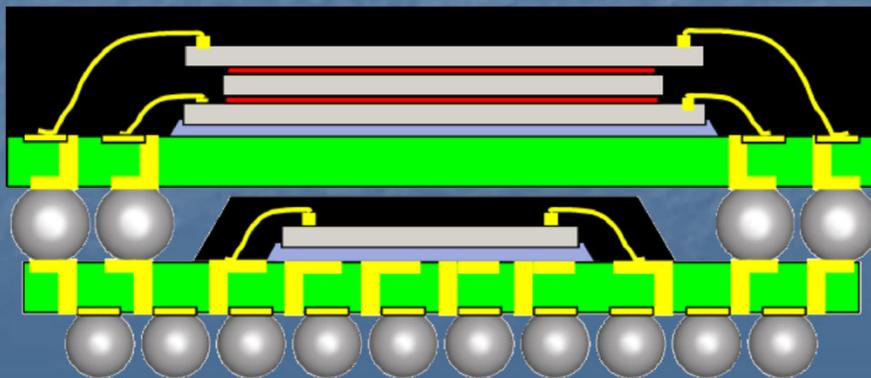
Oh, Yes, and Cell Phones



Extremely small form factors

Extremely low power

Extremely poor cooling



Package on Package popular

Edge bonded die required

Mix of technologies



To Fight the Unbeatable Foe

We engineers continually face the impossible math...

Smaller + Denser + Faster

= the same price



Changing the Rules

Increasing gap between DRAM generations

- ⊕ Increasing demands for high capacity
- ⊕ Frequency increases still quite aggressive
- ⊖ Rethinking module & chip architectures



Moving the Roadblocks

Reduce loading

Reduce stub wire length

Resistors to buffer & terminate signals



Moving the Roadblocks

Socket is the limiting factor, capacity requirements low?

→ Move to solder down

UDIMM loading prevents address bus from getting better?

→ Move applications to RDIMM

RDIMM loading prevents data bus from getting better?

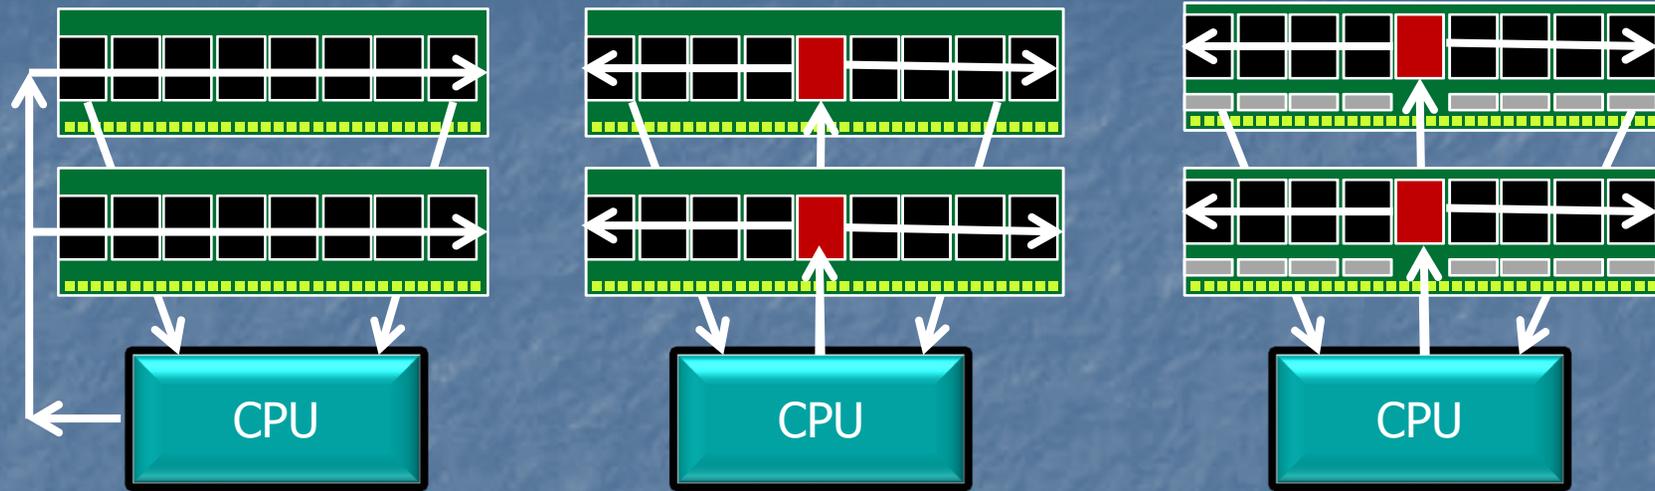
→ Move applications to LRDIMM

LRDIMM loading prevents capacity from getting better?

→ Move to memory hubs



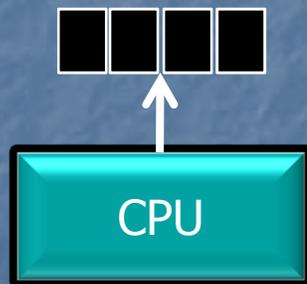
Current Lay of the Land



Unbuffered

Registered

Load Reduced

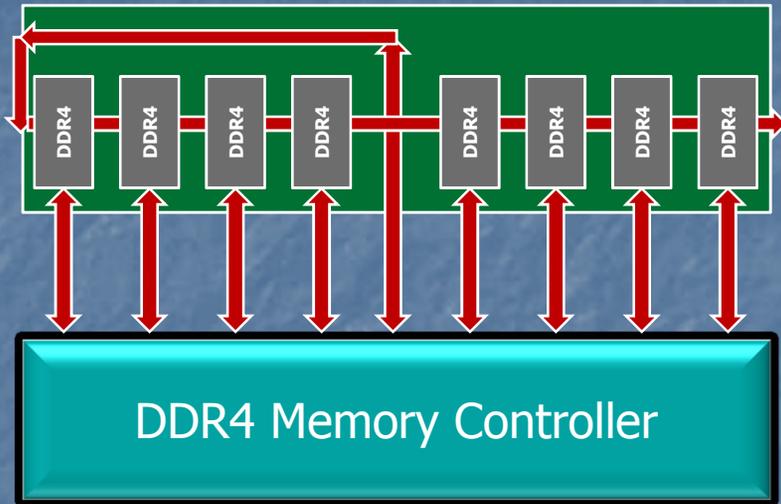


Solder Down

Each solution has tradeoffs regarding frequency, latency, capacity... and cost

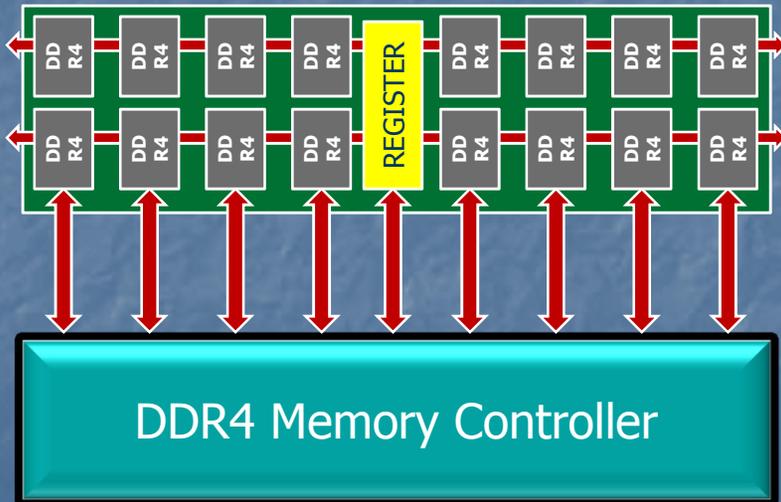


UDIMM vs RDIMM



UDIMM

Address bus sees 0 to 18 loads per slot
Data bus sees 0 to 2 loads per slot

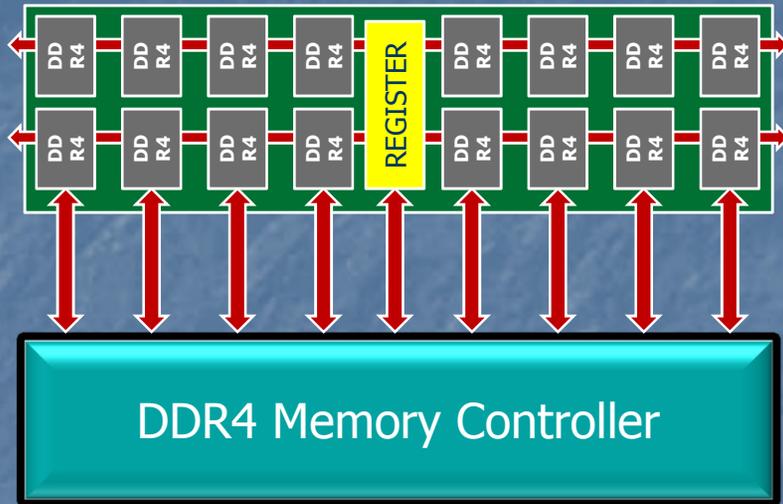


RDIMM

Address bus sees 0 to 2 loads per slot
Data bus sees 0 to 4 loads per slot

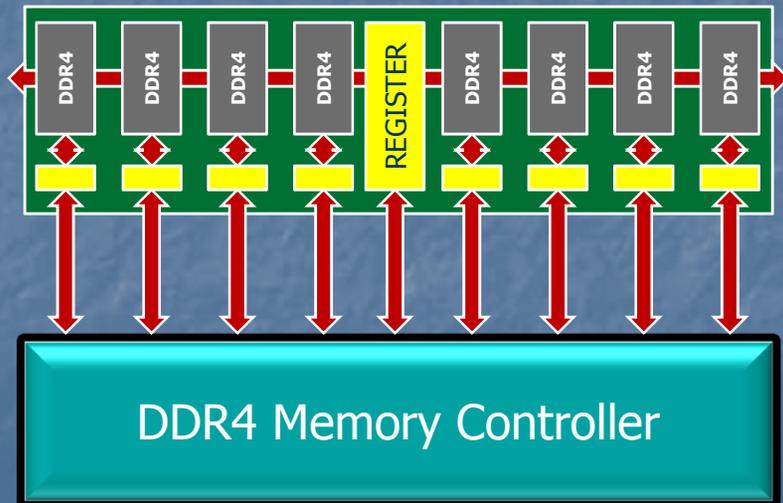


RDIMM vs LRDIMM



RDIMM

Data bus sees 0 to 4 loads per slot
Matrix of 63 combinations to test
for a 3 slot system

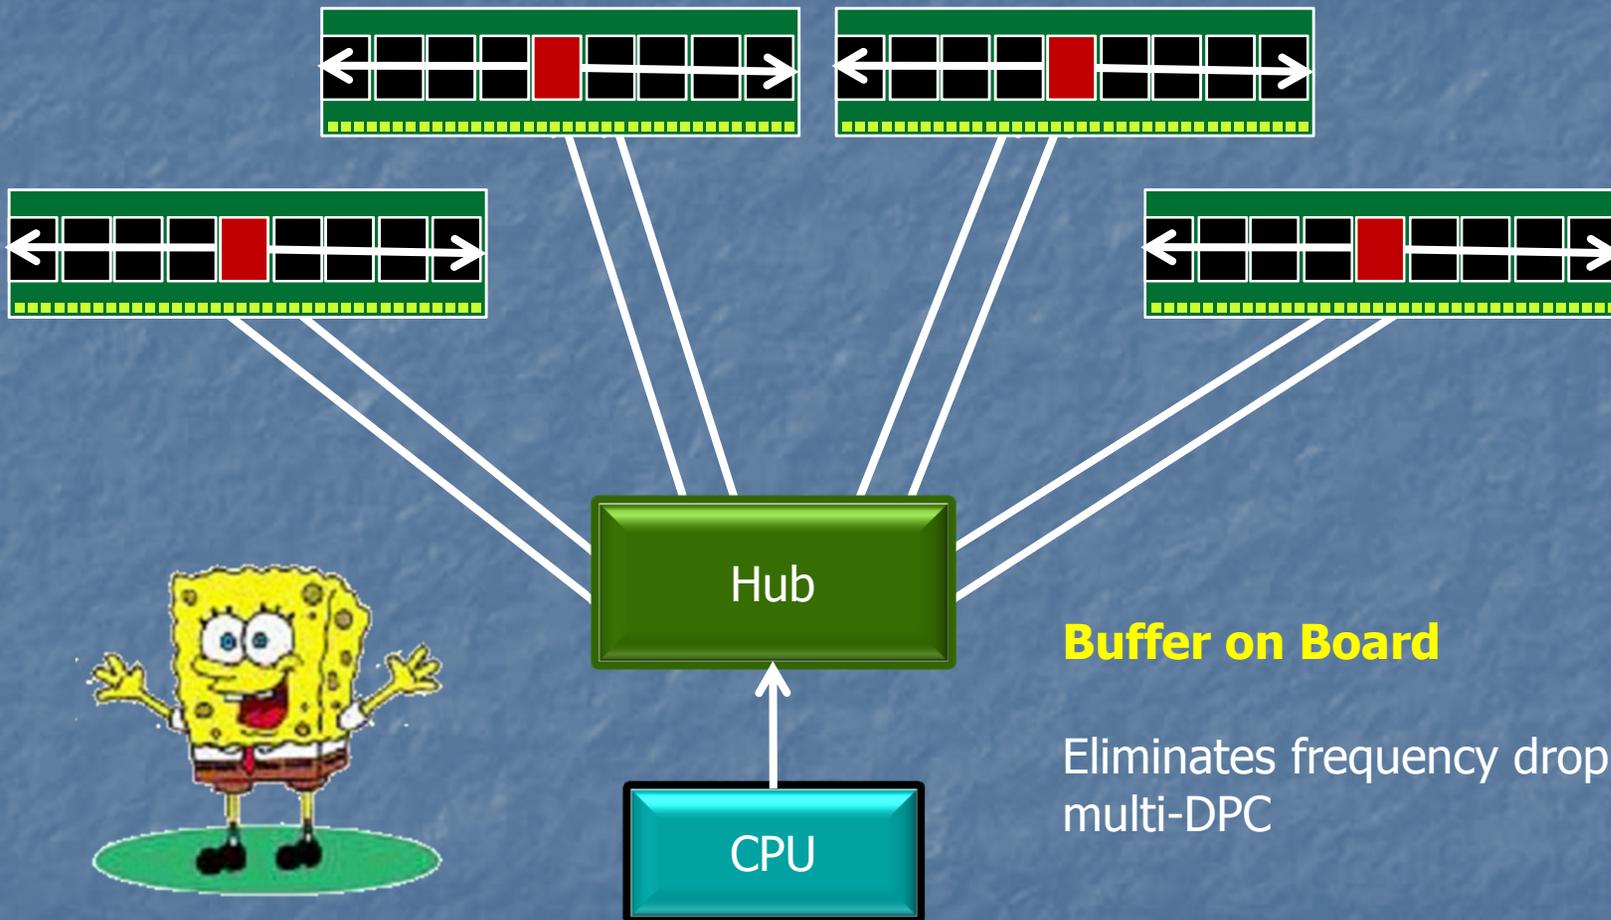


LRDIMM

Data bus sees 0 or 1 loads per slot
Matrix of 7 combinations to test
for a 3 slot system



Memory Hub

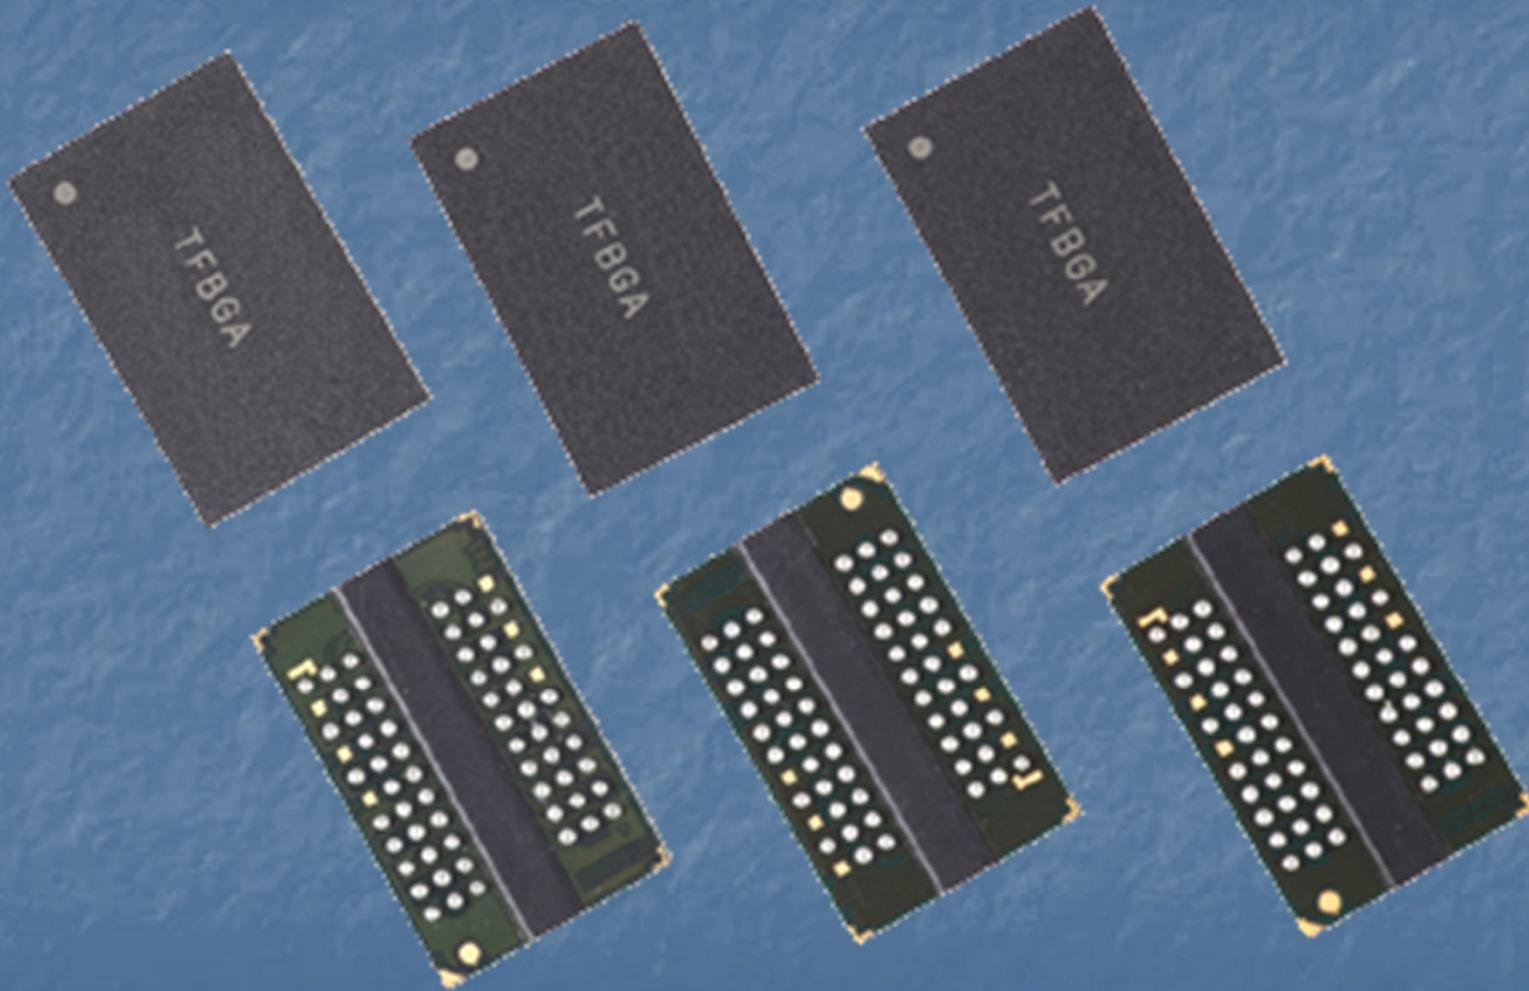


Buffer on Board

Eliminates frequency drop from multi-DPC



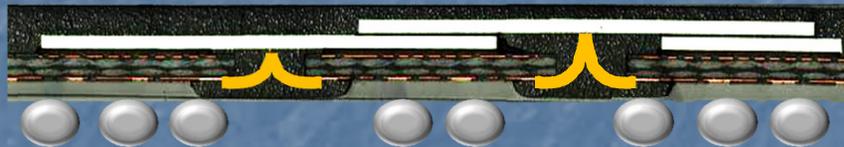
What About at the Chip Level?



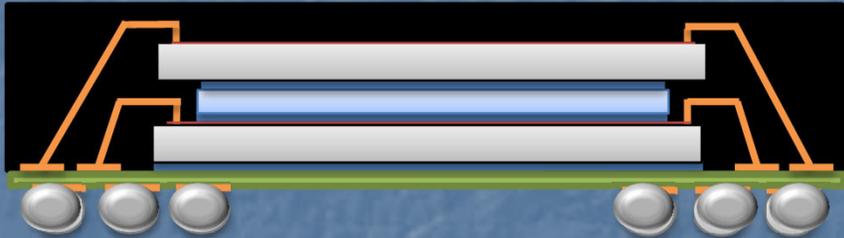
Stacking Technology Matters



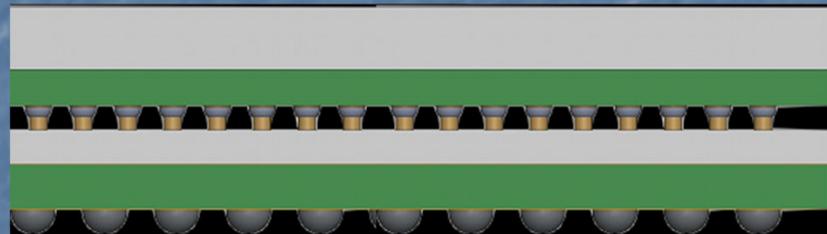
Classic window BGA



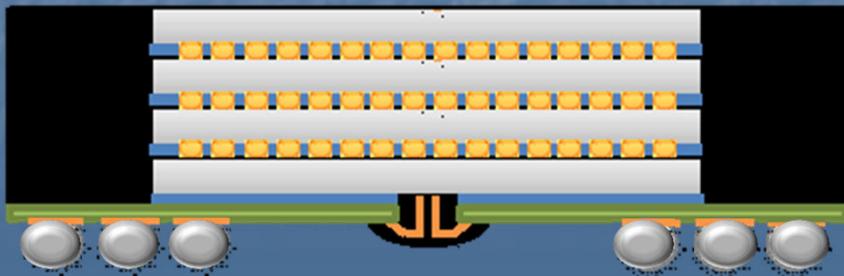
Dual Face Down DDP



RDL DDP



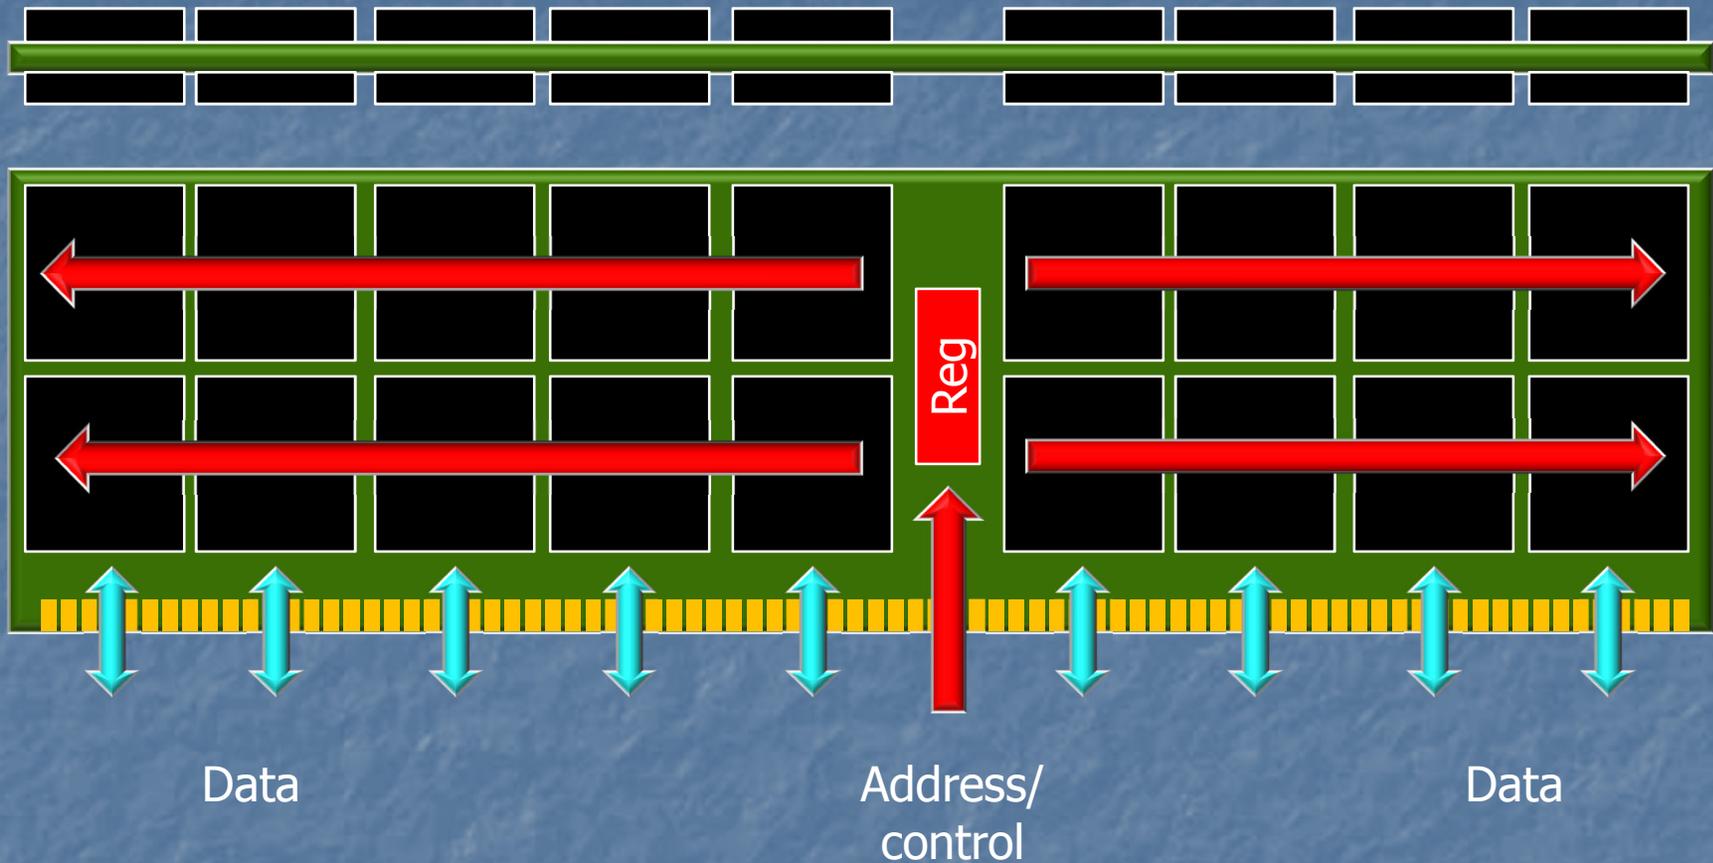
Bond Via Array PoP



3DS with
through silicon vias



Example: RDIMM Architecture



Data

Address/
control

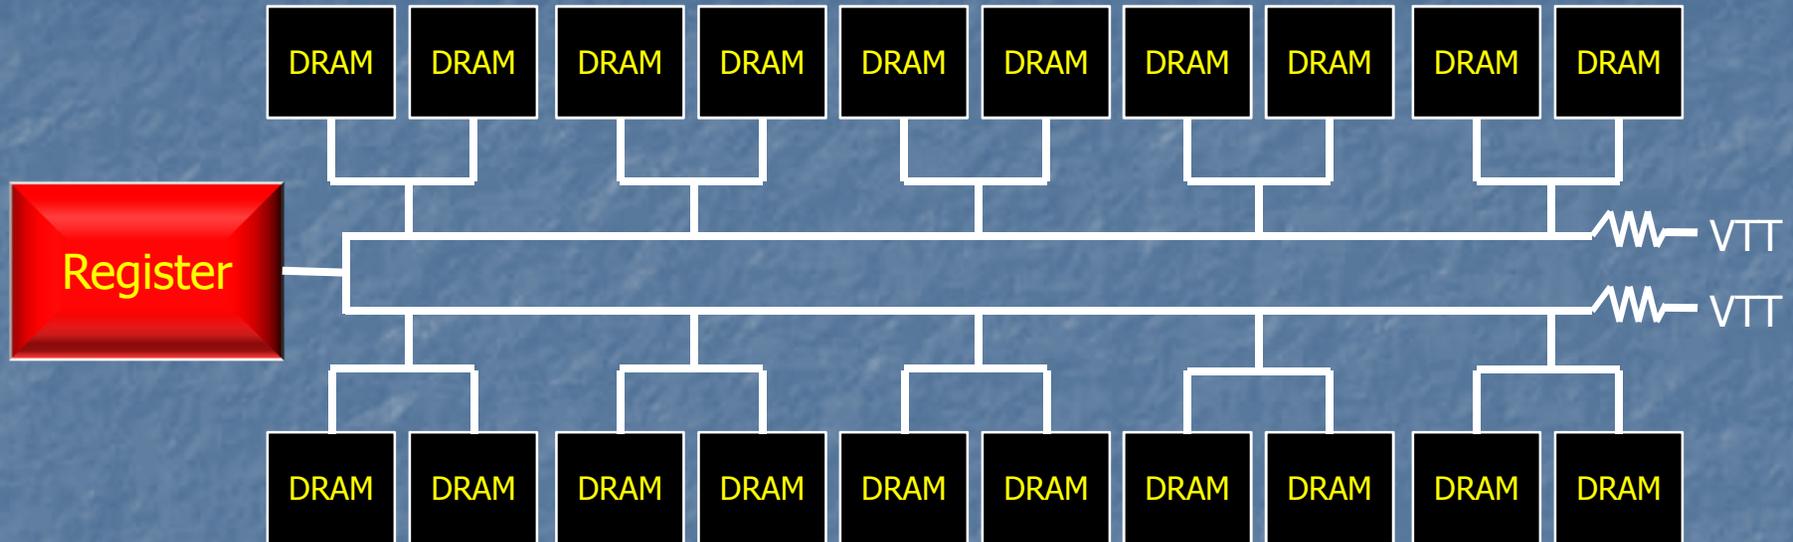
Data

Address and control: one load per slot

Data: one to four loads per slot



RDIMM Flyby Topology

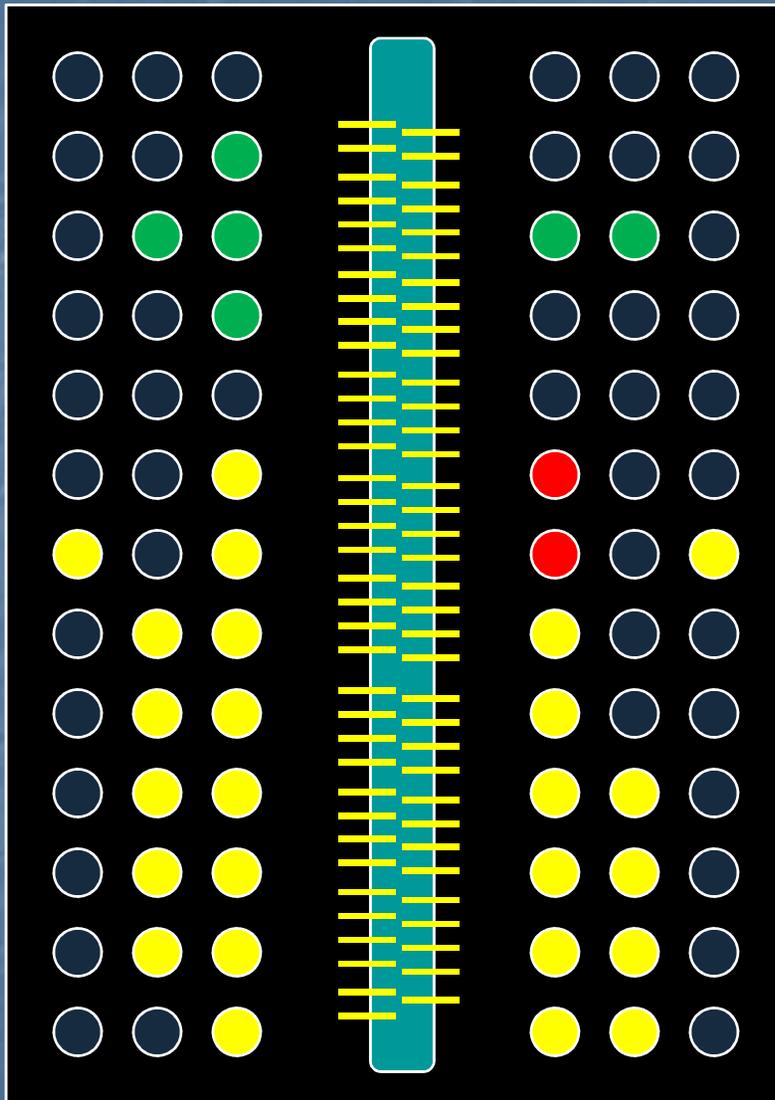


Typical two rank design

External resistor termination to VTT



Standard Monolithic DRAM Package

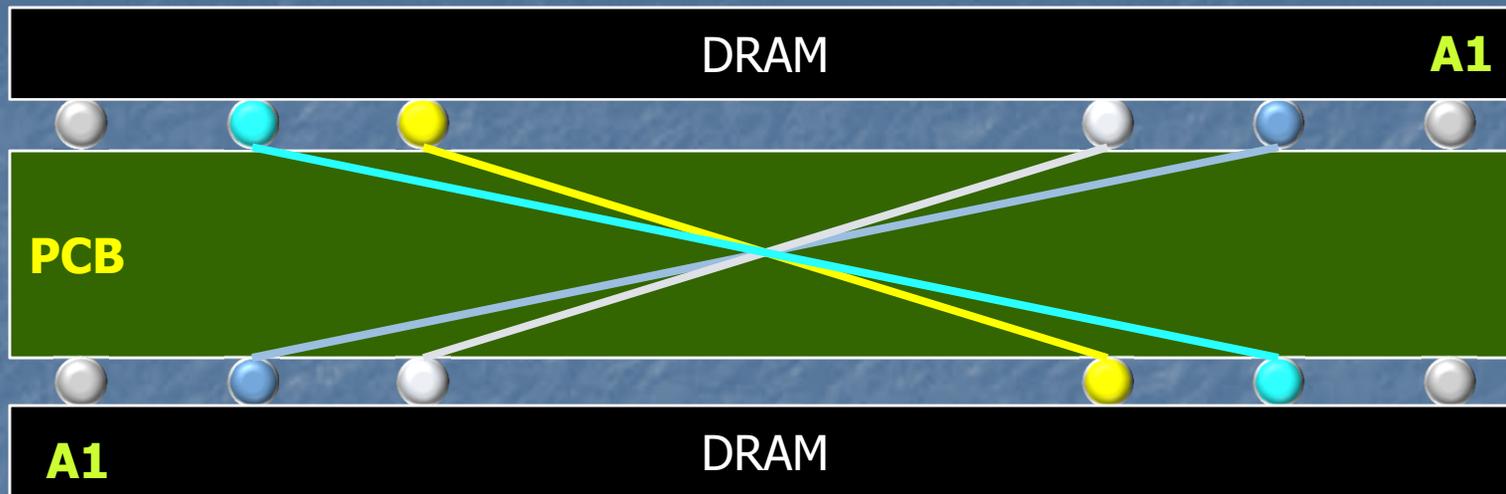


- Data signals
- Address/control signals
- Clocks
- Gold wire bonds

**Designed around needs
of monolithic DRAM**



Double-Sided Board Problem



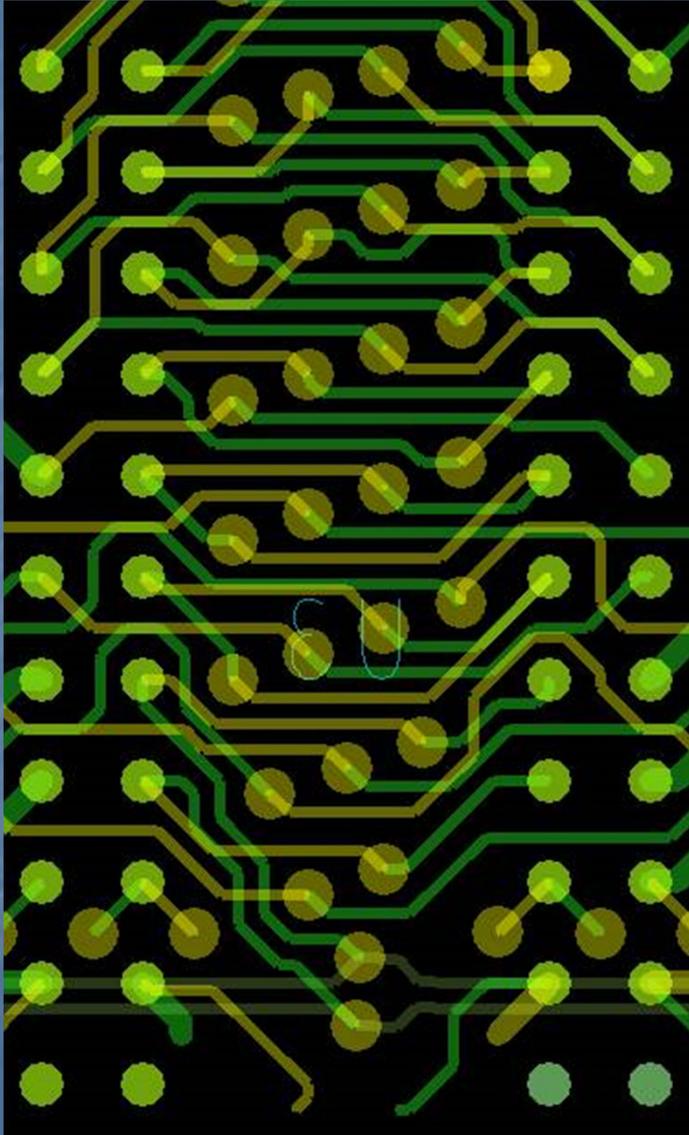
The "Bowtie" problem

**Matching address signals
are diagonally opposed**

Results in long stubs



Bowtie Impact on Layout



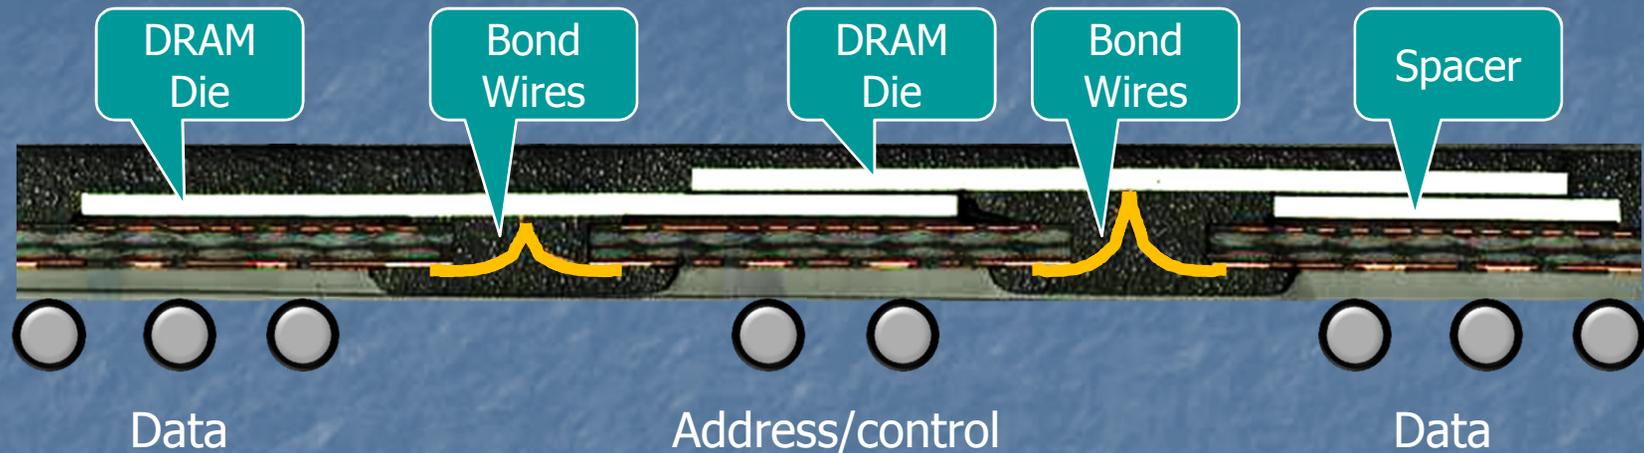
Two routing layers just for bowtie

More crowded if stub length matching is done

Signal quality issues if stub length matching is not done



Rethinking Stacking



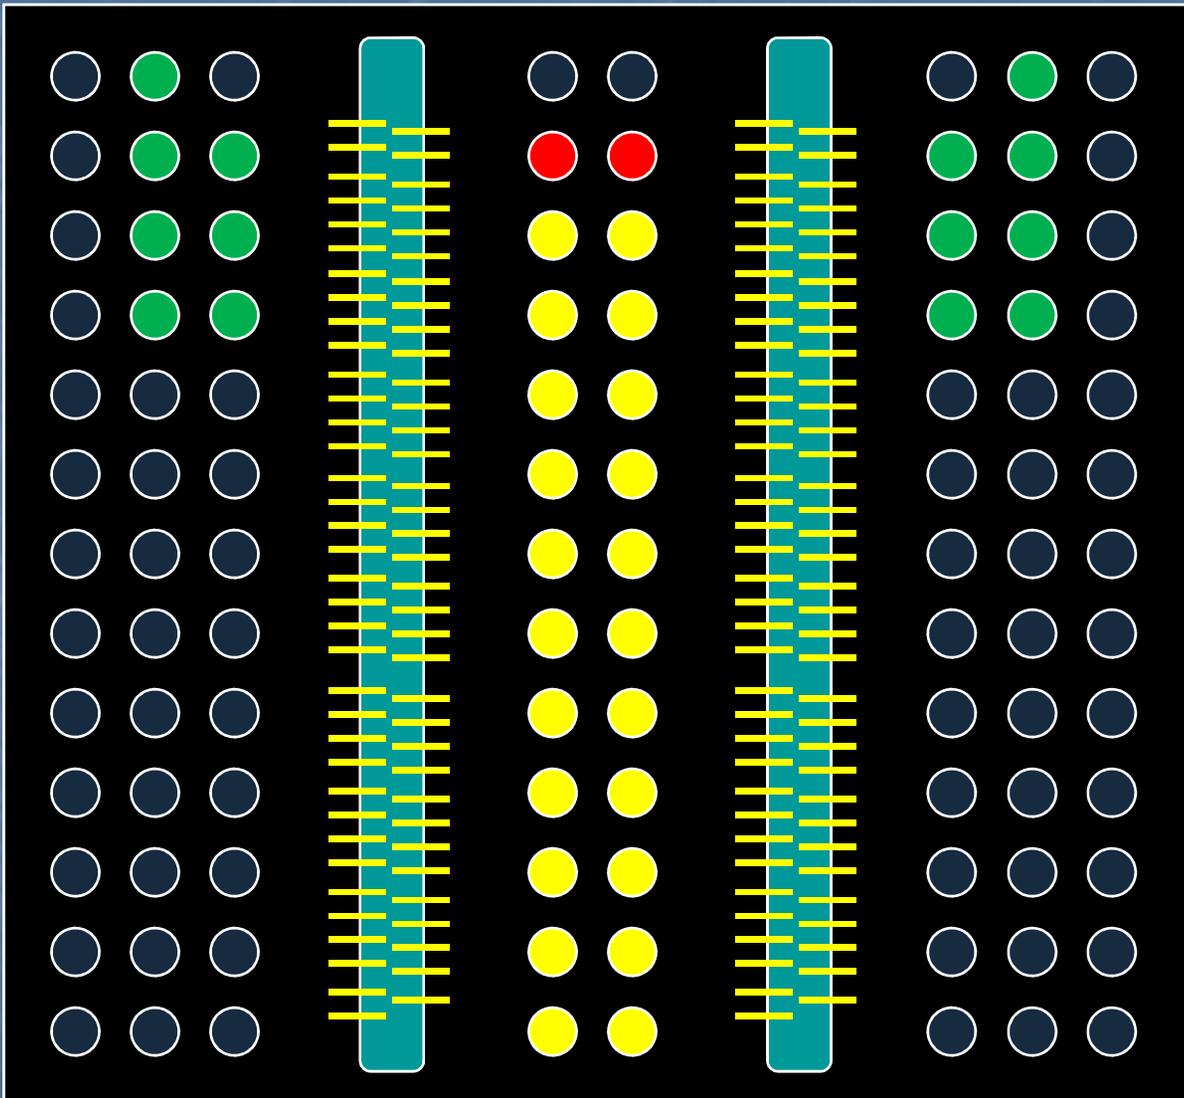
Matched parasitics between DRAMs

Excellent power delivery

Reduced thermal envelope



DFD Ballout

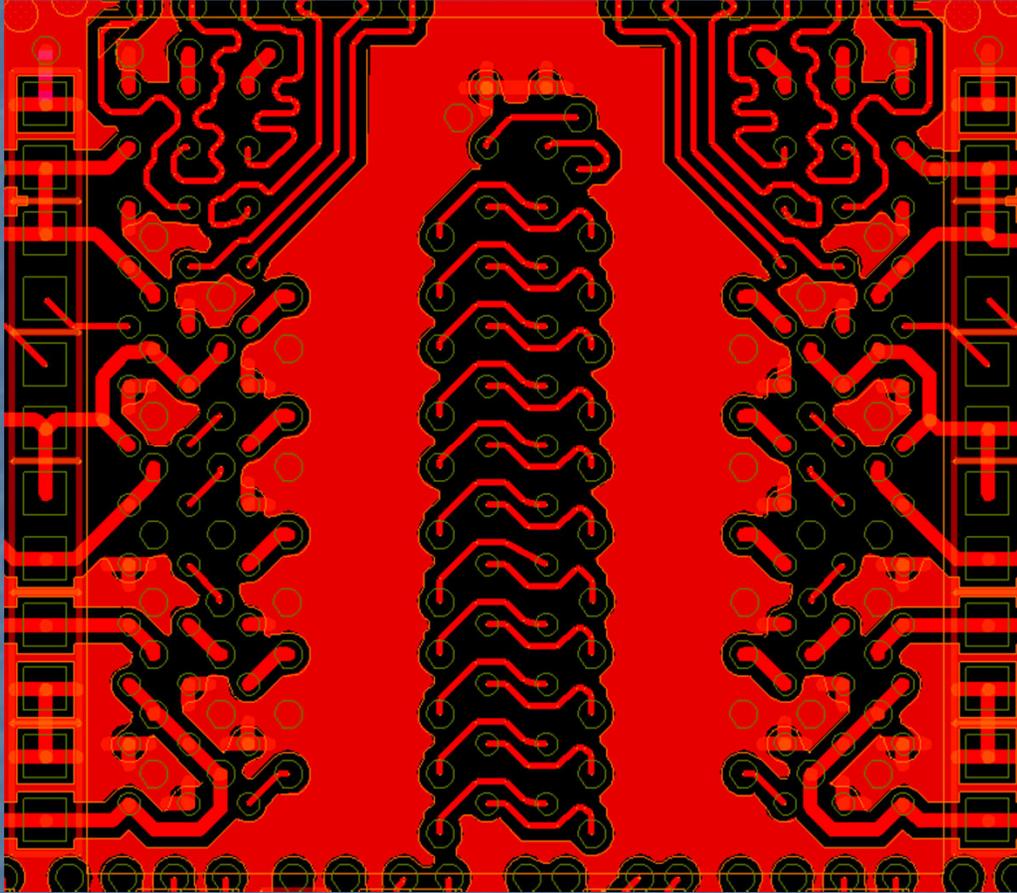


- Data signals
- Address/control signals
- Clocks
- Gold wire bonds

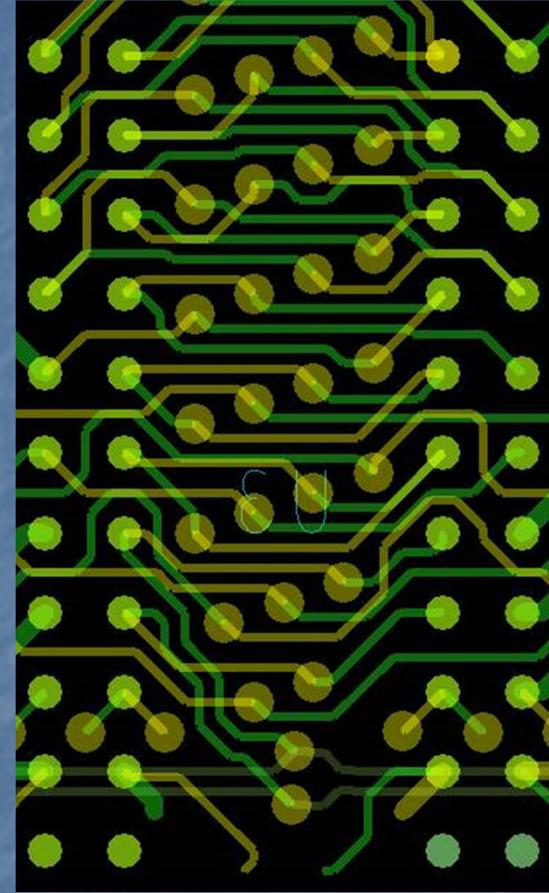
Package size =
11.5 x 11.5 mm



Compare DFD to DDP Routing



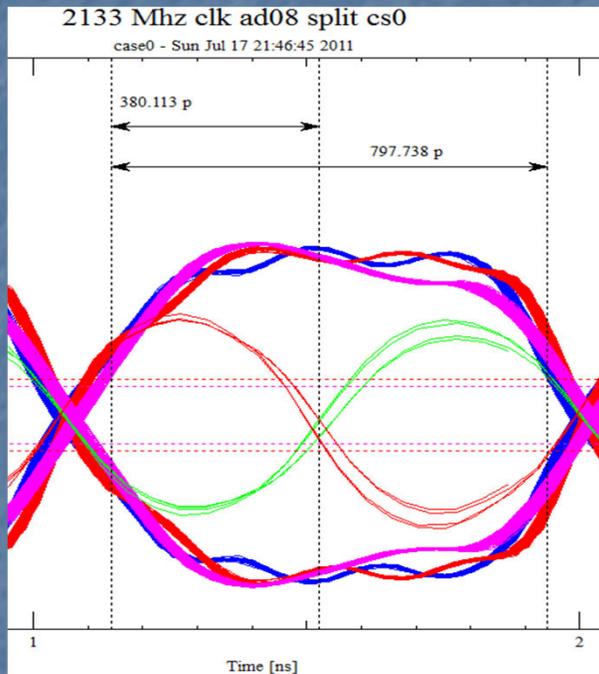
DFD



DDP

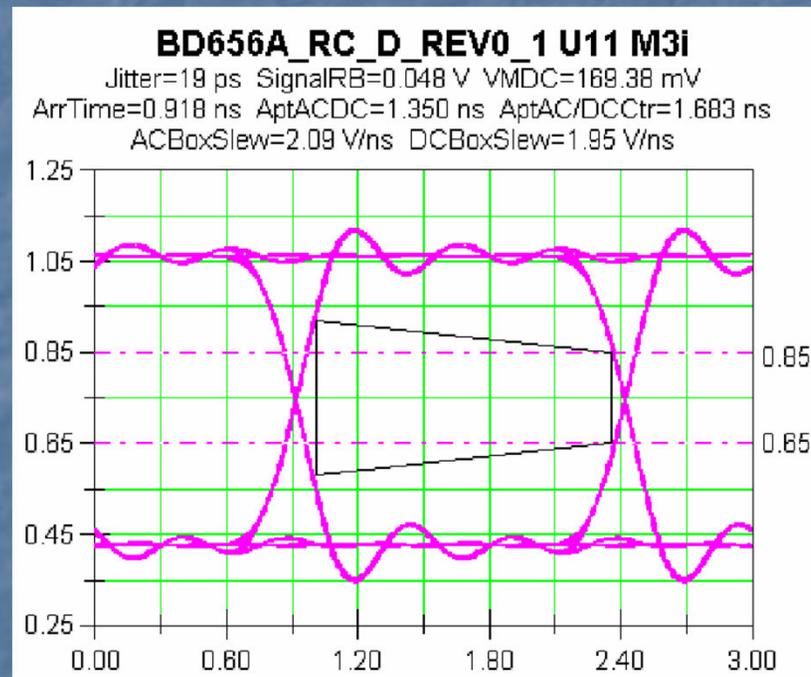


Results from Reducing Stub Length



DFD @ 2133

Window = 798 ps
Ideal tCK = 938 ps
85% of a tCK



DDP @ 1333

Window = 1093 ps
Ideal tCK = 1500 ps
73% of a tCK



Miniaturization

10 pounds of #\$\$&@

Miniaturization forcing everything into smaller and smaller form factors

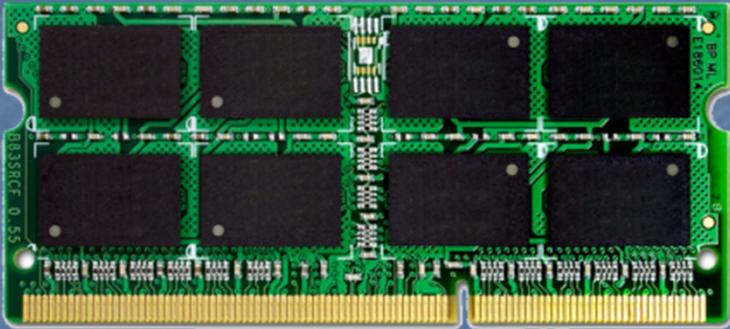
Increased frequency requiring more use of termination resistors

What's an engineer to do???

5 lb



Embedded Resistors



Traditional SMT



With Embedded Resistors

Eliminates between 35 and 200 placed parts

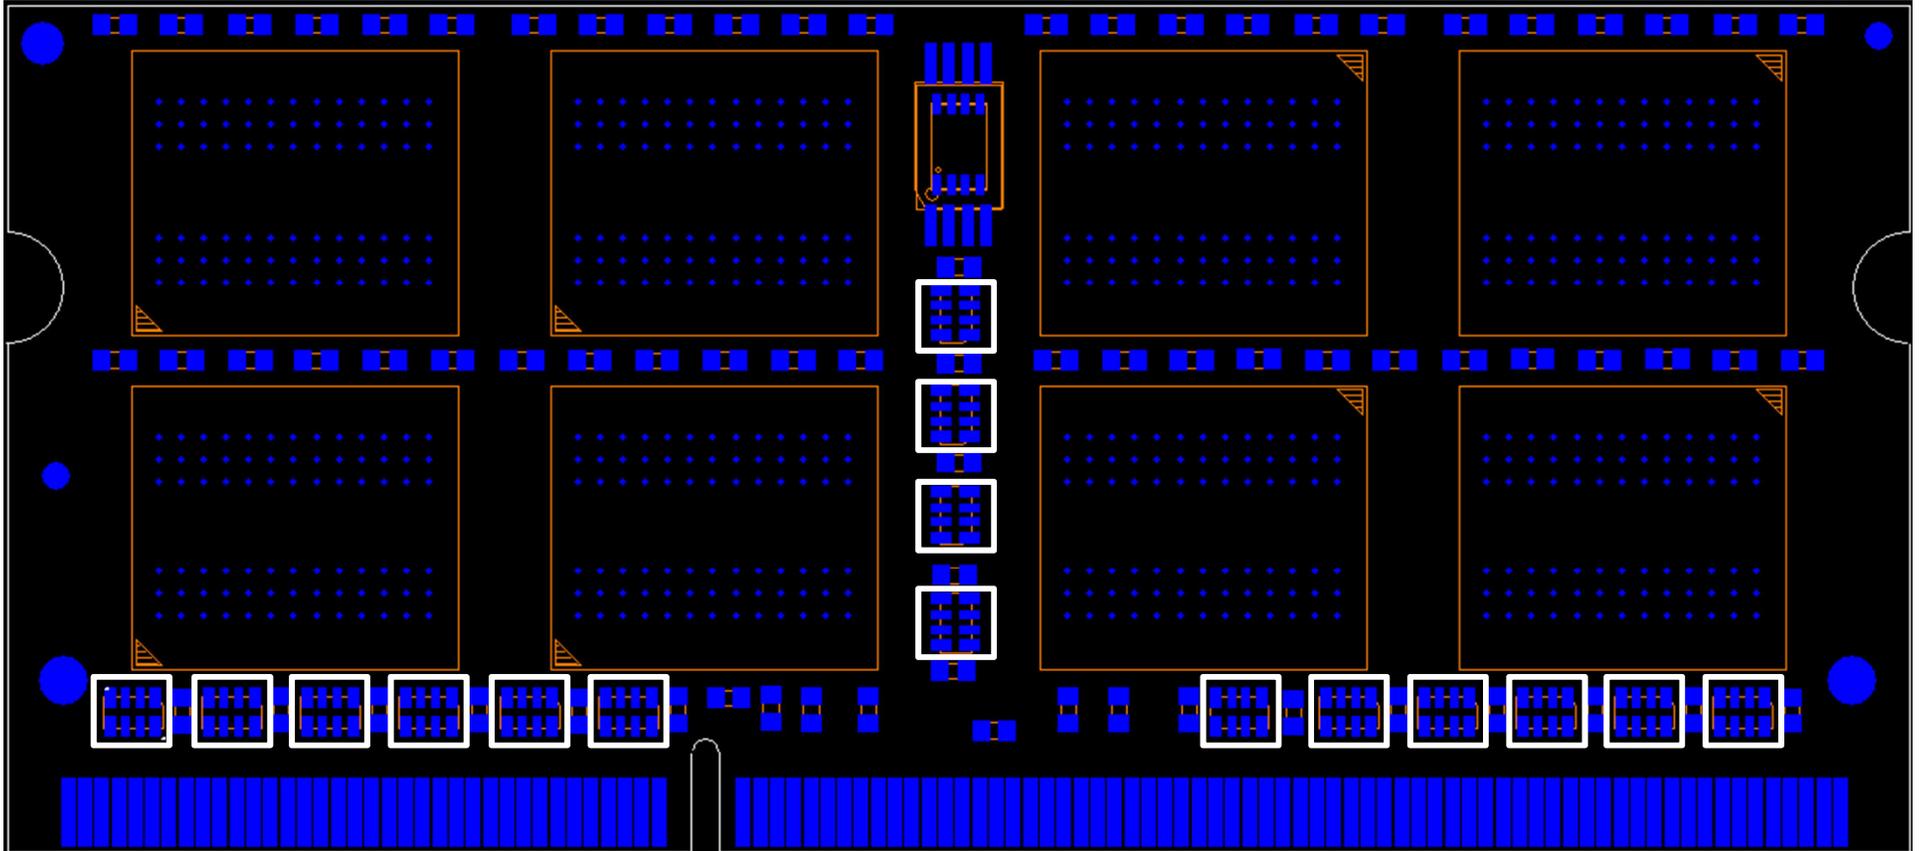
Eliminates the need for 0201 or smaller parts

More room for what customers actually pay for... GB/cm³

Signal integrity improvements enabled with design tricks



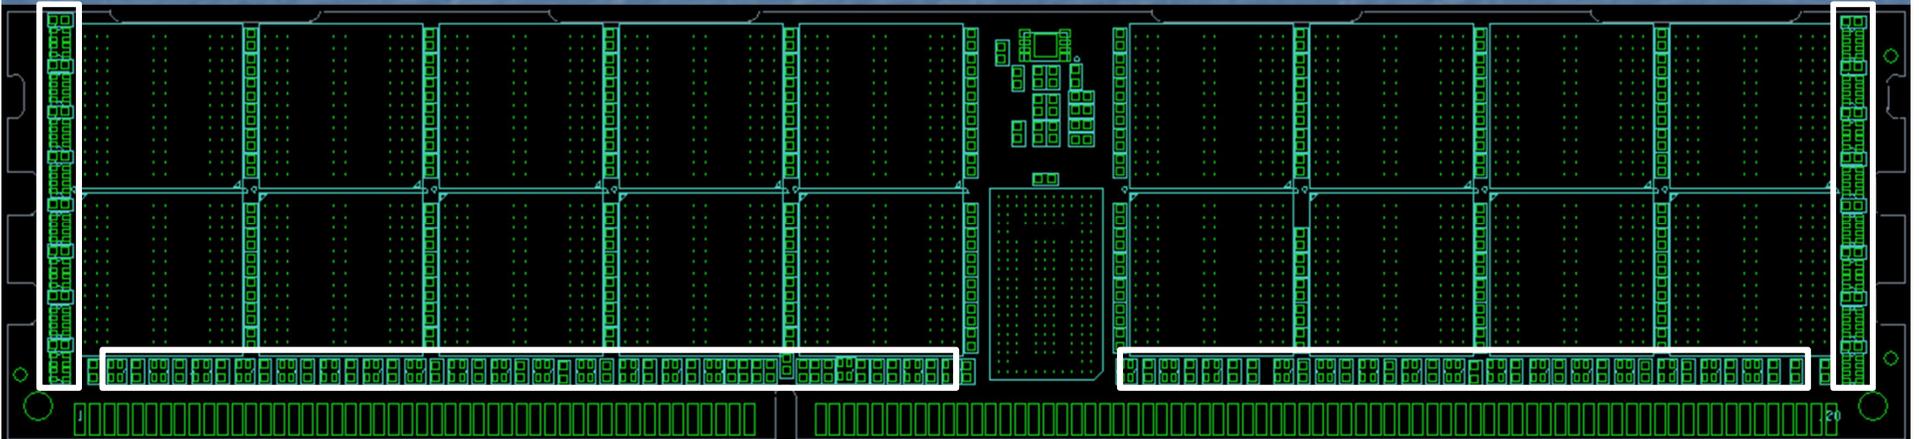
Example of Space Saved



**64 mm² of layout space saved
from 1500 mm² total component area
→ 4.3% saved using embedded resistors**



More Extreme Example

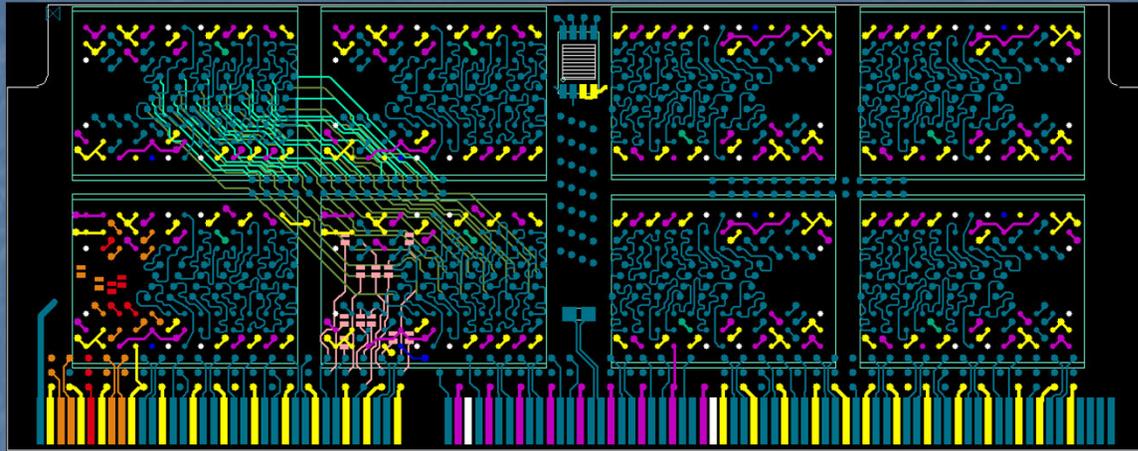


RDIMM

600 mm² of layout space saved
from 3180 mm² total component area
→ 19% saved using embedded resistors



Enable the “Impossible”



56 x 22 mm

MicroDIMM concept drawing

Single sided assembly

8 DRAMs + SPD

Series termination on every data bit

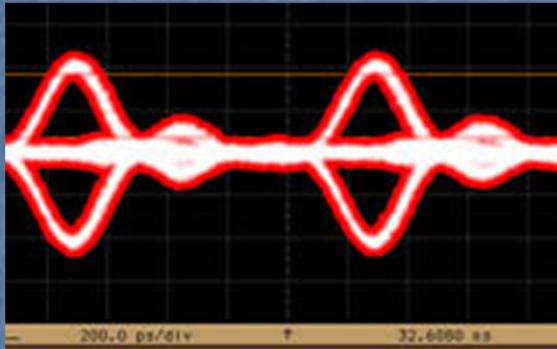
Two parallel terminations on every address

SMT resistors would consume 7% of layout



Improving Signal Quality

50Ω SMT resistor

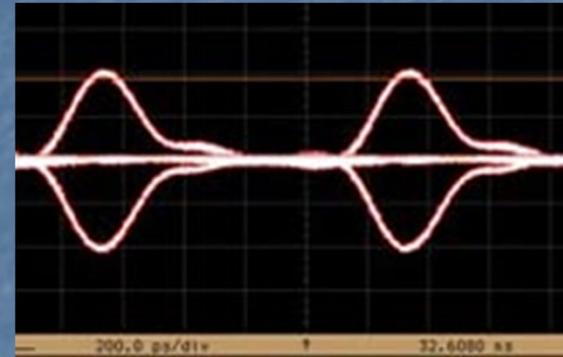


Mechanical structure has impedance problems



Signal reference suffers through body, solder joints, physical spacing

50Ω embedded resistor



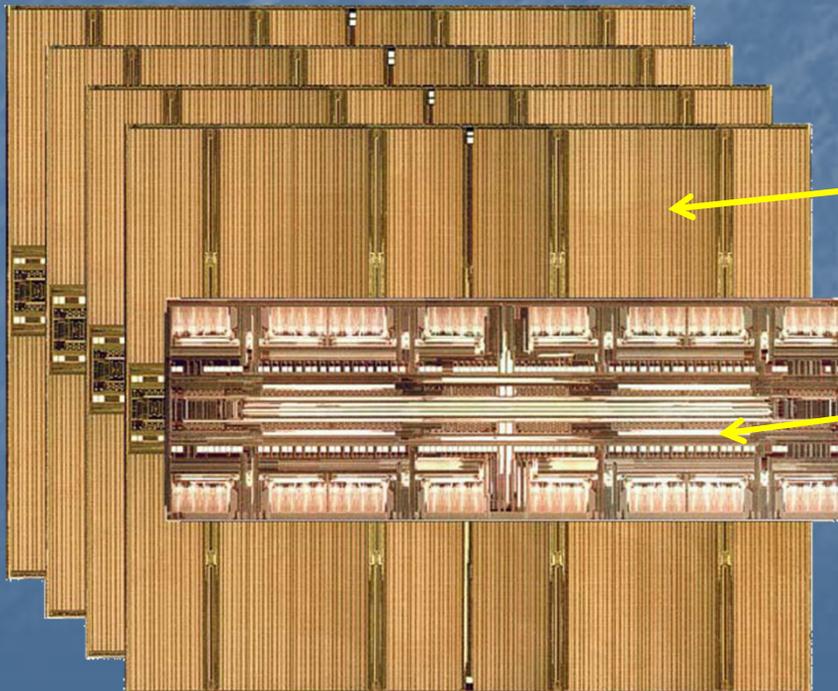
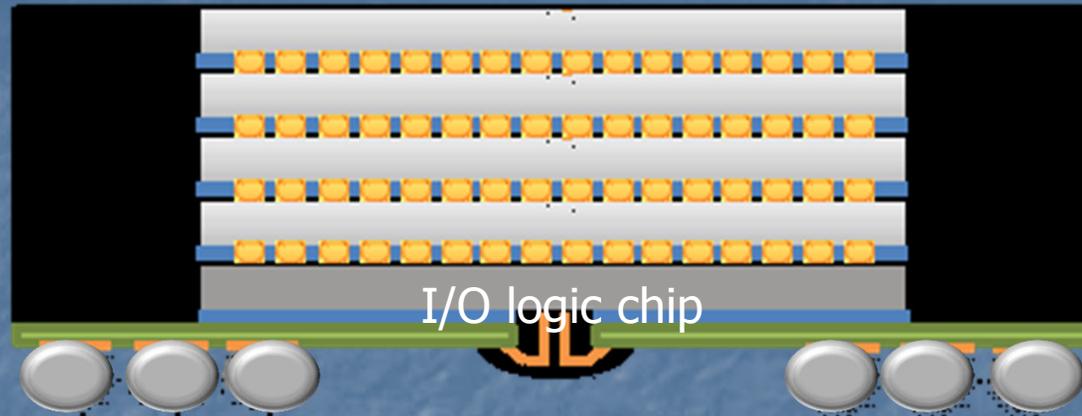
Impedance held through routing



Signal reference maintained throughout



3DS Memory Stacks



DRAM dies

I/O logic die



Current Status of TSV

Manufacturability of TSV getting better

Die thinning increases yield fallout

Known good die with speed binning not possible at this time



Compound Yield Problem



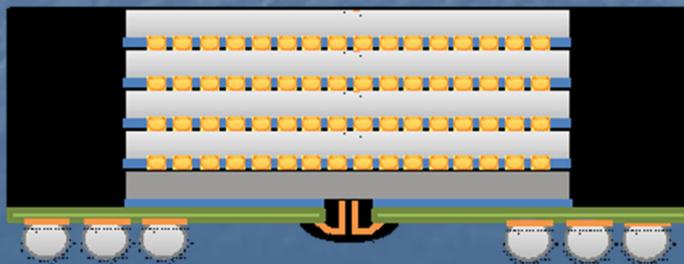
Yield 90% to go to market



Yield 99% for I/O chip x
90% for DRAM = **89% yield**



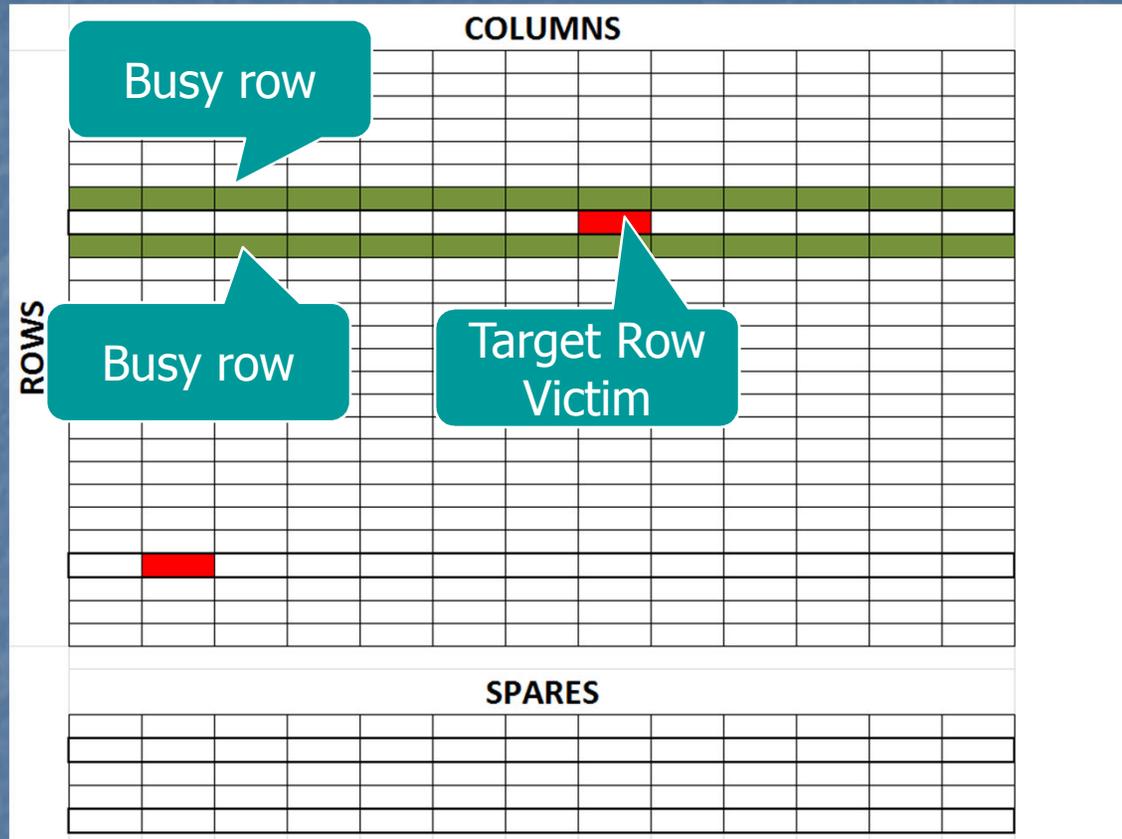
Yield 99% for I/O chip x
90% for DRAM = **80% yield**



Yield 99% for I/O chip x
90% for DRAM = **65% yield**



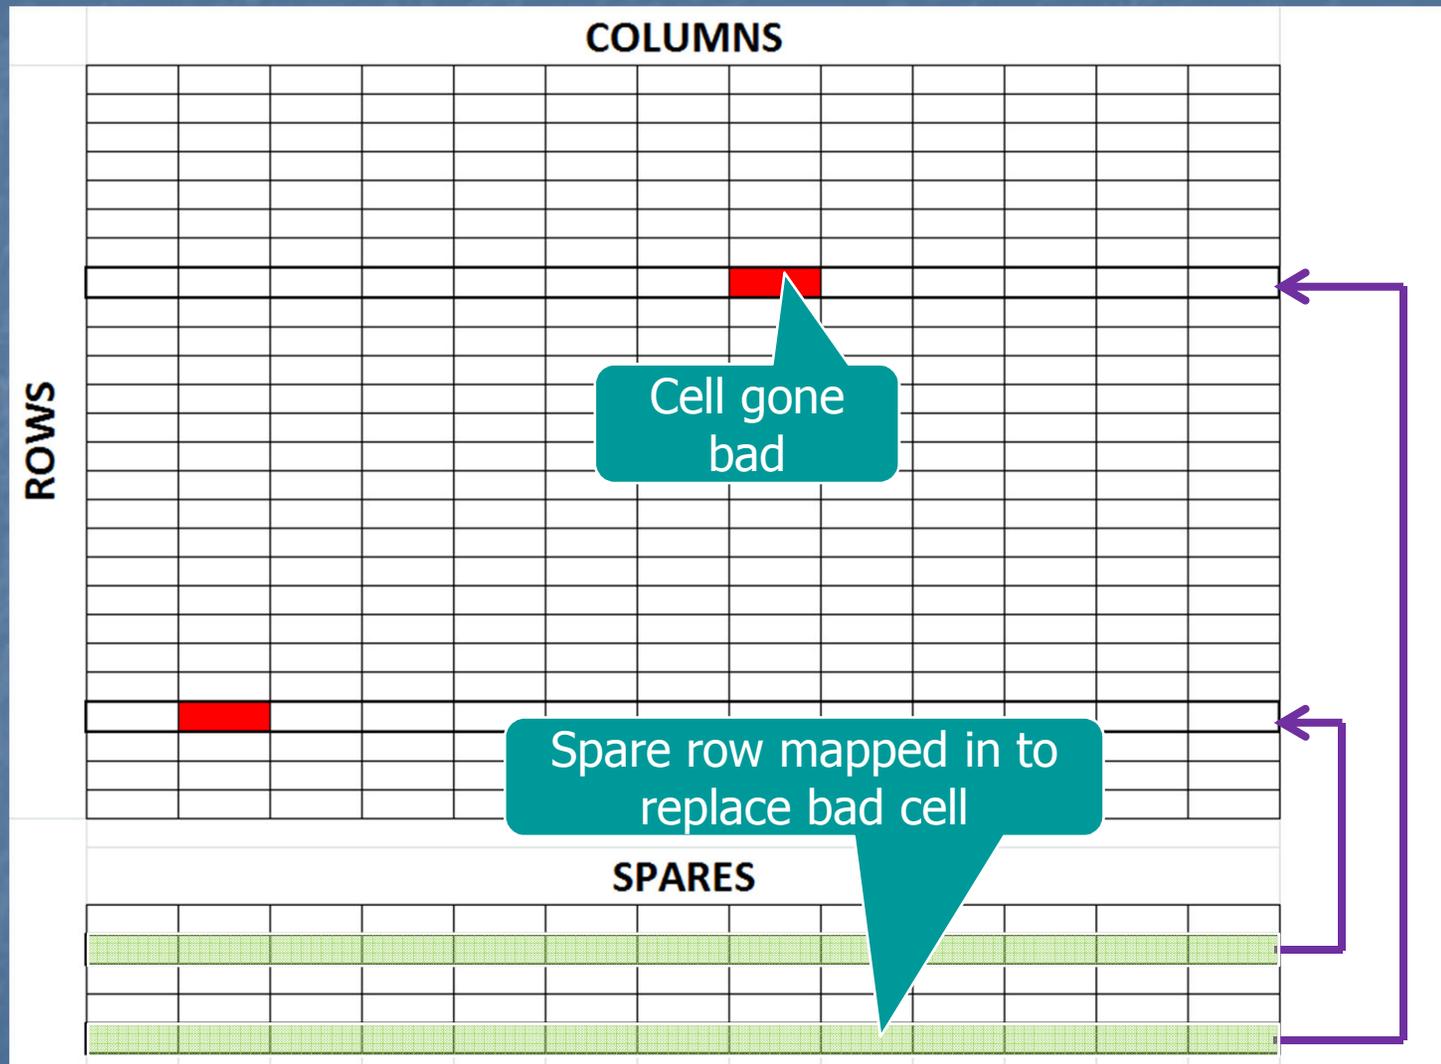
Target Row Errors



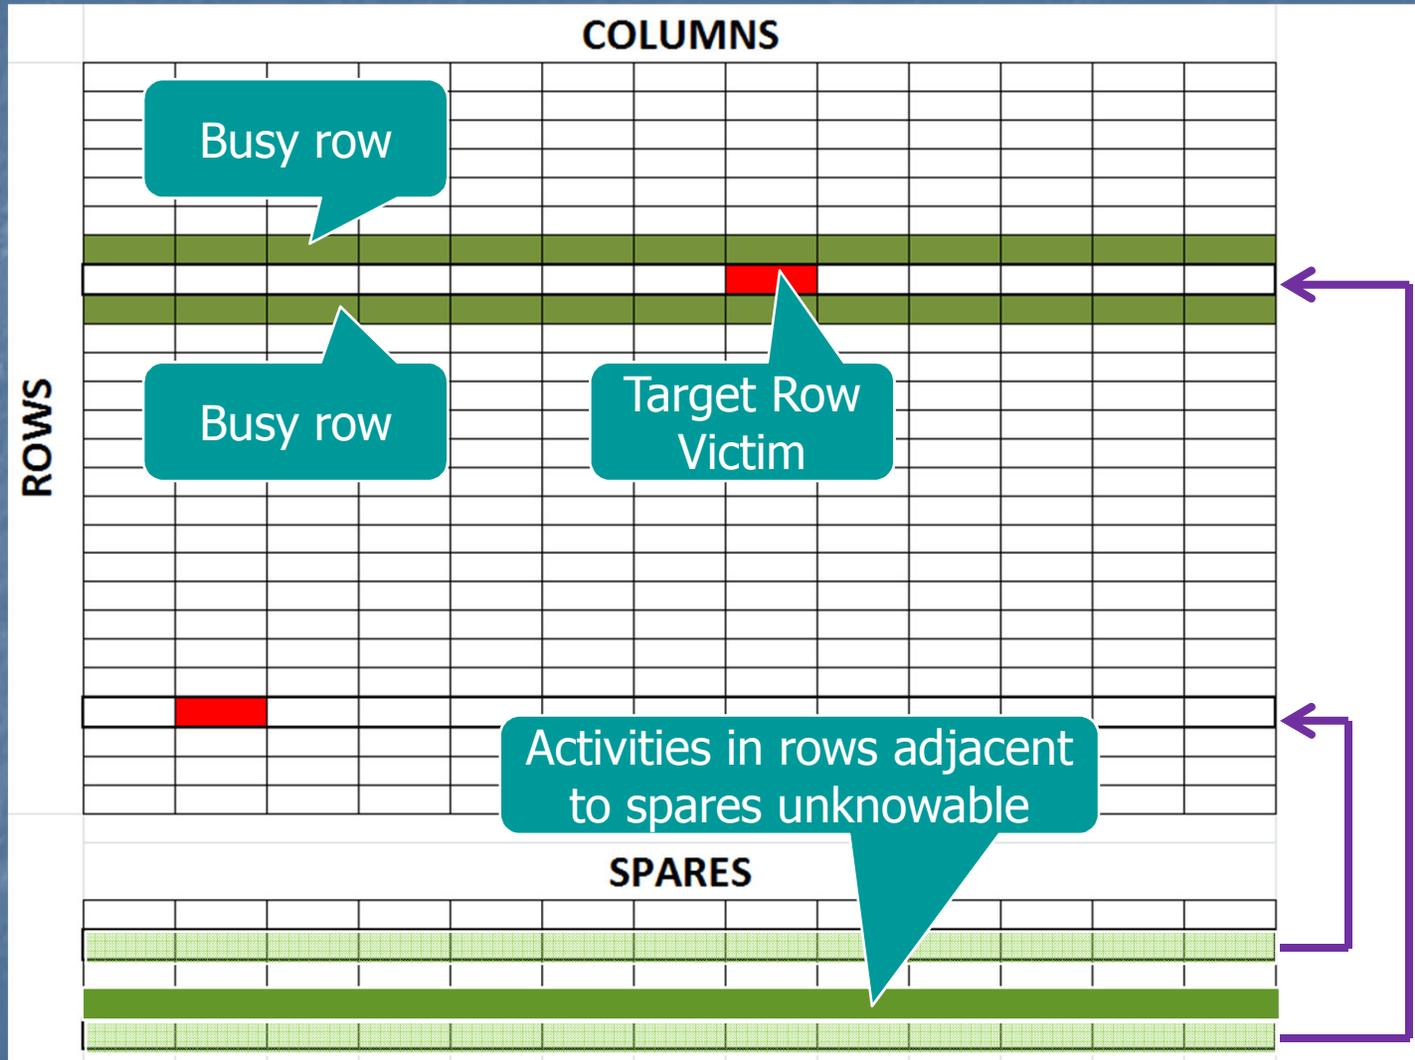
Controllers limit accesses to adjacent rows to avoid disrupting target rows



Does Sparing Fix the Yield Problem?



Makes Target Row Errors Worse



TSV Challenge

DRAM sparing is already done – no target row solution defined yet

3DS DRAM will need to increase use of sparing to raise yields

Thermal trapping will make it worse

Expect increased error rates in the field



3DS Summary

Complex construction makes 3DS with 2 DRAMs cost ineffective

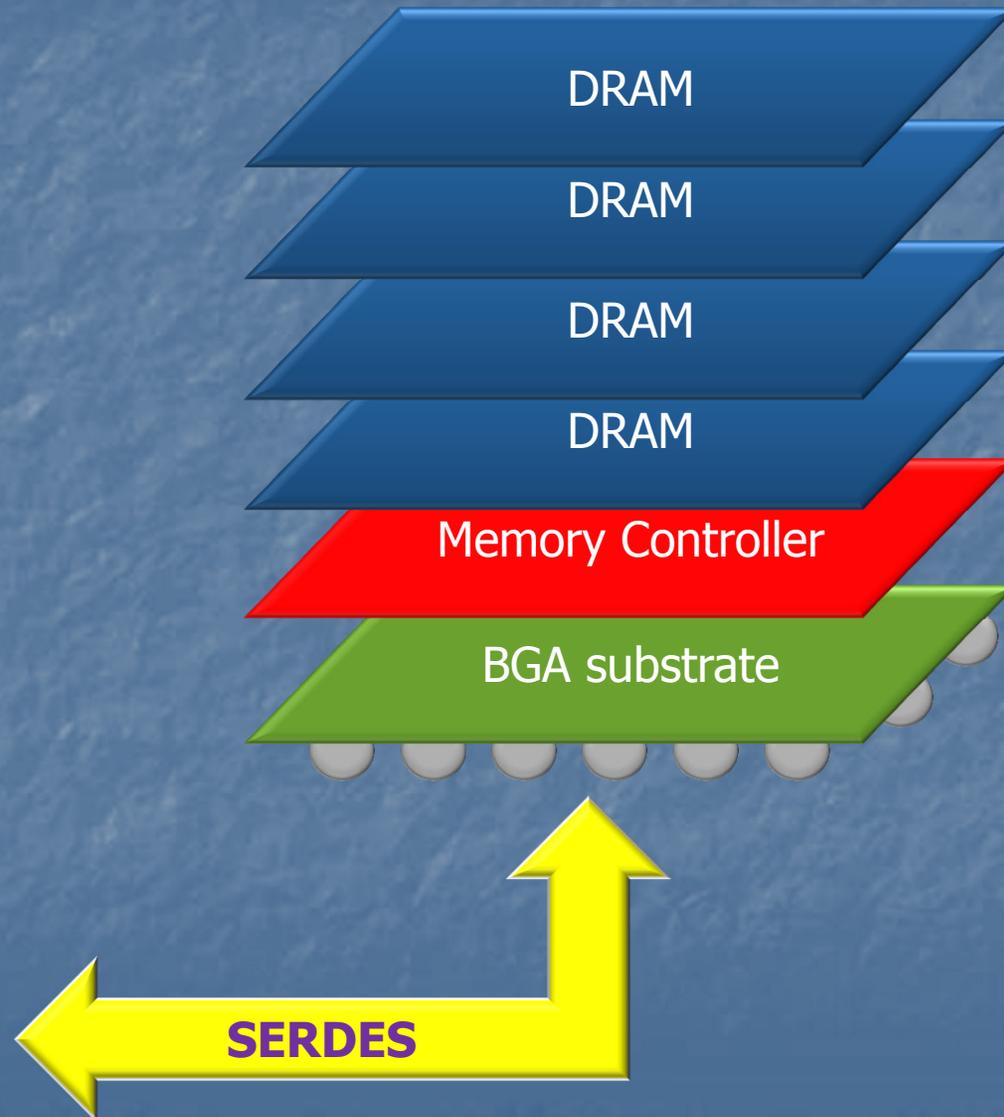
4 DRAMs per 3DS minimum to justify cost

Probably not viable until 2018-2020

DRAM speed by then will be 3200 MTps

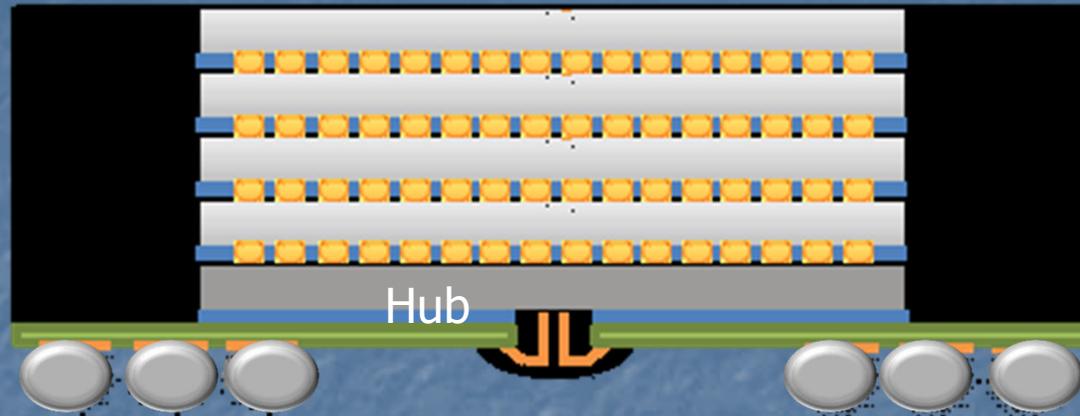


What if you Merged Controller+DRAM?



HMC = Combined Hub & 3DS

Hybrid
Memory
Cube



Inherits problems from the FB-DIMM generation

SERDES runs hot

SERDES increases latency

Minimum silicon size for hub is pretty large

Must incorporate the entire controller

ECC, sparing, refresh, etc.

Political battle to hand over control to suppliers?



HMC Advantages

Can incorporate knowledge of sparing into target row problem solution

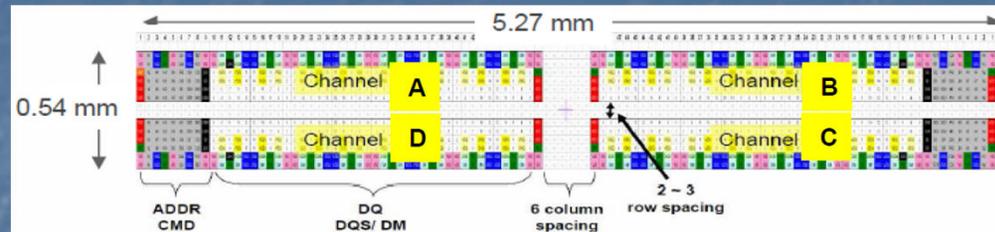
ECC on board to deal with increased error rates

Has potential to make it to market faster than 3DS

Need not be DRAM specific; adding NVM is simple



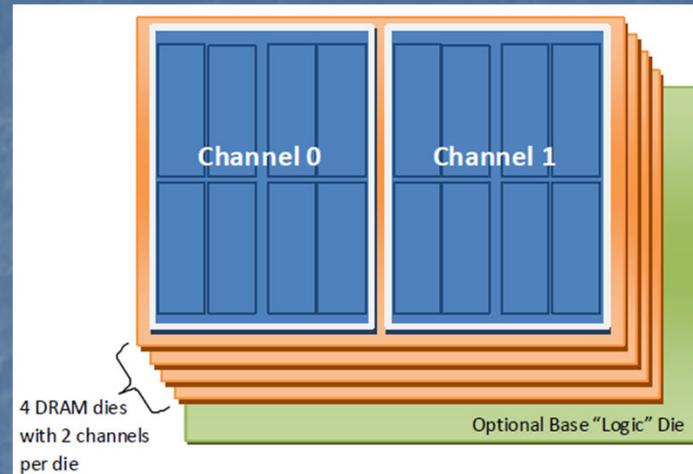
Wide I/O



Low power focus SDR @ 200 MHz
128 bits x 4 channels → 12.8 GB/s
Target market mobile devices
Power reduction is critical
Fairly low capacity required



High Bandwidth Memory



GDDR family running out of steam

128 bit channel, 2n prefetch = 256bit access

Target is graphics

High bandwidth

Low capacity



Migration of Other Memories

DRAM used to be the
“only” memory



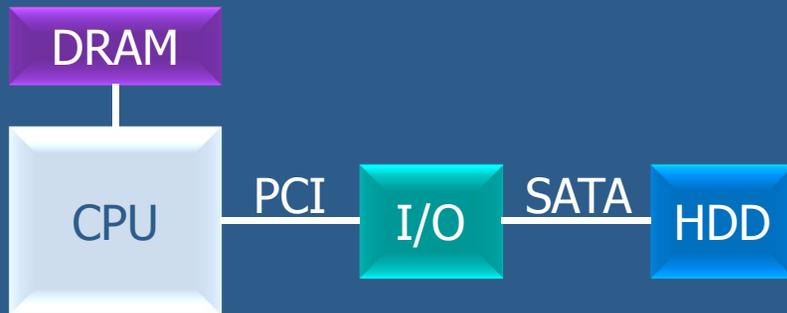
then Flash dropped to
a fraction of the price
of DRAM

So naturally people attempt to work
around Flash’s limitations

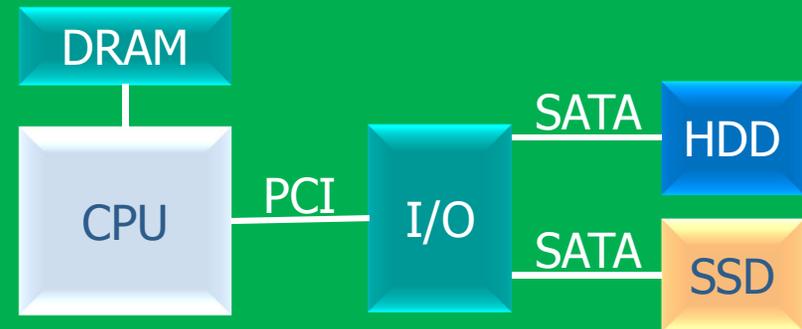


Evolution of Heterogeneous Memory

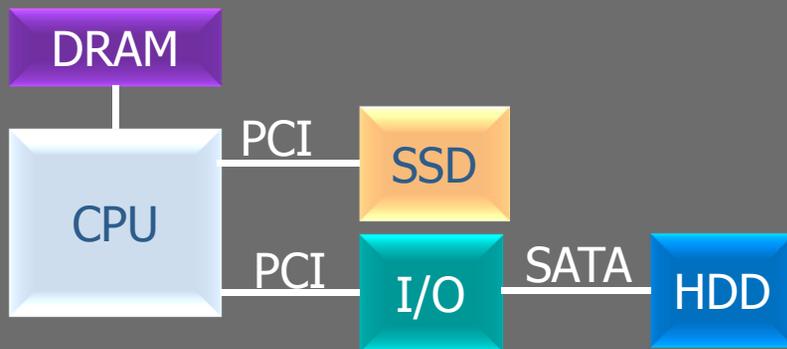
Gen 1



Gen 2



Gen 3



Gen 4



Leveraging Hetero Memory

Initial use of NVM on DRAM bus is for content persistence

Evolutionary improvement is SSD emulation

Eventual change is architectural support of multiple memory types



Summary

Market evolution continues

Market fragmentation occurring

Segments must achieve volume & price

There will be successes

There will be failures

Miniaturization continues to drive innovation

Heterogeneity requires new ideas



I'm Excited to be
a Part of It!

Aren't you?



Thank You

Bill Gervasi

Discobolus Designs

bilge@discobolusdesigns.com



Author of "Nerd Story", available on amazon.com