

CHAPTER 5



The Intel Mobile Processor

If GM had kept up with technology like the computer industry has, we would all be driving \$25 cars that got 1,000 MPG.

—Bill Gates

Intel is the original microprocessor designer: the first commercially available microprocessor was the Intel 4004 in 1971. Intel processors currently dominate the high-performance market, capturing almost all modern high-end servers. Intel's Pentium line of processors were ubiquitous in personal computing during the 1990s, and their Core i-series are the most popular central processing units for laptops and ultrabooks today. In addition, Intel has entered the mobile marketplace with mobile specific microprocessor products competitive with ARM Ltd., the market leader in the mobile space. Combined with the robustness and flexibility of the Android OS, the power and compatibility of the x86 series of processors is bringing a competitive new device family to the mobile market.

Intel's x86 Line

x86 forms the base architecture of an enormous family of Intel processors, ranging from the earliest Intel 8086 to the Pentium line, the i-series, recent virtualization hypervisor-equipped server processors, low-power Atom and Haswell microprocessors designed for mobile and embedded use, and the tiny Quark system-on-chip aimed at wearable computing. Originally, the 8086 architecture was designed for embedded systems. But Intel's early implementations of the 8086 architectures were wildly successful and led to a long line of revisions and upgrades adding power and rich features.

The x86 architecture is a Complex Instruction Set Computing (CISC) system, built with more complex instructions facilitating ease of use and simpler implementations. ARM, Intel's main competitor for mobile processors, is a Reduced Instruction Set Computing (RISC) system, without these features. For example, in a RISC system it might take three to four instructions to load a given value into memory, whereas in a CISC system there is one single instruction specifically written to do this. The x86 architecture is also register-to-memory based, meaning instructions can affect both registers and memory.

History

Intel is one of the oldest semiconductor manufacturing companies in the world, and is known for building innovative and functional technologies in the computer hardware and related industries. The company was started by Bob Noyce and Gordon Moore in 1968. Venture Capitalist Arthur Rock solidified this company with an initial investment of \$10,000 and a later contribution of \$2.5 million, resulting in his position as chairman.

Intel released their first two products in 1969: the 3101 Schottky bipolar random access memory, and the 1101, the world's first metal oxide semiconductor (MOS).

As mentioned, the first Intel processor was released in 1971, and it was called the 4004.

In 1978, Intel first released the 8086 series of processor architectures that would change the world. Only five years after that, Intel could officially call itself a billion dollar company. Intel is the largest semiconductor company in the world, and per the 2012 year-end reports, holds a market share of 15.7% and a revenue of 47.5 billion dollars. The original x86 architecture has split, diversified, added new specifications, and been reshaped into smaller form factors, continuing to be used in products around the world. The incorporation of Android on x86 is just another step forward for Intel.

Because the x86 architecture has been used in so many technologies, from servers to personal computers, mobile phones, laptops, and tablets, compiling a complete list of its devices would be prohibitively difficult. Its wide use has resulted in the creation of tools, applications, frameworks, and libraries specific to x86 platforms for developer use.

It all started in 1978 with the Intel 8086, originally built as an experimental 16-bit extension to the Intel 8080 8-bit microprocessor. The 8086 was the processor that drove the "IBM PC" and all of its clones. The term x86 was derived from the successors to the 8086, all of which ended in "86." In 1985, Intel continued the x86 architecture with the Intel 80386, the first 32-bit processor. It wasn't until 2005, with the release of the Pentium 4, that x86 64-bit processors hit the market.

Intel's latest home computing processor series based on x86 architectures is nicknamed the Intel Core i-series. This series supports 64-bit operations and is focused on performance and speed. All of the processors support hyperthreading and have multiple cores, which allow for concurrent processing. Running parallel to the personal computing Core i-series of microprocessors is the x86-based Atom series for mobile devices.

Strengths and Weaknesses

As industry leaders in the semiconductor market, Intel processors have distinct advantages. First and foremost, Intel processors have the highest performance of any other processors. This performance consists of both processor speed and the number of cores and virtual cores. The x86 architecture also grants developers access to the largest collection of software available. The last major advantage is the scalability of high-end systems that use Intel CPUs; the addition of processors gives direct performance improvements.

Table 5-1 highlights some of the differences of the Intel Atom processor family, which implement the x86 instruction set. The Intel Atom processor family consists of many different varieties for each platform type, including tablets, smartphones, netbooks, and other mobile consumer electronics, and Table 5-1 represents comparable high-end models for each.

Table 5-1. Intel Atom Processor Family Comparison

Name (Code Name)	System Type	Processor Speed	Cores (Threads)
Intel Atom Z2480 (Medfield)	Mobile	2GHz	1 (2)
Intel Atom D2700 (Cedarview)	Desktop	2.13GHz	2 (4)
Intel Atom C2750 (Avoton)	Server	2.40GHz	8 (8)

Note *These are the top-of-the-line Atom processors for each system type.*

There are situations where the x86 family is not the right choice of microprocessor. The Intel family is physically much larger than other brands of CPUs, taking more than an inch of space in the Core series. Up until the Intel Atom series, the power consumption of Intel's processors was too demanding for embedded devices; however, leading Atom processors compete with ARM for battery life. Lastly, Intel processors cost significant amounts and there are systems in which a 4-core 3GHz processor is overkill. In these situations using an ARM or some other lower performance CPU may be desirable.

Business Model

In its home computing efforts, Intel has continued to produce powerful and energy-efficient processors for laptops, ultrabooks, and desktop platforms. The closest competitor is the Semiconductor company Advanced Micro Devices, Inc. There was a point in 2006 when the desktop market was close to being split between AMD and Intel, but that is no longer the case. As of November 2012, Intel CPU's hold roughly 71% of the market, to AMD's 28%.

As laptops and tablets grew in popularity, Intel released the Atom series of processors. The atom processor balances heat and power with performance aimed specifically at items that will run on batteries for extended periods. The Atom series can be found in over 100 million devices, and is now expanding to the mobile market place.

Clash of the Mobile Titans: ARM versus Intel

ARM entered the microprocessor market in 1983, and emerged as a strong competitor in certain areas. Intel and their x86-based processors have managed to capture a majority of the desktop and home computing market. ARM on the other hand is the current leader in the mobile and embedded device market. This next section discusses in detail the properties of each company's processors, including its benefits and weaknesses.

ARM

If you look at per-unit sales, ARM is the current mobile-arena winner. With over 30 billion units in the current market, and 16 million sold per day, ARM is generating revenues of well over 900 million dollars a year. ARM's history, business strategies, and future plans are all relevant to the reason for ARM's success.

History

The ARM story begins at the British personal computer company Acorn. The original Acorn RISC machine was developed between the years of 1984 and 1985. In 1982, prior to ARM, the British Broadcasting Company (BBC) signed with Acorn to develop a home computer that would be later known as the BBC microcomputer. The *BBC micro* was wildly successful, and led to the growth of Acorn as a small company with a handful of employees, to a medium-sized business with hundreds of employees.

Around the end of the BBC micro era, Acorn began looking for the next processor to carry their new personal computer forward. Acorn tried a variety of 16- and 32-bit processors, including the 65C816 used in the Apple IIGS, but couldn't find one with the performance that Acorn needed. Acorn's solution to this problem was simply to develop a new processor, the ARM1.

Despite its incredible power and performance, the ARM1 wasn't largely used until the release of Archimedes, the first true ARM-based platform. Archimedes was a desktop computer, released in mid-1987, which primarily was used in schools and other educational environments. Even with the mild success and response from consumers, the ARM team pushed forward and created the ARMv3, focusing on increasing performance to compete with Intel and Motorola workstations.

In 1990, the Acorn RISC machine became the Advanced RISC Machine, and Advanced RISC Machines Ltd was created. With help of founding partners Apple, Acorn, and VLSI Technology, the company was founded for the sole purpose of continuing development of the ARM processor. From this foundation came the birth of the ARMv6, released to licensees in October of 2002. The ARMv6 architecture is used widely in embedded and mobile devices today, along with its more recent relatives, ARMv7 and ARMv8.

Strengths and Weaknesses

ARM processors have some very attractive qualities. To start, they are incredibly small. In fact the most modern ARM11 family of processors is under 2mm². Because of the small form factor, the heat generated from use is generally low enough to avoid any sort of heat sink or cooling system. Even as small as it is, ARM chips can contain many core system components inside of a single piece of silicon. These components include CPUs, GPUs, and DSPs. The last major advantage is the minimal amount of power consumed relative to competitors; some reports claim as much as a 66% savings. The less power used, the longer batteries last and the cheaper the electricity bill.

Table 5-2 showcases some of the more popular ARM processors currently used in the mobile market. This table is only a sample of the many options that ARM provides for mobile devices, but comparison with Table 5-1 demonstrates the significant differences from the Intel processor family. The comparable ARM mobile processors offer significantly less processor speeds, even at the high-end of the spectrum with the A15.

Table 5-2. ARM Cortex-A Series Comparison

Name (Architecture)	System Type	Processor Speed	Cores
A15 (ARMv7)	Mobile	1-2.5GHz	1-4 + clustering*
A7 (ARMv7)	Mobile	.8-1.5GHz	1-4
A5 (ARMv7)	Mobile	.3-.8GHz	1-4

Note *The A15 supports up to four clusters of processors with 1-4 cores each.

For all of the strengths that ARM chips have, there are plenty of weaknesses. To start with, ARM chips lack the serious performance required for any sort of heavy processing situations. The ARM processors are also inherently less scalable, especially in comparison to modern Intel CPUs. Software for ARM needs to be specifically created for the architecture; luckily, some of the more common tools and utilities already exist for ARM.

Business Model

An analysis of ARM's corporate decisions can help reveal their focus in the processor market. It is obvious the primary values that are held—RISC architecture, high performance, low power consumption, and low price point. These differentiations set ARM up perfectly for the mobile market, and are the key reasons that ARM processors are almost exclusively used in smartphones.

ARM processors however are not sold or manufactured by ARM Ltd. Instead, the processor architecture is licensed to interested parties. ARM Ltd. offers a variety of terms, and varying costs. With all its licensees, ARM provides in-depth documentation, a complete software development toolset, and the right-to-sell manufactured silicon with the licensed CPU.

This business model has done well for the company; in the second quarter of 2013, ARM reports categorized 51% of their income from royalties, with 39% from licensing. The report went on to detail the number of units for both royalties and licenses. The average cost of royalties per unit was roughly \$.07 cents, with over 2.6 billion units. On the other hand, there were 25 new licenses signed that quarter, averaging about \$1.84 million dollars per license.

Future

The latest processor that ARM has publically released is the ARM7, with various modified implementations. The ARM7 is used widely in the modern smartphone market. There have been rumors in corporate waters that ARM will be pursuing additional directions with its processors.

With the release of Windows 8 for x86, Microsoft has created a version of Windows called Windows RT for ARM processors. Windows RT was written almost entirely from scratch, and has managed to eliminate many, but not all, of the bottlenecks of modern backward-compatible Windows versions. Tests have concluded that RT applications are running as much as 20% faster than the same applications on competing Intel chips.

Experts have also projected ARM's entry into the server and datacenter markets. With support from ARM-based Linux server operating systems, this is becoming more of a reality. The ability to run performance systems on ARM conceivably means lower power costs. This has yet to be seen in the current performance of high-end ARM systems versus the performance of high-end Intel systems.

Intel's Atom Line of Microprocessors

Atom processors are featured in devices that are used on the go. Typical devices include small laptops, netbooks, tablet computers, televisions, and new smartphones. The Atom balances performance and power usage to enable much longer battery life for the device.

With over 100 million Atom CPUs shipped, the outreach of the Atom is apparent. As with all Intel processors, the Atom is a member of the Intel Architecture (IA) family. The distinct and cross portability of the IA family allows for quick and effortless transitions between processors.

Intel Atom Evolution

The Intel Atom is the successor of the Intel A100 and A110, low-power processors primarily used in notebook computers. The A100 and A110 were code-named Stealey and originally built at the size of 90nm. Tables 5-3 and 5-4 highlight some of the iterations of Atom, for tablets and for smartphones, from the processor family's infancy in April of 2008 through its modern releases.

Table 5-3. Intel Atom Smartphone Processors

Code Name	GPU* Speed	Processor Speed	Cores (Size)	Release
Silverthorne	200MHz	800MHz – 2.13GHz	1 (45nm)	April 2008
Lincroft	400MHz	800MHz – 1.9GHz	1 (45nm)	May 2010
Penwell	400MHz	1.2GHz – 2.0GHz	1-2 (32nm)	2012

Note *GPU stands for Graphics Processing Unit, the hardware responsible for producing graphics for a given device.

Table 5-4. Intel Atom Tablet Processors

Code Name	GPU* Speed	Processor Speed	Cores (Size)	Release
Lincroft	400MHz	1.2GHz – 1.5GHz	1 (45nm)	May 2010
Cloverview	300MHz – 533MHz	1.2GHz – 2.0GHz	2 (32nm)	2012
Bay Trail-T	311MHz – 688MHz	1.33GHz – 2.41	4 (22nm)	2013*

Note *Bay-Trail T devices are expected in the first half of 2014.

At first glance, the processors listed in Table 5-3 only seem to be getting marginally better, but in order to truly understand what's going on, you need to take into account all of the variables. The Penwell is the forerunner for processors that Intel produces for smartphones today, with a size of only 32nm, multi-core support, and top-of-the-line operating frequency with embedded GPU support. It is the obvious choice from Intel for modern device manufacturers.

In comparison to the existing processors in Table 5-3, the tablet processors listed in Table 5-4 are much more capable. These tablet processors support even more cores, and have faster GPU speeds, which help accommodate much larger and often high-resolution display components.

Intel Atom Security

With technology in this modern day and age, security is always a concern. The Intel Atom processor offers support for many security features. These include Secure Boot, Intel Platform Trust Technology, hardware-enhanced encryption, and operating system-level key storage. Secure Boot is part of the current Unified Extensible Firmware Interface (UEFI) specifications, and is best described in Intel's own words:

When enabled and fully configured, Secure Boot helps a computer resist attacks and infection from malware. Secure Boot detects tampering with boot loaders, key operating system files, and unauthorized option ROMs, by validating their digital signatures. Detections are blocked from running before they can attack or infect the system.

The Intel Platform Trust Technology, or PTT for short, is a virtual smartcard reader on tablets that allows for certificate-based authentication through the CPU.

Intel Atom Features

The Intel Atom processor supports a significant amount of the features that exist in other Intel processors. Energy efficiency is a new idea in the Intel world, and the Atom brings this to the forefront. The Atom processor can be custom-tailored to bring the correct balance of incredibly low power with varying performance scalability options. Performance-wise, the Atom supports both Intel Hyper-Threading and Intel Burst Technology to help deal with required performance and power efficiency. The last major feature that Intel promotes with the Atom is the concept of mobility, supporting NFC, advanced camera imaging, 3G, and 4G LTE.

Android and the Atom

The Atom processor is the current x86 processor of choice for Android platforms. The Atom Android team brings a wardrobe packed with top-of-the-line features. This includes 3D graphics with full 1080p HD support for multiple formats, screen sharing and device pairing, optimized web page rendering, and simple cross computability. Android SDK applications are supported out-of-the-box on Atom Android platforms. Android NDK applications

require only a recompile, in most cases, to be fully supported. More information about the compatibility and conversion process can be found the following section titled **Application Compatibility** and in **Chapter 7: Creating and Porting NDK-Based Android Applications**.

Inside the Medfield System-on-Chip

Intel's Medfield platform is intended for smartphones and tablets running the Android operating system. One model of Medfield, the Intel Atom Z2610 System-on-Chip (SOC), will be discussed in greater detail shortly (in Figure 5-1). As stated earlier, Intel has recently started producing standalone mobile processors, including one codenamed Penwell. Although the Penwell processor contains some of the same segments as the Medfield SoC, namely Saltwell-family microprocessor architectures, Penwell is a standalone processor primarily targeted at smartphones as opposed to Medfield's multiple-part and higher-performance system targeting both smartphones and tablets.

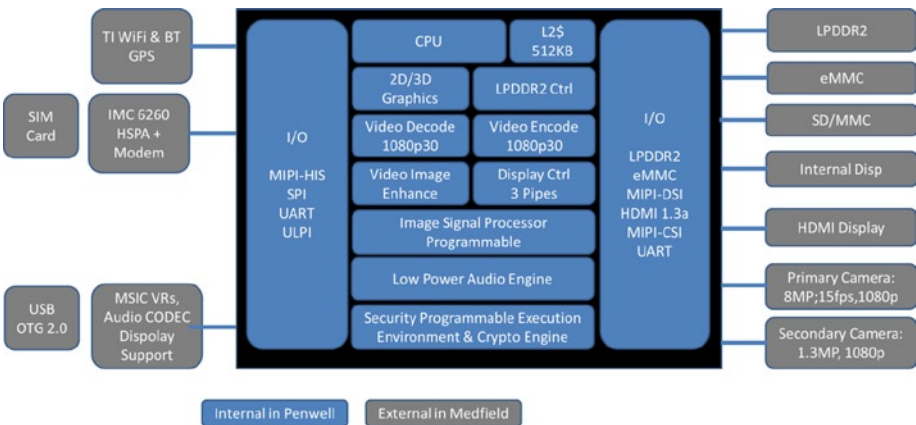


Figure 5-1. Medfield Block Diagram

This Medfield model, the Z2610, is physically divided into two complexes, the North Complex and the South Complex. The North Complex consists of a Saltwell-family single-core processor, a 32-bit dual channel LPDDR2 memory controller, a 3D graphics core, video decode and encode engines, a 2D display controller that is capable of supporting up to three displays, and an image processor for camera input. The South Complex consists of all the necessary I/O interfaces to complete a smartphone design, such as a security engine, a storage controller supporting SD/eMMC storage cards, a USB OTG controller, a 3G modem, Complimentary Wireless Solution (CWS) interfaces, SPI, and UART. See Figure 5-1.

Zooming In on the Saltwell CPU Architecture

The Saltwell CPU architecture is fairly simple. The idea of the design is to create a processor with a balance between optimized performance and efficient power consumption. The processor uses in-order architecture, which is different from most of the other processors in the market, which instead use out of order execution. The processor has a 64-KB L1 cache and a 512-KB L2 cache. This processor supports Intel Burst Performance technology, which lets the processor dynamically increase the CPU speed. There are three frequency modes in Saltwell: Low Frequency Mode (LFM) runs at 600MHz, High Frequency Mode (HFM) runs at 900MHz, and Burst Frequency Mode (BFM) runs at 1.6GHz. Among the power optimization features, Saltwell has an ultra-low power smart L2 cache that keeps data while the CPU is in C6 states, in order to lower the latency during the resumption of C states. In addition, Saltwell has separate power planes and clock inputs for the core and the rest of the SoC, which makes power and clock gating easily configurable through Intel Smart Idle Technology (Intel SIT). This technology enables the CPU to be switched off completely while the SoC is still in the ON state (S0 state).

Architecture Differences between Intel's Saltwell and ARM's Cortex A15

As listed in the book, *Break Away with Intel Atom Processors: A Guide to Architecture Migration*,¹ the Intel Atom architecture is very different from the ARM architecture in every way. Table 5-5 shows a list of high-level differences between Saltwell and ARM Cortex architecture.

Table 5-5. High-Level Differences Between Saltwell and ARM (Cortex A15)

Feature	Saltwell	ARM Cortex
Technology	32nm	28nm
Architecture	In-order	Out-of-order
Integer pipelines	16	15
L1 cache	64KB	Configurable up 64KB
L2 cache	512KB	Max 4MB
Instruction set	IA32, Intel Streaming SIMD Extensions, Intel Supplemental Streaming SIMD Extensions 3	ARM, Thumb
Multi core/thread support	Single core with Intel Hyper-Threading Technology	Multi-core
Security technology	Intel Smart & Secure Technology (Intel S&ST)	TrustZone* Technology

¹Matassa, Lori and Max Domeika, *Break Away with Intel® Atom™ Processors: A Guide to Architecture Migration*. Intel Press, 2010.

Architecture

As mentioned, Saltwell has an architecture similar to other processors in the Intel Atom series. It uses an in-order execution design. With an in-order processor, all the instructions are executed according to the order they are fetched, whereas out-of-order processors are capable of executing multiple instructions simultaneously and reordering them later in the pipeline. ARM processors use out-of-order architecture, which has the advantage of executing instructions with minimal latency. However, this increases the complexity of the core design. The elimination of the reordering logic is one of the power reduction initiatives of the Intel Atom processor.

Integer Pipelines

There are six phases in Intel Atom pipelines; the details are listed in Table 5-6.

Table 5-6. Intel Atom Instruction Phases and Pipeline Stages

Phase	Pipeline Stages
Instruction fetch	3
Instruction decode	3
Instruction issue	3
Data access	3
Execute	1
Write back	3

This instruction architecture results in a total of 16 integer pipelines in the Intel Atom processor, and three extra stages are required to execute floating point instructions. The latest ARM processor has 15 integer pipelines. The lengthy pipeline in the ARM processor trades energy over performance. Saltwell can decode up to two instructions per clock cycle while the latest ARM processor is a triple issue superscalar architecture.

Instruction Sets

ARM instruction sets are always 32-bit and aligned on a four-byte boundary, whereas IA32 instruction sets vary in size and do not require any alignment. Another difference between ARM instructions and IA32 instructions is how the instruction is executed. For ARM, all the instructions are conditionally executed to reduce branch overhead and mis-prediction during branching. There are condition flags that each instruction needs to fulfill in order to take effect, otherwise the instruction will act as NOP and get discarded. There are conditional instructions as well in Intel architecture; these are called conditional MOV instructions. Other instructions in IA32 are not conditionally executed.

Multi-Core/Thread Support

As mentioned previously, Saltwell supports Intel Hyper-Threading Technology (Intel HT Technology), where tasks are completed by using shared resources. The details of the technology are discussed further in the next section. ARM multi-core architecture has unique resources to perform its tasks on each core. The coherency of the cores is handled by AMBA 4 AXI, a compatible slave interface that is directly interfaced to the core.

Security Technology

There is a security subsystem in Medfield called Intel Smart & Secure Technology (Intel S&ST). It is a complete hardware and software security architecture. This subsystem is compliant with industry standards, supporting AES, DES, 3DES, RSA, ECC, SHA-1/2, and DRM. It also supports 1,000 bits of OTP and enables SecureBoot. The implementation in the ARM processor for a security system is different. There is no separate controller for the security subsystem as Intel implemented. The ARM processor uses TrustZone technology, where resources in the system such as processor and memory are divided into two worlds: the *normal world* and the *secure world*. There are three motivations for this Trust Zone architecture:

- To provide a security framework that allows designers to customize the functions needed depending on the use cases.
- To save silicon area and power where there will be no need to have a dedicated processor for secured tasks.
- To prevent intrusion during debug to security-sensitive tasks in the *secure world* or non-security-sensitive tasks in the *normal world*, by providing a single debug component.

Intel Hyper-Threading Technology

Intel Hyper-Threading Technology (Intel HT Technology) enables software to have a view of multiple logical processors in a physical processor package. The Saltwell CPU architecture uses Intel Hyper-Threading Technology as a boost to its performance. Having a second thread in a single in-order architecture processor enables Saltwell to execute multiple instructions within a clock cycle sharing the execution resources among the two threads, giving a 50% performance improvement compared to a single thread processor, as shown in Figure 5-2.




Figure 5-2. Benefits of the Intel Hyper-Threading Technology

In Intel HT technology, the processor has duplicates of the architecture state that consists of general purpose registers, control registers, the advanced programmable interrupt controller (APIC) registers, and some machine state registers. The duplication of architecture states is the reason software can view a single core processor as two logical processors. Caches, execution units, branch predictors, control logic, and buses are shared between the two threads. This created a concern where there might be resource contention and workload imbalance between the threads. However, most of the current development kits such as Dalvik and JavaScript already have the capability to support multi-threaded environments, giving developers an easy way to generate applications that utilize the advantage of Intel HT technology. Applications developers on Android can also utilize the Intel VTune performance tool to analyze the workload and perform resource tuning on their applications.

Application Compatibility: Native Development Kit and Binary Translator

Android has been ported to x86 and all further releases will be available in both the x86 and ARM architectures. It is not an issue to run the OS on an Intel Atom platform. However, in some cases, existing Android applications may need to be recompiled with or without source code modification.

It is believed that roughly 75–80% (commonly cited numbers) of Android applications in the Google Play Store run on top of the Dalvik VM and use the Android Framework (see Figure 5-3). The vast majority of Dalvik VM applications written in the Java language using the Android Software Development Kit (SDK) are processor agnostic. They run

“as-is” on Intel Atom platforms transparently without requiring porting efforts. For the subset of applications that include C and C++ code, developers need to recompile their code using the latest Android Native Development Kit (NDK).

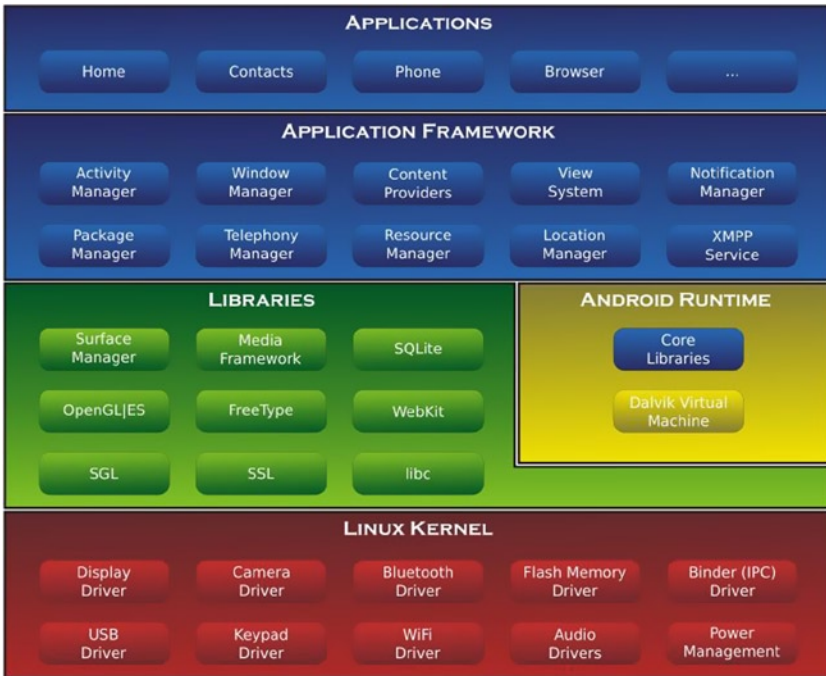


Figure 5-3. *Android Framework*

In most cases, NDK application developers simply need to recompile the project, which supports x86 (x86), ARMv5 (armeabi), and ARMv7 (armeabi-v7a). Compiling for x86 (by GCC with compiler flags `-march=i686 -msse3 -mstackrealign -mfpmath=sse`) will generate code that is tailored exactly to the Intel Atom CPU feature sets. Only applications that use an ARM vendor’s specific feature will require source code to be rewritten and then recompiled.

The resulting APK application package may comprise three versions of machine code for x86, ARMv5, and ARMv7. Upon installation, only the appropriate version of code is unpacked and installed onto the target platform.

The rest of the applications are either Dalvik VM applications that use Java Native Interface (JNI) libraries built for ARM, or Native Development Kit (NDK) applications that haven’t been compiled for x86. These applications cannot run without alteration on the Intel Atom platform due to the calls to native libraries (especially ARM-specific native libraries).

Intel and Google have worked together to ensure native application execution on Intel Atom platform “as-is” without porting efforts. Intel provides Binary Translation (BT) that translates ARM code to x86 code on-the-fly during execution, as shown in Figure 5-4.

This translation mitigates the inconvenience of JNI libraries and NDK applications that have not yet been ported to x86. It allows the device to expose itself as supporting two applications binary interfaces (ABIs): x86 and ARMv5. This could be observed from the **build.prop**, as shown in Figure 5-4.

```
ro.product.cpu.abi=x86
...
ro.product.cpu.abi2=armeabi
```




Figure 5-4. Binary Translation

If the NDK applications haven't been rebuilt for the x86 platform, the binary translator will locally translate the armeabi version into x86. The same applies for Dalvik VM applications that request ARM-based JNI libraries. The translation process is optimized and completely transparent to the end users.

The combination of all of these efforts should result in approximately 90% of applications in the Google Play working right away. The other 10% of applications may take some additional configuration and setup to be fully functional. In **Chapter 7: Creating and Porting NDK-Based Android Applications**, we will cover some more specifics about native code development with x86, and this should offer some general suggestions to help with any applications that fit into this bucket.

Overview

This chapter covered a brief history of both Intel and ARM from a company and a processor standpoint. You looked at some specific processors from Intel and Arm and read about what each brings to the table. After a brief overview of each company, we spent some time talking about the Intel Atom processor and which features and specifics exist with the modern versions. Finally we jumped into a technical discussion about Intel Atom's Medfield architecture, which is being featured in the newest x86 phones and tablets. We have discussed how the Integer pipeline flows, which security systems are in place, and even how the Intel Hyper-Threading optimizes performance. Binary translation was discussed at length, with an explanation of how NDK-based applications have to be prepared to later be ported to the Intel platform.