

Aaron Kobayashi

## Considerations for the Intel Itanium 2 Data Cache Design

- 1) Out-of-order microarchitectures – Most contemporary superscalar architectures require multiple memory pipelines (or M ports); The Itanium 2 provides 4 M-ports to support these demands.
- 2) Instruction Bundling – The EPIC design has a fundamental instruction packaging, called a “bundle”, which is made up of three instructions; the Itanium 2 processor can issue 2 bundles per processor cycle.
- 3) Predication – Code with predication can increase the number of memory ops requested increasing the need for a well designed and balanced caching system.
- 4) Control Speculation – Speculation can also increase the number of memory ops requested since speculation can create memory requests that will never be used.
- 5) Data Speculation – EPIC allows loads to be moved ahead of stores, but this increased performance doubles the demand of the M-port resources.
- 6) Compilers – The EPIC design structure is heavily dependent on the compilers ability to assist the processor at compile time; The Itanium 2 tries to reduce hardware restrictions visible to the compiler in an effort to assist the compiler developer in make the compiler better.
- 7) Load Latencies – The Itanium 2 processor provides a 1-cycle first level data cache to enable compilers to better realize the performance made possible by hardware capable of issuing 6 instructions per cycle.
- 8) Commercial Applications – Commercial database applications place extreme demands on a processor’s memory subsystems; The Itanium 2 processor data cache was structured to enable fast access even to the largest level of cache.
- 9) Data locality – Similarly data references in database applications show little locality resulting in many memory references that cannot be satisfied by any cache level creating the need for fast access to memory.
- 10) Technical Applications – Simulation codes common of technical applications may require streaming of instructions from a large cache; The Itanium 2 processor data cache is designed to support this requirement by providing the capability to execute a stream of two MMF bundles per cycle.

### References:

- T. Lyon, E. Delano, C. McNairy, and D Mulla, "Data Cache Design Considerations for the Intel Itanium 2 Processor," *International Conference on Computer Design (ICCD'02)*, September 2002, pp. 356-362.