Anatomy of contemporary GSM cellphone hardware

Harald Welte <laforge@gnumonks.org>

August 8, 2010

Abstract

Billions of cell phones are being used every day by an almost equally large number of users. The majority of those phones are built according to the GSM protocol specifications and interoperate with GSM networks of hundreds of carriers.

Despite being an openly published international standard, the architecture of GSM networks and its associated protocols are only known to a relatively small group of R&D engineers.

Even less public information exists about the hardware architecture of the actual mobile phones themselves, at least as far as it relates to that part of the phone implementing the GSM protocols and facilitating access to the public GSM networks.

This paper is an attempt to serve as an introductory text into the hardware architecture of contemporary GSM mobile phone hardware anatomy. It is intended to widen the technical background on mobile phones within the IT community.

Contents

1 Foreword

This document is the result of my personal research on mobile phone hardware and system-level software throughout the last six years.

Despite my past work for Openmoko Inc., I have never been professionally involved in any aspect of the actual GSM related hardware of any phone. Nevertheless I have the feeling that in the wider information technology industry, I am part of a very, very small group of people who actually understand mobile phones down to the lowest layer.

I hope it is useful for any systems level engineer with an interest in understanding more about how mobile phone hardware actually works.

There are no guarantees for accuracy or correctness of any part of the document. I happily receive your feedback and corrections.

2 Is your phone smart or does it have features?

Initially, for the first couple of years, GSM cell phones were actual phones with very little additional functionality. They provided everything that was required for voice calls, as well as SIM phone book editing features. The only additional non-features were simple improvements like the ability to use them as an alarm clock.

In the mid-1990s, a certain new type of devices became popular: The PDA (personal digital assistant). They pioneered handheld computing by introducing touch screen user interfaces and a wide range of application programs, ranging from calendar/scheduling applications, dictionaries, exchange rate and tip calculators, scientific calculators, accounting / finance software, etc.

While in mobile phones the actual cellphone aspect was becoming more and more commoditized, at some point the PDA features and functionalities were added to phones, coining the term *smartphone*. At that point there was a need to differentiate from those phones that were not-so-smart. Those phones were then called *feature phones*.

There has never been an industry-wide accepted definition of those terms, and especially in the late 2000s, feature phones started to inherit a lot of the functionality that was formerly only present in smartphones.

This document will define the terms (only for the purpose of this document) along a very clear border in hardware architecture, as will be described in the following sections:

2.1 Feature Phone

A feature phone is a phone that runs the GSM protocol stack (the software implementing the GSM protocol) as well as the user interface and all applications on a single processor. For historic reasons, this processor is known as the so-called *baseband processor* (BP). Some manufacturers also call it Cellular Processor (CP) or CMT.

The baseband processor often exposes a serial port (or today USB) over which the phone can be used as a terminal adapter, similar to old wireline modems. The industry standard protocol for this interface is an AT command set - extended and modified from how computers interfaced old wireline modems. The AT-command interface can be connected to a computer. The computer can then use the phone to establish data calls, send/receive short messages via SMS, and generally remote-control the phone.

2.2 Smartphone

There is no clear, industry-wide definition on the term "smartphone".

Originally, and for the scope of this paper, a smartphone is a phone that has one dedicated processor for the GSM protocol stack, and another (potentially multi-core) general purpose processor for the user interface and applications. This processor is known as the *application processor* (AP).

The baseband processor (BP) part in a smartphone is typically the same as in a feature phone. But instead of connecting it to a personal computer, a small PDA (personal digital assistant) is built into the same case.

We will later discuss smartphone hardware architecture in more detail, but let's first look at the GSM modem side of things.

3 GSM modem architecture

Every GSM phone, feature phone and smartphone alike, has a GSM modem interfacing with the GSM network.

This GSM modem consists of several parts:

- RF Frontend, responsible for receiving and transmitting on GSM frequencies
- Analog Baseband, responsible for modulation and demodulation

- Digital Baseband, responsible for digital signal processing and the GSM protocol stack

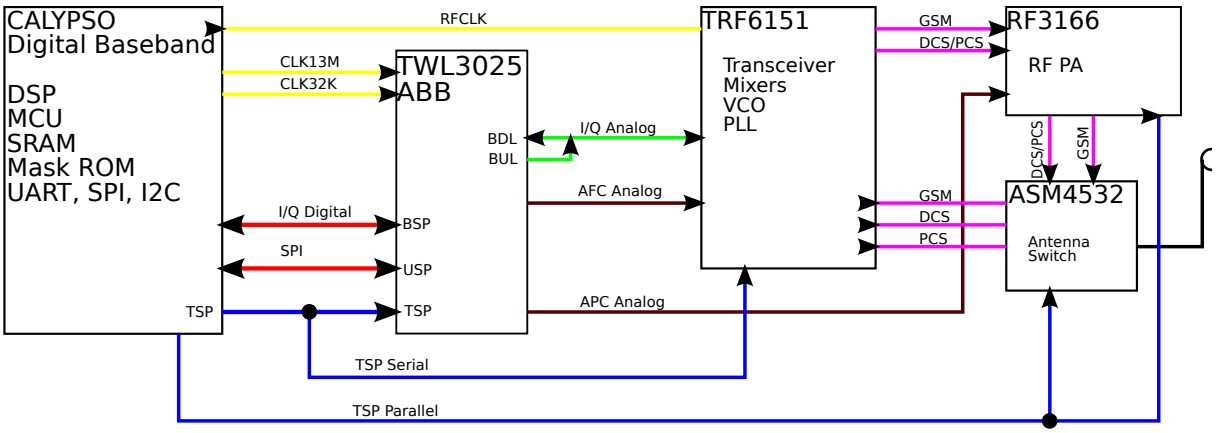

Figure 1: Block schematics of a TI Calypso/Iota/Rita GSM modem

3.1 The RF Frontend

The RF Frontend is tasked with the physical receive and transmit interface with the GSM air interface (sometimes called Um interface).

It minimally consists of an antenna switch, GSM band filters, low-noise amplifier (LNA) for the receive path, power amplifier for the transmit path, a local oscillator (LO) and a mixer.

By mixing the LO frequency with the received RF signal, it generates an analog baseband signal that is passed to the Analog Baseband (ABB) part of the modem. By mixing the output of the ABB with the LO frequency, it generates a RF signal that is to be transmitted in the GSM frequency band.

As the receive and transmit framing has an offset of 3 TDMA frames, there is no need for a frequency duplexer. Instead, an antenna switch is used. The switch typically is implemented using a MEMS or diode switch. For a quad-band phone, typically a single-pole 6-throw (SP6T) switch is used: 4 for the four Rx bands (one for each band), and 2 for Tx (where 850+900 and 1800+1900 each share one PA output, respectively).

3.1.1 RF Frontend receive path

The antenna picks up the GSM radio signal as it is sent from a GSM cell tower (properly called a Base Transceiver Station, or abbreviated as BTS). The antenna signal first hits the antenna switch, which connects the antenna with the Rx path for the GSM band of the to-be-received radio frequency. It is then filtered by a bandpass to block out-of-band signals before entering a low-noise amplifier for increasing signal amplitude.

After passing the LNA, the RF signal is mixed with a frequency generated by the LO. Depending on the LO signal, either an intermediate frequency (IF) or a direct baseband signal is produced. In modern GSM modems, zero-IF designs with immediate down-conversion to analog baseband signals are most common.

The baseband signal is then filtered to remove unwanted images and sent as analog I/Q signals (representing amplitude and phase) to the ABB.

3.1.2 RF Frontend transmit path

The ABB generates analog I/Q signals, which are filtered and passed into the mixer, where they are mixed with the LO frequency and thus up-converted to the GSM RF band. From there, they are sent to the transmit amplifier (RF PA) for amplification. After amplification, they traverse the antenna switch and are transmitted by the antenna.

3.1.3 Local Oscillator

The LO of a GSM modem has to be synchronized very closely to that of the cell (BTS). To achieve the required precision, a Voltage-Controlled, Temperature-Compensated Crystal Oscillator (VCTCXO) is used.

Common frequencies for this VCTCXO are 26MHz or 13MHz, as the GSM bit clock (270,833 Hz) is an integral division (/96 or /48, respectively) of those frequencies. The tuning range of the VCTCXO is several kHz to compensate for temperature drift.

3.2 The Analog Baseband (ABB)

The ABB part of a GSM modem is responsible to interface between the digital domain and the analog domain of the GSM modem.

3.2.1 ABB Receive path

The analog baseband I/Q signals are potentially filtered again and digitized by an Analog-Digital converter (ADC). The sample clocks used are typically integral multiples of the GSM bit-clock. The sample clock itself is derived by dividing the VCTCXO of the RF frontend.

The digital I/Q samples are passed to the Digital Signal Processor (DSP) in the Digital Baseband (DBB). To reduce the number of traces to be routed on the PCB, the samples are typically sent over some kind of synchronous serial link.

3.2.2 ABB Transmit path

There are multiple architectures found in the ABB transmit path.

The obvious architecture is to do the inverse of the receive operation: Transmit digital I/Q samples from the DSP to the ABB and convert them into an analog signal that is then to be sent to the mixer of the RF Frontend.

However, sending a GSM signal with its GMSK modulation is much simpler than receiving. So in order to reduce computational complexity (and thus cost as well as power consumption) inside the DSP, the modulation of the bits is often performed in hardware inside the ABB.

In this design, the unmodulated GSM burst bits are sent from the DBB to a burst buffer inside the ABB. From there, based on ROM tables and a Digital-to-Analog converter (DAC), an analog GMSK modulated signal is generated.

3.3 The Digital Baseband (DBB)

The digital baseband implements the actual GSM protocols from Layer1 up to Layer3 as well as higher layers such as a user interface in the case of the feature phone. In a smartphone, the DBB only implements a machine interface to be used by the AP.

A typical DBB design includes a Digital Signal Processor (DSP) for the lower half of Layer1, and a general-purpose processor (MCU) for the upper part of Layer1, as well as anything above.

3.3.1 Digital Signal Processor

The choice of DSP architecture largely depends on the DBB chipset vendor. Often they already have a line of DSP cores in-house and will of course want to reuse that in their DBB chip designs. Every major DSP architecture can be found (TI, Analog Devices, ...).

The DSP performs the primary tasks such as Viterbi equalization, demodulation, decoding, forward error correction, error detection, burst (de)interleaving.

Of course, if actual speech data is to be communicated over the GSM network, the DSP will also have the auxiliary task to perform the computation of the lossy speech codec used to compress the speech.

Communication between the DSP and MCU happens most commonly by a shared memory interface. The shared memory contains both actual data that is to be processed, as well as control information and parameters describing what to be done with the respective data.

For the receive side, the MCU will instruct the DSP to perform decoding for a particular GSM burst type, after which the DSP will receive I/Q samples from the ABB, perform detection/demodulation/decoding and report the result of the operation (including any decoded data) back to the MCU.

For the transmit path, the MCU will present the to-be-transmitted data and auxiliary information to the DSP, which then takes care of encoding and sends the corresponding burst bits to the ABB (remember, most ABB devices take care of the modulation to reduce DSP load).

The detailed programming information (API) of the DSP shared memory interface is a closely-guarded secret of the baseband chip maker and is not commonly disclosed even to their customers (the actual phone making companies).

In doing so, the baseband chip makers create a close dependency between the GSM Layer1 software (running on the MCU) driving/implementing this API and the actual baseband chip. Whoever buys their chip will also have to license their GSM protocol stack software.

It is thus almost impossible for an independent software vendor to get access to the DSP API documentation, which the author of this paper finds extremely anti-competitive.

3.3.2 DSP Peripherals

The specifications of the GSM proprietary On-air encryption A5/1 and A5/2 are only made available to GSM baseband chip makers who declare their confidentiality. Implementing the algorithm in software is apparently considered as breach of that confidentiality. Thus, the encryption algorithms are only implemented in hardware - despite them being reverse-engineered and publicly disclosed by cryptographers as early as 1996.

Thus, the DSP in a DBB commonly has a integrated peripheral implementing the A5 encryption.

Further integrated DSP peripherals may include a viterbi hardware accelerator, a DMA capable serial interface to the ABB and others.

3.4 Baseband Processor (MCU)

The MCU of almost all modern GSM DBBs is a System-on-a-Chip (SoC) utilizing a 32bit ARM7TDMI core. The only notable exception are low-cost Infineon chips like PM7870, who still use a version of their 16bit C166 core.

Baseband chips that support 3G cellular networks often use a more powerful ARM926 or ARM975 core as MCU.

3.5 MCU peripherals

The MCU cores have the typical set of peripherals of any ARM7 based microcontroller, such as RTC, UARTs for RS232 and IRDA, SPI, I2C, SD/MMC card controller, keypad scan controller, USB device, ...

However, there are some additional peripherals that are very GSM specific:

- A GPRS crypto unit for the proprietary GEA family of ciphers
- Extended power management facilities, including a timer that can calibrate the RTC clock based on the synchronized VCTCXO in order to wake-up the MCU ahead of pre-programmed events in the GSM time multiplex
- GSM TDMA timers that can synchronize to the on-air time frames and generate interrupts to MCU and DSP
- Software-programmable hardware state machines for sequencing GSM burst Rx or Tx in ABB and RF Frontend
- An ISO7816 compatible smart card reader interface for the SIM card

- Audio routing, i.e. selecting how audio is routed in the phone, considering integrated earpiece, ringtone speaker and microphone, as well as external analog headset, handsfree or even bluetooth-attached audio devices.

The programming of those peripherals is highly device-specific and there are no industry standards. Every DBB architecture of every supplier has its own custom register set and programming interface.

The register-level documentation for those proprietary peripherals is (like all documentation on DBB chipsets) closely guarded by NDAs, effectively preventing the development of Free Software / Open Source drivers for them, unless such documents are leaked by third parties.

However, as opposed to the DSP API documentation, the register-level documentation to the MCU peripherals is normally provided to the cellphone manufacturers.

4 Digital Baseband Software Architecture

This section provides an introductory reading in the typical software architecture as it is found on contemporary GSM digital baseband designs.

The MCU usually runs a very small realtime operating system (RTOS) such as Nucleus, VxWorks or the L4 microkernel. In some cases, no operating system is used at all, in order to save royalties or licensing fees that would occurr for proprietary RTOS.

4.1 GSM Layer 1

The Layer1 (L1) software is highly device-specific, as it closely interacts with the DSP using the shared memory DSP API, as well as the proprietary integrated peripherals controlling the ABB and RF Frontend.

However, there are some general observations that can be made about the L1:

4.1.1 L1 Synchronous part

The synchronous part is executed synchronously to the GSM TDMA frame clock. Both CPU and DSP are interrupted by some hardware GSM timer every TDMA frame.

The L1 synchronous part typically runs inside IRQ or FIQ context of the MCU, taking care of retrieving data from and providing data to the DSP API.

4.1.2 L1 Asynchronous part

The asynchronous part is scheduled as a normal task, potentially with high or even real-time priority. It picks up the information provided by the L1 Sync and schedules its next actions.

The L1 async typically communicates via a message queue with the Layer2 above. Common primitives for L1 control are described (as non-normative parts) of the GSM specifications.

4.2 GSM Layer 2

As opposed to L1, the GSM Layer 2 (L2) is already fully hardware independent. It implements the LAPDm protocol as specified in GSM TS 04.06. LAPDm is a derivative of the ISDN Layer 2 called LAPD, which in turn is a descendent of the HDLC family of protocols.

LAPDm takes care of providing communication channels for Layer3. Those channels are protected from frame loss by the use of sequence numbers and retransmissions.

The interface to Layer3 is typically implemented by means of a message queue.

Primitives (but no detailed protocol) for use of the Layer2 / Layer3 interface are provided in the GSM specifications.

4.3 GSM Layer 3

GSM Layer 3 (L3) consist of sublayers for Radio Resource (RR), Mobility Management (MM) and Call Control (CC).

There is sufficient treatment of the GSM L3 and its sublayers in existing texts, so there is no point in making a futile attempt repeating that here.

5 Synchronization and Clocking

The author of this paper has been quoted saying *GSM is a synchronous TDMA nightmare*. This is by no means intended as an insult to the technology itself or to its inventors. It merely serves as evidence how hard it is to get into the synchronous TDMA mindset, especially for engineers who have spent most of their career in the world of packet switched networks.

GSM is synchronous in multiple ways between cell (BTS) and phone (MS):

- Synchronization of the carrier clock to tune the receiver and transmitter to the correct frequency
- Synchronization of the bit clock in order to perform sampling at the most optimal sample intervals
- Synchronization of the frame clock (and thus timeslots) to know when a TDMA frame and its 8 timeslots start
- Synchronization of the TDMA multiplex to correctly (de)multiplex the logical channels that are sent over each timeslots

As all those clocks are related to each other, they can (and should) all be derived from the same master clock: The VCTCXO present in each GSM phone.

5.1 How to synchronize the VCTCXO

Every cell sends frequency correction bursts as part of the Frequency Correction CHannel (FCCH), which is itself part of the BCCH, which in turn is constantly transmitted by the BTS.

To acquire initial synchronization to the GSM network, the LO is tuned to the desired GSM RF channel (ARFCN) frequency. However, at this point, the LO frequency is a multiple of the VCTCXO frequency which in turn still has an undetermined error. This initial frequency error is as large as that of a regular crystal oscillator, potentially already with temperature compensation.

The resulting baseband signal thus can be shifted by a fairly large amount in our baseband spectrum. A specific DSP code is now using correlation and other techniques to identify the frequency correction burst. The DSP can then further calculate the actual frequency error of the LO by comparing the received FCCH burst with the FCCH burst as specified.

This computed frequency error can be fed into a (software) frequency control loop filter. The loop filter output is applied to an auxiliary DAC, which generates the control voltage for the VCTCXO.

After a number of FCCH bursts and corresponding frequency control loop iterations, the VCTCXO generated clock has only a residual error. Whenever the phone is receiving, the frequency control loop is continuously exercised in order to maintain synchronization.

5.2 How to synchronize the frame clock

When the DSP performs FCCH burst detection as described above, it identifies the exact position in the incoming sample stream when the FCCH burst was happening. By knowing from the specification that the FCCH burst is part of the BCCH, and that the BCCH is sent on timeslot 0, the Layer1 software can then synchronize the phone to the TDMA frame start.

Commonly, a hardware timer unit is clocked by a (divided) VCTCXO clock and thus counts in multiples of the GSM bit clock, wrapping/resetting at the TDMA duration of 1250 bits.

By scheduling events synchronously to this GSM bit-clock timer, the L1 can now trigger events (such as asking the DSP to demodulate incoming data) or instructing the LO to retune synchronously to every TDMA frame. From this timer the DBB typically also generates interrupts to the DSP and MCU.

5.3 How to synchronize the GSM TDMA multiplex

As part of the BCCH, the BTS not only sends the FCCH but also the Synchronization CHannel (SCH). The Synchronization channel indicates the current GSM time / frame number (skipping the 3 least significant bits). By using this received GSM time and incrementing it every time the GSM bit-clock timer wraps at the beginning of a new TDMA frame, the GSM time is synchronized.

Understanding the multiple layers of time multiplex such as the 26/51 multiframe, superframe and hyperframe, the L1 can multiplex and demultiplex all the logical channels of GSM.

6 Miscellaneous Topics

6.1 GPRS

GPRS was the first packet switched extension to GSM. In fact, it is much more its entirely own mobile network, independent of GSM. The only parts shared are the GSM modulation scheme (GMSK) and time multiplex, in order to ensure peaceful coexistence between them.

The L1 and L2 protocols are very different (and much more complex) than GSM.

So while the phone baseband hardware did not need any modifications for a basic GPRS enabled phone, the software needed to be extended quite a lot.

6.2 EDGE

EDGE is a very small incremental set of changes from GPRS. It reuses all of the time multiplex and protocol stack, but introduces a new modulation: Offset 8-PSK instead of GMSK to increase the bandwidth that can be transmitted. Offset 8-PSK is used (as opposed to simple 8-PSK) to avoid zero-crossings in the modulator output.

So while the software modifications from GPRS to EDGE are minimal, the 8PSK modulation scheme has a significant impact on the DSP, ABB and even RF PA design.

6.3 UMTS

UMTS (sometimes called WCDMA) is an entirely separate cellular network technology. Its physical layer, modulation schemes, encoding, frequency bands, channel spacing are entirely different, as is the Layer1.

UMTS Layer2 has some resemblance to the GPRS Layer2.

UMTS Layer3 for Mobility Management and Call Control are very similar to GSM.

Given the vast physical layer and L1 differences, a UMTS phone hardware design significantly differs from what has been described in this document.

Notwithstanding, all known commercial UMTS phone chipsets as of today still include a full GSM modem in hardware and software to remain backwards-compatible.

6.4 Dual-SIM and Triple-SIM phones

In recent years, a large number of so-called *Dual-SIM* or even *Triple-SIM* phones have entered the market, particularly in China and other parts of East Asia.

Those phones come in various flavours. Some of them simply have a multiplexer that allows electrical switching between multiple SIM card slots. This is similar to replacing the SIM card in a phone, just without

the manual process of mechanically removing/inserting the card. As a result, you can only use one of the two SIMs at any time.

The more sophisticated Dual-SIM phones have two complete phones in one case. Yes, that's right! They contain two full GSM phone chipsets, i.e. 2 antennas, 2 rf frontends, 2 analog basebands, 2 digital basebands, ...

However, they use the same trick as smartphones: One of the two basebands does not have keypad or display and is simply a GSM modem connected via serial line to the other baseband processor.

So if a smartphone (as defined in this document) is a GSM modem connected to a PDA in one case, a Dual-SIM phone is a GSM modem connected to a feature phone in one case.

Triple-SIM phones often combine the two approaches, i.e. they contain two complete GSM baseband chips, but three SIM slots that can be switched among the base bands. Only two SIMs can be active at the same time.

6.5 GSM baseband security features

There are several (sometimes conflicting) security requirements that apply to mobile phones. Interestingly, the security features are typically used to protect some industry interest against the interest of the customer. There are very few security features in a phone that are meant to protect the users or their interests.

6.5.1 IMEI - The hardware serial number

The International Mobile Equipment Identifier (IMEI) uniquely identifies a GSM phone. It is a globally unique serial number which is programmed into the phone non-volatile memory (Flash or EEPROM) during the production process. There are no particular security features used to store the IMEI. Once you are able to write to flash and have found it, it can be changed.

6.5.2 The SIM Card

The SIM card is a cryptographic smart card that stores the secret key used for authenticating the user to the GSM network (Ki). The Ki is never released by the card, and as such copying (cloning) of the SIM is prevented. Some early implementations of the SIM card had cryptographic weaknesses that inadvertently permitted cloning until the late 1990s.

Furthermore, the SIM stores the International Mobile Subscriber Identity (IMSI). The IMSI is not encrypted or protected in any way.

There is no security channel on the connection between the SIM card and the baseband MCU. Furthermore, there is no way how the MCU can securely identify/authenticate the SIM itself. Physical access to the SIM card slot allows sniffing and/or modification of the communications between the MCU and the SIM.

6.5.3 SIM or Operator Locking

GSM is an international standard. This ensures interoperability, i.e. any phone can be used on any network.

However, in some cases, the vendors of a GSM phone want to restrict this interoperability and lock a phone to one specific network, or even lock it to a particular SIM.

Those locks are implemented by software only, i.e. the MCU software will instruct the GSM protocol stack not to register with a network unless its network operator code is a certain factory-programmed network number.

As such, techniques for circumventing those locks have become commonplace. It's somewhat of an ongoing race between the phone makers and the phone-unlockers. The industry invents ever more complex methods of obfuscating their locks in the software, while the phone-unlockers reverse engineer those bits for each and every phone model after some time.

6.5.4 DBB firmware signatures

In order to protect the operator and phone manufacturers peculiar business models based on SIM or operator locking, some vendors extended their baseband software with cryptographic signatures. Only if the correct signature is present in a software update, the bootloader program will accept the new software.

However, there are more or less invasive hardware-related approaches in circumventing those signature checks, such as hardware debugging interfaces like JTAG, or replacing the external flash memory components.

More recently, GSM chipset vendors introduced features such as hardware-assisted software signature checks. In this case a master key hash might be present in DBB-internal fuses, together with a signature-verifying boot loader in DBB-internal mask ROM. As the root of the chain of trust is moving deeper into the hardware, it becomes more difficult for anyone to make software modifications to the DBB.

Especially with tighter integration, where RAM and FLASH are no longer present as discrete components but part of a multi-chip-package, the number of options are becoming more limited.

On the other hand, an ever more complex baseband software stack is opening up many more options for exploiting software vulnerabilities. Given the lack of a proper/modern operating system with privilege separation and virtual memory, such exploits immediately give away full access to all aspects of the respective DBB.

7 Smartphone hardware archtiecture

A smartphone is a phone that has a dedicated processor for the GSM protocol stack, and another (potentially multi-core) general purpose processor for the user interface and applications. This processor is known as the *application processor* (AP).

The purpose of the application processor is to run a general-purpose operating system (OS) driving a rich user interface (UI) software stack, which provides a platform for running application programs. Such applications migh be developed by the original phone vendor or by third parties.

There is no specific technical reason for splitting the functionality among two independent processors. It is more likely that business and political issues played a role in this. The baseband processor vendors are typically very reluctant to give up any amount of control over "their" processor. They also don't usually provide sufficient informtaion or a general-purpose enough operating system on it.

So the logical choice is to keep the "phone part" (aka GSM modem) just like it was (and is) in feature phones, but add an entirely separate PDA-like embedded system with an application processor into the same device.

It is common to use existing feature phone GSM modem designs in smartphones. The same BP chipset and peripherals are used.

7.1 Fully separate AP and BP

The first hardware generations of smartphones did nothing else but to put the feature phone and a PDA into one case. The keypad and display connection to the BP is removed. What remains of the feature phone is a *GSM modem*, controlled by AT commands sent from the AP.

Each processor has its own memory (RAM and Flash), peripherals, clocking, etc. So this setup is not to be confused with the symmetric multi-processor setups like those seen in the personal computer industry.

The interface between AP and BP originally was a simple serial (UART) line, but ever since there has been a growing variety of electrical-level interfaces. For more information, see the section below on the Interface between AP and BP

7.2 Integrated Smartphone-on-a-chip Solutions

Due to market pressure for ever smaller phones with ever more functions, the industry has produced highly integrated products, uniting the AP and BP inside one physical package. The first popular example was the Qualcomm MSM7200 as used in the first generation of Android and many Windows Mobile phones.

More recently, other manufacturers such as ST-Ericsson (a merger of the cellular chipset business of NXP, ST Micro and Ericsson Mobile Platforms) have been shipping similar products.

Such integrated chips typically combine the

- Application processor (typically ARM11, Cortex-A8 or A9)
- AP peripherals such as RAM-controller, display controller, I2C, SPI, SDIO, etc.
- Digital Baseband (DBB), typically including an ARM9 core and a DSP
- Integrated peripherals of a the BP, including ADC and DAC
- A GPU (Graphics Processing Unit) for 3D and/or video codec acceleration

Sometimes, even a second DSP is added for signal processing tasks of the AP side.

Further pressure on reducing cost and PCB footprint has led to products where there is no need to have independent RAM and Flash chips for AP and BP. Rather, a single RAM and Flash chip is divided by assigning portions of the RAM and Flash to each of the two processors.

In such systems, some integrated peripheral logic is separating the physical RAM and flash into portions that are accessible from the AP and portions accesible from the BP. The division ratio as well as the access levels might be configurable by software, eFuses or bootstrap pins of the package.

However, the fundamental separation between the AP and BP, each with their own memory address space and software, remains present in all smartphones until today.

7.3 Control + Data Interface between AP and BP

7.3.1 Serial Line

The interface between AP and BP originally was a simple serial line (UART), over which AT commands compliant with GSM TS 07.05 / 07.07 are spoken. A serial line with a standard speed of 115200bps is sufficient for the control of GSM voice calls, SMS, circuit switched data (CSD), as well as most GPRS data speeds. However, for concurrent data and voice services, a serial multiplexor protocol according to GSM TS 07.10 was used. It provides multiple virtual channels with each their own instance of an AT command parser on the BP.

As the data speeds of the cellular networks were increased with EDGE (both ECSD and EGPRS), an asynchronouse serial connection at standard speeds became to narrow as a communications channel.

7.3.2 Universal Serial Bus (USB)

The EDGE capable GSM modems that were once again coming from the feature phone designs typically included a USB device mode controller for attaching those feature phones to personal computers.

While many of the USB-device-capable BPs use the standardized CDC-ACM protocol to emulate one or multiple serial ports over USB, there never was any standard or even any recommendation in the GSM/3GPP specifications.

So a number of smartphone designs such as the Motorola EZX platform (A780, A1200, ROKR E6, etc.) simply used that existing USB device-mode controller and connected it to a USB host controller inside the AP. However, USB is far from being a good protocol for this application, mostly due to power management issues. If the phone is idle, the AP switches in some kind of deep-sleep state. To do this, it has to disable the USB host controller, whcih in turn means that the BP has no way how to actually issue a wake-up to the AP in case of an incoming call. The solution to the problem was connecting some general purpose output signal of the BP to a wakeup-capable general-purpose input of the AP. However, this means that the system is no longer fully USB compatible, and that the BP software has to be specifically modified.

7.3.3 Serial Peripheral Interface

Some smartphone designs, most notably those of E-TEN corporation (now Acer) have started to use SPI-class electrical interfaces between the AP and BP.

However, as SPI normally is a master/slave type of protocol, additional handshaking was needed to allow the slave to request an outgoing data transfer from the master.

Modern application processors support SPI with speeds of up to 25 or sometimes 50 MHz, providing more than sufficient bandwidth for even the fastest available cellular transfer speeds over the air interface.

The second-layer protocol on top of this SPI link is vendor-specific and proprietaty. One of them is known as CAIF by Ericsson Mobile Products (EMP).

7.3.4 Shared Memory / Dual Ported RAM

Another method for interfacing wth AP with the BP is by using some form of shared memory. The clear advantage is speed, as access to parallel RAM is typically several orders of magnitude faster than any serial link. Furthermore, there is no need for serializer/deserializer, the use of DMA controllers and the like. The data is available without any copying (zero-copy).

Management of shared memory is a complex problem though, and there has to be some kind of mutual exclusion mechanism to prevent coherency/concurrency problems like race conditions.

Depending on the chipset architecture, this is either an actual external dual-ported RAM (DPRAM) that provides separate address and data busses for AP and BP. Sometimes that DPRAM is built into the BP - or simulated by the BP using some internal arbitration logic.

In the latest Smarphone-on-a-Chip systems, the shared memory is simply one portion of the phyiscal RAM which is mapped into the address space of both AP and BP parts - while the remaining RAM is mapped exclusively to either the AP or the BP.

7.4 Audio Interface between AP and BP

In feature phones, the audio architecture is quite simple: Microphone and earpiece speaker are all connected to an audio codec which is co-located with the analog baseband (ABB) chip. A digital serial audio interface connects this codec with the DSP inside the BP. As the DSP does both the demodulation/decoding of the baseband signal as well as the actual voice codec processing, speech data never needs to leave the DSP. The MCU is purely handling control, but no voice data.

With a bluetooth headset things get slightly more complex, but not much. Bluetooth chipsets for mobile phones have a control interface (typically UART), as well as a synchronous serial PCM interface. That PCM interface then needs to be hooked up either to the ABB or the DSP.

In a smartphone however, things are getting more complicated. The Application Processor wants to play-back music, both via the ringtone speaker as well as the headset. Ringtones are no longer played back by the BP, but are typically mp3 files on the AP. Voice recording / memo features require the microphone to be routed to the AP. Some advanced phones allow you to record an actual voice call. While that call is handled by the BP, recording is done by the AP, which is storing it to mass storage memory such as a (micro)SD card.

There are different audio architectures on smartphones, all of them are complex.

7.4.1 Analog audio interface

This is the most "logical" interface, looking at the idea of a smartphone being a feature phone and a PDA in one box: The AP gets an audio codec chip not different to what a "sound card" used to be for the PC.

Using proper analog impedance matching networks, you connect the analogue output of the ABB to a line input of the codec chip. One of the codec outputs is connected to the microphone input of the codec chip.

The actual micrphone is connected to the microphone input of the codec chip, while the headphone jack and ringtone speakers are connected to corresponding outputs of the codec.

The digital (PCM/IIS) interface of the codce is driven by the AP.

So all connections between ABB and codec are analog, while the AP-codec connection is digital.

If you add a bluetooth interface for wireless headsets, the codec chip will need a second IIS/PCM interface which is then connected to that bluetooth chip.

Analog audio signals on an otherwise completely digital device can be cumbersome. They will likely catch noise from power supply or digital signals.

7.4.2 Digital audio interface

The solution to this problem is to use digital audio interfaces. This will require some cooperation/integration with either the ABB or the DBB of the baseband processor and was not possible with re-purposed BP chipsets that were not built with smartphones in mind.

One possible architecture is to have an ABB that offers a secondary PCM/IIS interface for the AP. Another solution is to use the PM/IIS as a multi-master bus, which is either driven by the AP or the BP, depending on the current use case.

The third option is to no longer use any voice band DAC/ADC that might be present in the ABB and use a codec chip that has at least two (three with bluetooth) PCM/IIS interfaces, and a DBB that has a compatible digital PCM interface.

8 Powerful feature phones

Feature phones are becoming more and more powerful. However, their comparatively lower market price cannot afford a full-blown smartphone design with its two independent processors and the associated design complexity.

Thus, more and more hardware peripherals are added to the only processor left in the phone: The baseband processor. Such peripherals include sophisticated camera interfaces, high-resolution color display controllers, TV output, touchscreen controllers, audio and video codecs and even interfaces for mobile TV reception.

However, all of those features are still implemented on a fairly weak ARM7 or ARM9 CPU core (compared to ARM11 and Cortex-A8 in the smartphone market). They also lack a real operating system and still run on top of a real-time microkernel intended for much less complex systems. They almost always lack any form of memory protection or multiple address spaces. This makes them more prone to security issues as there is no privilege separation between the GSM protocol stack and the applications, or between the applications themselves.

The chipset vendor most associated with this strategy is Mediatek (MTK) in Taiwan, who bough the cellular chipset business from Analog Devices (ADI).

In MTK chipsets , you can find unusual combinations of an ARM7 core with Jazelle (ARM7EJS), which has H.264 hardware codecs attached to it. Most of the competition doesn't have Jazelle or advanced video codecs in anything smaller than an ARM9.

9 Personal rant on the closedness of the GSM industry

The GSM industry is one of the most closed areas of computing that I've encountered so far. It is very hard to get any hard technical information out of them. All they like to spread is high-level marketing information, but they're very reluctant when it comes down to hard technical facts on their products.

If you want to build a phone, you need to buy a GSM chipset for your product. There are only very few companies that offer such chipsets. The classic suppliers are Infineon, Texas Instruments, ST/Ericsson, ADI (now MediaTek) and Freescale.

The GSM handset products they sell are not generally available and distributed like other electronic components they manufacture. If you need a Microcontroller/SoC, a power management IC, a Wifi or Bluetooth chip, RFID reader ASIC, you simply approach the respective distributors and order them. You get your samples directly from Digikey.

This is impossible for GSM (or other cellphone) chipsets. For some reason those chips are sold only

to hand-picked manufacturers. If you want to qualify, you have to subscribe to at least six-digit annual purchasing quantities. And in order for them to believe you, you have to cough up a significant NRE (non-refundable engineering fee). This has been reported as high as a seven-digit US$ amount and is to make sure that even if you end up buying less chips than you indicate, the chipset maker will still have your upfront NRE money and keep it.

And if you buy your way into that select club of cellphone makers, what you get from the chipset maker is typically not all too impressive either. The documentation you get is incomplete, i.e. it alone would not enable you as a cellphone maker to make any use of the hardware, unless you license the software (drivers, GSM protocol stack, ...) from the chipset maker, too.

On the software side, most of the technologically interesting bits (like the protocol stack) are provided as binary-only libraries, you only get source code to some parts of the systems, i.e. some hardware drivers that might need modification for your particular phone electrical design.

That GSM protocol stack was not written by the chipset maker either. They simply license a stack from one of the estimated 4 or 5 organizations who have ever implemented a commercial GSM protocol stack.

It is not like the GSM protocols were some kind of military secret. They are a published international standard, freely accessible for anyone. So why does everybody in that industry think that there is a need to be so secretive?

Having spent a significant part of the last 6 years with reverse engineering of various aspects of mobile phones in order to understand them better and to write software tools for security analysis, I still don't understand this secrecy.

All the various vendors do more or less the same. The fundamental architecture of a GSM baseband chip is the same, whether you buy it from TI, Infineon or from MediaTek. *They all cook with water*, like we Germans tend to say. The details like the particular DSP vendor or whether you use a traditional IF, zero-IF or low-IF analog baseband differ. But from whom do they want to hide it? If people like myself with a personal interest in the technical aspects of mobile phones can figure it out in a relatively short time, then I'm sure the competition of those chipset makers can, too. In much less time, if they actually care.

This closedness of the cellular industry is one of the reasons why there has been very little innovation in the baseband firmware throughout the last decades. Innovation can only happen by very few players. Source code bugs can only be found and fixed by very few developers at even fewer large corporations. There is little to no chance for a small start-up to innovate, like they can in the sphere of the internet.

It is fundamentally also the reason why the traditional phone makers have been losing market share to newcomers to the mobile sphere like Apple with its iPhone or Google with its Android platform.

Those innovations really only happened on the application processor on high-end smartphones. The closed GSM baseband processor had to be accompanied by an independent application processor running a real operating system, with real processes, memory management, shared libraries, memory protection, virtual memory spaces, user-installable applications, etc.

They still don't happen on the baseband processor, which is as closed as it was 15 years ago.