


A PowerPC compatible
processor[®] supporting high-
performance
3-D graphics

PETER SANDON
IBM Microelectronics
Processor Architect



System Characteristics

➤ Processor

- PowerPC® derivative
- Support for high-performance 3D graphics and game play
 - ✓ Floating point computation
 - ✓ Data movement
- Integrated L2 cache
- 485 MHz

➤ System controller


- Geometry, lighting and rendering
- Memory and I/O controller
- 162 MHz

➤ Memory

- 24 MB, 1T-SRAM




NINTENDO GAMECUBE™ System



Gekko's Role

Programmers utilize Gekko to customize games

- **Game scripting, artificial intelligence (AI)**
 - Planning for animated agents
- **Physics and collision detection**
 - Avoiding walking through walls
- **Custom effects, custom geometry**
 - Smoothing via interpolated transformations
- **Custom and close-up lighting**
 - Detail on characters' faces
- **Moving graphics data through the system**



Gekko Tradeoffs

➤ Processor Organization

- *Options*: Single issue, superscalar, multi-core, multiprocessor
- *Choice*: Superscalar for high performance, low complexity

➤ On-chip memory


- *Options*: L1 cache, L2 cache, embedded DRAM, SRAM
- *Choice*: separate L1 caches, unified 256KB L2, L1 cache locking and DMA for efficient data movement

➤ Floating point support

- *Options*: Single FP pipeline, dual pipeline, vector engine
- *Choice*: Dual pipeline for high performance while maintaining register architecture


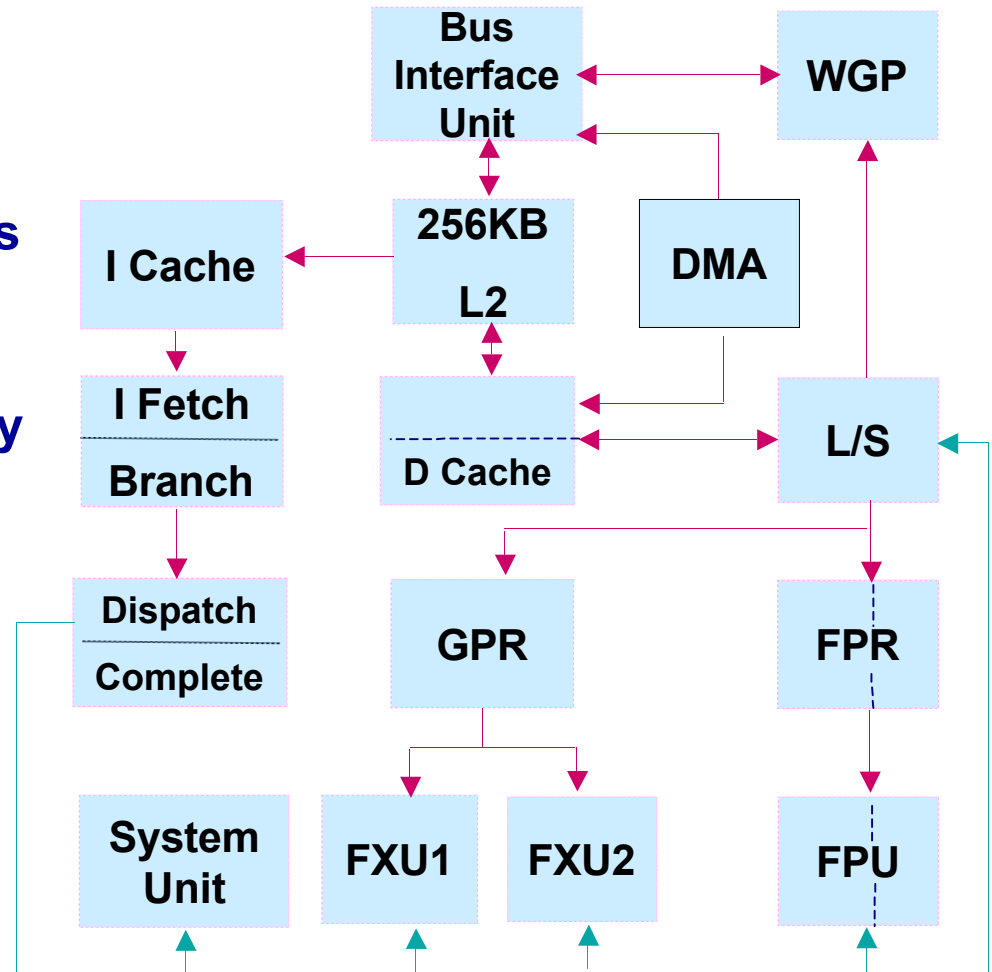
➤ Volume production

- Various process, package and test choices (described later)



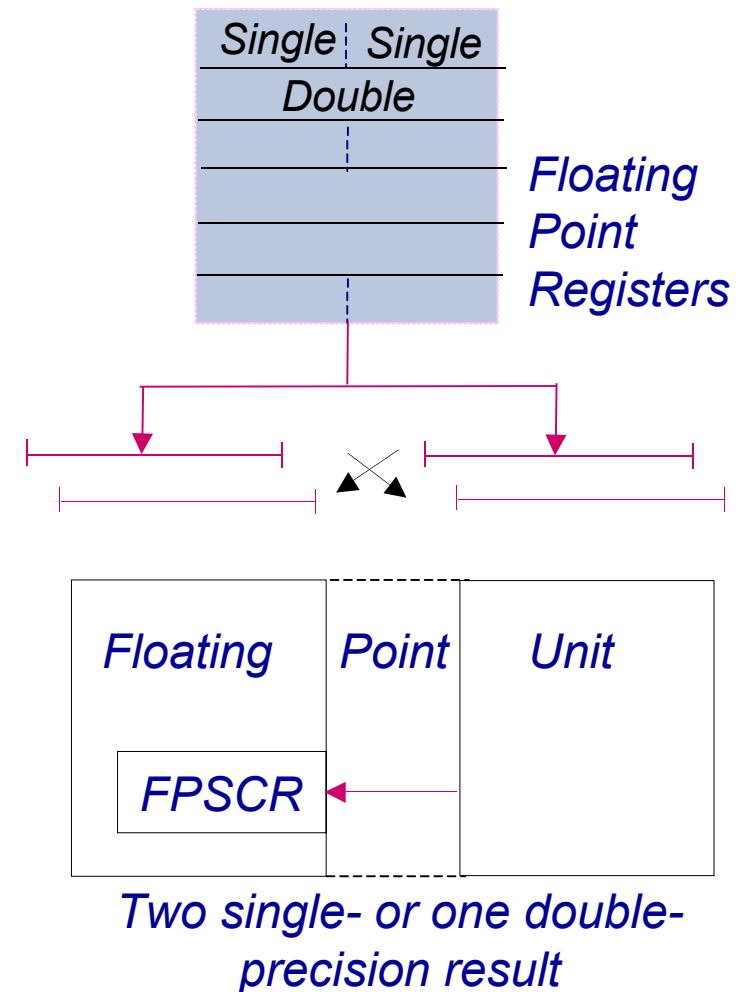
Gekko Block Diagram

- **Superscalar**
 - ✓ 3 issue (2 plus branch folding)
 - ✓ BHT, BTIC
- **L1 Data and Instruction Caches**
 - ✓ 32KB, 8-way set associative
 - ✓ 16KB of lockable Dcache
- **Unified L2 cache: 256KB, 2-way**
- **60x bus interface unit**
 - ✓ DMA, Write gather pipe
- **Load-store unit**
 - ✓ Data quantization
- **Dual fixed-point units**
- **IEEE 754 floating point**
 - ✓ Single-, double-precision
 - ✓ Paired-single operations



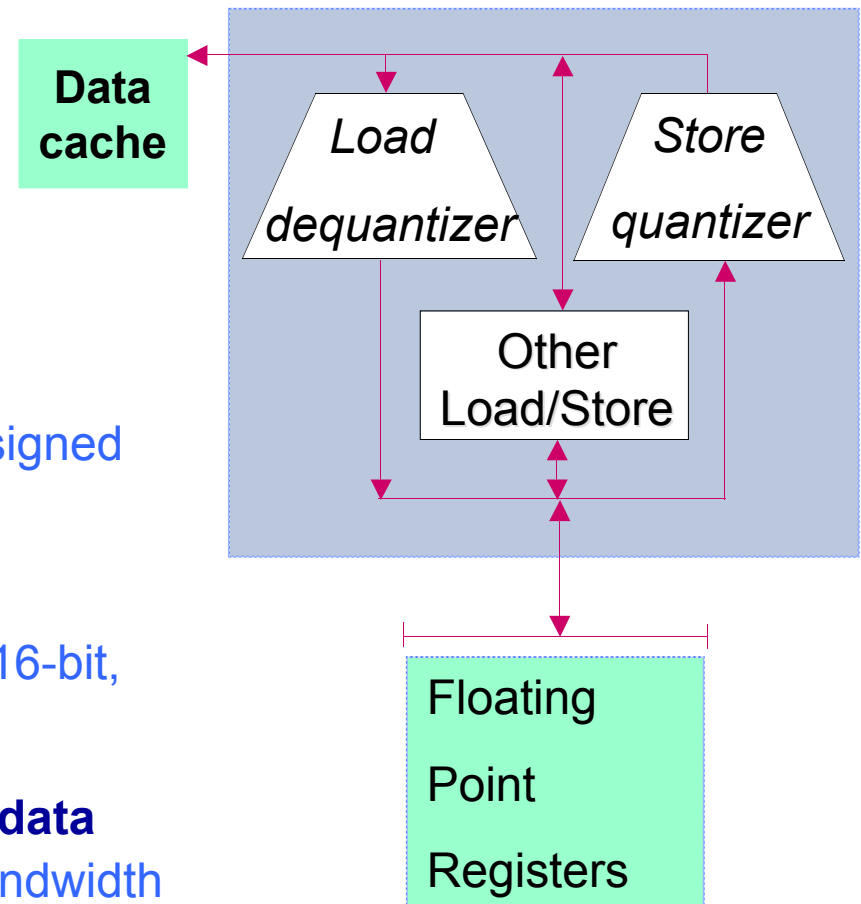
Paired-single operations

- FPRs each hold one double-precision (DP) 64-bit operand or two single-precision (SP) 32-bit operands
- FPU performs one DP operation or two SP operations per cycle
- Two reservation stations / internal busses allow full pipelining
- Finish synchronization to allow mixed instruction programming
- Operand lane crossing for merge and sum instructions
- Single FPSCR merges status from pair of single-precision computations
- Two multiply-adds per cycle yields 1.9 GFLOPS peak throughput


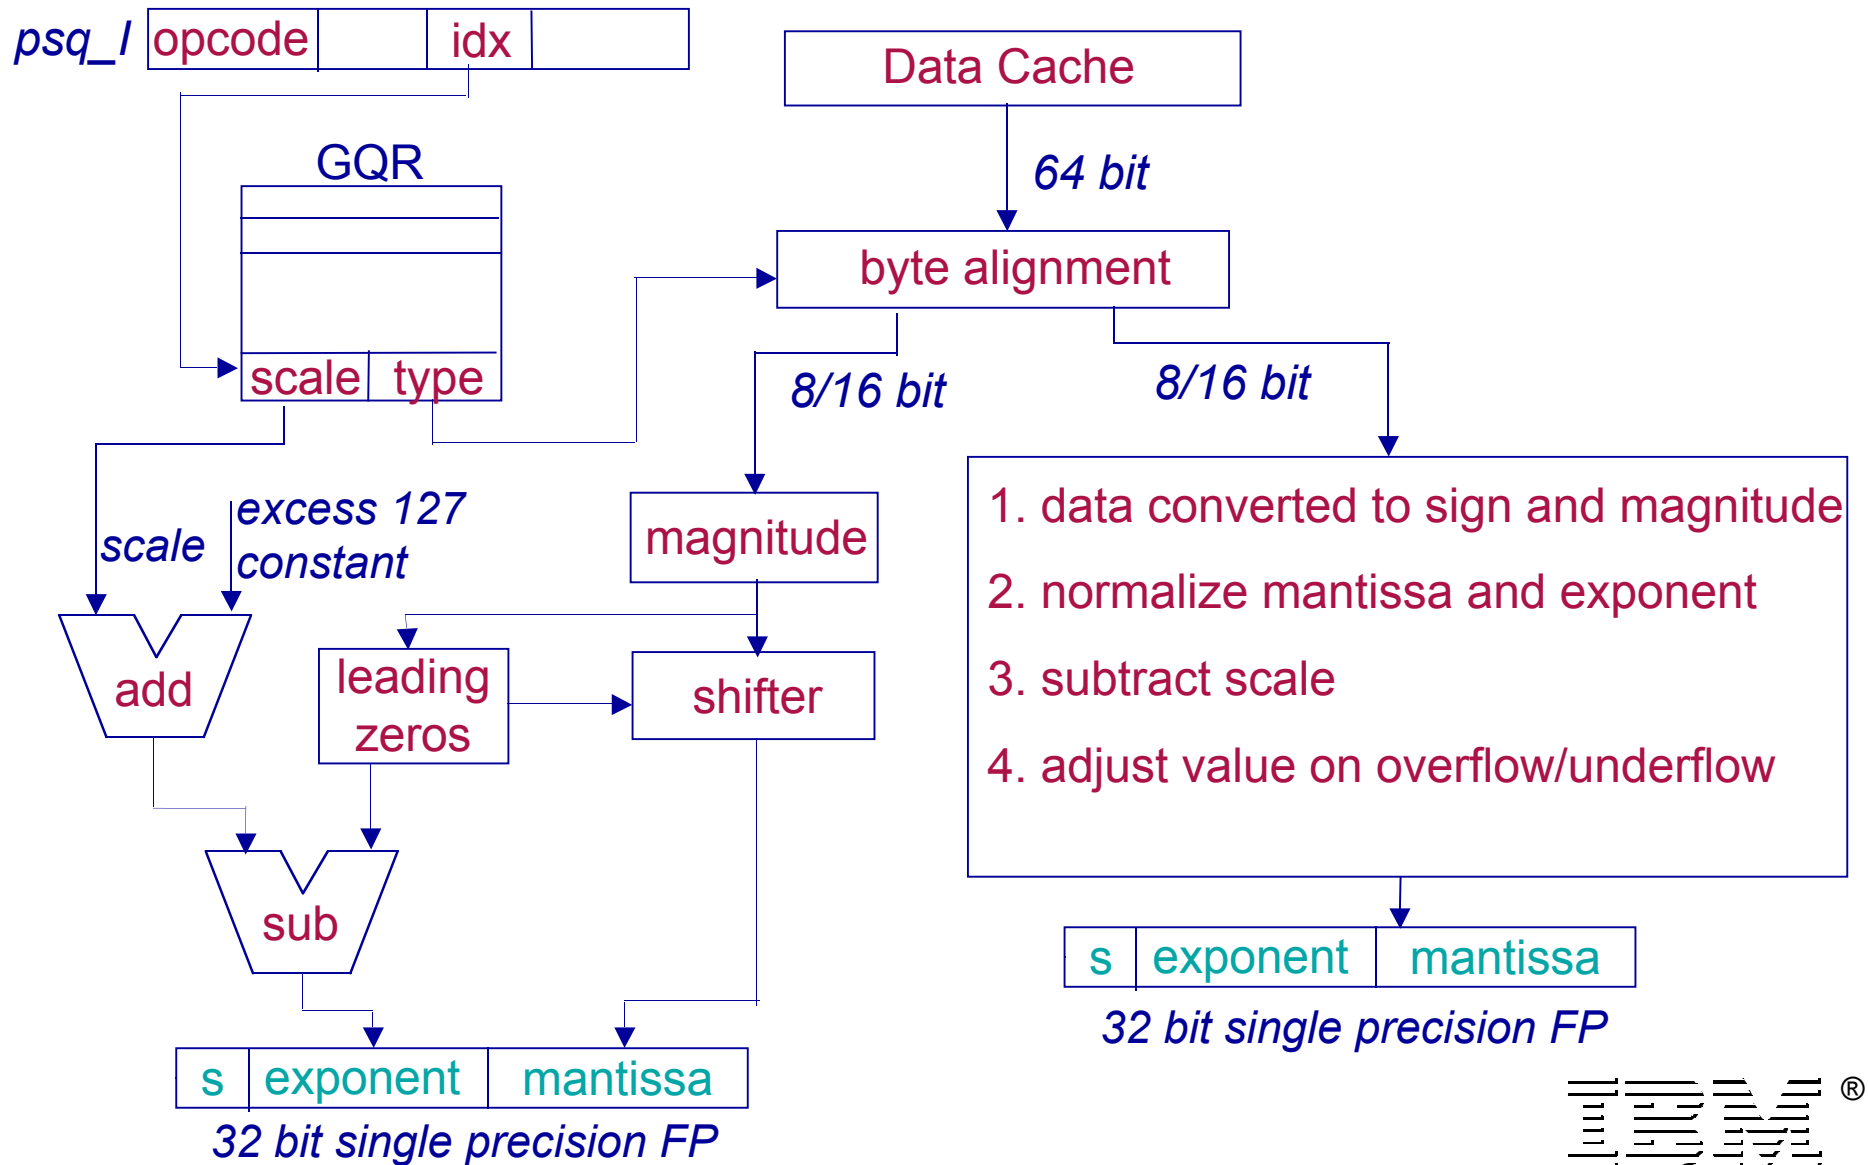


Data Compression

- **Precision required for graphics elements might be:**
 - Position : 16 bits
 - Surface normal : 8 bits
 - Weighting coefficient : 8 bits
 - Transformation matrix : 24 bits
- **Load Q instruction:**
 - converts 8-bit or 16-bit, signed or unsigned integers to SP floating point
- **Store Q instruction:**
 - converts SP floating point to 8-bit or 16-bit, signed or unsigned integer
- **2:1 and 4:1 compression for graphics data**
 - Yields 5.2GB/s peak effective bus bandwidth

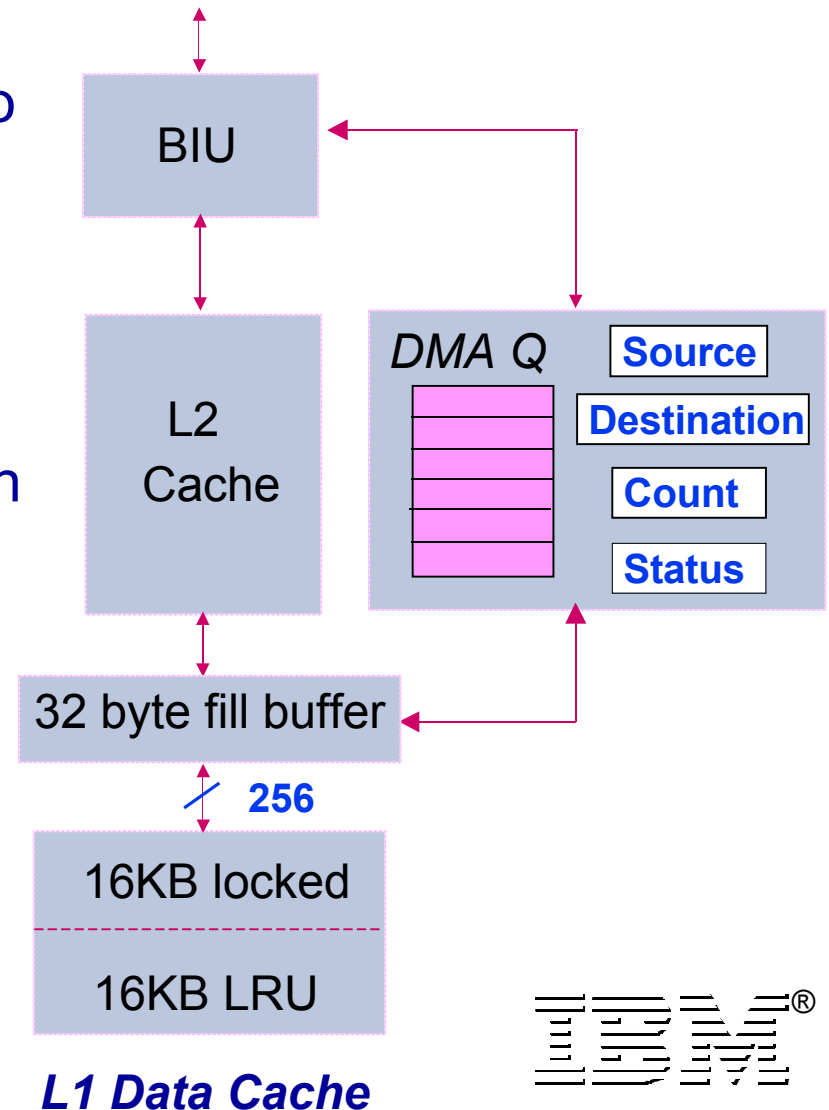


Load Dequantization



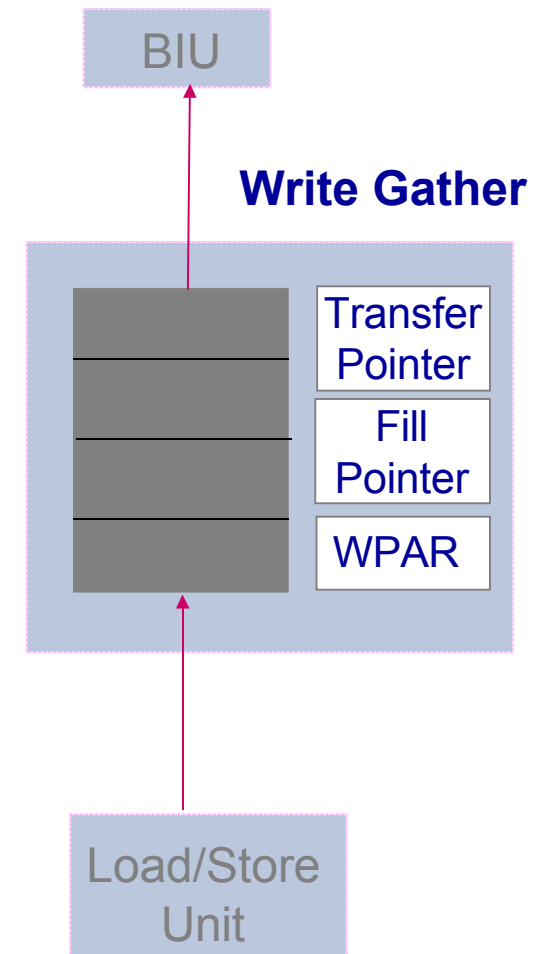
Data Cache Locking and DMA

- Data cache block clear to zero - locked (dcbz_I) instruction used to allocate lines in locked cache
- 15 entry queue, each DMA command can move up to 4KB
- Data in locked half of D-cache can be transferred in parallel with instruction execution
- Status register is polled for completion
- Transient data do not displace persistent data in L2 and unlocked L1




Write Gather Pipe

- Write Pipe Address Register (WPAR) defines address of stores to be gathered
- 128-byte FIFO gathers sequential non-cacheable stores to be transferred in parallel with instruction execution
- Handles packing of byte, halfword, word and double word operands
- 32-byte block transfers over memory bus
- FIFO structure is specifically designed for streaming graphics data



Summary of Graphics Features

- SIMD FPU accelerates computation for custom lighting and geometry, supporting more realistic (or fantastic) visual effects
- Inline data compression allows more complex scenes to be represented within a given space and bandwidth, while maintaining the efficiency of floating point computation
- Asynchronous DMA of large, transient graphics datasets reduces latency and preserves persistent code and game state in caches.
- The write gathering FIFO collects the sequentially generated components of a graphics command stream and “bursts” the data (in four consecutive 8 byte transfers) to the graphics chip




Design-for-Test Features

➤ **Level Sensitive Scan Design**

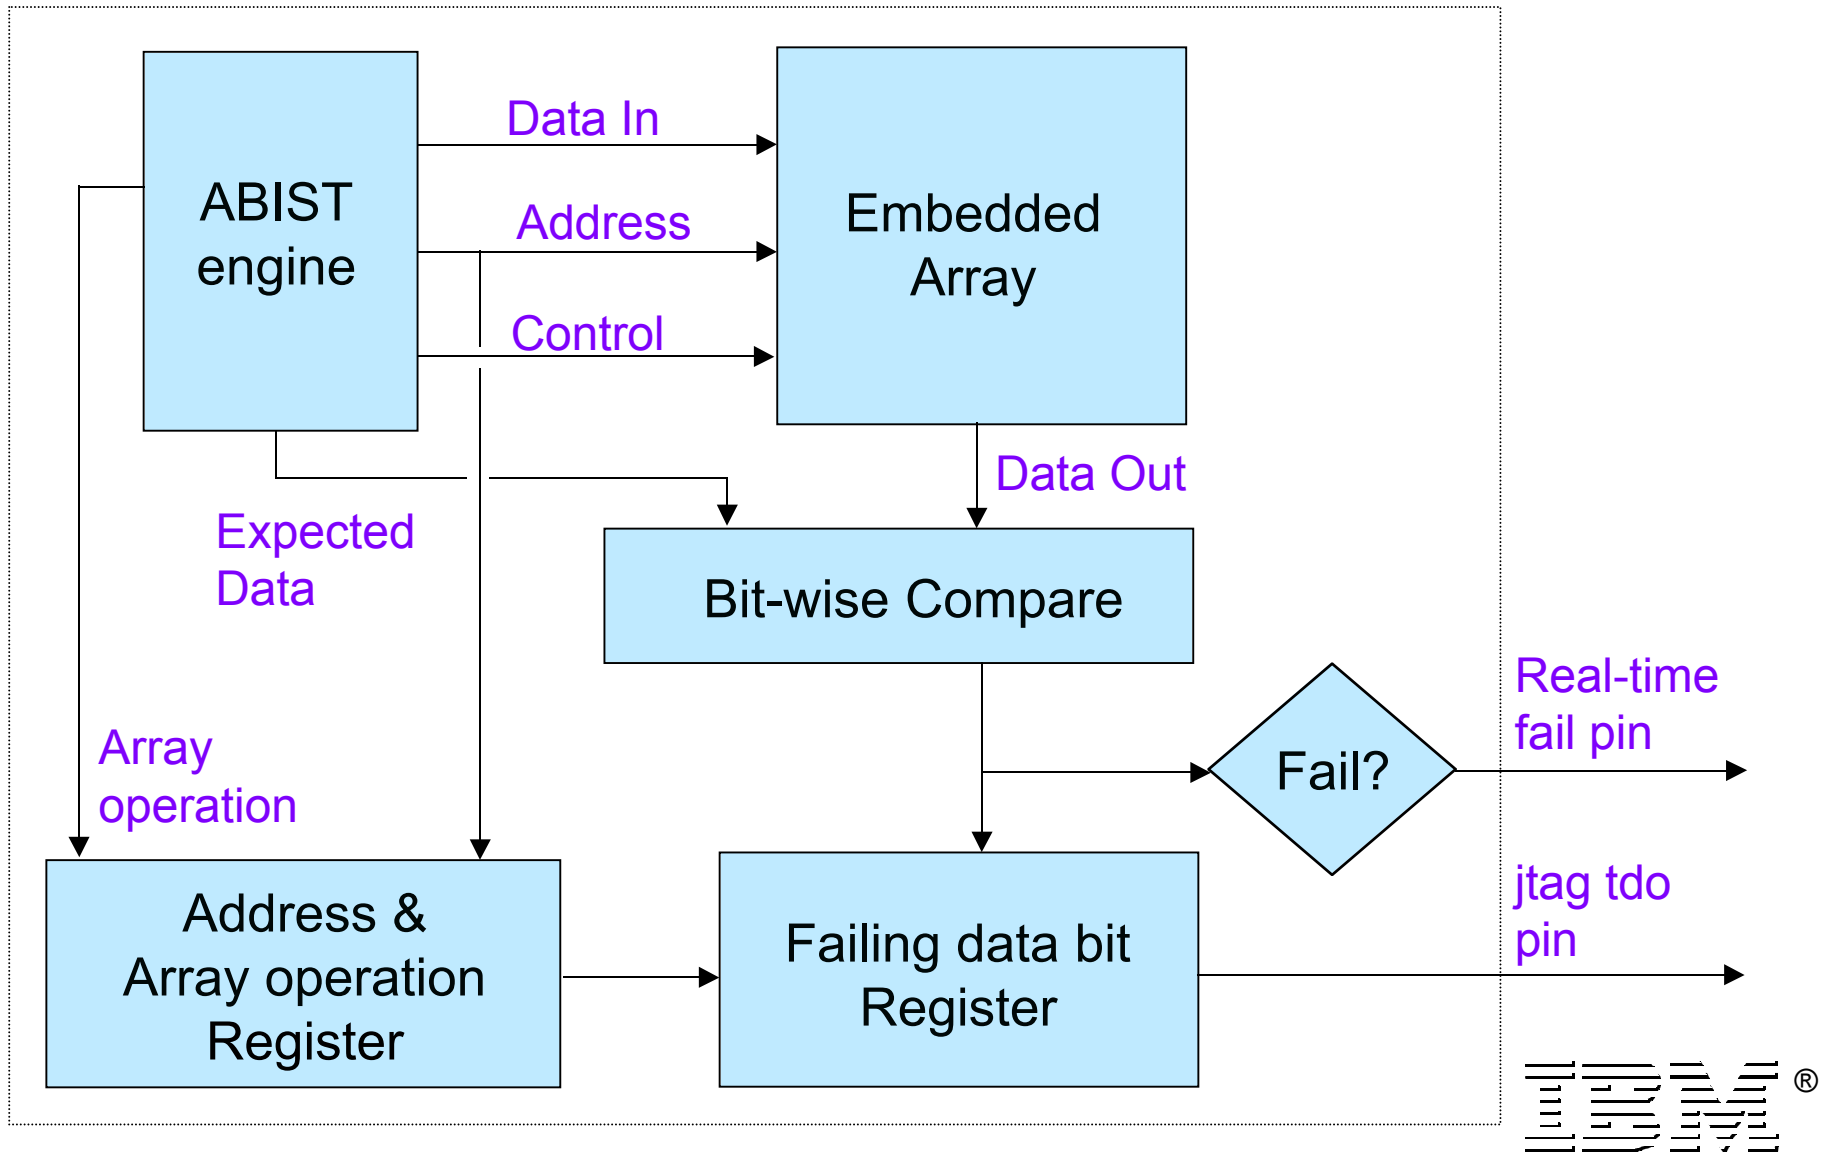
- ❑ Provides test coverage of latches, combinatorial logic, and array interface logic
- ❑ 100 Mhz scan load/unload operation
- ❑ On product clock generation (OPCG) circuitry to detect at-speed (>400 Mhz) transition faults

➤ **Array Built-In Self Test (ABIST)**

- ❑ Provides test coverage of embedded arrays (e.g. caches)
- ❑ Array diagnostics circuitry for continuous technology process learning




Array diagnostics block diagram



Gekko Specifications

Frequency – CPU	485MHz
Performance	1125 DMIPS (Dhrystone 2.1)
Power Dissipation	4.9W (typical)
Caches	L1: 32/32KB, 8-way set associative L2: 256KB, 2-way set associative
System Bus	1.3GB/s peak bandwidth (162MHz, 32-bit address, 64-bit data) 5.2GB/s effective with compression
Package	256 I/O, thermally enhanced 27x27mm PBGA
Technology	0.18 μ m CMOS copper technology, 6 levels of metal
Power Supply	1.8V logic and I/O



Summary

- **Gekko is a custom-designed, superscalar processor that implements the PowerPC® instruction set architecture with extensions to support the high computation and bandwidth requirements of a game console**
- **In the end, the goal is to enable compelling games. The processor's role is to provide general-purpose processing for game play, FP compute power for custom lighting and geometry, and data movement facilities, all at a high level of performance**
- **The SIMD floating point unit, along with clock speed and system design, support peak computation and data transfer rates that exceed the customer's original guidelines**
- **Additional processor facilities, including the DMA, write gather, and data compression facilities, allow the system to sustain computation and data movement such that performance exceeds goals, as measured by internal benchmarks**



Questions?

