(intel®)

# Introduction to Intel® Architecture

**The Basics**

## Executive Summary

The term Intel® architecture encompasses a combination of microprocessors and supporting hardware that creates the building blocks for a variety of computing systems.

Although the architecture is straightforward and remarkably well-supported, the workings of these components may not be obvious to engineers, programmers, or product developers with no previous Intel architecture experience. This paper describes the basic operations and functions of the relevant components, using three example systems. Specifically, the paper will focus on the Intel® Core™ i7 processor (high-performance) and the Intel® Atom™ processor (low-power) implementations. In each case, the paper will walk the reader through the operation of the microprocessor's communication with memory and peripheral I/O devices, the interaction between different types of components, and related design criteria.

The various system components are described along with the services they provide. This paper will also define common terms used when describing Intel architecture designs and their operation. The final section highlights design aids, support, and collateral provided by Intel and its partner provides to help create successful products.

*The goal of this paper is to describe the basic operation and function of three classes of hardware platforms based on the Intel architecture and the components used.*

**Jim Turley**
Principal Analyst, Silicon Insider
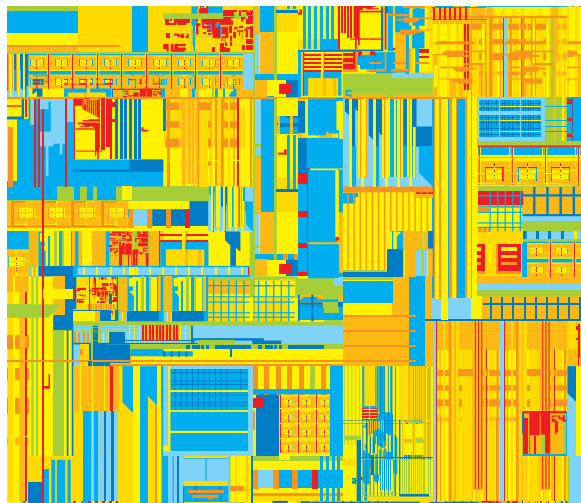
## Table of Contents

## What is Intel® Architecture?

Intel® is the world's oldest and most-established microprocessor company, producing the world's most popular microprocessor chips. Although perhaps best known for its PC processors, Intel devices are used in virtually every field of electronics, including automotive, industrial, automation, robotics, consumer electronics, image processing, networking, encryption, military, construction, medical, energy, and other industries.

Designers unfamiliar with the Intel architecture may have concerns about the architecture's fundamental concepts, its inner workings, or its complexity. The goal of this paper is to educate skilled developers with no previous exposure to the Intel architecture and to provide guidance regarding system components and concepts.

## Introduction to Intel® Architecture

Since the first tiny Intel 4004 microprocessor chip was made in 1971, Intel has produced an unbroken series of upgrades and improvements to the world's best known microprocessor family. From its early 8-bit beginnings, the Intel architecture now encompasses a range of 32-bit and 64-bit microprocessors that address a range of applications, performance requirements, power levels, and price points.

The cornerstone of Intel architecture's popularity is its compatibility. Each new generation of Intel architecture microprocessor is a superset of its predecessors, providing backward compatibility with older chips and older software, while also adding new or enhanced features. This compatibility allows engineers, programmers, and development teams to reuse the software and software-development tools from earlier projects, protecting their investment in time and talent. It also makes developing new Intel-based systems easier by leveraging developers'

experience. Software can be reused across generations of products, and product teams can protect their investment (in both hardware and software) in a cost-efficient manner. Although original work may be required to take advantage of the newest microprocessor features, the old software will still work as-is.

Intel architecture chips have obviously undergone many changes over the past 40+ years. A list of currently available devices is available here.[1] Early chips were given technical part numbers, such as 8086, 80386, or 80486. This led to the commonly used shorthand of "x86 architecture," in reference to the last two digits of each chip's part number. Beginning in 1993, the "x86" naming convention gave way to more memorable (and pronounceable) product names such as Intel® Pentium® processor, Intel® Celeron® processor, Intel® Core™ processor, and Intel® Atom™ processor.

Although every branch of the broad Intel architecture (or x86) family tree retains the same basic features and functionality as the earlier chips, and retains backward compatibility with them, each new generation also adds its own unique features to the mix. For example, Intel Pentium processor added multimedia extensions (called MMX™ technology) that accelerated audio and video processing. Extended temperature Intel Pentium processor with MMX technology is with more streaming-media capabilities known as Intel® Streaming SIMD Extensions (Intel® SSE) and Intel® Streaming SIMD Extensions 2 (Intel® SSE2). Floating-point units (FPUs) went from optional upgrade to standard feature of Intel architecture processors, and today encryption/decryption extensions, power-management features, and multilevel caches are now found on most Intel architecture processors. Data paths have widened from 8 bits to 32 bits, 64 bits, and even 128 bits and more. Operating frequencies have jumped from a few megahertz to 2 GHz (two billion cycles per second) and beyond.

[1] http://ark.intel.com

**Some abbreviations
used in the diagrams below:**

BIOS   Basic Input/Output
       System; a boot ROM

DDR3   Double Data-Rate v3; a popular
       DRAM interface standard

DMI    Direct Memory Interface; a
       video graphics standard

FIVR   Fully Integrated Voltage
       Regulator

LPC    Low Pin Count; a simple interface
       to slower I/O devices

PCH    Platform Controller Hub;
       a companion chip

PCI    Peripheral Component
       Interconnect; a popular
       expansion bus

PCIe*  PCI Express*; an upgraded PCI
       standard

SATA   Serial ATA; a popular
       disk-interface standard

SPI    Serial Peripheral Interface;
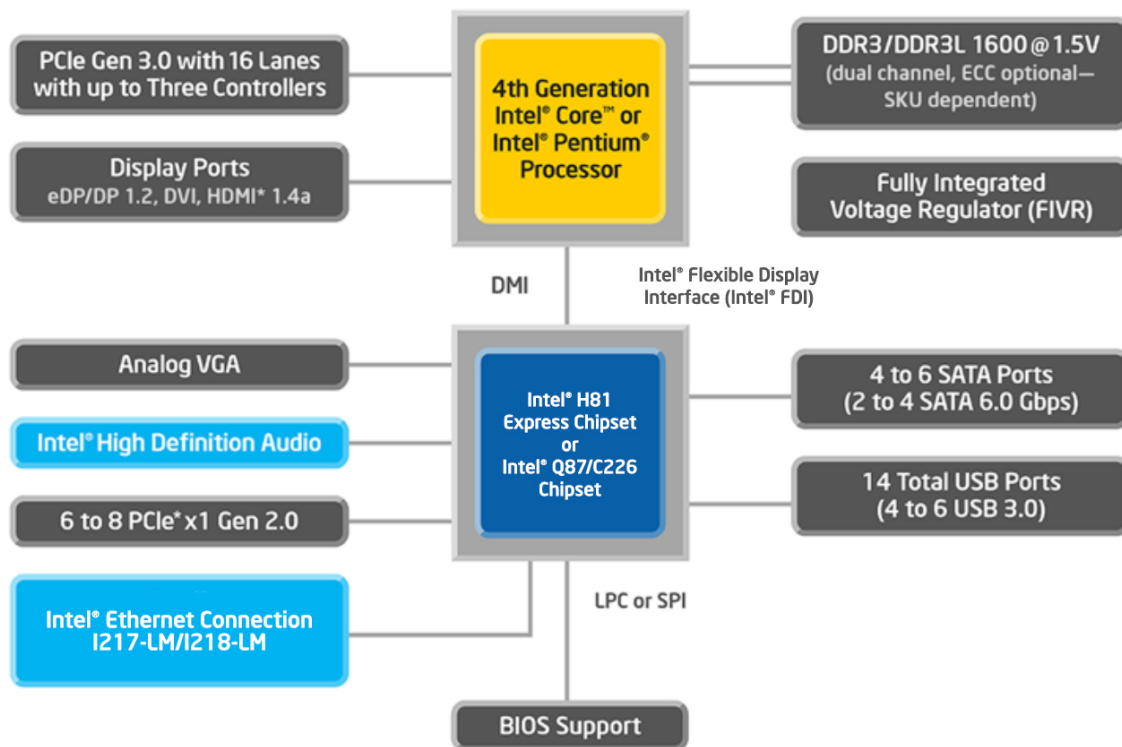       simple interface to
       slower devices



**Figure 1:** Typical system based on the Intel® Core™ i7 processor
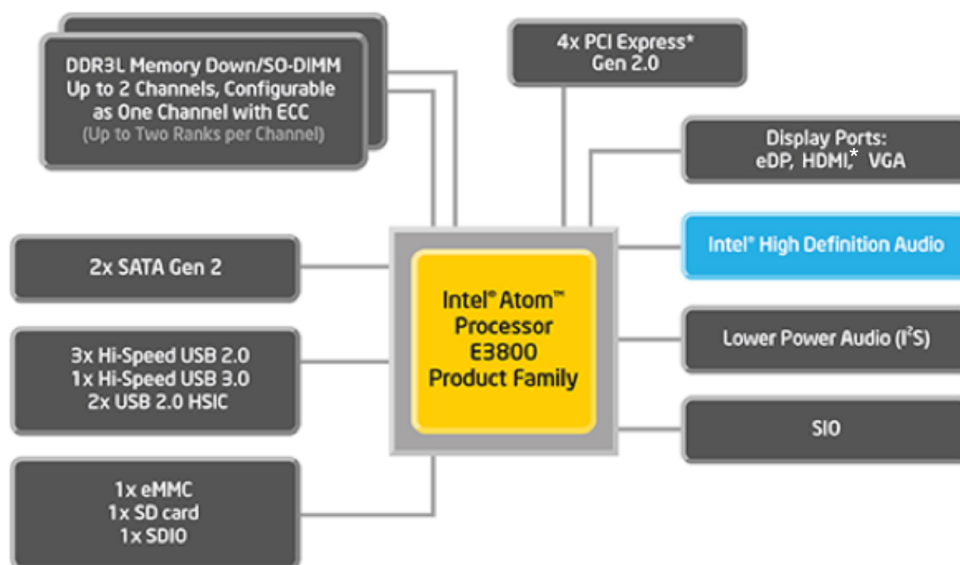
**Figure 2:** Typical system with Intel® Atom™ processor (SoC)

Similarly, many Intel® architecture chips now boast multicore performance, meaning that two or more Intel architecture processor cores, or "engines," operate within a single chip. Many also offer multithreading, a technique that is designed to improve performance by allowing a single Intel architecture core to perform multiple tasks. All the while, the power consumption and heat dissipation of these processor chips has been kept in check thanks to aggressive on-chip power-management hardware (some of which is adjustable through software) and industry-leading semiconductor-manufacturing technology. An in-depth description of the Intel architecture technical features can be found in the Intel Software Developer's Manual, available here.[2]

And while the processors have gotten ever more capable, there have been similar advances with the support logic, the development tools, and completely integrated hardware "platforms." Every Intel architecture processor is supported by one or more chip sets that provide needed system-level functions, and some Intel architecture processors include their own integrated functions such as memory controllers, graphics engines, or network interfaces. Sometimes the "chip set" is internal, and the processor becomes a standalone SoC – a system on a chip.

## Basics of an Intel® Architecture System

The hardware requirements for each customer application will be different, of course, but some basics apply to all. The processor chip itself is just the beginning. With few exceptions, the processor does not stand alone, but works in concert with compatible support chips. Different Intel architecture processors work with different support chips, and this paper outlines two representative examples.

The first example describes a typical hardware platform based on the high-performance Intel® Core™ i7 processor combined with similarly high-performance support logic (see Figure 1). The second example focuses on a small, low-cost Intel Atom processor–based system (see Figure 2).

Generally speaking, every Intel architecture hardware platform will include two major components: the microprocessor chip, and a companion chip known as the platform controller hub (PCH). In earlier times, Intel® processors were paired with two companion chips, often called the "north bridge" and

[2] http://download.intel.com/design/processor/manuals/253665.pdf

the "south bridge," or the MCH and the ICH. Nowadays, the functions of the north bridge are usually included in the processor itself, while the south bridge has been replaced by the much more capable PCH. In single-chip (SoC) configurations, there is no PCH; its functions are included in the processor itself.

Surrounding these two components (i.e., processor and PCH) will be other customer-supplied components including DRAM, a boot ROM, a power supply, and the peripheral interfaces appropriate for the system, such as a network or sensor connection. Most systems will also include some nonvolatile memory (e.g. flash or E2PROM), and perhaps some "glue logic" that is specific to the application.

As these diagrams show, one of the example systems is based on a two-chip set (processor and PCH), while the second example uses only a single chip (the processor) with integrated controllers. The former is designed for higher performance and more expansion capability, while latter is optimized for small size and low cost. These examples highlight just two of the many options available to designers of Intel architecture systems.

## Intel® Core™ i7 Processor–Based System

The example in Figure 1 illustrates a high-performance system based on the quad-core Intel Core i7 processor, the Intel® Q87 chipset (Intel® DH82Q87 PCH), two banks of external DDR3 DRAM, and several peripheral devices and interfaces. This configuration represents a high-end system with maximum performance with maximum capability and expandability.

## The Intel® Core™ i7 Processor

The heart of this system design is the Intel® Core™ i7-4770S processor, a high-end 64-bit implementation of the Intel architecture. The particular 4th generation, or "Haswell," Intel Core i7 processor shown in the diagram has several notable features, including:

- Four independent CPU cores
- Two-way multithreading per CPU core
- A built-in two-channel DDR3 DRAM controller
- Integrated L1, L2, and L3 caches
- Direct Media Interface (DMI) connection between the processor and the PCH

The Intel Core i7 processor achieves its high performance through its multiple CPU cores and its Simultaneous Multithreading (SMT) feature. Between the four cores (in this example) and the two-way multithreading per core, the Intel Core i7 processor appears to software as eight independent 64-bit CPUs.

Figure 3 shows a representation of the silicon die for Intel Core i7 processor, with its four independent CPU cores and other features highlighted. Figure 4 shows a conceptual diagram of the same processor, illustrating how the four CPUs each has its own L1 and L2 caches, and the shared L3 cache.

The processor's on-chip DRAM controller is responsible for cache coherence. If data at the address requested is not in one of the processor's caches, or if the data in external memory is newer than the cached copy, the memory controller is told to retrieve the data at the requested address. Data transfers between the processor and memory are always 64 bits wide, the full width of the L2 cache on the processor. If only a byte of data is requested, the full 64 bits are retrieved but the processor may use only 8 of those
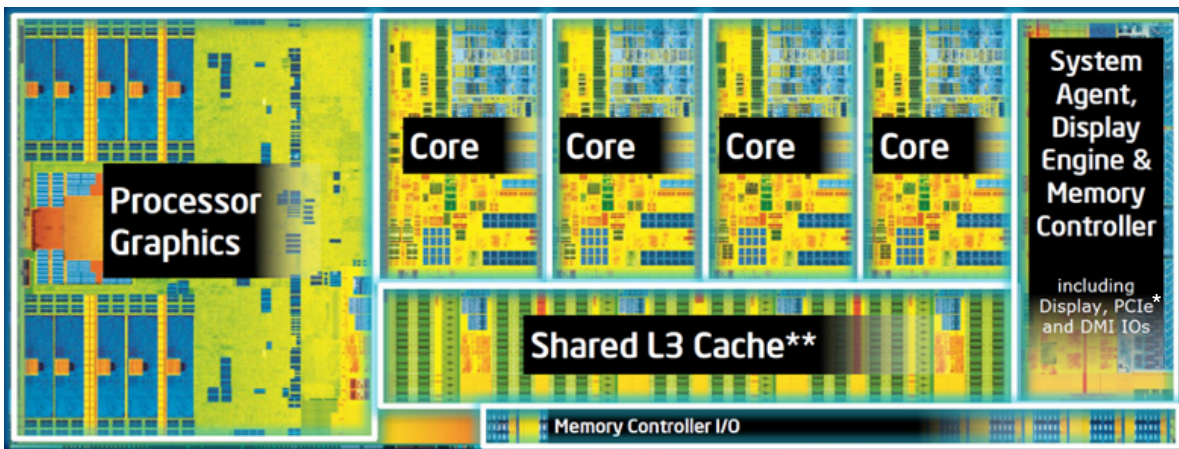


**Figure 3:** Intel® Core™ i7 processor internal die photograph

bits. The memory controller is configurable via the BIOS to support multiple speeds and/or sizes of memory. DRAM refresh is also handled by the memory controller, after it's initially configured. The specific type, size, and speed of memory supported varies by processor.

The PCI Express (PCIe) interface is the highest bandwidth I/O interface in the Intel architecture system. The number of PCIe lanes can vary depending on the processor used, but will usually be in multiples of four. A common width for PCIe is 16 lanes, as this is the maximum width for discrete PCIe graphics cards. The PCIe interface uses the same differential signaling as the DMI, although PCIe supports higher transfer rates. The original PCIe specification specifies data rates of 2.5 GT/sec (250 MB/sec) per lane, the same as DMI. The second generation of PCIe doubles the data rate to 5 GT/sec (500 MB/sec).

### Intel's Direct Media Interface (DMI)

DMI is the name given to the link between the Intel Core i7 processor and its companion chip, the Intel Q87 chipset (Intel® DH82Q87 PCH). The same DMI interface is used by several different processor and PCH implementations, including the Intel Core i7-4770S, Intel® Core™ i5-4570S, Intel® Core™ i3-4330 processors, Intel® Pentium® processor G3420, Intel Q87 chipset, Intel® H81 Express chipset, and Intel® C226 chipset. The DMI bus is a high-speed, point-to-point link between two chips, as shown in the diagram. It differs from Intel® QuickPath Interconnect (Intel® QPI) in that it supports only two chips: a processor and a PCH, whereas Intel QPI is used in multi-socket configurations with two or more processors. DMI supports transfers rates of 2 GB/second over each of two unidirectional lanes.

DMI is implemented as four serial links with dedicated transmit and receive pins. These serial links are referred to as "lanes"
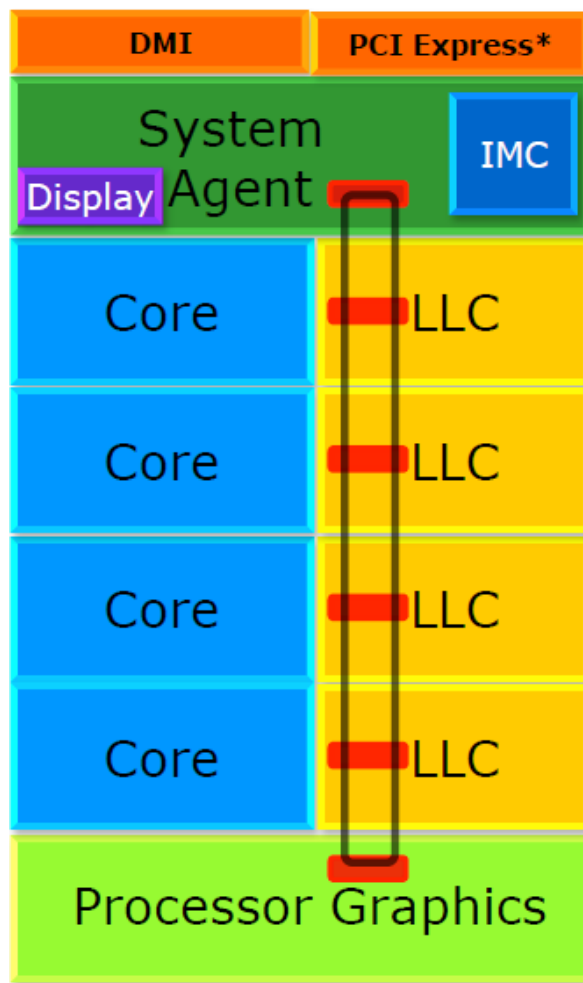


**Figure 4:** Intel® Core™ i7 processor internal block diagram.

and all use differential signaling. Thus, the DMI is 4 lanes × Transmit and Receive (2) × differential signaling (2) = 16 pins. DMI supports signaling of 2.5 billion transfers/sec. (Other chips implementing DMI 2.0 may achieve 5 GT/sec.)

DMI also integrates advanced priority-based servicing, allowing for concurrent traffic and isochronous transfer capabilities that improve performance over the earlier QPI interface. This ensures that the I/O subsystem (PCI Express, audio, SATA, USB, etc.) receives the bandwidth necessary for peak performance.

### Platform Controller Hub (PCH)

The Platform Controller Hub (PCH) chip is a highly integrated device intended to prove all the high-value features and interfaces required for a midrange to high-end system. Depending on the particular model of PCH chip, it may provide a controller for multiple banks of DRAM, a controller for multiple graphics displays (an accelerator for which may be present on the processor itself), several USB interfaces, SATA controller for disk drives and other storage media, an Ethernet LAN port, Intel® High Definition
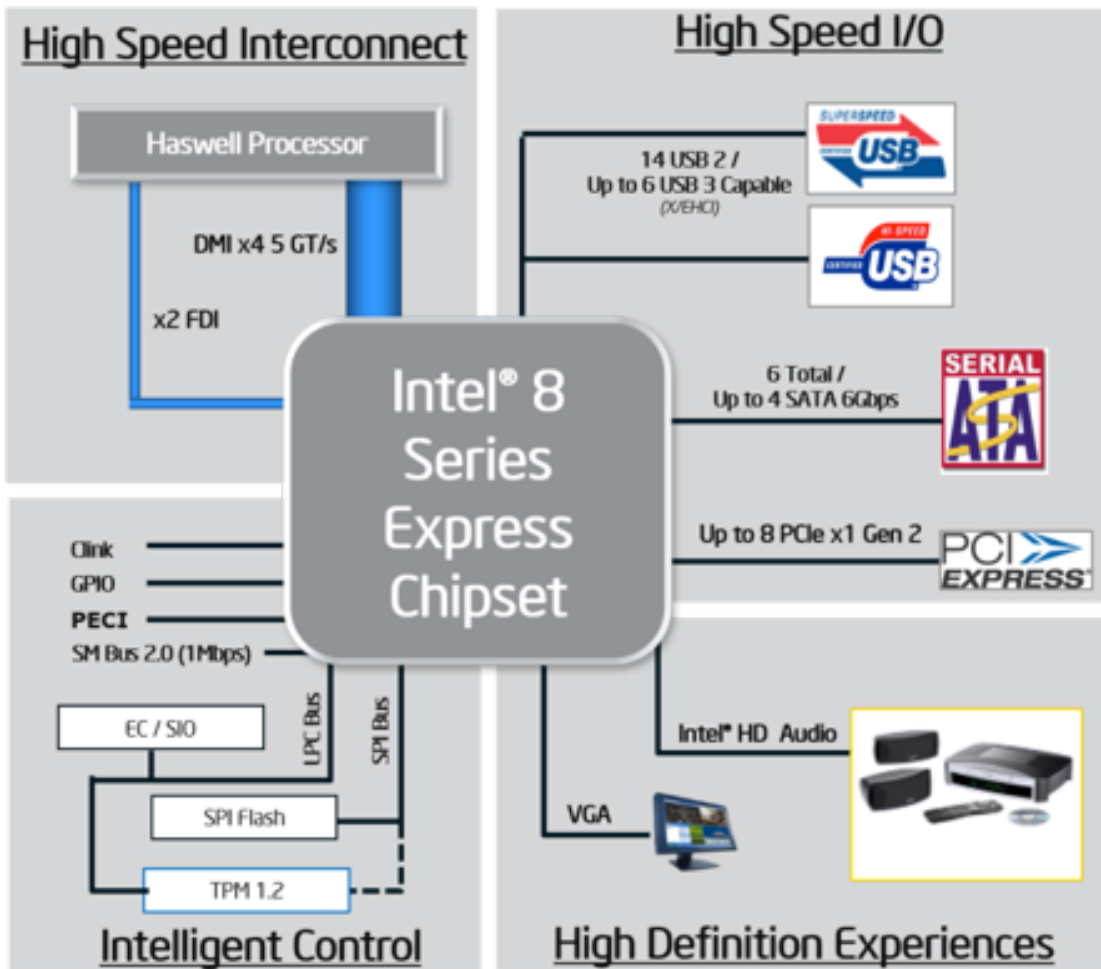
**Figure 5:** PCH internal block diagram

Audio (Intel® HD Audio), and more.

As the diagram in Figure 5 shows, the PCH connects directly to the processor's DMI interface, with no additional engineering required from the developer. Simply connect the appropriate pins and go.

The PCH itself is a complex and software-programmable component with several features to enhance system performance and reliability. Chief among them are:

- PCI Express Root Port Controllers. In the case of the Intel Q87 chipset (Intel® DH82Q87 PCH), the PCH supports up to eight PCI Express 2.0 root ports supporting up to 5 GT/second (five billion transfers per second).

- Up to six SATA ports, with integrated AHCI controller, and data rates of 6.0, 3.0, and 1.5 Gb/sec on all ports.

- An IEEE 802.3 Ethernet LAN MAC supporting 10/100/1000 Mbit/sec data rates and jumbo frames.

- Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d); Intel® Anti-Theft Technology (Intel® AT); Intel® Trusted Execution Technology (Intel® TXT)

- Clock controller, DMA controllers, event timers, real-time clock, JTAG boundary scan, low-pin-count interface for a trusted platform module (TPM), serial-peripheral interface (SPI), and much more.

## Modern I/O Interfaces

The PCH acts a bridge or controller for a variety of industry-standard interfaces allowing the system designer to choose from a wide range of peripherals, including:

- **PCI** interface operates at 33 MHz and allows for a number of external bus masters. The PCH acts as the central arbiter and root of the PCI bus.

- **PCI Express** root port controllers. The number of ports varies with the specific PCH component but is generally in the range of 1 to 4. Link widths of ×1 ×4 are support at speeds of 2.5 GT/sec.

- **Serial ATA (SATA)** controllers supporting both legacy operation using I/O space and the Advanced Host Controller Interface (AHCI) using memory-mapped I/O, as well as allowing advanced features such has hot-plug and native command queuing. SATA II supports data rates of 1.5 Gb/sec and 3 Gb/sec.

- **Integrated Drive Electronics (IDE**) controllers are also used to control hard disc drives and CD/DVD drives. They have been replaced in some platforms by the newer SATA interface since SATA offers better performance over a smaller interface.

- **Universal Serial Bus (USB)** supporting High Speed USB 2.0 (480 Mb/sec) operation as well as full-speed (12 Mb/sec) and low-speed signaling.

- **General Purpose I/O (GPIO)** pins for system customization. Many pins can also be configured to cause interrupts or wake events.

- **System Management Bus (SMBus 2.0)** The SMBus Host interface allows the processor to communicate with SMBus slaves. This interface is also compatible with most I2C devices. Slave functionality, including the Host Notify protocol is implemented.

- **Serial Peripheral Interface (SPI)** is used to interface to BIOS flash devices that contain boot firmware and initialization code. Up to two SPI flash devices operating at 33 MHz can be connected. Note that the flash devices connected to the LPC interface are quickly becoming obsolete and SPI is expected to be standard interface for BIOS flash in the future. The PCH is always a master on the SPI interface.

- **Low Pin Count Interface (LPC)** This interface replaces the ISA bus originally developed by IBM in the early 1980s, but uses only 7 signals plus a clock. It can be used to connect to a variety low speed devices that don't require the bandwidth of PCI or PCI Express. This interface is typically used to interface with Super I/O devices which contain many interfaces such as floppy driver controller, PS2 keyboard/mouse controls and serial ports.

- **JTAG Boundary Scan** allows testing of PCB board after assembly.

## Support Peripherals

The PCH integrates numerous support peripherals that replace many external components.

- **Real Time Clock (RTC):** The RTC is compatible with the popular MC146818A. It contains 256 bytes of RAM that can be maintained with a 3V battery. Of that space, 242 bytes are available for programmer use, while the remaining are dedicated to the clock function. The RTC can generate wake events up to 30 days in the future. An external 32.768-KHz crystal is required for operation.

- **High Precision Event Timers** These are high-resolution timers that can be used to generate periodic or one-shot interrupts. There are eight comparators, which share a common counter that is clocked from a 14.31818-MHz source.

- **Advanced Programmable Interrupt Controller (APIC)** is a more modern interrupt controller than the earlier 82C59 (see below). It supports multiprocessor/multicore interrupt management by allowing interrupts to be directed to a specific processor. The I/O APIC in the PCH can support up to 24 interrupt vectors and works in conjunction with I/O APICs in other devices to help eliminate the need for share interrupts among multiple devices.

## Compatibility Peripherals

The PCH contains peripherals compatible with those dating back to the earliest IBM PCs with an ISA bus. Modern systems replace the ISA bus with the Low Pin-Count (LPC) bus, but the peripherals that were once discrete components are now integrated into the PCH. One key strength of Intel architecture is its combination of backward compatibility and continuing innovation.

The PCH contains two 82C37 DMA controllers, two ISA-compatible 82C59 interrupt controllers, and three 82C54 programmable interval timer equivalents.

Intel Atom processor is a fairly recent addition to the Intel architecture family of 32-bit processors. It is intended for embedded systems where small size, modest power consumption, and low cost are important.

The 82C37 DMA controllers should not be confused with the DMA engines found in some earlier MCH (Memory Controller Hub) components. These DMA controllers are tied to the ISA/LPC bus and used mostly for transfers to/from slow devices such as floppy disk controllers.

The ISA-compatible 82C59 interrupt controllers have been largely supplanted by the Advanced Programmable Interrupt Controller (APIC, see above) since the letter offers support for more than 15 interrupt sources and supports multicore/multiprocessor systems. However, the 82C59 controllers are still used by some older operating systems that run only on single-processor (single CPU) systems.

### The BIOS

Like any computer, an Intel architecture system requires a boot ROM to bootstrap the processor and, optionally, load an operating system and configure components. In an Intel architecture system, this boot ROM has typically been known as the BIOS: the basic input/output system.

The BIOS controls the activity of the Intel architecture hardware until the operating system takes over. One job of the BIOS is to configure registers and components and set up the devices to the particulars of the system hardware into which the Intel architecture is designed. In a typical PC design, some of the hardware is dedicated by the design based on the motherboard design, but other hardware aspects vary based on what the end user may plug into the motherboard. As the BIOS executes, after the initial configuration is done, it will determine the type and amount of memory, then it goes through a discovery phase. Once all the devices and hardware are configured the BIOS will turn over control of the system to an operating system.

Many Intel architecture platforms come with the necessary boot firmware already installed. To create a boot ROM for a custom or updated system, Intel offers its

Intel® Boot Loader Development Kit (Intel® BLDK), which can be used to create a UEFI-compliant (Unified Extensible Firmware interface) boot loader compatible with many operating systems.

For a more minimally functional boot loader, Intel® Firmware Support Package (Intel® FSP) may be sufficient. Intel FSP includes royalty-free code that supports the most critical functions of Intel architecture processors and chipsets. The Intel FSP code may be included in a more fully featured UEFI-compliant boot loader, if desired.

## Intel® Atom™ Processor–Based System

Compared to the Intel Core i7 processor–based system described above, an Intel Atom processor-based system is almost trivially simple. This example is based on an Intel® Atom™ processor E3800 product family, which has all the necessary controllers and peripherals already integrated into it, doing away with the need for a companion chip. Refer back to Figure 2 for an overview of this system configuration.

### The Intel® Atom™ Processor

Intel Atom processor is a fairly recent addition to the Intel architecture family of 32-bit processors. It is intended for embedded systems where small size, modest power consumption, and low cost are important. For that reason, Intel Atom processor includes its own on-chip DRAM controller, PCI Express interface, optional display controller, USB controllers, real-time clock, timers, interfaces to system-management functions. The overall system thus enjoys small size with high integration.

With no DMI (nor any need for one), Intel Atom processor uses PCI Express as its primary means of expansion. PCIe has sufficient bandwidth for extensive expansion capability, and its standardized interface is compatible with many devices from multiple vendors.

Like the Intel Core i7 processor, Intel Atom processor E3800 product family has its own on-chip DRAM controller, able to handle two channels of DDR3 DRAM. One channel can optionally support ECC (error correction), further enhancing system reliability.

## Tools for New Designs

This section gives an overview of some reference documentation available to help with an Intel architecture design. Many of these documents are available at http://developer.intel.com/design/index.htm. Others may require a non-disclosure agreement (NDA) and will be made available through your sales representative.

Intel® Embedded Design Center site (http://edc.intel.com) also provides a wealth of information, documentation, and two-way communication with other Intel architecture developers.

**Platform Design Guides** are available under nondisclosure agreement and give detailed printed-circuit board design recommendation including PCB stack-up, impedance targets, material selection, and layout recommendations.

**Platform Reference Schematics** given an example system implantation.

**Thermal and Mechanical Design Guide** covers heat sink design recommendations and package/socket attachment to the printed circuit board.

**Datasheets (Processors)** contain electrical, thermal and package mechanical information for Intel processors.

**Datasheets (MCH/IOH/ICH/PCH)** contains a functional and register description of the device. Pinout and package mechanical information is also included. Electrical Electrical specifications may be in the datasheet or might require a non-disclosure agreement.

**Intel® 64 and IA-32 Architectures Software Developer's Manuals** These manuals describe the architecture and programming environment of the Intel® 64 and IA-32 processors.

**Specification Updates** contain lists of errata and the most recent changes to the other documents. Specification updates should always be consulted for the latest available information.

**Intel's Packaging Databook** is intended to serve only as a data reference guide to Intel package selection and availability. As the packaging landscape changes very rapidly, information can become outdated very quickly. Please refer to the product specifications on the products site for the most current detailed package information.

## Conclusion

The Intel architecture offers a complete and well-thought-out launching point for a variety of systems. It spans an array of features, performance, and power levels. The ability to reuse software and software tools across different generations and different product families is a great benefit to product teams and gives designers the ability to scale performance and features without new software rewrites. The basic building blocks of an Intel architecture system are highly integrated and make system design as straightforward as possible. Intel provides all of the resources necessary to design leading-edge products, so go to http://www.intel.com and start creating.

## References

Intel CPU History http://www.intel.com/pressroom/kits/quickreffam.htm

Current Intel® Product Information http://ark.intel.com

Intel Developers Centers http://developer.intel.com/design/index.htm

### Acronyms

| | |
|---|---|
| CPU | Central Processing Unit |
| BLDK | Boot Loader Development Kit |
| DMI | Direct Media Interface |
| FSB | Front Side Bus |
| FSP | Firmware Support Package |
| GLCI | Gigabit LAN Connection Interface |
| ICH | Input / Output Control Hub |
| IOH | Input / Output Hub |
| LAN | Local Area Network |
| LCI | LAN Connection Interface |
| LPC | Low Pin Count Interface |
| MCH | Memory Controller Hub |
| PCH | Platform Controller Hub |
| PCI | Peripheral Component Interconnect |
| PCIe | Peripheral Component Interconnect Express |
| QPI | QuickPath Interconnect |
| SPI | Serial Peripheral Interface |
| UEFI | Unified Extensible Firmware Interface |
| x86 | Common shorthand for Intel® architecture |