Intel® Itanium® Architecture

May 2004

Agenda

- Why migrate to Itanium® Processor Family?
- Itanium® Architecture
- Intel Software Development Tools

Why migrate to the Itanium® 2 Processor?

- Performance
- Performance/Price
- Reliability Accessibility Support
- Next Generation Architecture Features

*Other names and brands may be claimed as the property of others.

Itanium® 2 processor 6M vs. RISC

vs. SPARC

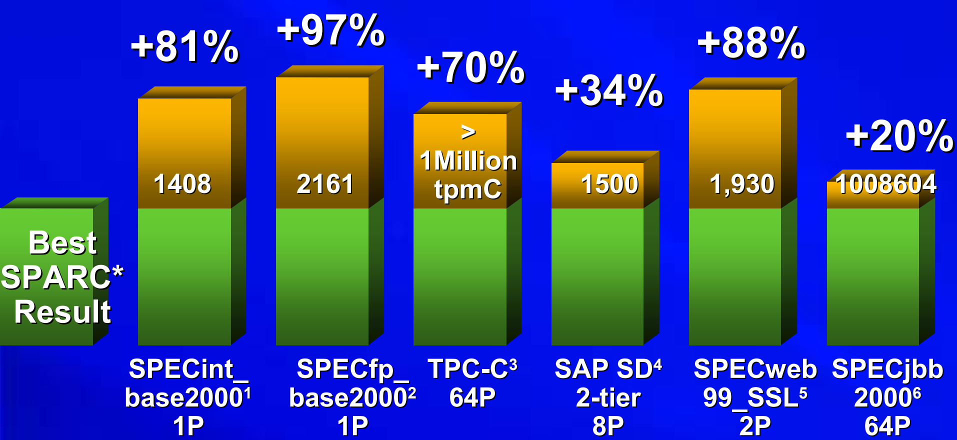

vs. IBM Power4

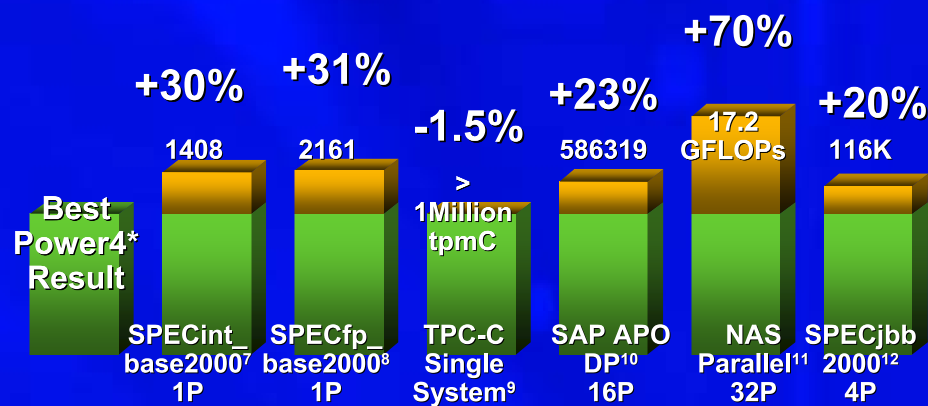

1 Source www.spec.org : SPECint_base2000 of 1408 on HP Integrity Server rx2600 with one Itanium® 2 processor 6M at 1.5GHz, 6M L3 cache, HP-UX OS and compilers. Best SPARC result of 776 on Fujitsu PRIMEPOWER650 using SPARC64 V 1.35GHz processor.

2 Source www.spec.org : SPECfp_base2000 of 2161 on HP Integrity Server rx4640 with one Itanium® 2 processor 6M at 1.5GHz, 6M L3 cache, RedHat Linux AS2.1 operating system. Best SPARC result of 1096 on Fujitsu PRIMEPOWER650 using SPARC64 V 1.35GHz processor.

3 Source: www.tpc.org : 1,008,144 tpmC at $8.33/tpmC on HP Integrity Superdome,, with 64 Intel Itanium 2 processors 6M, each at 1.5 GHz with 6MB of L3 cache, running HP UX 11.iv2 64-Bit with Oracle Database 10G Enterprise Edition, with 1024 GB RAM. TPC-C availability date: 4/14/2004.
Best SPARC: 595,702 tpmC, $12.43/tpmC on Fujitsu PRIMEPOWER 2500, with sixty four (64) Fujitsu SPARC64 V 1.3 GHz processors, running Sun Solaris 8 with Oracle Database 10g Enterprise Edition, 512GB RAM; TPC-C availability date: 12/31/2003.

4 Source: www.sap.com/benchmark: 1500 SD users on HP Server rx7620 with 8 Itanium® 2 processors 6M at 1.5GHz, each with 6MB L3 cache, 32GB of memory, HP-UX 11i, SAP rev 4.7, Oracle 9i database.
Best SPARC result of 1120 SD users on Fujitsu Siemens Computers PRIMEPOWER with 8 SPARCTM V, 1.35 GHz proceesors, 32GB of memory, Solaris 8, SAP rev 4.6C, Oracle 9i database.

5 Source: www.spec.org. Itanium® 2 processor result of 1930 on HP Server rx2600 using 2 Itanium® 2 processors 1.5GHz with 6MB L3 cache, 12GB memory, HP-UX, Zeus 4.2r2 and submitted to SPEC.
Best RISC result of 1022 on Sun Fire* V240 result with 2 UltraSPARC* IIIi processors at 1.28GHz, each with 8MB L2 cache (off chip), Solaris* 9 8/03, Sun ONE Web Server 6.0 SP5, 8GB RAM, published 2/04.

6 Source: www.spec.org for Best published SPARC result of 835,479 on Fujitsu PRIMEPOWER2500 using 64 SPARC64 V 1.35GHz processors, 262144MB memory, Solaris8 02/02, JVM HotSpot Server VM on Solaris/SPARC, version1.4.1_02.
Itanium® 2 processor 6M result of 1008,604 measured by HP on HP Integrity Superdome using 64 Itanium® 2 processors 6M at 1.5GHz with integrated 6MB L3 cache, 128GB of memory, HP-UX 11i v2.0, JVM Hotspot 1.4.2.00 and submitted to www.spec.org. SPECjbb* is a trademark of SPEC at www.spec.org.

7 Source www.spec.org : SPECint_base2000 of 1408 on HP Integrity Server rx2600 with one Itanium® 2 processor 6M at 1.5GHz, 6M L3 cache, HP-UX OS and compilers. Best RISC result of 1077 on eServer pSeries IBM 690 with Power4+ 1.7GHz processor.

8 Source www.spec.org : SPECfp_base2000 of 2161 on HP Integrity Server rx4640 with one Itanium® 2 processor 6M at 1.5GHz, 6M L3 cache, RedHat Linux AS2.1 operating system. Best RISC result of 1642 on eServer pSeries IBM 655 with Power4+ 1.7GHz processor.

9 Source: www.tpc.org. 1,008,144 tpmC at $8.33/tpmC on HP Integrity Superdome,, with 64 Intel Itanium 2 processors 6M, each at 1.5 GHz with 6MB of L3 cache, running HP UX 11.iv2 64-Bit with Oracle Database 10G Enterprise Edition, with 1024 GB RAM. TPC-C availability date: 4/14/2004.Best single system RISC: 1,025,486 tpmC, $5.43/tpmC on IBM eServer pSeries 690 Turbo 7040-681, with thirty two (32) IBM Power4+ processors at 1.9GHz, running IBM AIX 5L V5.2 with IBM DB2UDB 8.1, 1024 GB RAM; TPC-C availability date: 8/16/2004.

10 Source: www.sap.com/benchmark. Itanium® 2 processor result of 586,319 measured on Unisys ES7000 Model 400, 16-way SMP, Intel Itanium 2, 1.5 GHz, 6 MB L3 cache, 64 GB main memory, Windows Server 2003 Datacenter Edition, SQL Server 2000, SAP APO Release 3.1, LiveCache 7.4.2. Best Power4 result of 474,162 from www.sap.com/benchmark on IBM eServer pSeries 690, 16-way, POWER4 1.3 GHz, 256MB L3 cache, AIX 5.1, Oracle 9i, SAP APO Release 3.1A, LiveCache 7.4.1.18.

11 Source: www.sgi.com Itanium® 2 Processor 6M measurements on a SGI Altix 3000 with 32 Itanium® 2 Processors at 1.5GHz. Best Power4 results from NAS Parallel Benchmark, Class A www.nas.nasa.gov/Software/NPB/ http://www.csm.ornl.gov/~dunigan/sp4/index.html#mpi on IBM p690 with 32 Power4 processors at 1.3GHz. Itanium 2 Processor 6M: 17.2 GFLOPs, Best Power4: 9.5 GFLOPs.

12 Source: www.spec.org 116,466 ops/sec on HP Server rx5670 with 4 Itanium® 2 processors 6M at 1.5GHz, each with 6MB L3 cache, 16GB of memory, HP-UX 11i v2.0, Hotspot JVM 1.4.2.00.
Best published Power4 result of 96,377 on eServer pSeries IBM 655 with 4 Power4+ 1.7GHz processors, 16GB memory, AIX 5L V5.2 APAR IY43549, JVM J2RE 1.4.1 IBM AIX build cadev-20030410.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance on Intel products, visit: http://www.intel.com/performance/resources/limits.htm

Up to ~2X the performance of SPARC, ~20-30% higher than Power

Results as of 4/23/04

Itanium® 2 & Xeon™ processors Compared

Performance

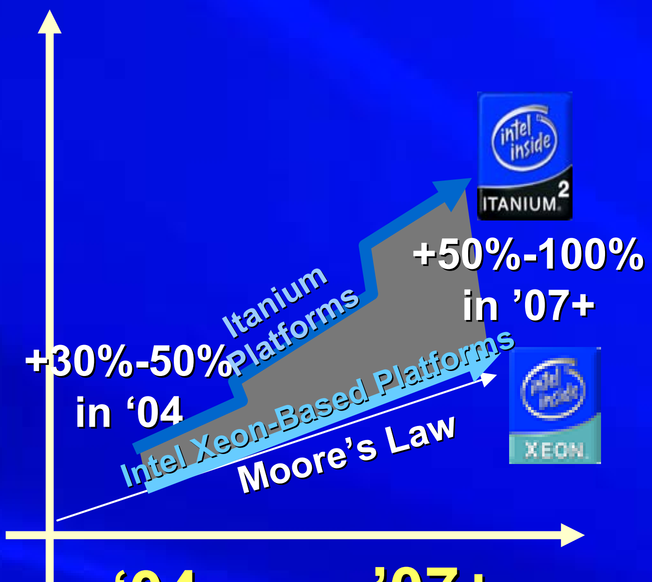

** For Enterprise & Technical Computing Application Segments*

Up to 2X performance at same price expected by '07

Reliability Features

Characteristic	Itanium® 2	Xeon™ MP
Error recovery on data bus-ECC	✓	
Internal soft error logic check	2005	
Machine Check Architecture	✓	
Bad data containment	✓	
Cache Reliability	2005	
Lockstep support	✓	✓
Memory SDEC, retry double-bit	✓	✓
Memory spares	✓	✓
Partitioning	✓ *node*	✓ *node*

High-end "RISC"-level RAS

Scalability

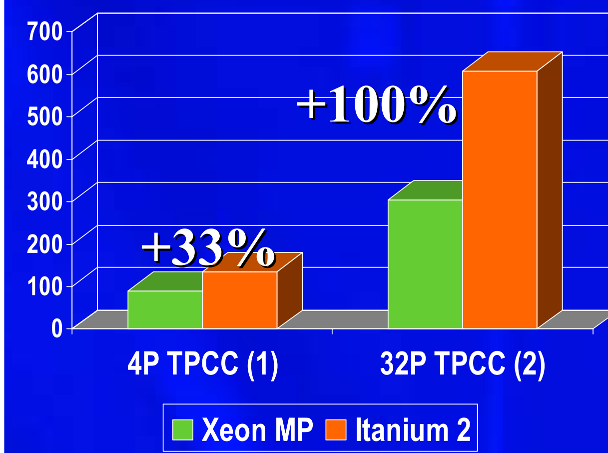

(1) www.tpc.org. Itanium 2 1.5GHz 6M = 136K on HP Integrity rx5670, Red Hat Enterprise Linux AS 3, Oracle 10g Standard Ed; Xeon MP 3.0GHz 4M = 103K on IBM x-Series 365, Windows Server 2003 Enterprise Edition, SQL Server 2000 Enterprise Ed.

(2) www.tpc.org. Itanium 2 1.5GHz 6M = 609K on NECExpress 5800, Suse Linux Enterprise Server v9, Oracle 10g Enterprise Ed; Xeon MP 3.0GHz 4M = 304K on Unisys ES7000, Windows Server 2003 Enterprise Ed., SQL Server 2000 Enterprise Ed.

~2x Higher Scalability

All dates, product, comparisons and features are preliminary and subject to change without notice

Itanium's EPIC Architecture: Intel's Highest Performance, Reliability Features, Scalability

Intel Enterprise Technologies

TECHNOLOGY	DESCRIPTION	BENEFIT
POWER: DBS *Demand Based Switching*	OS changes CPU power based upon utilization	Reduced energy consumption with lower utilization. Estimate most users would save up to 30% in power & cooling.
POWER: ACPC *Automatic Control of Power Consumption*	User defined maximum power threshold	Eliminates ambiguity in allocating power & cooling. 2-3 times more servers can be deployed in a rack.
I/O: PCI-EXPRESS	Serial I/O, removing bottlenecks caused by arbitration of parallel I/O	Increased performance, RAS and lower power. Headroom to accommodate LAN upgrade opportunities.
VIRTUALIZATION: Silvervale Technology	Processor includes "hooks". VMM talks to hooks.	Easier and more robust way for virtualized partitions to support consolidation
RELIABILITY: Pellston Technology	Capability to disable L3 cache lines that exhibit failures	Brings Itanium family to/beyond RISC & mainframe reliability features. Helps to reduce risk of ECC errors in L3
Multi-Core	Cores, cache, interconnect and arbitration logic within single die.	Increased performance, more throughput, better price / performance; no recompilation
Multi-threading	Run in parallel on Xeon processors (Hyper-Threading); On Itanium processor family, time share resources switching on long latency events.	Increased performance, more throughput, better price / performance

Comparisons referenced are in relation to Intel product offerings today (as of May 2004)

Itanium® 2 Processors Deliver Exceptional Performance!

How Does it do it?

intel

Agenda

- Why migrate to Itanium® Processor Family?
- Itanium® Architecture
- Intel Software Development Tools

EPIC Architecture Principles

1. “The *compiler* should play the key role in designing the plan of execution, and the architecture should provide the requisite support for it to do so successfully;
2. The Architecture should provide features that assist the *compiler* in exploiting statistical ILP; And
3. The Architecture should provide mechanisms to communicate the *compiler’s* plan of execution to the hardware.”

   Schlansker, Michael S. and Rau, B. Ramakrishna (HP Labs),

   “EPIC: Explicitly Parallel Instruction Computing”

   Computer, Vol 33 Issue: 2 , Feb. 2000 pp 37-45

Key Architectural Features

- Instruction-level parallelism
- Large number of registers
- Speculation
- Predication
- Register stack
- **Software pipelining**
- Advanced floating point
- Large Caches

Instruction Level Parallelism

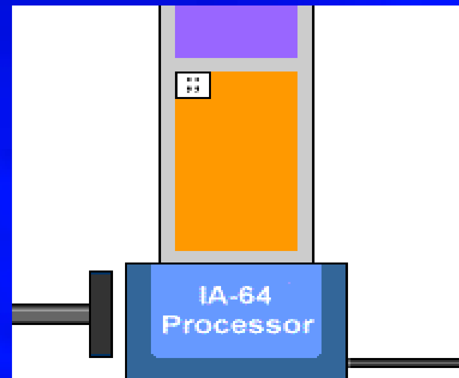

Instruction Groups

- An instruction group is one or more independent instructions that may be executed concurrently
- Read after write (RAW) or write after write (WAW) dependencies block concurrence
- Instruction groups issue in parallel, depending on available resources

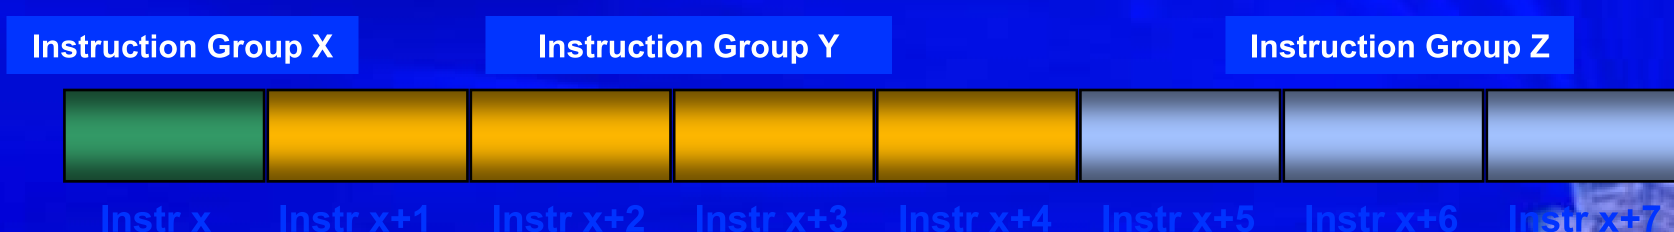

Instruction Level Parallelism

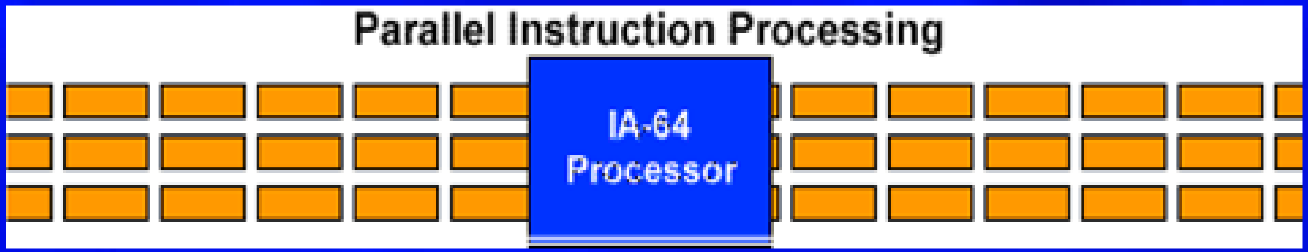

Bundles

- Three instructions and a template form a 128-bit bundle
- The template specifies the types of instructions in each bundle
- Instruction groups can span multiple bundles

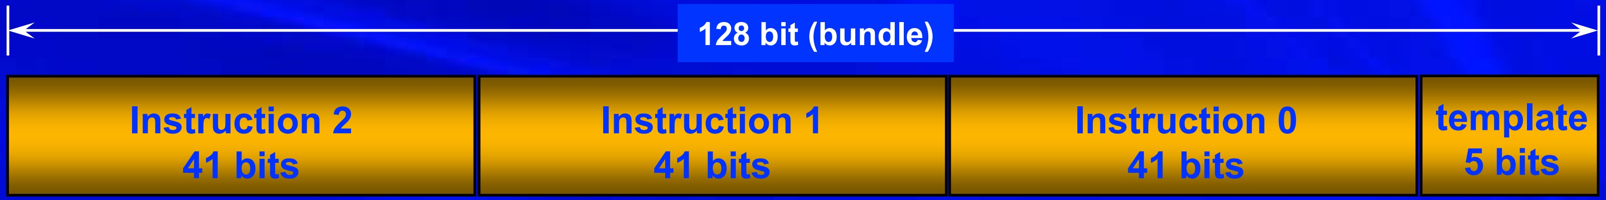

Massive Register Set

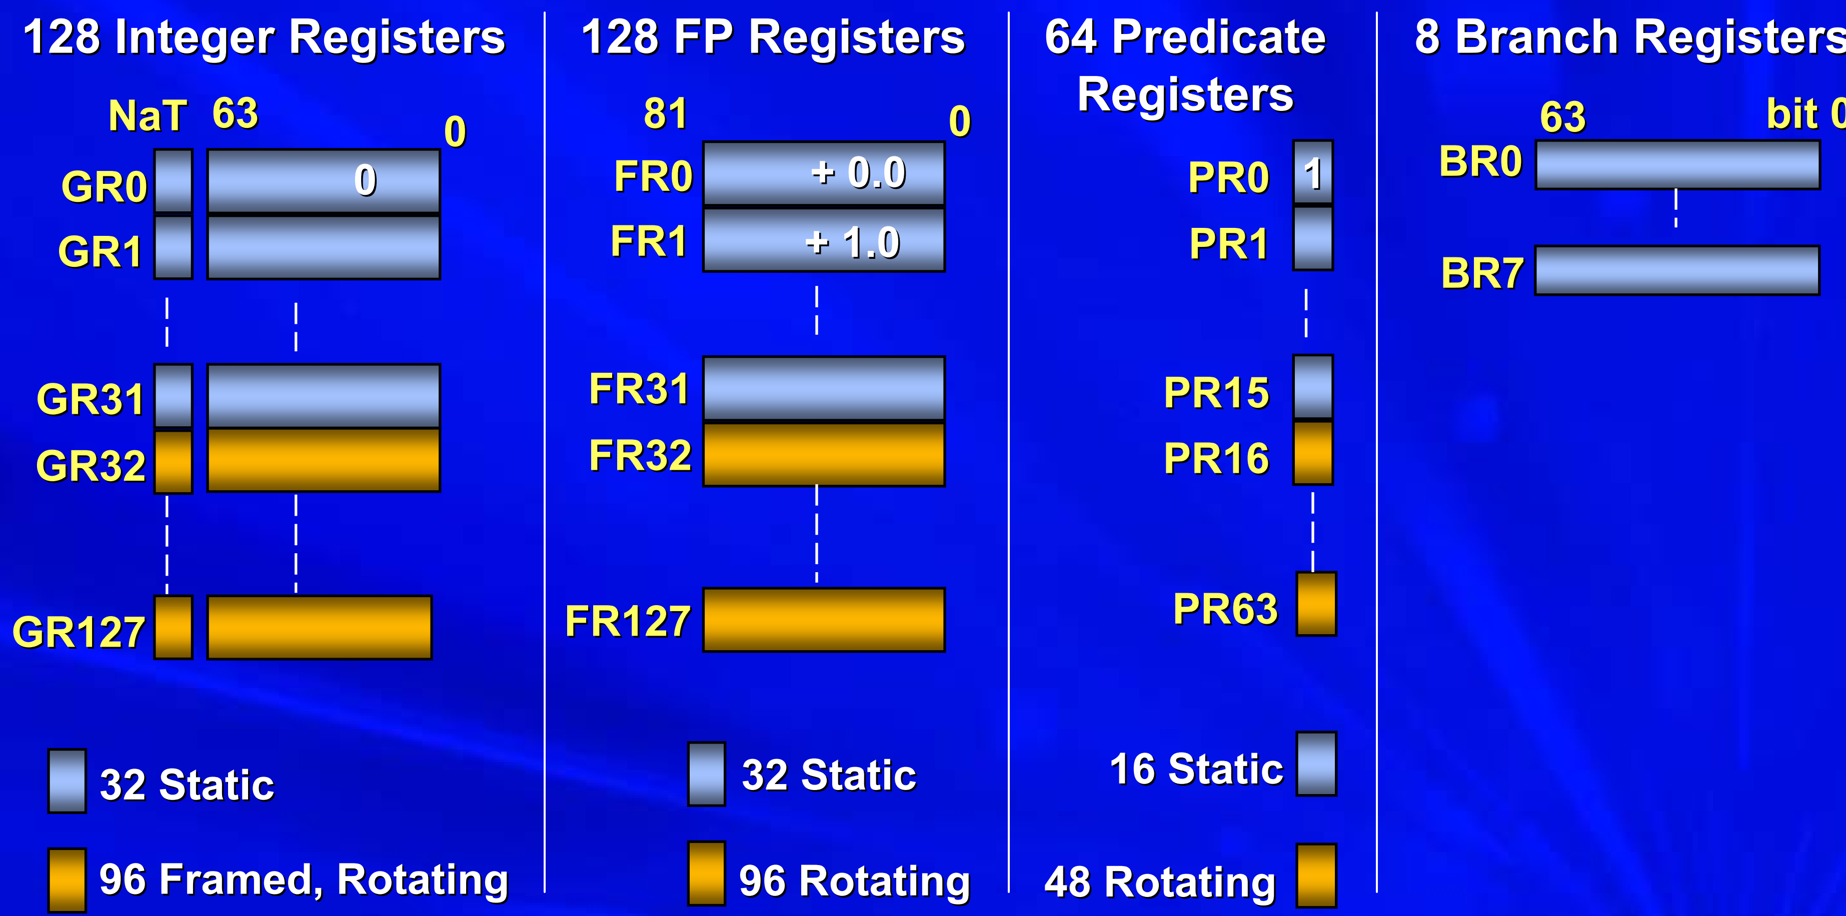

Large number of registers enables flexibility and performance

Register Stack

- GRs 0-31 are global to all procedures
- GRs 32-127 stacked registers, local to each procedure
- Only GRs implement a register stack
  - The FRs, PRs, and BRs are global to all procedures
- Register Stack Engine (RSE)
  - Upon stack overflow/underflow, registers are saved/restored to/from a backing store transparently

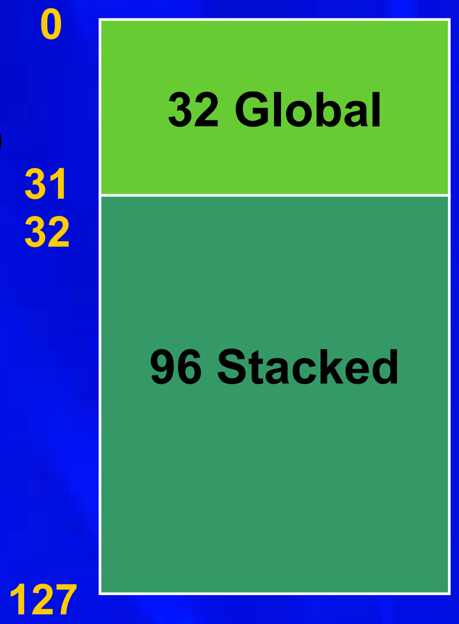

Optimizes the Call/Return Mechanism

Register Stack

- Call changes frame to contain only the caller's output
- Alloc sets the frame region to the desired size
  - Three architecture parameters: local, output, and rotating
- Return restores the stack frame of the caller

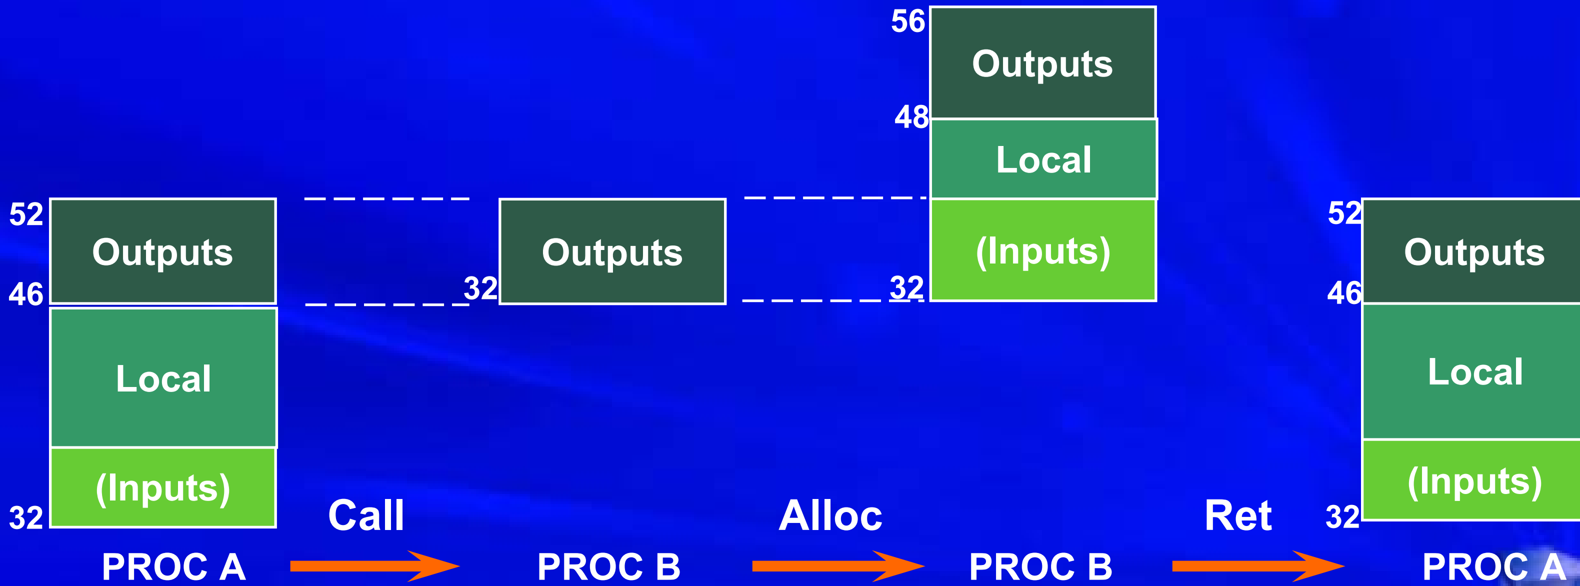

Avoids register spill/fill upon procedure call/return

Itanium® 2 Processor Software Pipelining

- Itanium 2 Processors invoke software loop pipelining through rotating registers
  - Absorbing latency through multiple loop iterations
  - Creating Parallelism through multiple iterations
- Predication creates pipeline prologues and epilogues
  - Single instruction stream for
    - Pipeline loading (prologue)
    - Loop body (kernel)
    - Pipeline draining (epilogue)

Rotating Registers

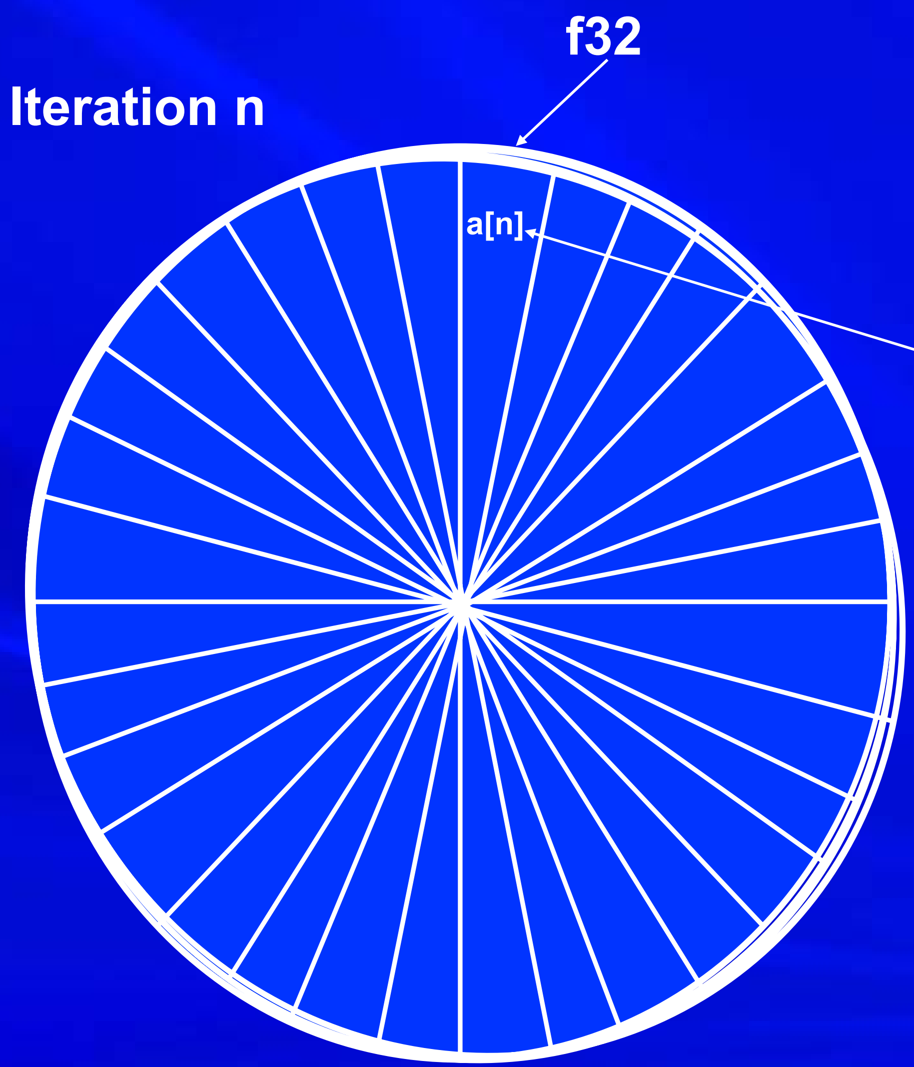

For(I=0;I<MAX;I++)sum += a[I];

Loop:

```
ldfs f32 = [r32],4
fma f46 = f50,f1,f38
br.ctop.sptk    loop;;
```

Rotating Registers

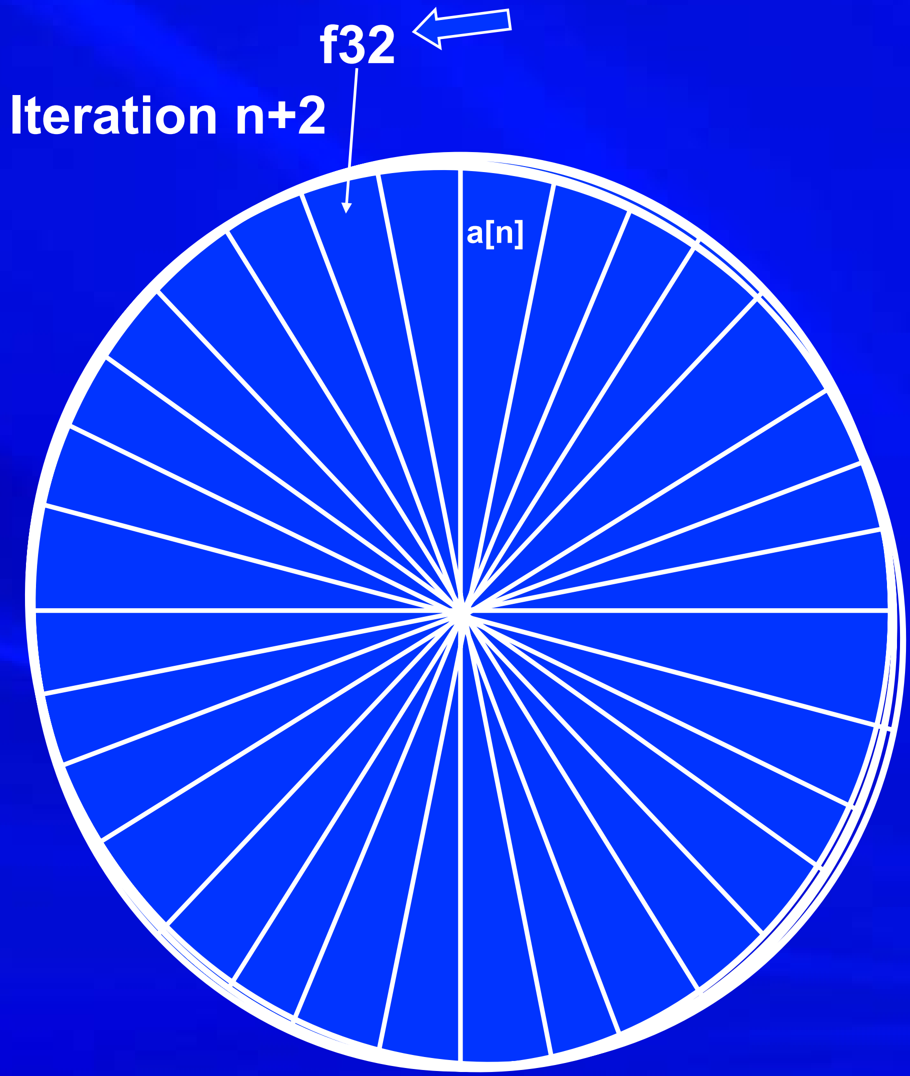


```
For(I=0;I<MAX;I++)sum += a[I];
Loop:
        ldfs f32 = [r32],4
        fma f46 = f50,f1,f38
        br.ctop.sptk    loop;;
```

Rotating Registers

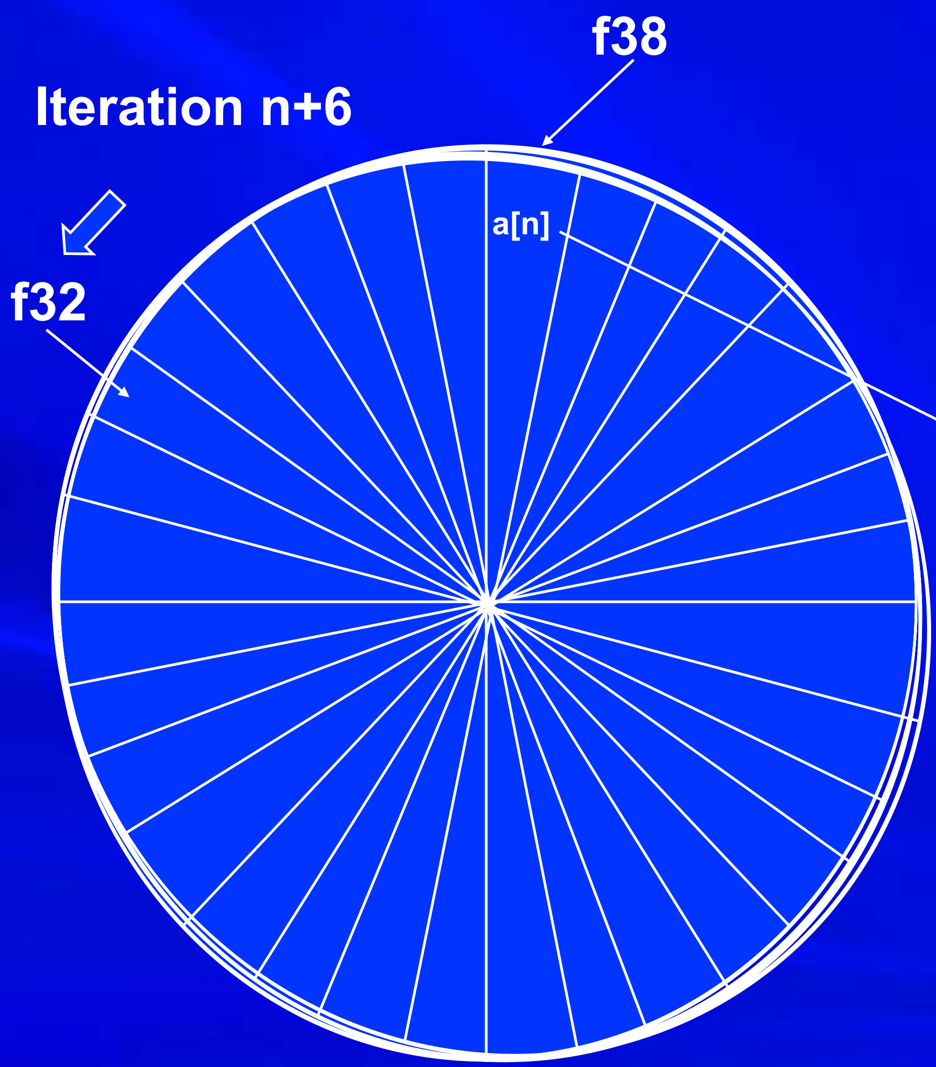

For(I=0;I<MAX;I++)sum += a[I];

Loop:

```
ldfs f32 = [r32],4
fma f46 = f50,f1,f38
br.ctop.sptk    loop;;
```

Executing Multiple iterations in Parallel

Predication Creates Prologue and Epilogue

Iteration 0

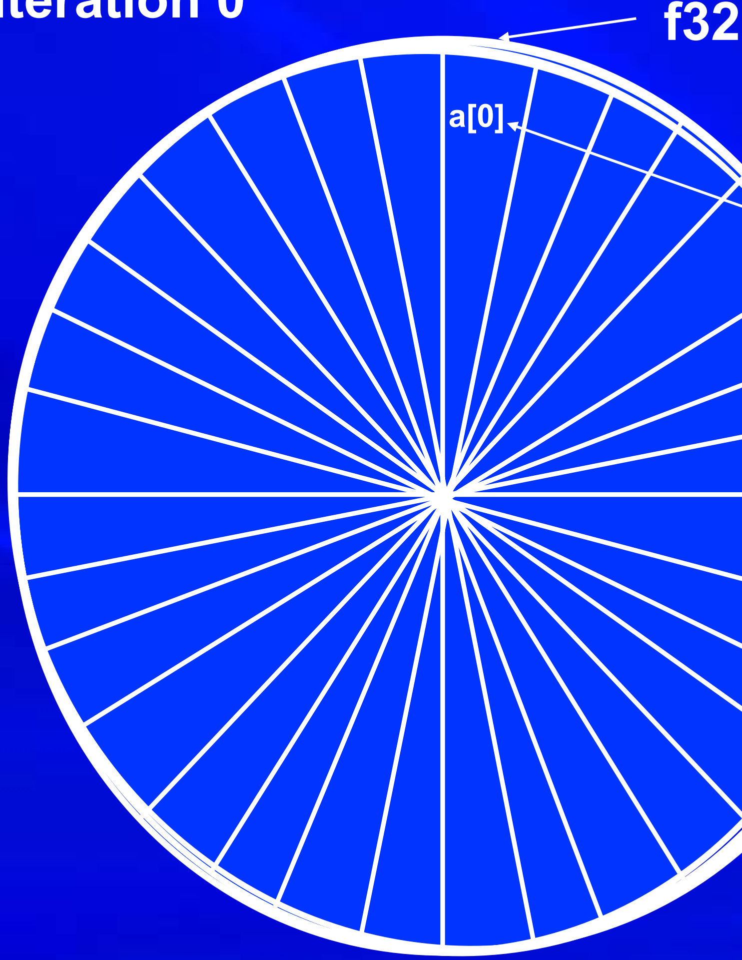


```
For(I=0;I<MAX;I++)sum += a[I];
Loop:
 (p16) ldfs f32 = [r32],4
 (p22) fma f46 = f50,f1,f38
       br.ctop.sptk    loop;;
```

p16	1
p17	0
p18	0
p19	0
p20	0
p21	0
p22	0
p23	0

Predication Creates Prologue and Epilogue

Iteration 2

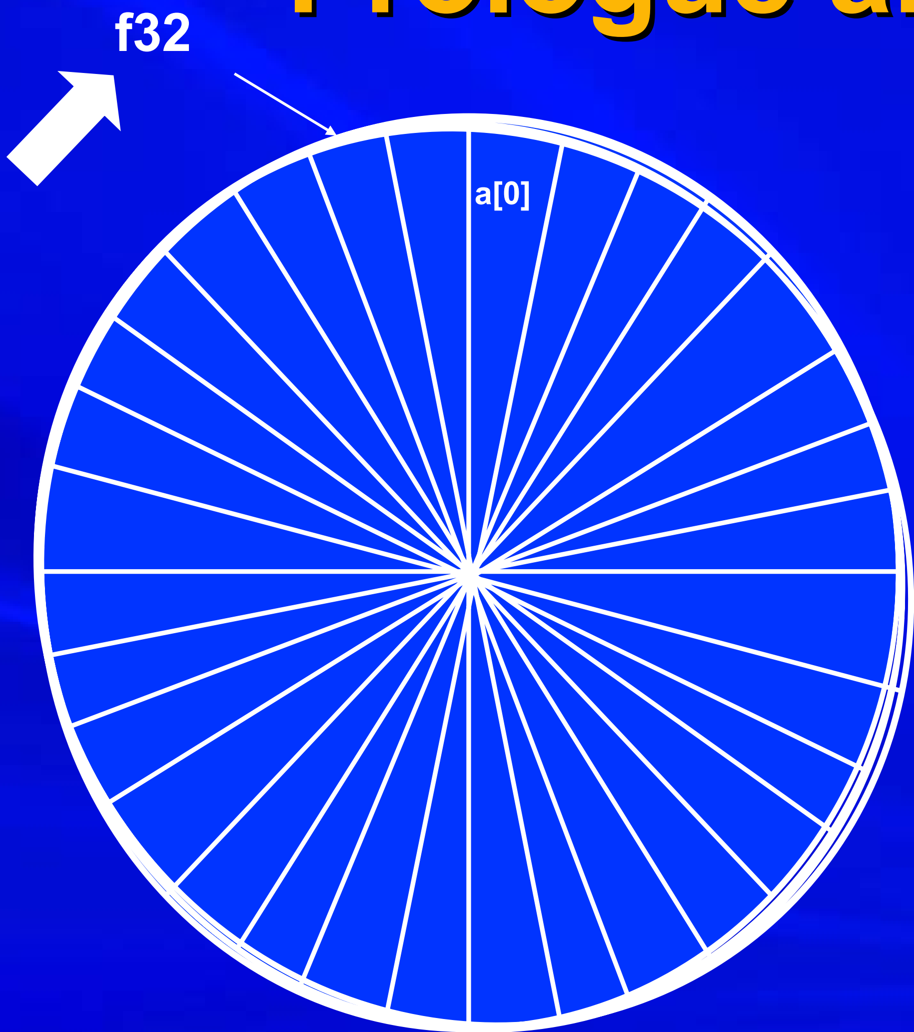


```
For(I=0;I<MAX;I++)sum += a[I];
Loop:
  (p16) ldfs f32 = [r32],4
  (p22) fma f46 = f50,f1,f38
        br.ctop.sptk    loop;;
```

p16	1
p17	1
p18	1
p19	0
p20	0
p21	0
p22	0
p23	0

Predication Creates Prologue and Epilogue

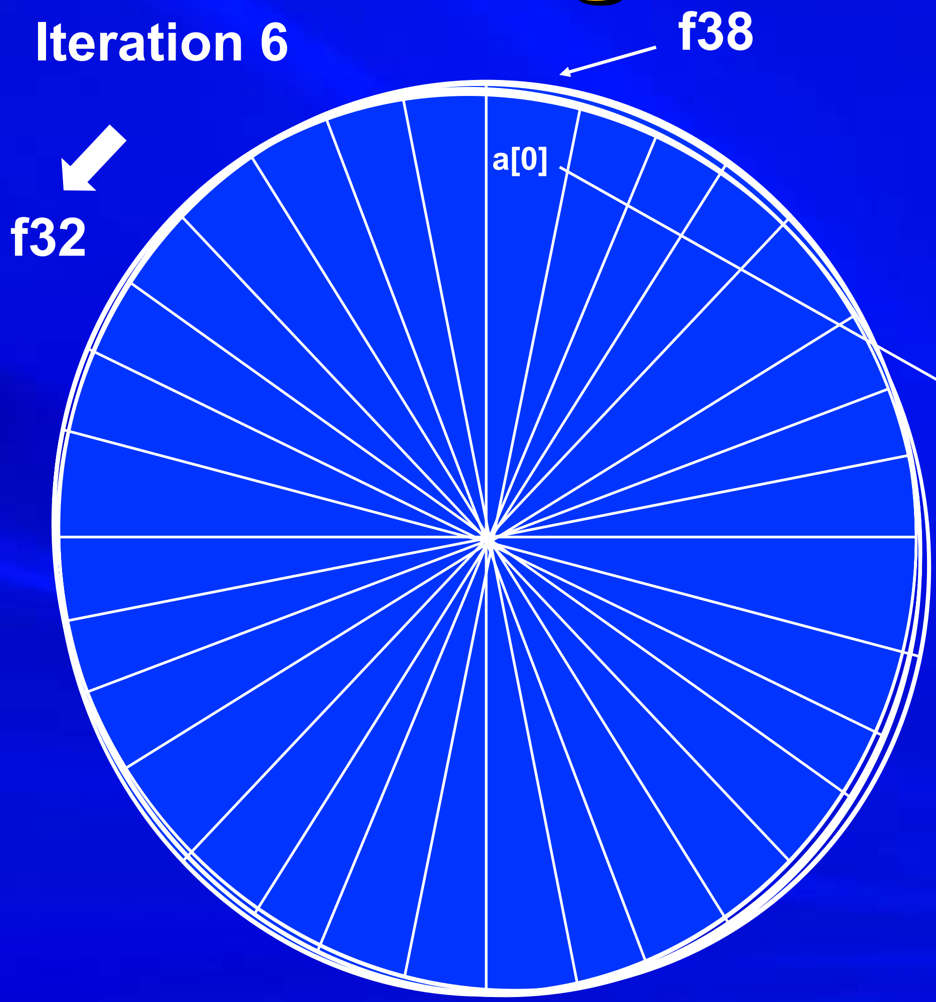

For(I=0;I<MAX;I++)sum += a[I];

```
Loop:
  (p16) ldfs f32 = [r32],4
  (p22) fma f46 = f50,f1,f38
        br.ctop.sptk    loop;;
```

p16	1
p17	1
p18	1
p19	1
p20	1
p21	1
p22	1
p23	0

Itanium® Architecture contains many advanced architectural features which enable performance

How can you maximize your use of those features to enable more performance?

Agenda

- Why migrate to Itanium® Processor Family?
- Itanium® Architecture
- Intel Software Development Tools

EPIC Architecture Principles

1. “The ***compiler*** should play the key role in designing the plan of execution, and the architecture should provide the requisite support for it to do so successfully;
2. The Architecture should provide features that assist the ***compiler*** in exploiting statistical ILP; And
3. The Architecture should provide mechanisms to communicate the ***compiler’s*** plan of execution to the hardware.”

   Schlansker, Michael S. and Rau, B. Ramakrishna (HP Labs), “EPIC: Explicitly Parallel Instruction Computing” Computer, Vol 33 Issue: 2 , Feb. 2000 pp 37-45

Intel® Software Development Products

- **Intel® Compilers**

  Best way to get application performance on Intel processors

- **Intel® VTune™ Analyzers**

  Quickly identify "hot spots" and how to fix them

- **Intel® Performance Libraries**

  Highly optimized, ready to use building-block functions

- **Intel® Threading Tools**

  Speeds, simplifies development & maintenance of threaded apps

- **Intel® Cluster Tools**

  Create, analyze, optimize and deploy cluster-based applications

Intel Software Development Products for
Intel® Personal Internet Client Architecture processors,
Pentium® M, Pentium® 4, Intel® Xeon™
and Itanium® 2 Processors

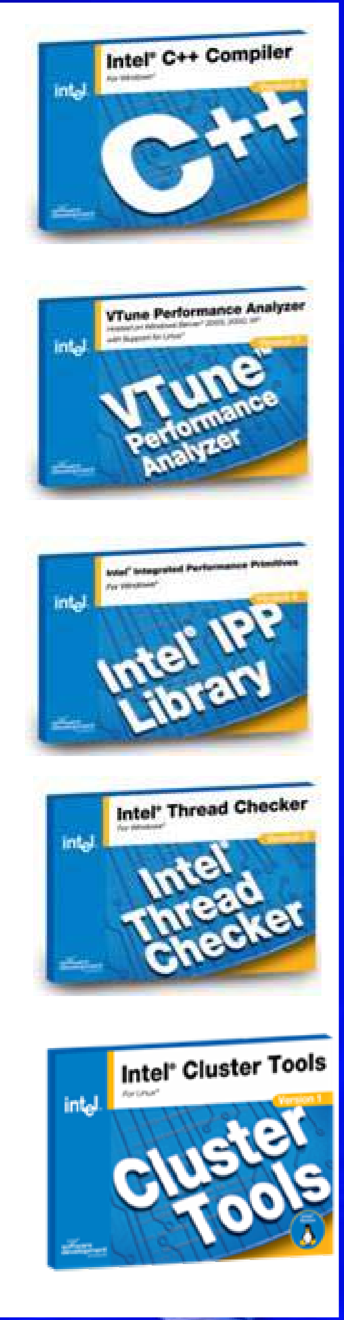

Lawrence Livermore National Laboratory

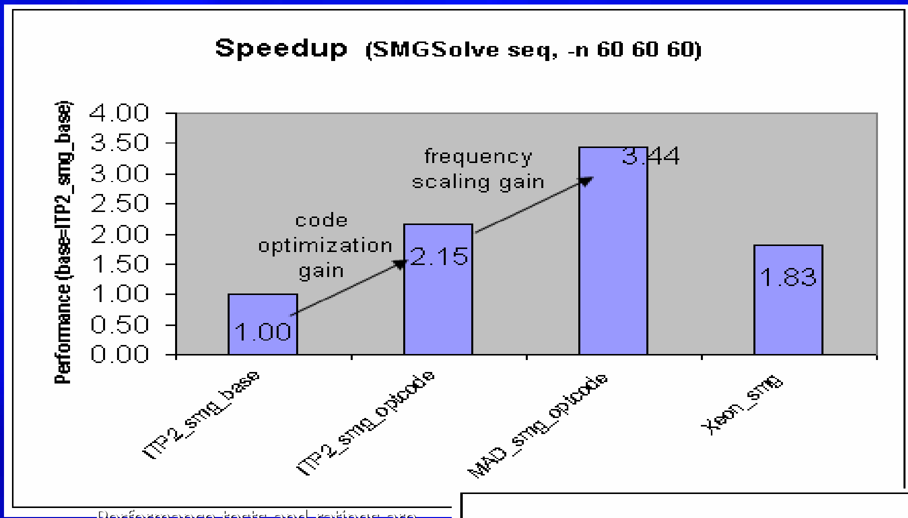

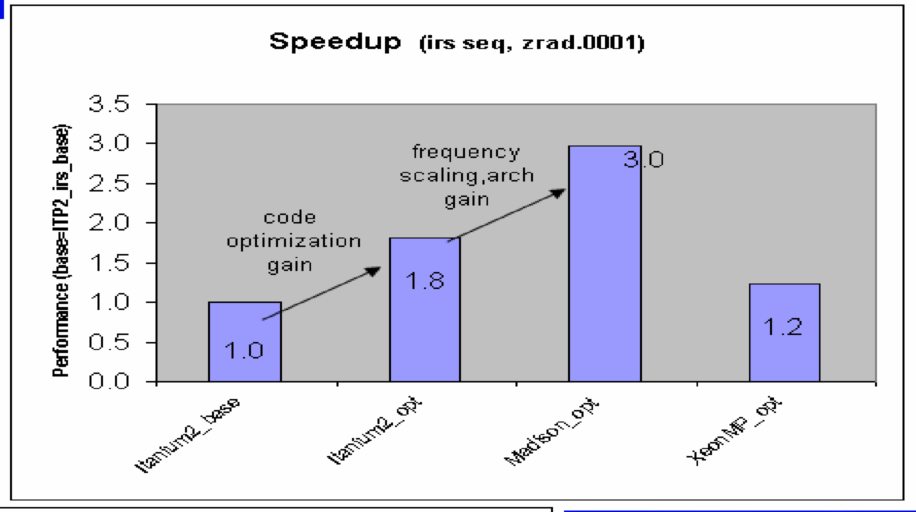

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/procs/perf/limits.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

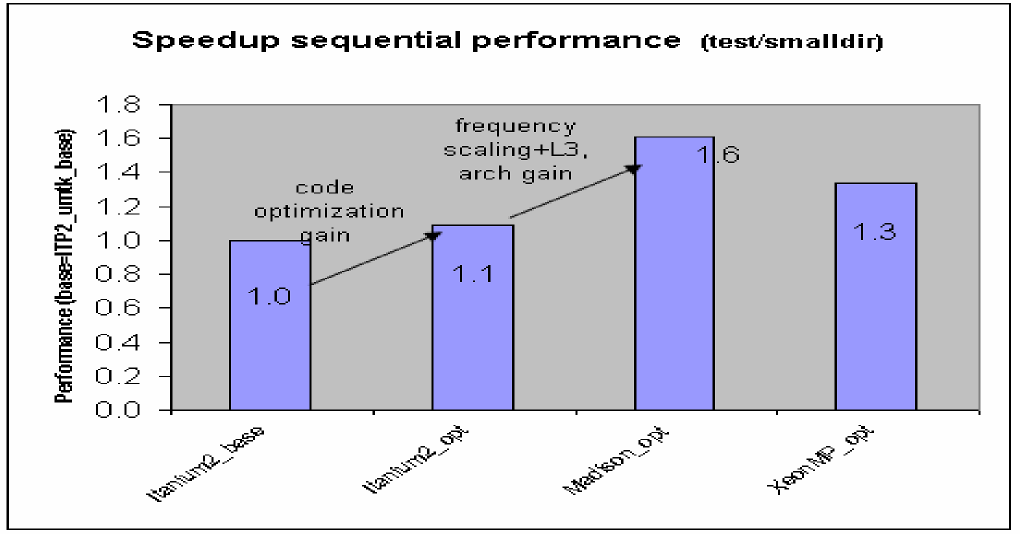

NOTE: Itanium® 2 processor @900 MHz with 3MB cache; Itanium® 2 processor @1.5 GHz with 6MB cache; Xeon™ processor @ 2 GHz with 512KB cache

Los Alamos National Laboratory

- Results
  - After a 2 week engagement, achieved 1.55x performance improvement (used Cell Updates per Cycle per second as the performance metric)
  - Most of the speed-up came from applying optimization flags specific for Itanium® 2 processor
  - Additional performance improvement came from prefetching and loop unrolling to relieve processor stalls

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/procs/perf/limits.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

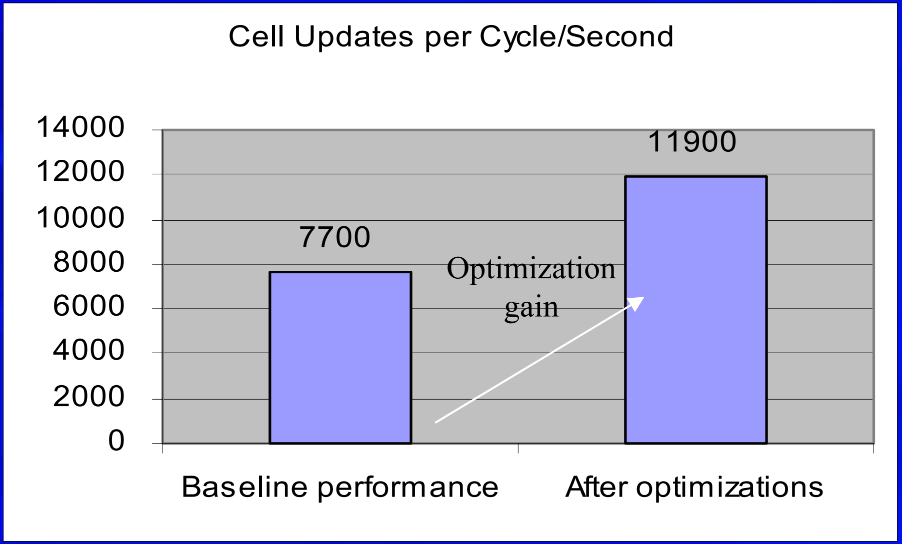

NOTE: Itanium 2 processor @1.2 GHz with 6MB cache

Summary

- Intel® Itanium® Architecture Delivers Excellent Performance
- Contains Many Next Generation Processor features
  - i.e. H … ware Pipe …
- The Performance Features are Uncovered via The Software Development Tools

Up next – Technical Details on How to Optimize for Itanium® 2 Processor

Back-up

Multi-Core Technology

2004	2005	2007
Single Core	Dual Core	Multi-Core

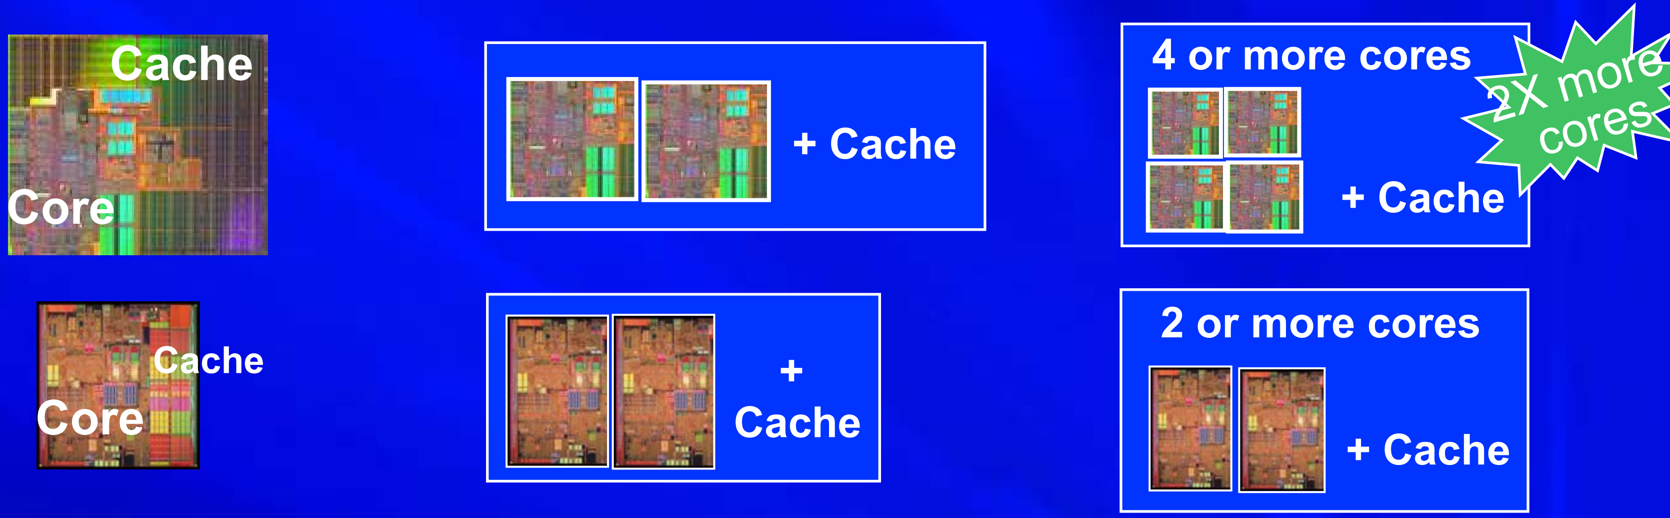

- Intel's manufacturing leadership (90nm, 65nm) enables leading multi-core
  - Xeon and Itanium multi-cores on smaller die sizes (lower cost) without stripping functionality (which sacrifices performance)
- Itanium® architecture has smaller core size – enabling up to 2x more cores per die than IA-32 for higher performance at same cost

All products, dates and features are preliminary and subject to change without notice

Itanium® architecture expected to enable up to 2x more cores per processor than Xeon processors by 2007

Cache Structure

- Very high speed memory for data that gets reused
- Organized into “cache lines”
  - Access of a single element brings in enough adjacent elements to fill the line (64/128 consecutive bytes)
  - Underlying assumption that if you need one element you will need its neighbors soon
- Cache lines are organized into “associative sets” or “ways”
  - Greater associativity allows the hardware more flexibility in cache line replacement algorithms

Cache Structure

- Very high speed memory for data that gets reused
- Organized into “cache lines”
  - Access of a single element brings in enough adjacent elements to fill the line (64/128 consecutive bytes)
  - Underlying assumption that if you need one element you will need its neighbors soon
- Cache lines are organized into “associative sets” or “ways”
  - Greater associativity allows the hardware more flexibility in cache line replacement algorithms

L1D Cache Structure

Itanium® 2 Processor 16 KB 4-way associative L1 Data Cache (Integer Data)

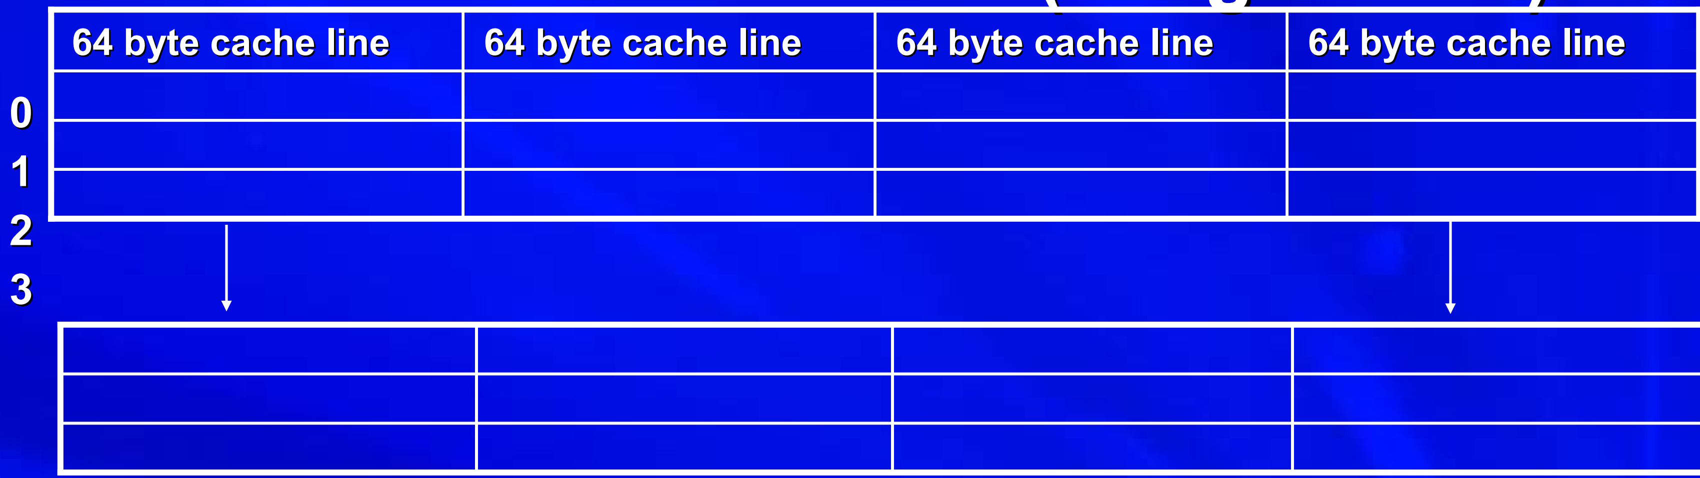

- L1D has one cycle latency for integer loads:
  - All cacheable integer loads go through L1D (write through)
  - Use L1D micro pipeline to access general register file
  - Cache line row determined by address bits 6 through 11
  - 64 associative sets

L2 Unified Cache Bank Structure

- 256KB, 128 byte cache lines, 8 way associativity
- Each associative set is 1KB, 256 associative sets
- Bank structure allows fast transfers from/to large Itanium® 2 Processor L2 Cache
- 16 banks each 16 bytes wide

16 banks cover 256 bytes = 2 cache lines

		Line	1	Of	8					Line	2	Of	8		
		Line	3	Of	8					Line	4	Of	8		
		Line	5	Of	8					Line	6	Of	8		
		Line	7	Of	8					Line	8	Of	8		

cache lines

cache lines

L2 Unified Cache Bank Structure

- 256KB, 128 byte cache lines, 8 way associativity
- Each associative set is 1KB, 256 associative sets
- Bank structure allows fast transfers from/to large Itanium® 2 Processor L2 Cache
- 16 banks each 16 bytes wide

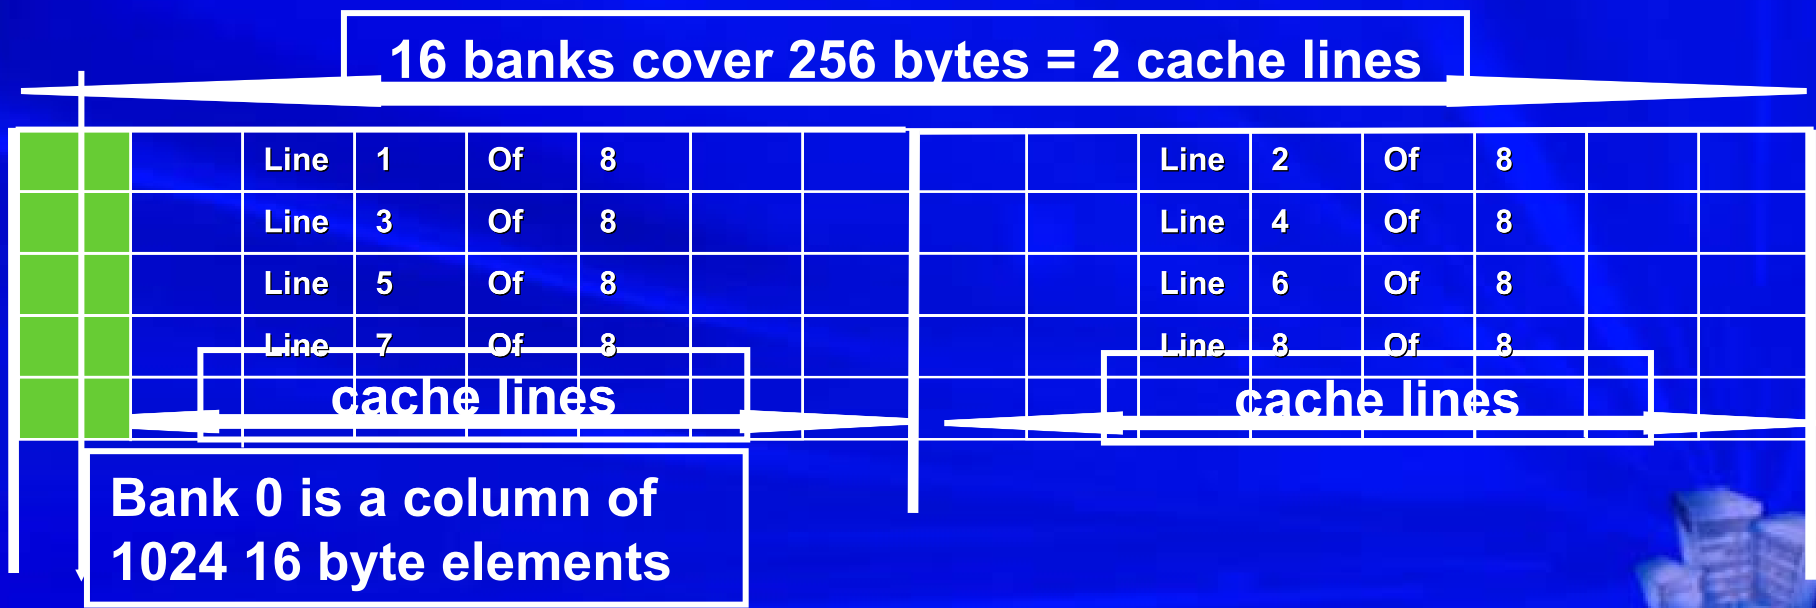

Itanium® 2 Processor L2 Cache Access

- L2 data access controlled by 32 entry queue (OzQ) and allows out of order data return
  - FP data loaded to FP register file directly from L2
- Minimum integer latency is 5 cycles
- Minimum floating point latency is 6 cycles
- Latency is increased by:
  - Cache miss
  - Bank conflicts cause OzQ cancels (measured to add 6 cycles)
  - Multiple misses and misses to lines being updated will cause OzQ recirculates (measured to add ~17 cycles)
    - Only one data access is escalated to L3 and the system bus, the others recirculate

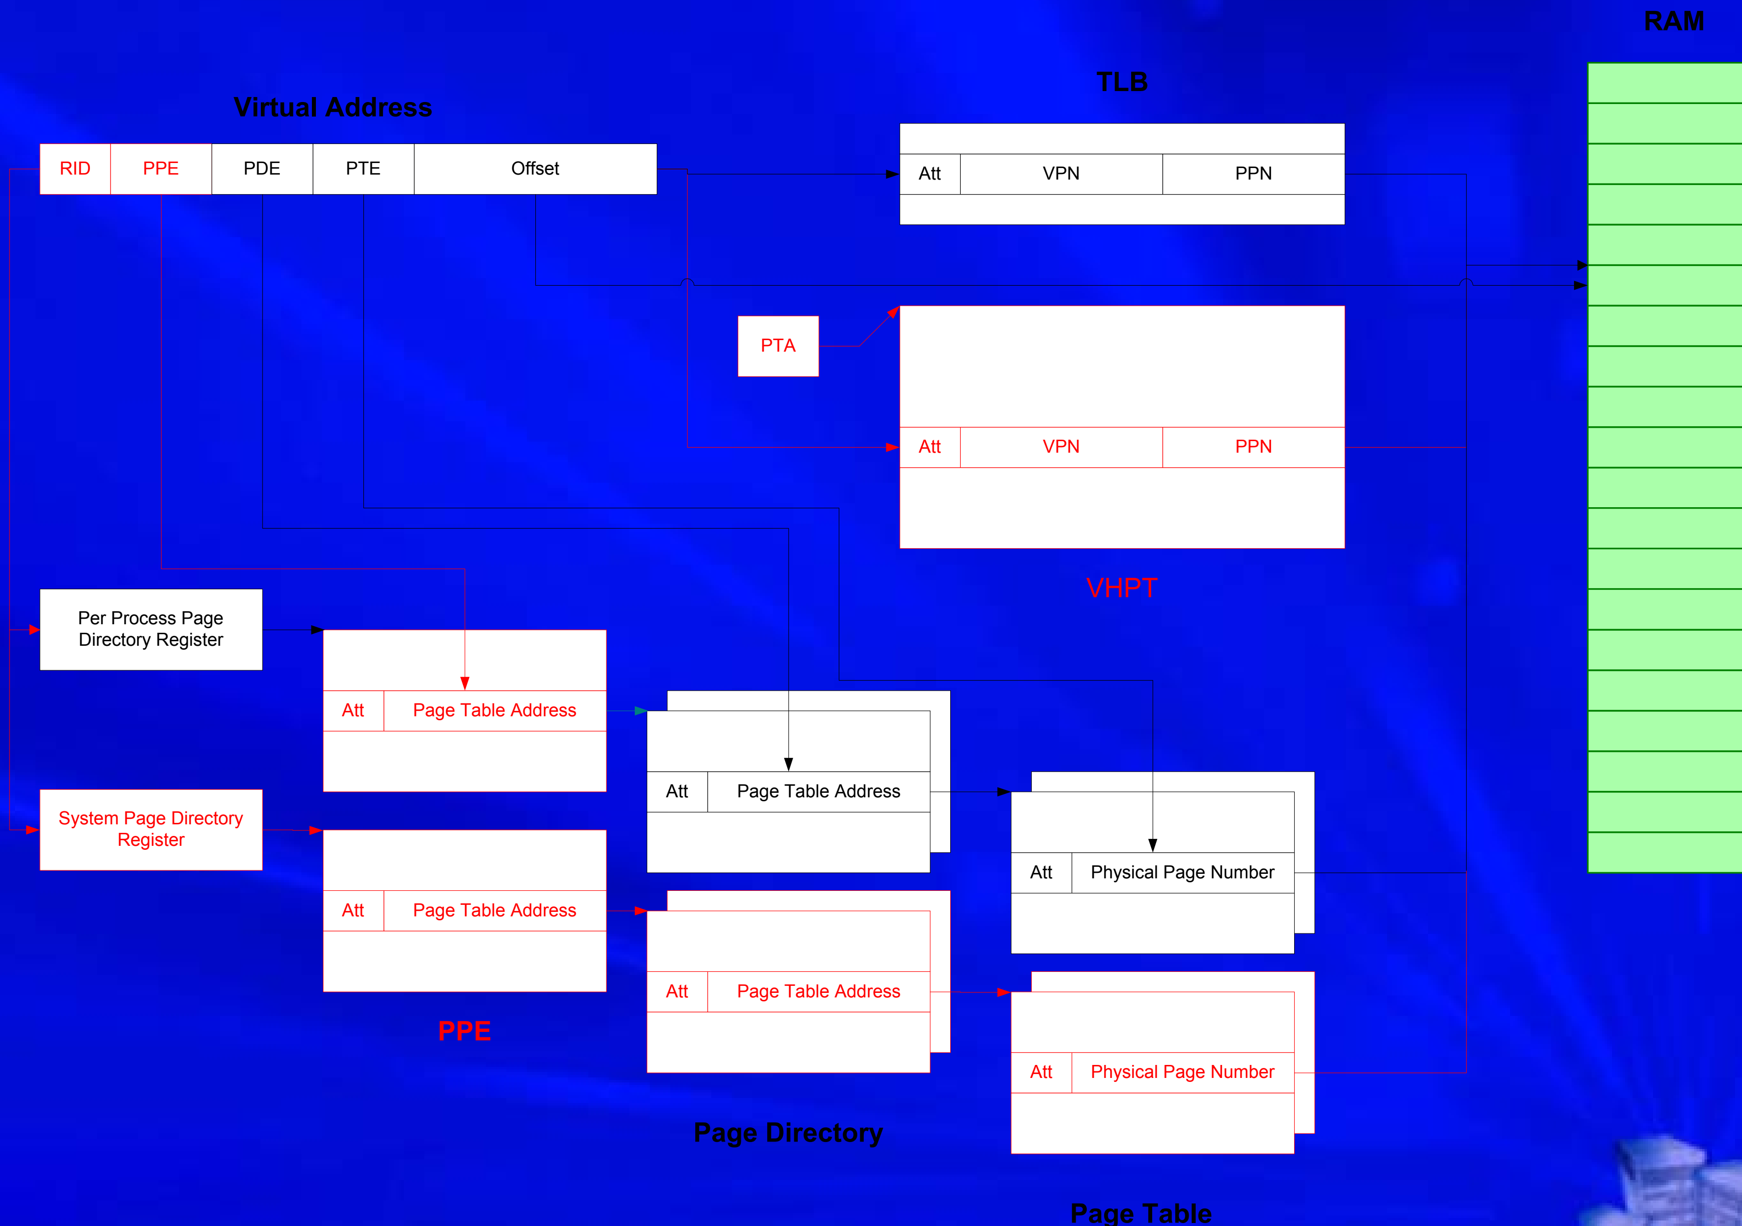
RAM
TLB
Virtual Address
RID
PPE
PDE
PTE
Offset
Att
VPN
PPN
PTA
Att
VPN
PPN
VHPT
Per Process Page Directory Register
Att
Page Table Address
Att
Page Table Address
System Page Directory Register
Att
Page Table Address
Att
Physical Page Number
Att
Page Table Address
Att
Physical Page Number
PPE
Page Directory
Page Table