Fully Buffered DIMM (FB-DIMM) Server Memory Architecture:

Capacity, Performance, Reliability, and Longevity

Pete Vogt
Principal Engineer
Intel Corp

February 18, 2004

Agenda

- Server Capacity Problem and Solution
- Performance
- Reliability
- Longevity

Intel Platform Memory Technology Roadmap

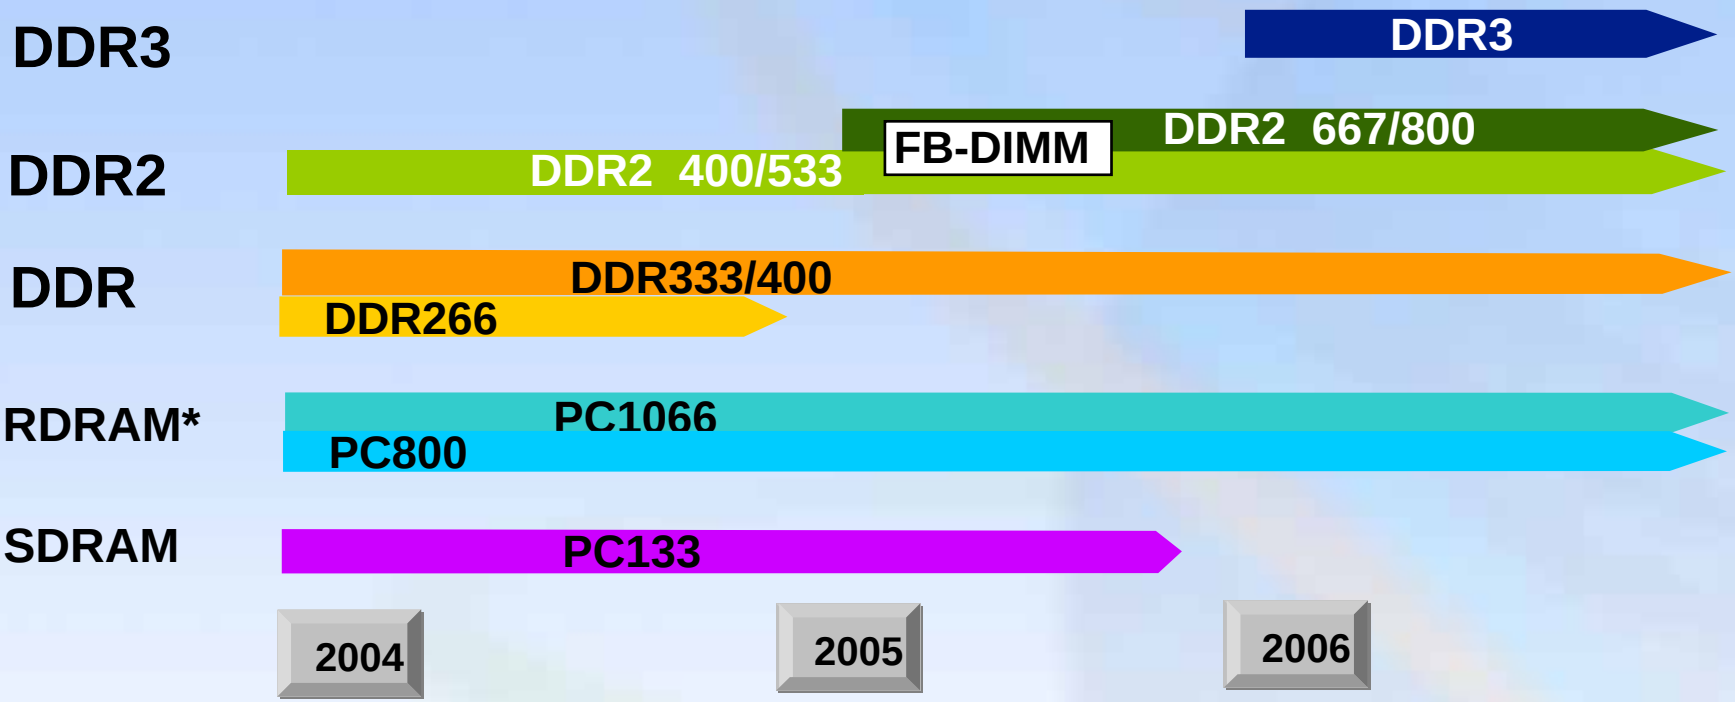

- DDR2 400/533 support in all main IA segments in 2004
  - with DDR flexibility
- FB-DIMM new server interconnect in 2005
- Low Power SDRAM moving to low power DDR
- RDRAM still used in specific applications

FB-DIMM New Server Interconnect in 2005

Why is FB-DIMM Needed in the Future?

DRAM Device Trends

- DRAM bit density follows Moore's law and increases 2X every ~2 years

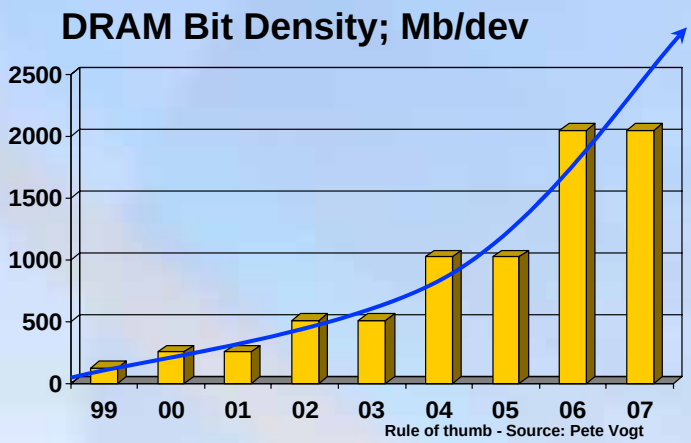

- DRAM data rates driven by CPU performance increase 2X every generation (3.35 years)

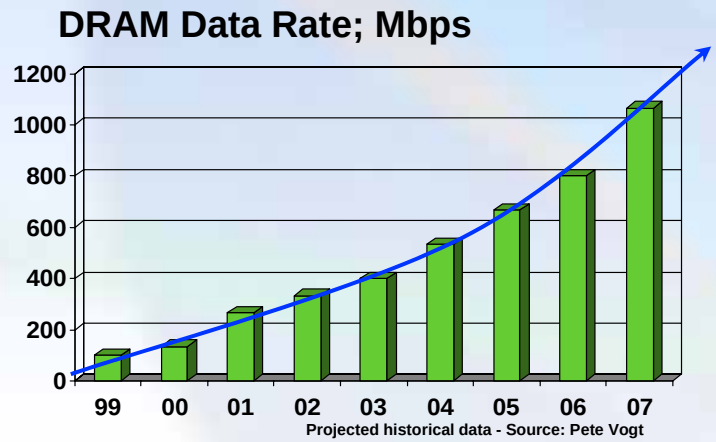

DRAM Devices Continue to Scale

Devices per Channel Trend

- Existing “stub-bus” architecture has impedance discontinuities that effect signal integrity

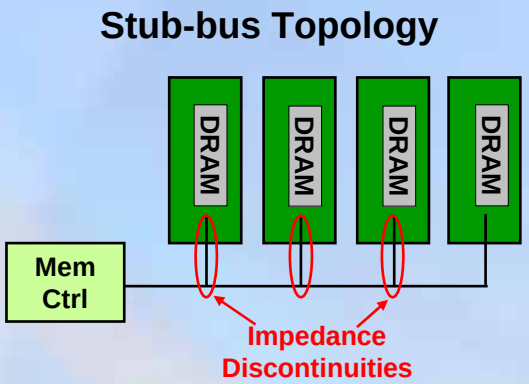

- As the data rate goes up the number of devices on the channel is going down

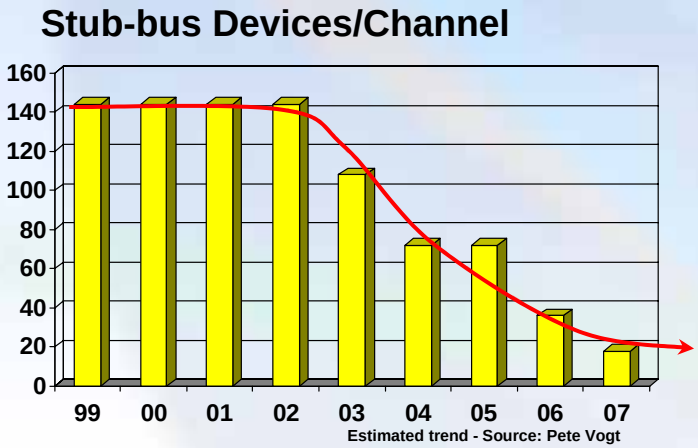

Devices per Channel is Decreasing

Memory Capacity Trends

- Server performance drives 2X capacity increase demand every ~2 years

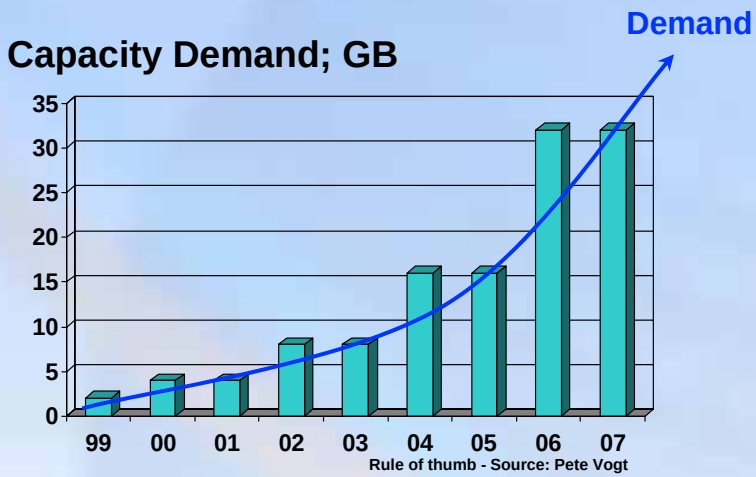

- Stub-bus channel capacity (density * devices/ch) supply has hit a ceiling

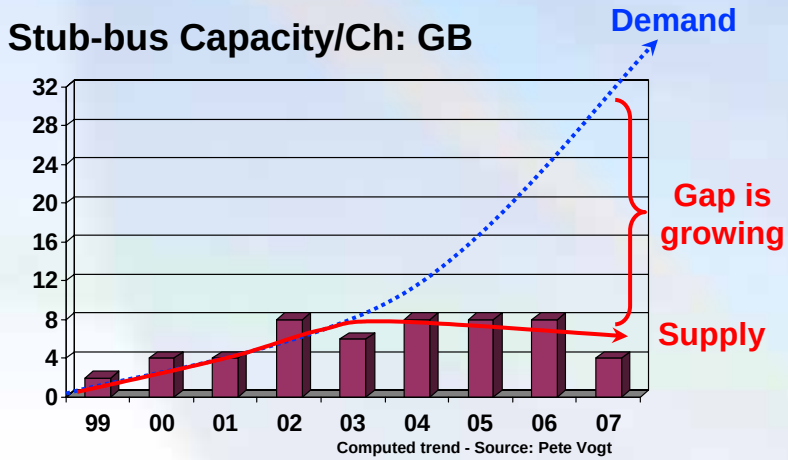

Meeting Demand Requires a Change

FB-DIMM Eliminates the "Stubs"

- FB-DIMM buffers the DRAM data pins from the channel and uses point-to-point links to eliminate the stub bus

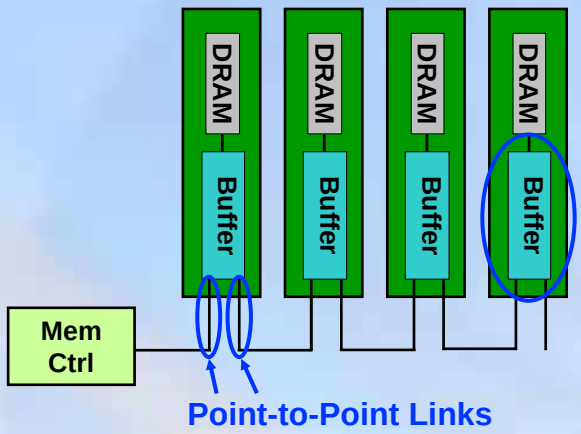

- FB-DIMM capacity scales throughout DDR2 & DDR3 generations

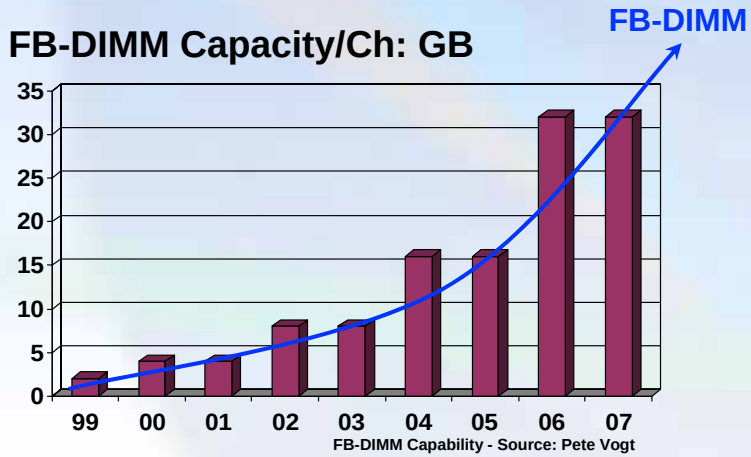

FB-DIMM Meets the Capacity Demand

FB-DIMM Solution Details

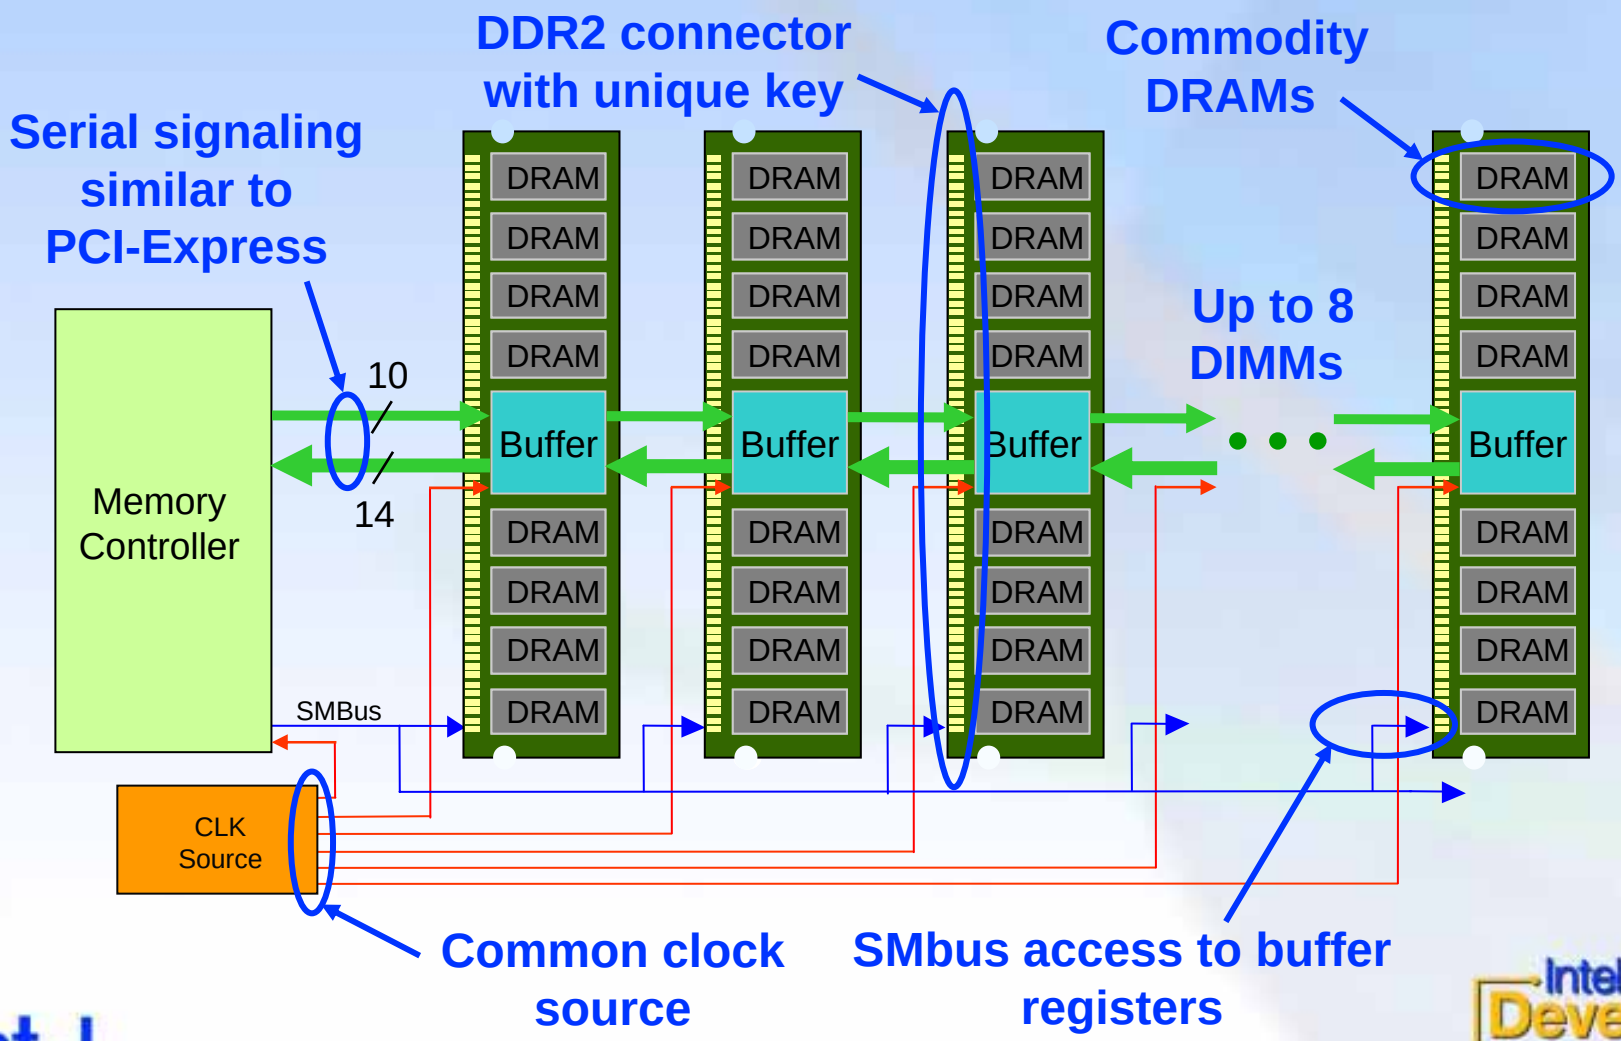

FB-DIMM Channel Pin Count

	Diff Signals	Pins
Data Path to DIMMs	10	20
Data Path from DIMMs	14	28
Total High-Speed Signals	48	
Power	6	
Ground	12	
Shared Pins (clocks, calibration, PLL pwr, test)	~ 3	
Total Pins	**~ 69**	

- **Compare with ~240 pins for DDR2 channel**

FB-DIMM Pin Count is 1/3rd of DDR2

Routing Comparison

Direct DDR2 Registered DIMMs:
1 Channel, 2 Routing Layers with 3rd layer required for power

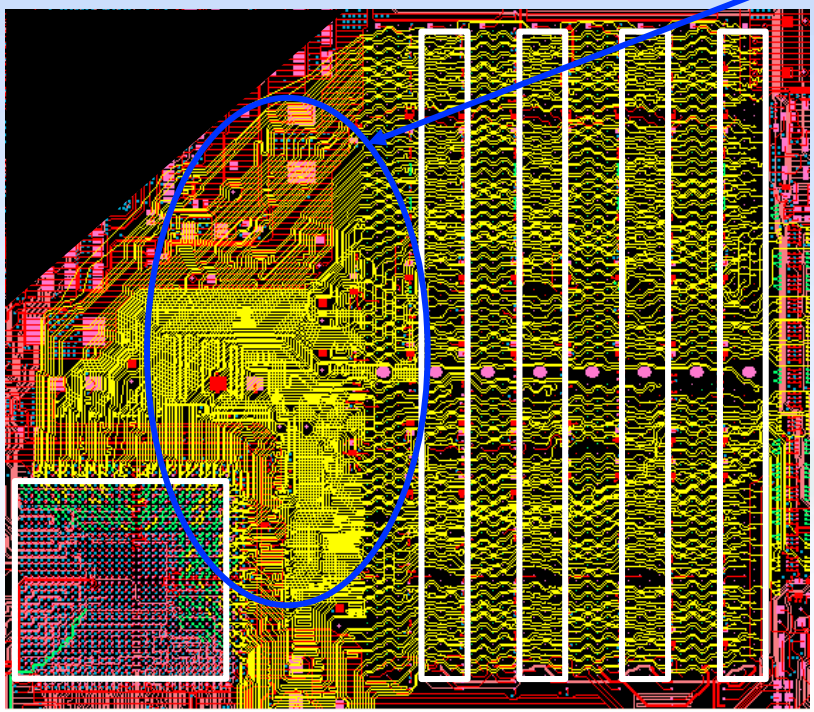

Serpentine routing is complicated and uses up a lot of board area

FB-DIMMs:
2 Channels, 2 Routing Layers (*includes* power delivery)

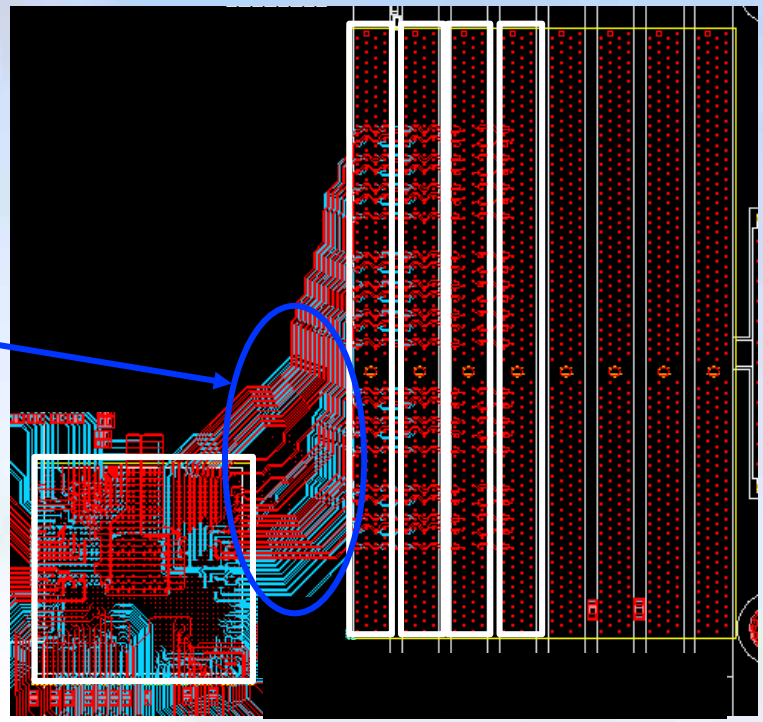

Fewer signals and no trace length matching minimizes board area

FB-DIMM: Fewer Layers, Less Routing Area

Source: Intel Enterprise Architecture Group

Expansion and Visibility

- **Repeaters support flexible system packaging and memory riser cards**

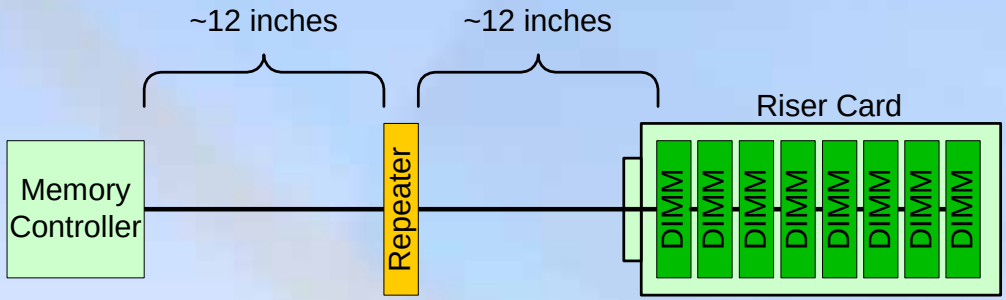

- **Logic Analyzer Interface provides visibility of high-speed bus activity for system analysis**

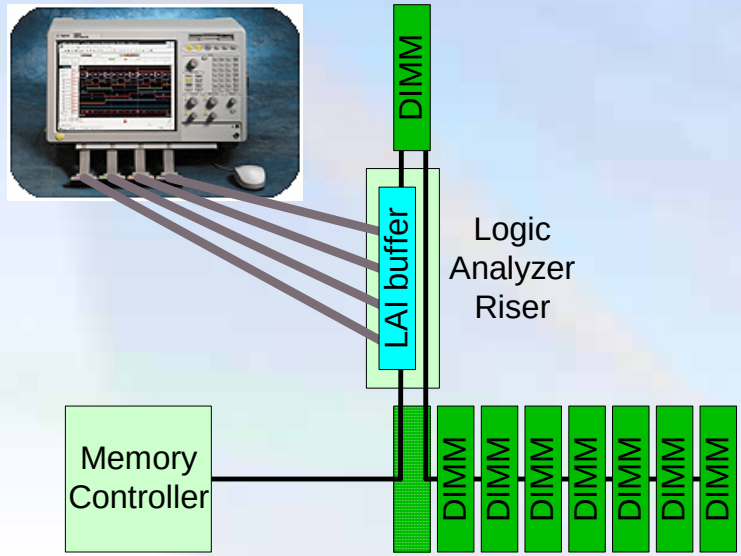

FB-DIMM is a Versatile Solution

Capacity Comparison

- 24x capacity
  - 8GB vs. 192GB
- ~4x bandwidth
  - ~10GB/s vs. ~40GB/s
- ~Lower pin count
  - ~480 vs. ~420

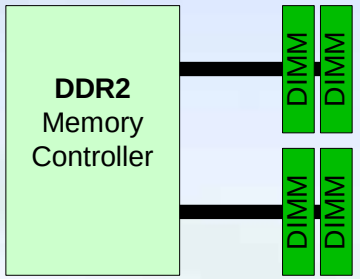

8GB with 1Gb x4 DRAMs
~10GB/s of BW w/DDR2-800
(only 2 ranks per channel)

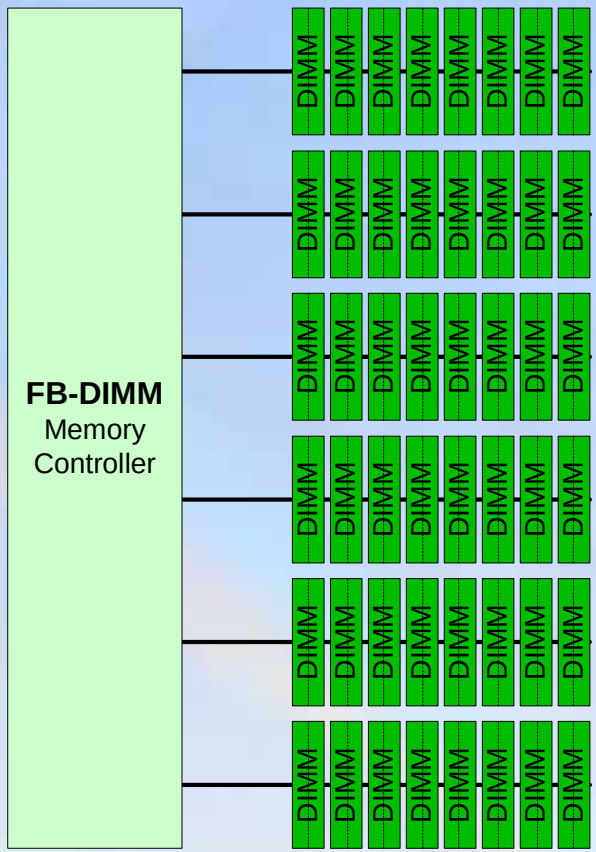

192GB with 1Gb x4 DRAMs
~40 GB/s of BW w/DDR2-800
(2 ranks per DIMM)

FB-DIMM Solves the Server Memory Capacity Problem

FB-DIMM Performance

Entry Configuration Comparison

- Equal capacity
  - 8GB vs. 8GB
- Better throughput
  - ~6.5GB/s vs. ~8.2GB/s
- Lower pin count
  - ~480 vs. ~280

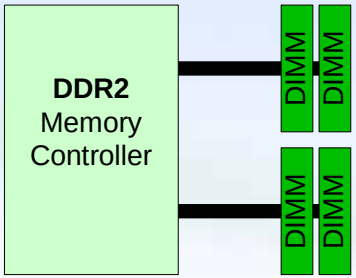

8GB with 1Gb x4 DRAMs
~6.5GB/s of throughput w/DDR2-800
(only 2 ranks per channel)

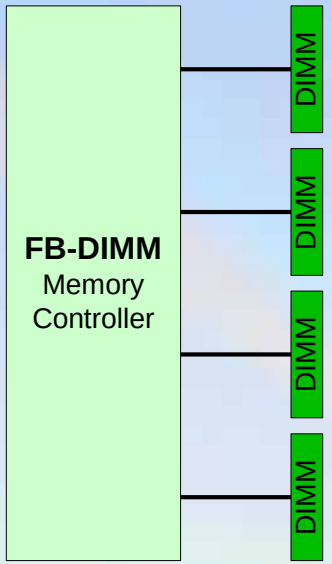

8GB with 1Gb x4 DRAMs
~8.2GB/s of throughput w/DDR2-800
(1 rank per DIMM)

Better Throughput with 200 Fewer Pins

Entry Configuration Comparison

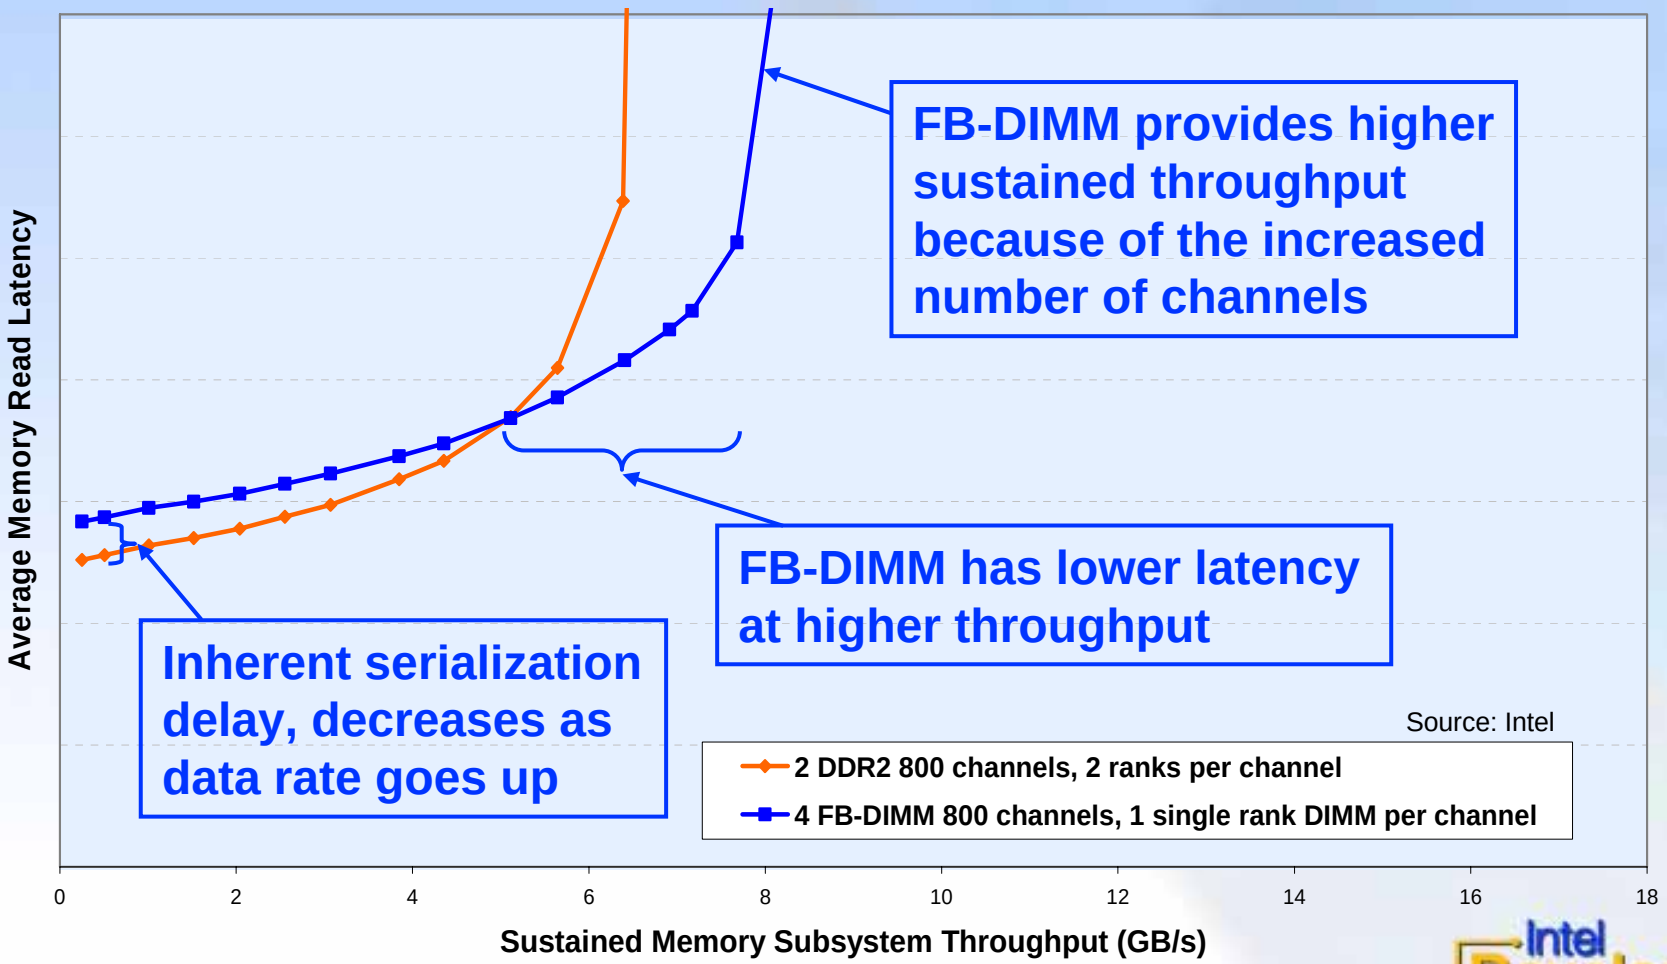

Mid-range Config Comparison

- 4x capacity
  - 8GB vs. 32GB
- ~2.5x throughput
  - ~6.5GB/s vs. ~16.5GB/s
- Lower pin count
  - ~480 vs. ~280

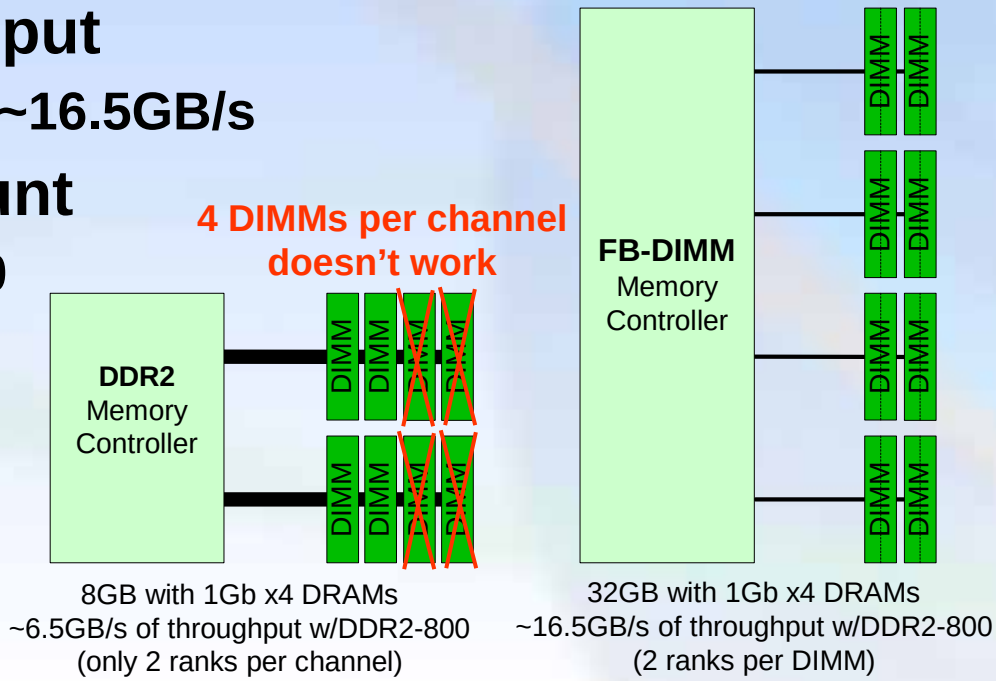

FB-DIMM Provides 4x the Capacity

Mid-range Config Comparison

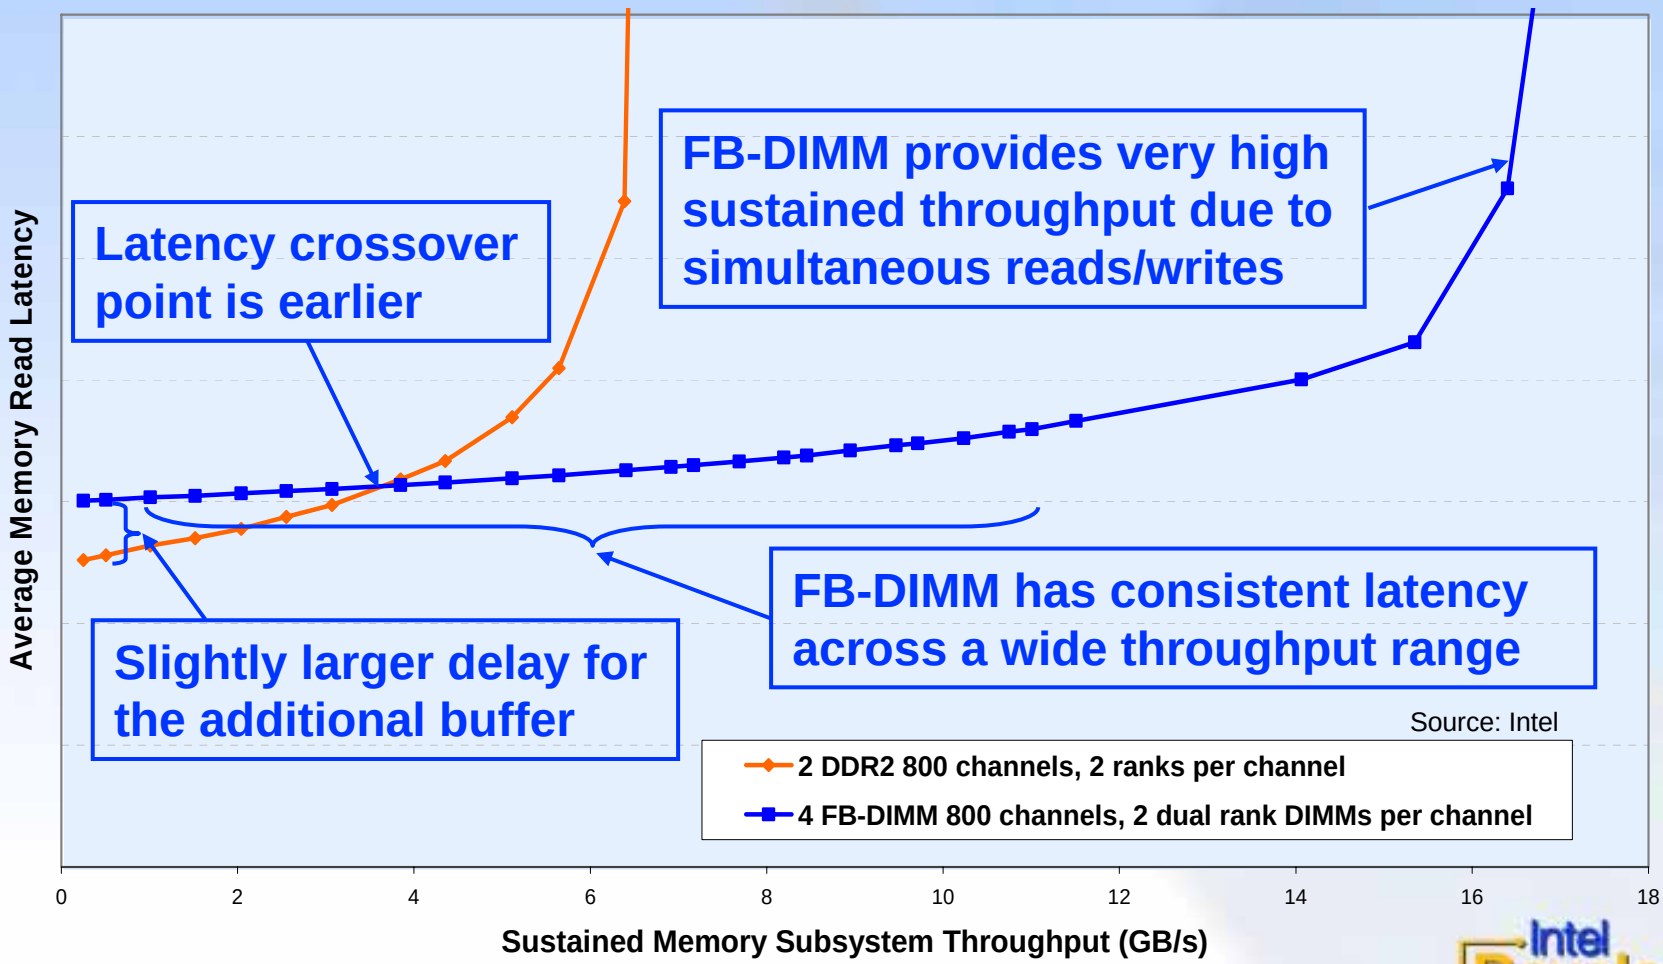

Features That Reduce Latency

- Fast pass-through-path in the IO cells
  - The data is not stored and forwarded
- DRAM synchronized to channel clock
  - No asynchronous clock boundary crossings
- Simultaneous Reads and Writes
  - Writes do not block read accesses
- Elimination of read-to-read dead time
  - Reads from different DIMMs have no dead time between data transfers

New Channel Features Mitigate Buffer Latency

FB-DIMM RAS Features

Channel Reliability Targets

- Intel's server goal is >100 years (1142 FITs) per Silent Data Corruption
  - The memory channels receive a small portion of this total FIT budget
- FB-DIMM is architected for exceptional Silent Data Corruption prevention
  - Channel data is protected by a strong CRC
  - Per channel segment SDC FIT rate <0.10 (1,142,000 years) to support even the highest-RAS servers

Improved RAS Features

- CRC protection for commands and data
- Optional bit widths and CRC coverage to cover wide range of applications
- Transient bit error detection and retry
- Bit lane fail-over “correction”
- Pass-through-path for high availability
- Hot Add while the channel is active
- Error registers in the buffer for improved fault isolation

Comprehensive RAS Features Deliver a Reliable Solution

FB-DIMM Longevity

FB-DIMM Momentum

- Industry is aligned around FB-DIMM
  - Server OEM suppliers recognize the need for a long-term buffered DRAM memory solution
  - DRAM & DIMM vendors committed to delivering FB-DIMM products: Samsung, Elpida, Infineon, Micron, Hynix, Nanya, Kingston, Smart
  - Complete set of tools for system development
- FB-DIMM is a long-term strategic direction
  - Smooth transition from DDR2 to DDR3 using the same connectors and topology
  - Industry investment in the technology leveraged over multiple generations

FB-DIMM Standardization

- Intel working with JEDEC to establish FB-DIMM as an industry standard
- Seven specifications being developed
  - FB-DIMM: Architecture & Protocol
  - FB-DIMM: High-speed Signaling
  - FB-DIMM: Connector (variation of DDR2)
  - FB-DIMM: Module
  - FB-DIMM: Advanced Memory Buffer (AMB)
  - FB-DIMM: SPD EPROM
  - FB-DIMM: Test & Verification

FB-DIMM Technology Roadmap

- DDR2 is the first technology intercept
  - Needed part-way through the DDR2 lifetime
- All DDR3 data rates supported

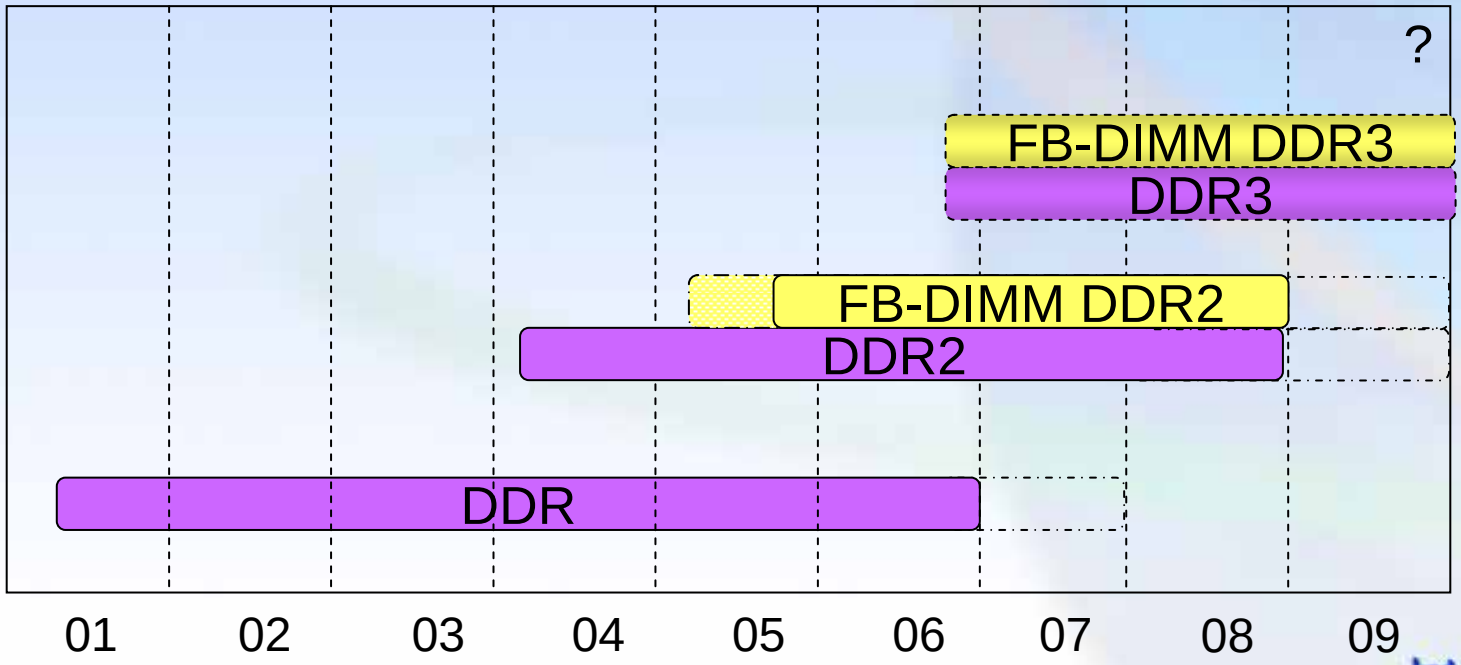

FB-DIMM Cost Effectiveness

- Cost competitive environment
  - Multiple buffer vendors and DIMM suppliers
- Eliminates memory expansion hubs
  - Motherboard not burdened with hubs
- Reduces motherboard routing area
  - Only 24 pairs, no serpentine trace matching
- Enables reduced PCB layer count
  - FB-DIMM can be routed in one signal layer

FB-DIMM is a cost effective long-term technology

Summary

- FB-DIMM solves the long-term server memory capacity problem
- Buffer latency is managed with new channel features
- Comprehensive RAS features deliver a reliable solution
- FB-DIMM is a cost effective long-term technology

FB-DIMM is The Server Memory DIMM of the Future

Fully Buffered DIMM Server Memory Architecture

Pete Vogt

Intel Corporation

Any Questions?

Please remember to turn in your session survey form.

intel

Definitions

- AMB: Advanced Memory Buffer – used to implement FB-DIMM
- FB-DIMM: Fully Buffered DIMM
- FIT: Failure in Time (failures in 1 billion hours)
- JEDEC: Joint Electronic Devices Engineering Council
- RAS: Reliability/Availability/Serviceability
- RDIMM: Registered DIMM
- SDC: Silent Data Corruption

Fully Buffered DIMM Architecture

- **Session Outline:**
- **Server Memory Capacity Problem & Solution**
  - Description of the capacity problem facing servers as data rates go up. Graph of DRAMs/channel vs. data rate. Description of the high-level FB-DIMM architecture that addresses this problem utilizing commodity DRAM devices.
  - Key point: FB-DIMM solves the Server Memory Capacity problem
- **FB-DIMM Performance Features**
  - Description of architecture features to mitigate the buffer latency; synchronous operation, fast pass through path, simultaneous reads/writes, & elimination of read-to-read dead time. Loaded latency graph illustrates result.
  - Key point: Buffer latency is managed with new channel features
- **FB-DIMM Reliability/Availability/Serviceability (RAS) Features**
  - Highlight RAS features in FB-DIMM including bit lane fail-over, reduced pin counts, high-availability pass through path, deterministic behavior, strong channel CRC protection.
  - Key point: Comprehensive RAS features deliver a reliable solution
- **Long-term Server Memory Solution**
  - Describe the longevity of FB-DIMM from DDR2-533 to DDR3-1600, indicate multiple buffer vendor sources, working within JEDEC to establish as a standard.
  - Key point: FB-DIMM is a cost effective long-term technology

Title:

Fully Buffered DIMM (FB-DIMM) Server Memory Architecture: Capacity, Performance, Reliability, and Longevity

Abstract:

What you'll get from this session:

- Trend data of the server memory capacity problem
- Loaded latency curves illustrating FB-DIMM architecture features to mitigate inherent buffer delay
- List of mechanisms that address FB-DIMM reliability concerns
- Roadmap of multiple generations of DRAM support