



# P6 Family of Processors

Hardware Developer's Manual

---

*September 1998*



Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Pentium® II processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>

Copyright © Intel Corporation 1998.

\* Third-party brands and names are the property of their respective owners.

---

<b>1</b>	<b>Introduction</b> .....	1-1
	1.1 P6 FAMILY OF PROCESSORS OVERVIEW .....	1-1
	1.2 TERMINOLOGY .....	1-2
	1.3 SPECIFIC PRODUCT REFERENCES .....	1-2
<b>2</b>	<b>Micro-Architecture Overview</b> .....	2-1
	2.1 FULL CORE UTILIZATION .....	2-1
	2.2 THE P6 FAMILY PROCESSOR PIPELINE .....	2-2
	2.2.1 The Fetch/Decode Unit .....	2-3
	2.2.2 The Dispatch/Execute Unit .....	2-4
	2.2.3 The Retire Unit .....	2-6
	2.2.4 The Bus Interface Unit .....	2-6
	2.3 ARCHITECTURE SUMMARY .....	2-7
<b>3</b>	<b>System Bus Overview</b> .....	3-1
	3.1 SIGNALING ON P6 FAMILY SYSTEM BUS .....	3-1
	3.2 SIGNAL OVERVIEW .....	3-2
	3.2.1 Execution Control Signals .....	3-2
	3.2.2 Arbitration Signals .....	3-3
	3.2.3 Request Signals .....	3-4
	3.2.4 Snoop Signals .....	3-4
	3.2.5 Response Signals .....	3-5
	3.2.6 Data Response Signals .....	3-6
	3.2.6.1 LINE TRANSFERS .....	3-6
	3.2.6.2 PART LINE ALIGNED TRANSFERS .....	3-7
	3.2.6.3 PARTIAL TRANSFERS .....	3-7
	3.2.7 Error Signals .....	3-7
	3.2.8 Compatibility Signals .....	3-9
	3.2.9 Diagnostic Signals .....	3-9
<b>4</b>	<b>Data Integrity</b> .....	4-1
	4.1 ERROR CLASSIFICATION .....	4-1
	4.2 P6 FAMILY PROCESSOR SYSTEM BUS DATA INTEGRITY ARCHITECTURE .....	4-2
	4.2.1 Bus Signals Protected Directly .....	4-2
	4.2.2 Bus Signals Protected Indirectly .....	4-3
	4.2.3 Unprotected Bus Signals .....	4-3
	4.2.4 Hard-Error Response .....	4-3
	4.2.5 P6 Family Processor System Bus Error Code Algorithms .....	4-3
	4.2.5.1 PARITY ALGORITHM .....	4-3
	4.2.5.2 P6 FAMILY SYSTEM BUS ECC ALGORITHM .....	4-4
<b>5</b>	<b>Configuration</b> .....	5-1
	5.1 DESCRIPTION .....	5-1
	5.1.1 Output Tristate .....	5-2
	5.1.2 Built-in Self-test .....	5-2
	5.1.3 Data Bus Error Checking Policy .....	5-2

5.1.4	Response Signal Parity Error Checking Policy .....	5-2
5.1.5	AERR# Driving Policy .....	5-3
5.1.6	AERR# Observation Policy .....	5-3
5.1.7	BERR# Driving Policy for Initiator Bus Errors .....	5-3
5.1.8	BERR# Driving Policy for Target Bus Errors.....	5-3
5.1.9	BERR# Driving Policy for Initiator Internal Errors .....	5-3
5.1.10	BINIT# Driving Policy .....	5-3
5.1.11	BINIT# Observation Policy.....	5-3
5.1.12	In-Order Queue Pipelining .....	5-4
5.1.13	Power-On Reset Vector .....	5-4
5.1.14	FFRC Mode Enable .....	5-4
5.1.15	APIC Mode .....	5-4
5.1.16	APIC Cluster ID .....	5-4
5.1.17	Symmetric Agent Arbitration ID.....	5-4
5.1.18	Low Power Standby Enable.....	5-7
5.2	CLOCK FREQUENCIES AND RATIOS .....	5-8
5.3	POWER-ON CONFIGURATION REGISTER .....	5-8
5.4	INITIALIZATION PROCESS .....	5-10
<b>6</b>	<b>Test Access Port (TAP).....</b>	<b>6-1</b>
6.1	INTERFACE .....	6-1
6.2	ACCESSING THE TAP LOGIC .....	6-2
6.2.1	Accessing the Instruction Register.....	6-4
6.2.2	Accessing the Data Registers.....	6-5
6.3	INSTRUCTION SET .....	6-6
6.4	DATA REGISTER SUMMARY .....	6-7
6.4.1	Bypass Register.....	6-7
6.4.2	Device ID Register.....	6-7
6.4.3	BIST Result Boundary Scan Register.....	6-8
6.4.4	Boundary Scan Register.....	6-8
6.5	RESET BEHAVIOR .....	6-8
<b>7</b>	<b>Integration Tools.....</b>	<b>7-1</b>
7.1	IN-TARGET PROBE (ITP) FOR P6 FAMILY PROCESSORS .....	7-1
7.1.1	Primary Function.....	7-1
7.1.2	Debug Port Connector Description .....	7-2
7.1.3	Debug Port Signal Descriptions .....	7-2
7.1.4	Debug Port Signal Notes .....	7-2
7.1.4.1	SIGNAL NOTE 1: DBRESET#.....	7-3
7.1.4.2	SIGNAL NOTE 5: TDO AND TDI.....	7-3
7.1.4.3	SIGNAL NOTE 7: TCK.....	7-6
7.1.5	Debug Port Layout.....	7-6
7.1.5.1	SIGNAL QUALITY NOTES .....	7-8
7.1.5.2	DEBUG PORT CONNECTOR .....	7-8
7.1.6	Using Boundary Scan to Communicate to the Processor.....	7-8
<b>A</b>	<b>Signals Reference .....</b>	<b>A-1</b>
A.1	ALPHABETICAL SIGNALS LISTING .....	A-1
A.1.1	A[35:3]# (I/O) .....	A-1
A.1.2	A20M# (I) .....	A-1
A.1.3	ADS# (I/O) .....	A-1



A.1.4	AERR# (I/O)	A-2
A.1.5	AP[1:0]# (I/O)	A-2
A.1.6	ASZ[1:0]# (I/O)	A-2
A.1.7	ATTR[7:0]# (I/O)	A-2
A.1.8	BCLK (I)	A-3
A.1.9	BE[7:0]# (I/O)	A-3
A.1.10	BERR# (I/O)	A-3
A.1.11	BINIT# (I/O)	A-4
A.1.12	BNR# (I/O)	A-4
A.1.13	BP[3:2]# (I/O)	A-4
A.1.14	BPM[1:0]# (I/O)	A-4
A.1.15	BPRI# (I)	A-4
A.1.16	BREQ[3:0]# (I/O)	A-5
A.1.17	D[63:0]# (I/O)	A-5
A.1.18	DBSY# (I/O)	A-6
A.1.19	DEFER# (I)	A-6
A.1.20	DEN# (I/O)	A-6
A.1.21	DEP[7:0]# (I/O)	A-6
A.1.22	DID[7:0]# (I/O)	A-6
A.1.23	DRDY# (I/O)	A-7
A.1.24	DSZ[1:0]# (I/O)	A-7
A.1.25	EXF[4:0]# (I/O)	A-7
A.1.26	FERR# (O)	A-7
A.1.27	FLUSH# (I)	A-8
A.1.28	FRCERR (I/O)	A-8
A.1.29	HIT# (I/O), HITM# (I/O)	A-8
A.1.30	IERR# (O)	A-8
A.1.31	IGNNE# (I)	A-9
A.1.32	INIT# (I)	A-9
A.1.33	INTR(I)	A-9
A.1.34	LEN[1:0]# (I/O)	A-9
A.1.35	LINT[1:0] (I)	A-10
A.1.36	LOCK# (I/O)	A-10
A.1.37	NMI(I)	A-11
A.1.38	PICCLK (I)	A-11
A.1.39	PICD[1:0] (I/O)	A-11
A.1.40	PRDY# (O)	A-11
A.1.41	PREQ# (I)	A-11
A.1.42	PWRGOOD (I)	A-11
A.1.43	REQ[4:0]# (I/O)	A-12
A.1.44	RESET# (I)	A-13
A.1.45	RP# (I/O)	A-13
A.1.46	RS[2:0]# (I)	A-13
A.1.47	RSP# (I)	A-14
A.1.48	SLP# (I)	A-14
A.1.49	SMI# (I)	A-14
A.1.50	SMMEM# (I/O)	A-14
A.1.51	SPLCK# (I/O)	A-14
A.1.52	STPCLK# (I)	A-14
A.1.53	TCK (I)	A-15
A.1.54	TDI (I)	A-15

A.1.55	TDO (O) .....	A-15
A.1.56	THERMTRIP# (O).....	A-15
A.1.57	TMS (I).....	A-15
A.1.58	TRDY# (I).....	A-15
A.1.59	TRST# (I) .....	A-15
A.2	SIGNAL SUMMARIES .....	A-16
<b>Index</b>	.....	INDEX-1

## Figures

2-1	Three Engines Communicating Using an Instruction Pool .....	2-1
2-2	A Typical Pseudo Code Fragment .....	2-1
2-3	The Three Core Engines Interface with Memory via Unified Caches .....	2-3
2-4	Inside the Fetch/Decode Unit .....	2-4
2-5	Inside the Dispatch/Execute Unit.....	2-5
2-6	Inside the Retire Unit .....	2-6
2-7	Inside the Bus Interface Unit.....	2-7
3-1	Latched Bus Protocol.....	3-1
5-1	Hardware Configuration Signal Sampling .....	5-1
5-2	BR[1:0]# Physical Interconnection with Two Symmetric Agents .....	5-5
5-3	BR[1:0]# Physical Interconnection with Four Symmetric Agents .....	5-6
6-1	Simplified Block Diagram of Processor TAP Logic .....	6-2
6-2	TAP Controller Finite State Machine .....	6-3
6-3	Processor TAP Instruction Register.....	6-4
6-4	Operation of the Processor TAP Instruction Register .....	6-5
6-5	TAP Instruction Register Access .....	6-5
7-1	Hardware Components of the ITP .....	7-2
7-2	GTL+ Signal Termination .....	7-3
7-3	Generic DP System Layout for Debug Port Connection .....	7-7
7-4	Debug Port Connector on Thermal Plate Site of Circuit Board.....	7-8
7-5	Hole Positioning for Connector on Thermal Plate Side of Circuit Board.....	7-8
7-6	Processor System Where Boundary Scan Is Not Used.....	7-9

## Tables

3-1	Execution Control Signals.....	3-2
3-2	Arbitration Signals.....	3-3
3-3	Request Signals.....	3-4
3-4	Snoop Signals.....	3-5
3-5	Response Signals.....	3-5
3-6	Data Phase Signals .....	3-6
3-7	Burst Order Used for P6 Family Processor Bus Line Transfers .....	3-6
3-8	Error Signals .....	3-7
3-9	PC Compatibility Signals .....	3-9
3-10	Diagnostic Support Signals.....	3-10
4-1	Direct Bus Signal Protection .....	4-2
5-1	APIC Cluster ID Configuration .....	5-4
5-2	P6 Family Processor Bus BREQ[1:0]# Interconnect (2-Way MP Processors).....	5-5



5-3	P6 Family Processor Bus BREQ[3:0]# Interconnect (4-Way MP Processors).....	5-6
5-4	Arbitration ID Configuration (Two Agents) .....	5-7
5-5	Arbitration ID Configuration (Four Agents).....	5-7
5-6	System Bus To Core Frequency Multiplier Configuration .....	5-8
5-7	Processor Power-On Configuration Register .....	5-9
5-8	Power-On Configuration Register APIC Cluster ID Bit Field.....	5-9
5-9	Power-On Configuration Register Arbitration ID Configuration.....	5-10
5-10	Power-On Configuration Register Bus Frequency to Core Frequency Ratio Bit Field.....	5-10
6-1	1149.1 Instructions in the Processor TAP.....	6-6
6-2	TAP Data Registers .....	6-7
6-3	Device ID Register .....	6-7
6-4	TAP Reset Actions.....	6-8
7-1	Debug Port Pinout Description and Requirements .....	7-4
A-6	LEN[1:0]# Signals Data Transfer Lengths .....	A-10





## 1.1 P6 FAMILY OF PROCESSORS OVERVIEW

The P6 family of processors is the generation of processors that succeeds the Pentium® line of Intel processors. This processor family implements Intel's dynamic execution microarchitecture, which incorporates a unique combination of multiple branch prediction, data flow analysis, and speculative execution. This enables P6 family processors to deliver higher performance than the Pentium family of processors, while maintaining binary compatibility with all previous Intel Architecture processors.

The first processor designed from the P6 family was the Pentium Pro processor which was followed by the Pentium II processor. As new products are designed, new technologies are utilized. For example, features were added to some P6 family processor products to aid in the design of energy efficient computer systems by offering multiple low-power states such as AutoHALT, Stop-Grant, Sleep and Deep Sleep, to conserve power during idle times.

The targeting of specific markets is another differentiator of products belonging to the P6 family including the Server and Workstation Market, Performance PC Market, Mobile Market and the Basic PC Market. All of these market segments demand specific features and performance. While all P6 family products have the benefits of Intel's dynamic execution microarchitecture, there are also product specific differentiators.

For example, the P6 family offers products with larger cache sizes and support for up to four processors to meet the higher performance demand of the server and workstation markets. Additionally, the Pentium II Xeon™ processor provides manageability requirements of the server and workstation environment by incorporating a System Management Bus (SMBus) interface. This interface can be used in conjunction with system hardware and software to provide more manageability options than any previous P6 family product. Memory is cacheable for 64 GB of addressable memory. This SMBus interface and larger L2 cache sizes, enables these products to provide higher performance and manageability for the server and workstation environment.

For high end desktop and business applications, the Pentium II processor can deliver the necessary computing power. Memory is cacheable for up to 4 GB of addressable memory space, allowing significant headroom for applications. It also incorporates Intel's MMX™ technology for enhanced media and communications performance.

The Intel P6 family also contains processors which are specifically designed and manufactured for the mobile market. These processors can operate under much more restrictive power and size constraints than the previously mentioned products, while still maintaining a high level of performance.

The Intel Celeron™ processor is designed for Basic PC desktops. It provides the same benefits of the P6 family architecture and adds the capabilities of Intel's MMX technology to bring a balanced level of performance and price to Basic PC consumers.

## 1.2 TERMINOLOGY

In this document, a '#' symbol after a signal name refers to an active low signal. This means that a signal is in the active state (based on the name of the signal) when driven to a low level. For example, when FLUSH# is low, a flush has been requested. When NMI is high, a non-maskable interrupt has occurred. In the case of signals where the name does not imply an active state but describes part of a binary sequence (such as address or data), the '#' symbol implies that the signal is inverted. For example, D[3:0] = 'HLHL' refers to a hex 'A', and D#[3:0] = 'LHLH' also refers to a hex 'A' (H= High logic level, L= Low logic level).

The term "system bus" refers to the interface between the processor, system core logic (a.k.a. the core logic components) and other bus agents. The system bus is a multiprocessing interface to processors, memory and I/O. The term "cache bus" refers to the interface between the processor and the L2 cache components. The cache bus does NOT connect to the system bus, and is not visible to other agents on the system bus.

When signal values are referenced in tables, a 0 indicates inactive and a 1 indicates active. 0 and 1 **do not** reflect voltage levels. A # after a signal name indicates active low. An entry of 1 for ADS# means that ADS# is active, with a low voltage level.

## 1.3 SPECIFIC PRODUCT REFERENCES

The reader of this document should also reference product datasheet specific details. Datasheets for Intel processors are located at <http://developer.intel.com>.

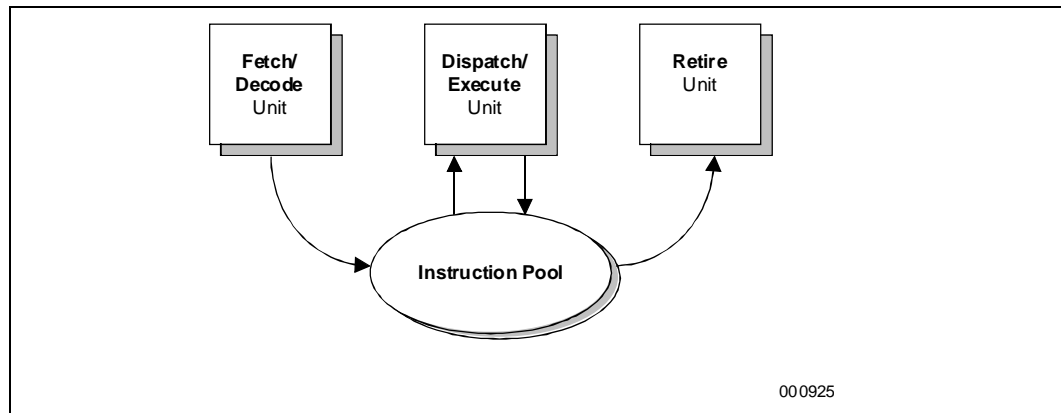
The task of providing all of the technical details of each product in one comprehensive manual is not the goal of this reference manual. The goal of this manual is to provide a reference of commonality between all P6 family products. It is the role each processor's datasheet to provide the specific differentiating details of each product. In the event that the datasheet and this reference manual contradict one another, please use the datasheet as the correct reference.

The P6 family of processor's may contain design defects known as errata. All characterized errata are available on-line at <http://developer.intel.com>.

The P6 family of processors use a dynamic execution micro-architecture. This three-way superscalar, pipelined micro-architecture features a decoupled, multi-stage superpipeline, which trades less work per pipestage for more stages. A P6 family processor, for example, has twelve stages with a pipestage time 33 percent less than the Pentium processor, which helps achieve a higher clock rate on any given manufacturing process.

The approach used in the P6 family micro-architecture removes the constraint of linear instruction sequencing between the traditional “fetch” and “execute” phases, and opens up a wide instruction window using an instruction pool. This approach allows the “execute” phase of the processor to have much more visibility into the program instruction stream so that better scheduling may take place. It requires the instruction “fetch/decode” phase of the processor to be much more efficient in terms of predicting program flow. Optimized scheduling requires the fundamental “execute” phase to be replaced by decoupled “dispatch/execute” and “retire” phases. This allows instructions to be started in any order but always be completed in the original program order. Processors in the P6 family may be thought of as three independent engines coupled with an instruction pool as shown in Figure 2-1.

**Figure 2-1. Three Engines Communicating Using an Instruction Pool**



## 2.1 FULL CORE UTILIZATION

The three independent-engine approach was taken to more fully utilize the processor core. Consider the pseudo code fragment in Figure 2-2:

**Figure 2-2. A Typical Pseudo Code Fragment**

```

    r1 <= mem [r0] /* Instruction 1 */
    r2 <= r1 + r2 /* Instruction 2 */
    r5 <= r5 + 1 /* Instruction 3 */
    r6 <= r6 - r3 /* Instruction 4 */
  
```

000922

The first instruction in this example is a load of `r1` that, at run time, causes a cache miss. A traditional processor core must wait for its bus interface unit to read this data from main memory and return it before moving on to instruction 2. This processor stalls while waiting for this data and is thus being under-utilized.

To avoid this memory latency problem, a P6 family processor “looks-ahead” into the instruction pool at subsequent instructions and does useful work rather than stalling. In the example in [Figure 2-2](#), instruction 2 is not executable since it depends upon the result of instruction 1; however both instructions 3 and 4 have no prior dependencies and are therefore executable. The processor executes instructions 3 and 4 out-of-order. The results of this out-of-order execution cannot be committed to permanent machine state (i.e., the programmer-visible registers) immediately since the original program order must be maintained. The results are instead stored back in the instruction pool awaiting in-order retirement. The core executes instructions depending upon their readiness to execute, and not on their original program order, and is therefore a true dataflow engine. This approach has the side effect that instructions are typically executed out-of-order.

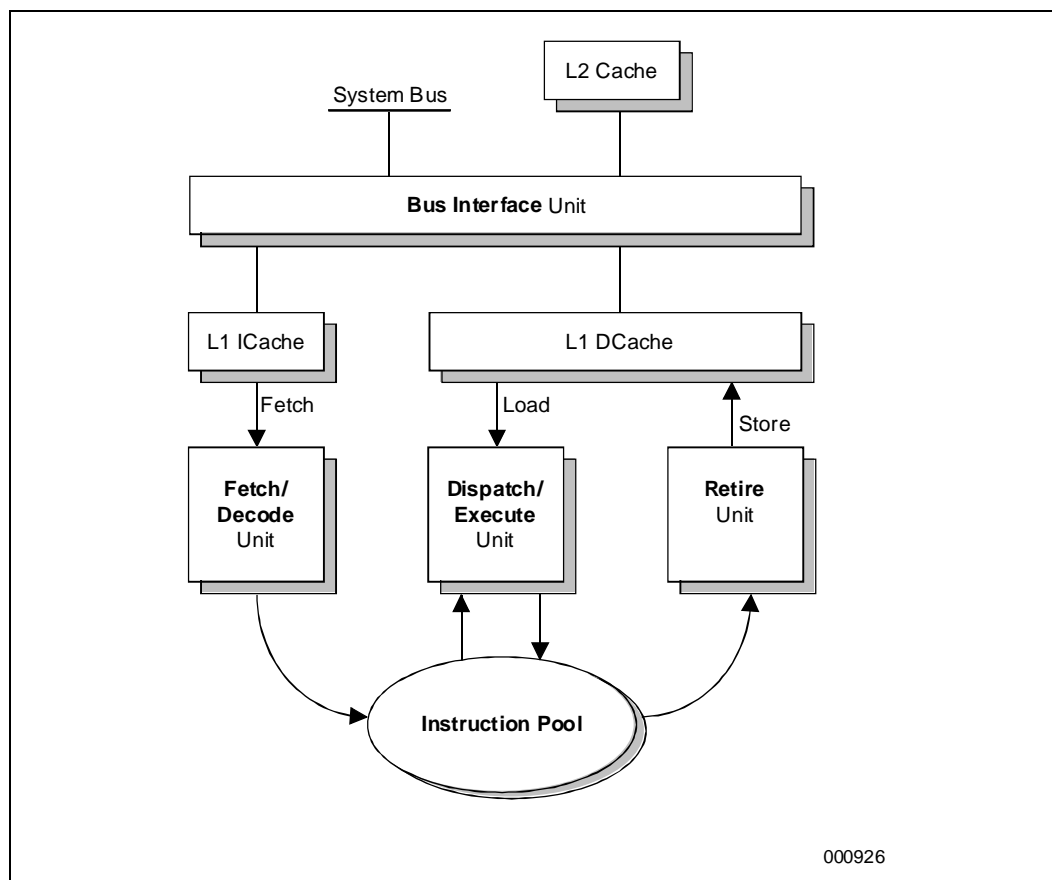
The cache miss on instruction 1 will take many internal clocks, so the core continues to look ahead for other instructions that could be speculatively executed, and is typically looking 20 to 30 instructions in front of the instruction pointer. Within this 20 to 30 instruction window there will be, on average, five branches that the fetch/decode unit must correctly predict if the dispatch/execute unit is to do useful work. The sparse register set of an Intel Architecture (IA) processor will create many false dependencies on registers so the dispatch/execute unit will rename the Intel Architecture registers into a larger register set to enable additional forward progress. The Retire Unit owns the programmer’s Intel Architecture register set and results are only committed to permanent machine state in these registers when it removes completed instructions from the pool in original program order.

Dynamic Execution technology can be summarized as optimally adjusting instruction execution by predicting program flow, having the ability to speculatively execute instructions in any order, and then analyzing the program’s dataflow graph to choose the best order to execute the instructions.

## 2.2 THE P6 FAMILY PROCESSOR PIPELINE

In order to get a closer look at how the P6 family micro-architecture implements Dynamic Execution, [Figure 2-3](#) shows a block diagram of the P6 family of processor products including cache and memory interfaces. The “Units” shown in [Figure 2-3](#) represent stages of the P6 family of processors pipeline.

Figure 2-3. The Three Core Engines Interface with Memory via Unified Caches

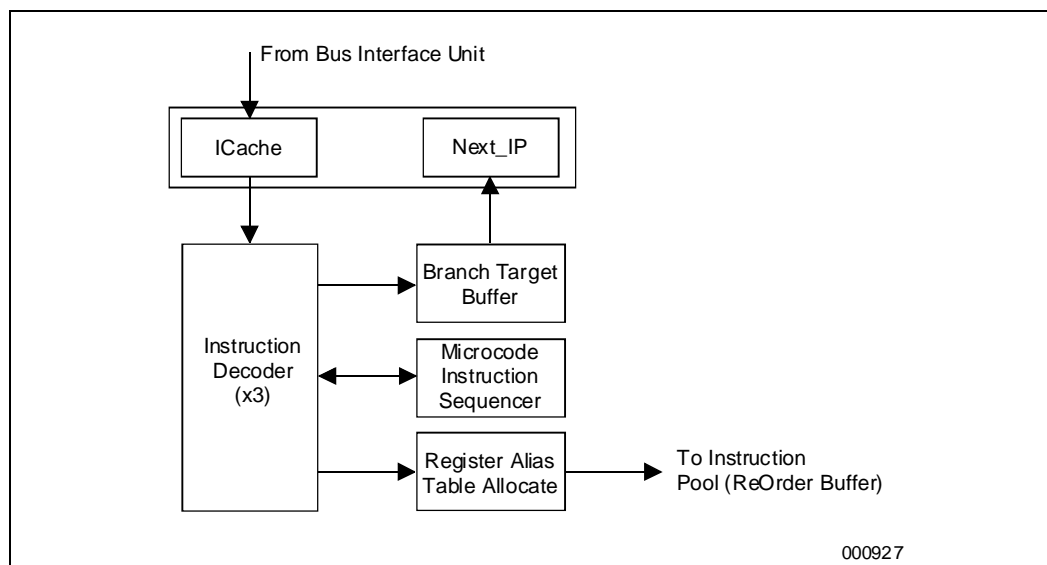


- The **FETCH/DECODE** unit: An in-order unit that takes as input the user program instruction stream from the instruction cache, and decodes them into a series of  $\mu$ operations ( $\mu$ ops) that represent the dataflow of that instruction stream. The pre-fetch is speculative.
- The **DISPATCH/EXECUTE** unit: An out-of-order unit that accepts the dataflow stream, schedules execution of the  $\mu$ ops subject to data dependencies and resource availability and temporarily stores the results of these speculative executions.
- The **RETIRE** unit: An in-order unit that knows how and when to commit (“retire”) the temporary, speculative results to permanent architectural state.
- The **BUS INTERFACE** unit: A partially ordered unit responsible for connecting the three internal units to the real world. The bus interface unit communicates directly with the L2 (second level) cache supporting up to four concurrent cache accesses. The bus interface unit also controls a transaction bus, with MESI snooping protocol, to system memory.

## 2.2.1 The Fetch/Decode Unit

Figure 2-4 shows a more detailed view of the Fetch/Decode unit.

Figure 2-4. Inside the Fetch/Decode Unit



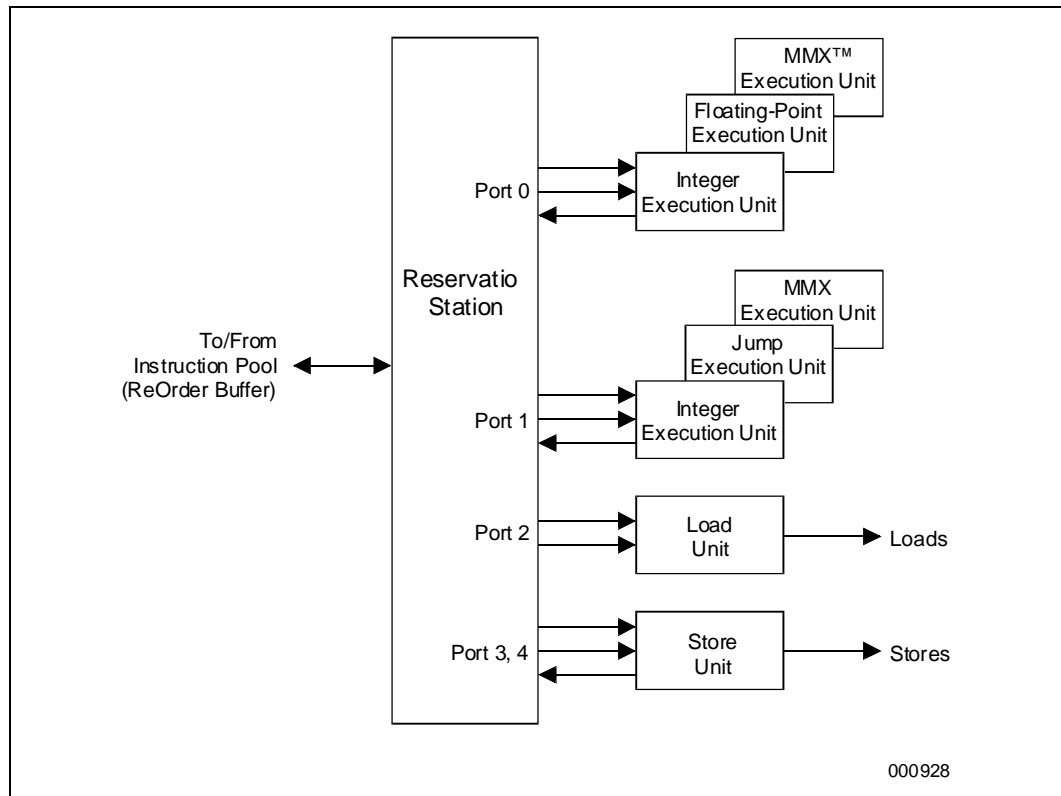
The L1 Instruction Cache is a local instruction cache. The Next\_IP unit provides the L1 Instruction Cache index, based on inputs from the Branch Target Buffer (BTB), trap/interrupt status, and branch-misprediction indications from the integer execution section.

The L1 Instruction Cache fetches the cache line corresponding to the index from the Next\_IP, and the next line, and presents 16 aligned bytes to the decoder. The prefetched bytes are rotated so that they are justified for the instruction decoders (ID). The beginning and end of the Intel Architecture instructions are marked.

Three parallel decoders accept this stream of marked bytes, and proceed to find and decode the Intel Architecture instructions contained therein. The decoder converts the Intel Architecture instructions into triadic  $\mu$ ops (two logical sources, one logical destination per  $\mu$ op). Most Intel Architecture instructions are converted directly into single  $\mu$ ops, some instructions are decoded into one-to-four  $\mu$ ops and the complex instructions require microcode (the box labeled Microcode Instruction Sequencer in Figure 2-4). This microcode is just a set of preprogrammed sequences of normal  $\mu$ ops. The  $\mu$ ops are queued, and sent to the Register Alias Table (RAT) unit, where the logical Intel Architecture-based register references are converted into references to physical registers in P6 family processors physical register references, and to the Allocator stage, which adds status information to the  $\mu$ ops and enters them into the instruction pool. The instruction pool is implemented as an array of Content Addressable Memory called the ReOrder Buffer (ROB).

## 2.2.2 The Dispatch/Execute Unit

The Dispatch unit selects  $\mu$ ops from the instruction pool depending upon their status. If the status indicates that a  $\mu$ op has all of its operands then the dispatch unit checks to see if the execution resource needed by that  $\mu$ op is also available. If both are true, the Reservation Station removes that  $\mu$ op and sends it to the resource where it is executed. The results of the  $\mu$ op are later returned to the pool. There are five ports on the Reservation Station, and the multiple resources are accessed as shown in Figure 2-5.

**Figure 2-5. Inside the Dispatch/Execute Unit**


The P6 family of processors can schedule at a peak rate of 5  $\mu$ ops per clock, one to each resource port, but a sustained rate of 3  $\mu$ ops per clock is more typical. The activity of this scheduling process is the out-of-order process;  $\mu$ ops are dispatched to the execution resources strictly according to dataflow constraints and resource availability, without regard to the original ordering of the program.

Note that the actual algorithm employed by this execution-scheduling process is vitally important to performance. If only one  $\mu$ op per resource becomes data-ready per clock cycle, then there is no choice. But if several are available, it must choose. The P6 family micro-architecture uses a pseudo FIFO scheduling algorithm favoring back-to-back  $\mu$ ops.

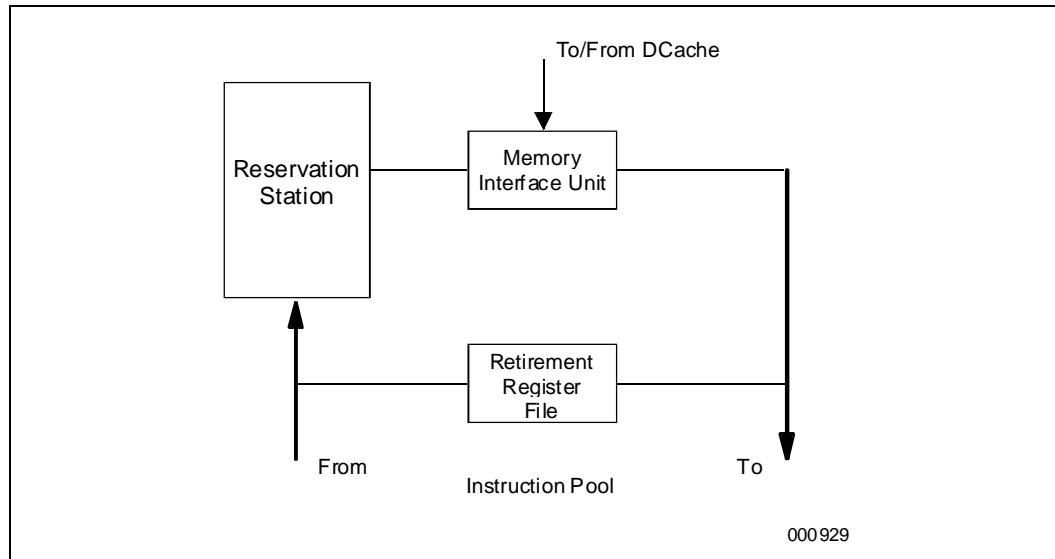
Note that many of the  $\mu$ ops are branches. The Branch Target Buffer (BTB) will correctly predict most of these branches but it can't correctly predict them all. Consider a BTB that is correctly predicting the backward branch at the bottom of a loop; eventually that loop is going to terminate, and when it does, that branch will be mispredicted. Branch  $\mu$ ops are tagged (in the in-order pipeline) with their fall-through address and the destination that was predicted for them. When the branch executes, what the branch actually did is compared against what the prediction hardware said it would do. If those coincide, then the branch eventually retires and the speculatively executed work between it and the next branch instruction in the instruction pool is good.

But if they do not coincide, then the Jump Execution Unit (JEU) changes the status of all of the  $\mu$ ops behind the branch to remove them from the instruction pool. In that case the proper branch destination is provided to the BTB which restarts the whole pipeline from the new target address.

## 2.2.3 The Retire Unit

Figure 2-6 shows a more detailed view of the Retire Unit.

**Figure 2-6. Inside the Retire Unit**



The Retire Unit is also checking the status of  $\mu$ ops in the instruction pool. It is looking for  $\mu$ ops that have executed and can be removed from the pool. Once removed, the original architectural target of the  $\mu$ ops is written as per the original Intel Architecture instruction. The Retire Unit must not only notice which  $\mu$ ops are complete, it must also re-impose the original program order on them. It must also do this in the face of interrupts, traps, faults, breakpoints and mispredictions.

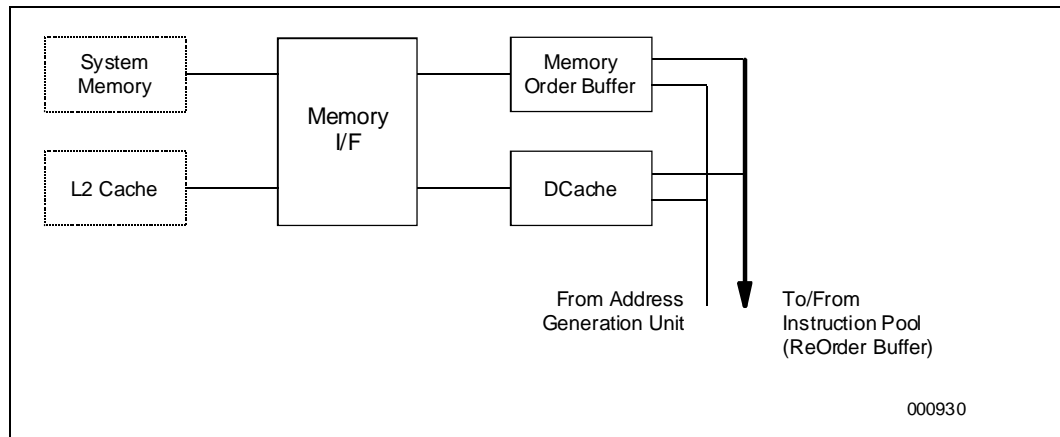
The Retire Unit must first read the instruction pool to find the potential candidates for retirement and determine which of these candidates are next in the original program order. Then it writes the results of this cycle's retirements to the Retirement Register File (RRF). The Retire Unit is capable of retiring 3  $\mu$ ops per clock.

## 2.2.4 The Bus Interface Unit

Figure 2-7 shows a more detailed view of the Bus Interface Unit.



**Figure 2-7. Inside the Bus Interface Unit**



There are two types of memory access: loads and stores. Loads only need to specify the memory address to be accessed, the width of the data being retrieved, and the destination register. Loads are encoded into a single  $\mu$ op.

Stores need to provide a memory address, a data width, and the data to be written. Stores therefore require two  $\mu$ ops, one to generate the address and one to generate the data. These  $\mu$ ops must later re-combine for the store to complete.

Stores are never performed speculatively since there is no transparent way to undo them. Stores are also never re-ordered among themselves. A store is dispatched only when both the address and the data are available and there are no older stores awaiting dispatch.

A study of the importance of memory access reordering concluded:

- Stores must be constrained from passing other stores, for only a small impact on performance.
- Stores can be constrained from passing loads, for an inconsequential performance loss.
- Constraining loads from passing other loads or stores has a significant impact on performance.

The Memory Order Buffer (MOB) allows loads to pass other loads and stores by acting like a reservation station and re-order buffer. It holds suspended loads and stores and re-dispatches them when a blocking condition (dependency or resource) disappears.

## 2.3 ARCHITECTURE SUMMARY

All products within the P6 family of processors are based upon this architectural summary. Dynamic Execution is the combination of improved branch prediction, speculative execution and data flow analysis that enable P6 family processors to deliver superior performance.