

Faculté des Sciences et Ingénierie
Master Informatique
Systèmes Électroniques, Systèmes Informatiques
Laboratoire d'Informatique Paris 6 - CIAN

Hardware initialization of modern computers

A review on the importance of firmware in modern computing and a documentation on the Asus KGPE-D16 RAM initialization

August, 2024

Adrien 'neox' Bourmault (neox@gnu.org)

Under the supervision of Franck WAJSBÜRT (franck.wajsburt@lip6.fr)

This is Edition 0.0.

Copyright (C) 2024 Adrien 'neox' Bourmault <neox@gnu.org>

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Acknowledgments

Thanks, I guess ? (TODO)

Abstract

The global trend is towards the scarcity of free software-compatible hardware, and soon there will be no computer that will work without software domination by big companies, especially involving firmware like BIOSes.

A Basic Input Output System (BIOS) was originally a set of low-level functions contained in the read-only memory of a computer's mainboard, enabling it to perform basic operations when powered up. However, the definition of a BIOS has evolved to include what used to be known as Power On Self Test (POST) for the presence of peripherals, allocating resources for them to avoid conflicts, and then handing over to an operating system boot loader. Nowadays, the bulk of the BIOS work is the initialization and training of RAM. This means, for example, initializing the memory controller and optimizing timing and read/write voltage for optimal performance, making the code complex, as its role is to optimize several parallel buses operating at high speeds and shared by many CPU cores, and make them act as a homogeneous whole.

This document is the product of a project hosted by the *LIP6 laboratory* and supported by the *GNU Boot Project* and the *Free Software Foundation*. It delves into the importance of firmware in the hardware initialization of modern computers and explores various aspects of firmware, such as Intel Management Engine (ME), AMD Platform Security Processor (PSP), Advanced Configuration and Power Interface (ACPI), and System Management Mode (SMM). Additionally, it provides an in-depth look at memory initialization and training algorithms, highlighting their critical role in system stability and performance. Examples of the implementation in the ASUS KGPE-D16 mainboard are presented, describing its hardware characteristics, topology, and the crucial role of firmware in its operation after the mainboard architecture is examined. Practical examples illustrate the impact of firmware on hardware initialization, memory optimization, resource allocation, power management, and security. Specific algorithms used for memory training and their outcomes are analyzed to demonstrate the complexity and importance of firmware in achieving optimal system performance. Furthermore, this document explores the relationship between firmware and hardware virtualization. Security considerations and future trends in firmware development are also addressed, emphasizing the need for continued research and advocacy for free software-compatible hardware.

Contents

Acknowledgments	3
Abstract	4
List of Figures	7
List of Listings	8
1 Introduction to firmware and BIOS evolution	9
1.1 Historical context of BIOS	9
1.1.1 Definition and origin	9
1.1.2 Functionalities and limitations	10
1.2 Modern BIOS and UEFI	11
1.2.1 Transition from traditional BIOS to UEFI (Unified Extensible Firmware Interface)	11
1.2.2 An other way with <i>coreboot</i>	11
1.3 Shift in firmware responsibilities	13
2 Characteristics of ASUS KGPE-D16 mainboard	14
2.1 Overview of ASUS KGPE-D16 hardware	15
2.2 Chipset	16
2.3 Processors	18
2.4 Baseboard Management Controller	19
3 Key components in modern firmware	21
3.1 General structure of coreboot	21
3.1.1 Bootblock stage	22
3.1.2 Romstage	24
3.1.3 Ramstage	25
3.1.3.1 Advanced Configuration and Power Interface	25
3.1.3.2 System Management Mode	26
3.1.4 Payload	26
3.2 AMD Platform Security Processor and Intel Management Engine	27
4 Memory initialization and training	29
4.1 Importance of DDR3 Memory Initialization	29
4.1.1 General steps for DDR3 configuration	30
4.2 Memory initialization techniques	32
4.2.1 Memory training algorithms	32
4.2.2 BIOS and Kernel Developer Guide (BKDG) recommendations	33
4.2.2.1 DDR3 initialization procedure	34
4.2.2.2 ZQ calibration process	34

4.2.2.3	Write leveling process	35
4.3	Current implementation and potential improvements [WIP]	36
4.3.1	Current implementation in coreboot on the KGPE-D16 [WIP]	36
4.3.2	Potential enhancements [WIP]	37
5	Virtualization of the operating system through firmware abstraction	38
5.1	ACPI and abstraction of hardware control	38
5.2	SMM as a hidden execution layer	39
5.3	UEFI and persistence	39
5.3.1	Memory Management	40
5.3.2	File System Management	40
5.3.3	Device Drivers	40
5.3.4	Power Management	40
5.4	Intel and AMD: control beyond the OS	40
5.5	The OS as a virtualized environment	41
	Conclusion	42
	Bibliography	43
	GNU Free Documentation License	50

List of Figures

1.1	The eight-striped wordmark of IBM (1967, public domain, trademarked)	9
1.2	An AMI BIOS chip from a Dell 310, by Jud McCranie (CC BY-SA 4.0, 2018)	10
1.3	The UEFI logo (public domain, 2010)	11
1.4	The <i>coreboot</i> logo, by Konsult Stuge & coresystems (coreboot logo license, 2008)	12
1.5	The <i>GNU Boot</i> logo, by Jason Self (CC0, 2020).....	12
2.1	The KGPE-D16 (CC BY-SA 4.0, 2021)	14
2.2	Basic schematics of the ASUS KGPE-D16 Mainboard, ASUS (2011)	15
2.3	The KGPE-D16, viewed from the top (CC BY-SA 4.0, 2024)	16
2.4	Functional diagram presenting the IOAPIC function of the SP5100, ASUS (2011)	17
2.5	Functional diagram of the KGPE-D16 chipset (CC BY-SA 4.0, 2024)	17
2.6	Annotated photography of an Opteron 6200 series CPU (2024), from a photography by AMD Inc. (2008).....	18
2.7	Functional diagram of an Opteron 6200 package (CC BY-SA 4.0, 2024)	19
3.1	<i>coreboot</i> 's stages timeline, by <i>coreboot</i> project (CC BY-SA 4.0, 2009)	21
3.2	<i>coreboot</i> ROM architecture (CC BY-SA 4.0, 2024).....	22
4.1	DDR3 fly-by <i>versus</i> T-topology (CC BY-SA 4.0, 2021).....	29
4.2	DDR3 controller state machine	31

List of Listings

4.1	<i>Example C code</i>	37
5.2	<i>How to estimate the impact of ACPICA in Linux</i>	38

Chapter 1

Introduction to firmware and BIOS evolution

1.1 Historical context of BIOS

1.1.1 Definition and origin

The BIOS (Basic Input/Output System) is firmware, which is a type of software that is embedded into hardware devices to control their basic functions, acting as a bridge between hardware and other software, ensuring that the hardware operates correctly. Unlike regular software, firmware is usually stored in a non-volatile memory like ROM or flash memory. The term "firmware" comes from its role: it is "firm" because it's more permanent than regular software (which can be easily changed) but not as rigid as hardware.

The BIOS is used to perform initialization during the booting process and to provide runtime services for operating systems and programs. Being a critical component for the startup of personal computers, acting as an intermediary between the computer's hardware and its operating system, the BIOS is embedded on a chip on the motherboard and is the first code that runs when a PC is powered on. The concept of BIOS has its roots in the early days of personal computing. It was first developed by IBM for their IBM PC, which was introduced in 1981 [43]. The term BIOS itself was coined by Gary Kildall, who developed the CP/M (Control Program for Microcomputers) operating system [96]. In CP/M, BIOS was used to describe a component that interfaced directly with the hardware, allowing the operating system to be somewhat hardware-independent.



Figure 1.1: The eight-striped wordmark of IBM (1967, public domain, trademarked)

IBM's implementation of BIOS became a de facto standard in the industry, as it was part of the IBM PC's open architecture [46][14], which refers to the design philosophy adopted by IBM when developing the IBM Personal Computer (PC), introduced in 1981. This architecture is characterized by the use of off-the-shelf components and publicly available specifications, which allowed other manufacturers to create compatible hardware and software. It was in fact a departure from the proprietary systems prevalent at the time, where companies closely guarded their designs to maintain control over the hardware and software ecosystem. For example, IBM used the Intel 8088 CPU, a well-documented and widely available processor, and also the Industry Standard Architecture (ISA) bus, which defined how various components like memory, storage, and peripherals communicated with the CPU. This open architecture allowed other manufacturers to create IBM-compatible computers, also known as "clones", which further popularized the BIOS concept. As a result, the IBM PC BIOS set the stage for a standardized method of interacting with computer hardware, which has evolved over the years but remains fundamentally the same in principle. IBM also published detailed technical documentation at that time, including circuit diagrams,

BIOS listings, and interface specifications. This transparency allowed other companies to understand and replicate the IBM PC's functionality [43].

1.1.2 Functionalities and limitations

When a computer is powered on, the BIOS executes a Power-On Self-Test (POST), a diagnostic sequence that verifies the integrity and functionality of critical hardware components such as the CPU, RAM, disk drives, keyboard, and other peripherals [114]. This process ensures that all essential hardware components are operational before the system attempts to load the operating system. If any issues are detected, the BIOS generates error messages or beep codes to alert the user. Following the successful completion of POST, the BIOS runs the bootstrap loader, a small program that identifies the operating system's bootloader on a storage device, such as a hard drive, floppy disk, or optical drive. The bootstrap loader then transfers control to the OS bootloader, initiating the process of loading the operating system into the computer's memory and starting it. This step effectively bridges the gap between hardware initialization and operating system execution. The BIOS also provides a set of low-level software routines known as interrupts. These routines enable software to perform basic input/output operations, such as reading from the keyboard, writing to the display, and accessing disk drives, without needing to manage the hardware directly. By providing standardized interfaces for hardware components, the BIOS simplifies software development and improves compatibility across different hardware configurations [14].




Figure 1.2: An AMI BIOS chip from a Dell 310, by Jud McCranie (CC BY-SA 4.0, 2018)

Despite its essential role, the early BIOS had several limitations. One significant limitation was its limited storage capacity. Early BIOS firmware was stored in Read-Only Memory (ROM) chips with very limited storage, often just a few kilobytes. This constrained the complexity and functionality of the BIOS, limiting it to only the most essential tasks needed to start the system and provide basic hardware control. The original BIOS was also non-extensible. ROM chips were typically soldered onto the motherboard, making updates difficult and costly. Bug fixes, updates for new hardware support, or enhancements required replacing the ROM chip, leading to challenges in maintaining and upgrading systems. Furthermore, the early BIOS was tailored for the specific hardware configurations of the initial IBM PC models, which included a limited set of peripherals and expansion options. As new hardware components and peripherals were developed, the BIOS often needed to be updated to support them, which was not always feasible or timely. Performance bottlenecks were another limitation. The BIOS provided basic input/output operations that were often slower than direct hardware access methods. For example, disk I/O operations through BIOS interrupts were slower compared to later direct access methods provided by operating systems, resulting in performance bottlenecks, especially for disk-intensive operations. This inflexibility restricts the ability to support new hardware and technologies efficiently[13]. Early BIOS implementations also had minimal security features. There were no mechanisms to verify the integrity of the BIOS code or to protect against unauthorized modifications, leaving systems vulnerable to attacks that could alter the BIOS and potentially compromise the entire system, such as rootkits and firmware viruses. Added to that, the traditional BIOS operates in 16-bit real mode, a constraint that limits the amount of code and memory it can address. This limitation hinders the performance and complexity of firmware, making it less suitable for modern computing needs [31]. Additionally, BIOS relies on the Master Boot Record (MBR) partitioning scheme, which supports a maximum disk size of 2 terabytes and allows only four primary partitions [38][93]. This constraint has become a significant drawback as storage capacities have increased. Furthermore, the traditional BIOS has limited flexibility and is challenging to update or extend. This inflexibility restricts the ability to support new hardware and technologies efficiently [13][1].

1.2 Modern BIOS and UEFI

1.2.1 Transition from traditional BIOS to UEFI (Unified Extensible Firmware Interface)

All the limitations listed earlier caused a transition to a more modern firmware interface, designed to address the shortcomings of the traditional BIOS. This section delves into the historical context of this shift, the driving factors behind it, and the advantages UEFI offers over the traditional BIOS.

The development of UEFI began in the mid-1990s as part of the Intel Boot Initiative, which aimed to modernize the boot process and overcome the limitations of the traditional BIOS. By 2005, the Unified EFI Forum, a consortium of technology companies including Intel, AMD, and Microsoft, had formalized the UEFI specification [38]. UEFI was designed to address the shortcomings of the traditional BIOS, providing several key improvements.



Figure 1.3: The UEFI logo (public domain, 2010)

One of the most significant advancements of UEFI is its support for 32-bit and 64-bit modes, allowing it to address more memory and run more complex firmware programs. This capability enables UEFI to handle the increased demands of modern hardware and software [31][95]. Additionally, UEFI uses the GUID Partition Table (GPT) instead of the MBR, supporting disks larger than 2 terabytes and allowing for a nearly unlimited number of partitions [32][93].

Improved boot performance is another driving factor. UEFI provides faster boot times compared to the traditional BIOS, thanks to its efficient hardware and software initialization processes. This improvement is particularly beneficial for systems with complex hardware configurations, where quick boot times are essential [31]. UEFI's modular architecture makes it more extensible and easier to update compared to the traditional BIOS. This design allows for the addition of drivers, applications, and other components without requiring a complete firmware overhaul, providing greater flexibility and adaptability to new technologies [1]. UEFI also includes enhanced security features such as *Secure Boot*, which ensures that only trusted software can be executed during the boot process, thereby protecting the system from unauthorized modifications and malware [13][22].

The industry-wide support and standardization of UEFI have accelerated its adoption across various platforms and devices. Major industry players, including Intel, AMD, and Microsoft, have adopted UEFI as the new standard for firmware interfaces, ensuring broad compatibility and interoperability [38].

1.2.2 An other way with *coreboot*

While UEFI has become the dominant firmware interface for modern computing systems, it is not without its critics. Some of the primary concerns about UEFI include its complexity, potential security vulnerabilities, and the degree of control it provides to hardware manufacturers over the boot process. Originally known as LinuxBIOS, *coreboot*, is a free firmware project initiated in 1999 by Ron Minnich and his team at the Los Alamos National Laboratory. The project's primary goal was to create a fast, lightweight, and flexible firmware solution that could initialize hardware and boot operating systems quickly, while remaining transparent and auditable[87]. As an alternative to UEFI, *coreboot* offers a different approach to firmware that aims to address some of these concerns and continue the evolution of BIOS.

One of the main advantages of *coreboot* over UEFI is its simplicity, as it is designed to perform only the minimal tasks required to initialize hardware and pass control to a payload, such as a bootloader or operating system kernel. This minimalist approach reduces the attack surface and potential for security vulnerabilities, as there is less code that could be exploited by malicious actors [92]. Another significant benefit of *coreboot* is its libre nature. Unlike UEFI, which is controlled by a consortium of hardware and software vendors, *coreboot*'s source code is freely available and can be audited, modified, and improved by anyone. This transparency ensures that security researchers and developers can review the code for potential vulnerabilities and contribute to its improvement, fostering a community-driven approach to firmware development[87]. This project also supports a wide range of bootloaders, called payloads, allowing users to customize their boot process to suit their specific needs. Popular payloads include SeaBIOS, which provides legacy BIOS compatibility, and Tianocore, which offers UEFI functionality within the *coreboot* framework. This flexibility allows *coreboot* to be used in a variety of environments, from embedded systems to high-performance servers [86].



Figure 1.4: The *coreboot* logo, by Konsult Stuge & coresystems (coreboot logo license, 2008)

Despite its advantages, *coreboot* is not without its challenges. The project relies heavily on community contributions, and support for new hardware often lags behind that of UEFI. Additionally, the minimalist design of *coreboot* means that some advanced features provided by UEFI are not available by default. However, the *coreboot* community continues to work on adding new features and improving compatibility with modern hardware or security issues [73]. For example, it provides a *verified boot* function, allowing to prevent rootkits and other attacks based on firmware modifications [85]. However, it's important to note that *coreboot* is not entirely free in all aspects. Many modern processors and chipsets require *proprietary blobs*, short for *Binary Large Object*, which is a collection of binary data stored as a single entity. These blobs are necessary for *coreboot* to function correctly on a wide range of hardware, but they compromise the goal of having a fully free firmware one day [67], since these blobs are used for certain functionalities such as memory initialization and hardware management.



Figure 1.5: The *GNU Boot* logo, by Jason Self (CC0, 2020)

To address these concerns, the GNU Project has developed GNU Boot, a fully free distribution of firmware, including *coreboot*, that aims to be entirely free by avoiding the use of proprietary binary blobs. GNU Boot is committed to using only free software for all aspects of firmware, making it a preferred choice for users and organizations that prioritize software freedom and transparency [68].

1.3 Shift in firmware responsibilities

Initially, the BIOS's primary function was to perform the POST, a basic diagnostic testing process to check the system's hardware components and ensure they were functioning correctly. This included verifying the CPU, memory, and essential peripherals before passing control to the operating system's bootloader. This process was relatively simple, given the limited capabilities and straightforward architecture of early computer systems [13]. As computer systems advanced, particularly with the advent of more sophisticated memory technologies, the role of firmware expanded significantly. Modern memory modules operate at much higher speeds and capacities than their predecessors, requiring precise configuration to ensure stability and optimal performance. Firmware now plays a critical role in managing the memory controller, which is responsible for regulating data flow between the processor and memory modules. This includes configuring memory frequencies, voltage levels, and timing parameters to match the specifications of the installed memory [38][9]. Beyond memory management, firmware responsibilities have broadened to encompass a wide range of system-critical tasks. One key area is power management, where firmware is responsible for optimizing energy consumption across various components of the system. Efficient power management is essential not only for extending battery life in portable devices but also for reducing thermal output and ensuring system longevity in desktop and server environments. Moreover, modern firmware takes on significant roles in hardware initialization and configuration, which were traditionally handled by the operating system. For example, the initialization of USB controllers, network interfaces, and storage devices is now often managed by the firmware during the early stages of the boot process. This shift ensures that the operating system can seamlessly interact with hardware from the moment it takes control, reducing boot times and improving overall system reliability [38]. Security has also become a paramount concern for modern firmware. UEFI (Unified Extensible Firmware Interface), which has largely replaced traditional BIOS in modern systems, includes features which prevents unauthorized or malicious software from loading during the boot process. This helps protect the system from rootkits and other low-level malware that could compromise the integrity of the operating system before it even starts [38]. In the context of performance tuning, firmware sometimes also plays a key role in enabling and managing overclocking, particularly for the memory subsystem. By allowing adjustments to memory frequencies, voltages, and timings, firmware provides tools for enthusiasts to push their systems beyond default limits. At the same time, it includes safeguards to manage the risks of instability and hardware damage, balancing performance gains with system reliability [13].

In summary, the evolution of firmware from simple hardware initialization routines to complex management systems reflects the increasing sophistication of modern computer architectures. Firmware is now a critical layer that not only ensures the correct functioning of hardware components but also optimizes performance, manages power consumption, and enhances system security, making it an indispensable part of contemporary computing.

This document will focus on *coreboot* during the next parts to study how modern firmware interact with hardware and also as a basis for improvements.

Chapter 2

Characteristics of ASUS KGPE-D16 mainboard




Figure 2.1: The KGPE-D16 (CC BY-SA 4.0, 2021)

2.1 Overview of ASUS KGPE-D16 hardware

The ASUS KGPE-D16 server mainboard is a dual-socket motherboard designed to support AMD Family 10h/15h series processors. Released in 2009, this mainboard was later awarded the *Respects Your Freedom* (RYF) certification in March 2017, underscoring its commitment to fully free software compatibility [41]. Indeed, this mainboard can be operated with a fully free firmware such as GNU Boot [69].

This mainboard is equipped with robust hardware components designed to meet the demands of high-performance computing. It features 16 DDR3 DIMM slots, capable of supporting up to 256GB of memory, although certain configurations may be limited to 192GB, with some reports suggesting the potential to support 256GB under specific conditions. In terms of expandability, the KGPE-D16 includes multiple PCIe slots, with five physical slots available, although only four can be used simultaneously due to slot sharing. For storage, the mainboard provides several SATA ports. Networking capabilities are enhanced by integrated dual gigabit Ethernet ports, which provide high-speed connectivity essential for data-intensive tasks and network communication [15]. Additionally, the board is equipped with various peripheral interfaces, including USB ports, audio outputs, and other I/O ports, ensuring compatibility with a wide range of external devices.




Figure 2.2: Basic schematics of the ASUS KGPE-D16 Mainboard, ASUS (2011)

The physical layout of the ASUS KGPE-D16 is meticulously designed to optimize airflow, cooling, and power distribution. All of this is critical for maintaining system stability, particularly under heavy computational loads, as this board was designed for server operations. In particular, key components such as the CPU sockets, memory slots, and PCIe slots are strategically positioned.



Figure 2.3: The KGPE-D16, viewed from the top (CC BY-SA 4.0, 2024)

2.2 Chipset

Before diving into the specific components, it is essential to understand the roles of the northbridge and southbridge in traditional motherboard architecture. These chipsets historically managed communication between the CPU and other critical components of the system [3].

The northbridge is a chipset on the motherboard that traditionally manages high-speed communication between the CPU, memory (RAM), and graphics card (if applicable). It serves as a hub for data that needs to move quickly between these components. On the ASUS KGPE-D16, the functions typically associated with the northbridge are divided between the CPU's internal northbridge and an external SR5690 northbridge chip. The SR5690 specifically acts as a translator and switch, handling the HyperTransport interface, a high-speed communication protocol used by AMD processors, and converting it to ALink and PCIe interfaces, which are crucial for connecting peripherals like graphics cards [11]. Additionally, the northbridge on the KGPE-D16 incorporates the IOMMU (Input-Output Memory Management Unit), which is crucial for ensuring secure and efficient memory access by I/O devices. The IOMMU allows for the virtualization of memory addresses, providing device isolation and preventing unauthorized memory access, which is particularly important in environments that run multiple virtual machines [3][126].

The southbridge, on the other hand, is responsible for handling lower-speed, peripheral interfaces such as the PCI, USB, and IDE/SATA connections, as well as managing onboard audio and network controllers. On the KGPE-D16, these functions are managed by the SP5100 southbridge chip, which integrates several critical functions including the LPC bridge, SATA controllers, and other essential I/O operations [3][129]. It is essentially an ALink bus controller and includes the hardware interrupt controller, the IOAPIC. Interrupts from peripheral always pass through the northbridge (fig. 2.4), since it translates ALink to HyperTransport for the CPUs and contains the IOMMU [11].




Figure 2.4: Functional diagram presenting the IOAPIC function of the SP5100, ASUS (2011)

In addition to the northbridge and southbridge, the KGPE-D16 also contains specialized chips for managing input/output operations and system health monitoring. The WINBOND W83667HG-A Super I/O chip handles traditional I/O functions such as legacy serial and parallel ports, keyboard, and mouse interfaces, but also the SPI chip that contains the firmware [133]. Meanwhile, the Nuvoton W83795G/ADG Hardware Monitor oversees the systems health by monitoring temperatures, voltages, and fan speeds, ensuring that the system operates within safe parameters [77]. On the KGPE-D16, access to the Super I/O from a CPU core is done through the SR5690, then the SP5100, as that can be observed on the functional diagram of the chipset (fig. 2.5) [11].

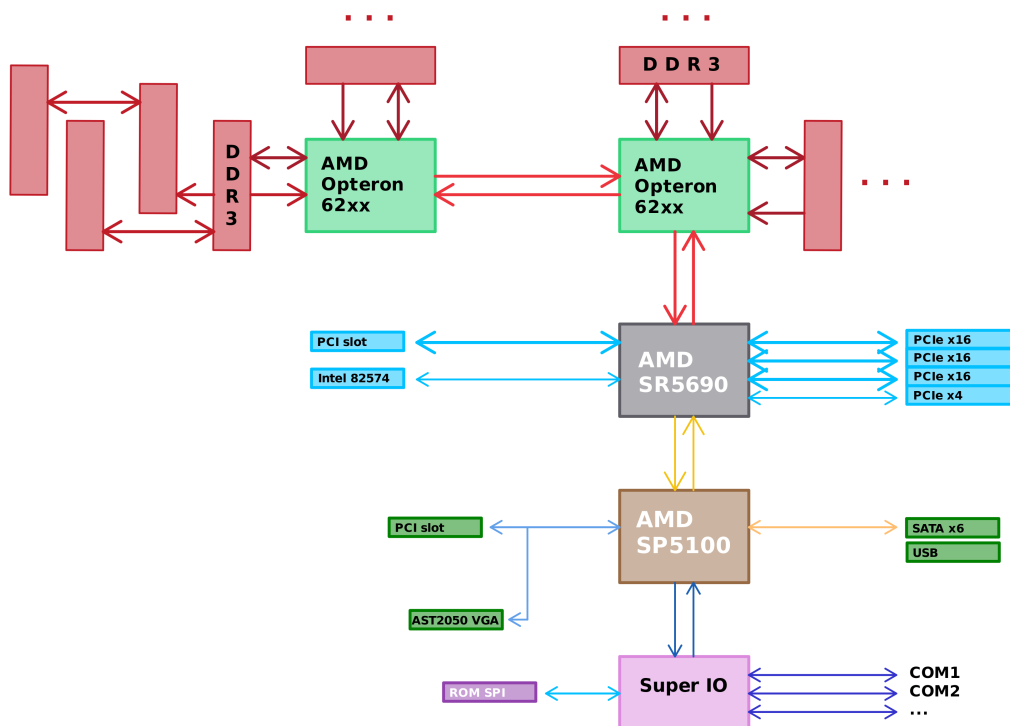


Figure 2.5: Functional diagram of the KGPE-D16 chipset (CC BY-SA 4.0, 2024)

2.3 Processors

The ASUS KGPE-D16 supports AMD Family 10h processors, but it is important to note that Vikings, a known vendor for libre-software-compatible hardware, does not recommend using the Opteron 6100 series due to the lack of IOMMU support, which is critical for security. Fortunately, AMD Family 15h processors are also supported. However, the Opteron 6300 series, while supported, requires proprietary microcode updates for stability, IOMMU functionality, and fixes for specific vulnerabilities, including a gain-root- via-NMI exploit. The Opteron 6200 series does not suffer from these problems and works properly without any proprietary microcode update needed [109].




Figure 2.6: Annotated photography of an Opteron 6200 series CPU (2024), from a photography by AMD Inc. (2008)

The Opteron 6200 series, part of the Bulldozer microarchitecture, was designed to target high-performance server applications. These processors feature 16 cores, organized into 8 Bulldozer modules, with each module containing two integer cores that shared resources like the floating-point unit (FPU) and L2 cache (fig. 2.6, 2.7) [7][94]. The architecture of the Opteron 6200 series is built around AMD's Bulldozer core design, which uses Clustered Multithreading (CMT) to maximize resource utilization. This is a technique where each processor module contains two integer cores that share certain resources like the floating-point unit (FPU), L2 cache, and instruction fetch/decode stages. Unlike traditional multithreading, where each core handles multiple threads, CMT allows two cores to share resources to improve parallel processing efficiency. This approach aims to balance performance and resource usage, particularly in multi-threaded workloads, though it can lead to some performance trade-offs in single-threaded tasks. In the Opteron 6272, the processor consists of eight modules, effectively creating 16 integer cores. Due to the CMT architecture, each Opteron 6272 chip functions as two CPUs within a single processor, each with its own set of cores, L2 caches, and shared L3 cache. Here, one CPU is made by four modules, each module in it sharing certain components, such as the FPU and L2 cache, between two integer cores. The L3 cache is shared across these modules. HyperTransport links provide high-speed communication between the two sockets of the KGPE-D16. Shared L3 cache and direct memory access are provided by each socket [7][50].

This architecture also integrates a quad-channel DDR3 memory controller directly into the processor die, which facilitates high bandwidth and low latency access to memory. This memory controller supports DDR3 memory speeds up to 1600 MHz and connects directly to the memory modules via the memory bus. By integrating the memory controller into the processor, the Opteron 6200 series reduces memory access latency, enhancing overall performance [7][6]. It is interesting to note that Opterons incorporate the internal northbridge that we cited previously. The traditional northbridge functions, such as memory controller and PCIe interface management, are partially integrated into the processor. This integration reduces the distance data must travel between the CPU and memory, decreasing latency and improving performance, particularly in memory-intensive applications [7].




Figure 2.7: Functional diagram of an Opteron 6200 package (CC BY-SA 4.0, 2024)

Power efficiency was a key focus in the design of the Opteron 6200 series. Despite the high core count, the processor includes several power management features, such as Dynamic Power Management (DPM) and Turbo Core technology. These features allow the processor to adjust power usage based on workload demands, balancing performance with energy consumption. However, the Bulldozer architecture's focus on high clock speeds and multi-threaded performance resulted in higher power consumption compared to competing architectures [94]. A special model of the series, called *high efficiency* models, solve a bit this problem by proposing a bit less performant processor but with a power consumption divided by a factor from 1.5 to 2.0 in some cases.

The processor connected to the I/O hub is known as the Bootstrap Processor (BSP). The BSP is responsible for starting up the system by executing the initial firmware code from the reset vector, a specific memory address where the CPU begins execution after a reset [4]. Core 0 of the BSP, called the Bootstrap Core (BSC), initiates this process. During early initialization, the BSP performs several critical tasks, such as memory initialization, and bringing other CPU cores online. One of its duties is storing Built-In Self-Test (BIST) information, which involves checking the integrity of the processor's internal components to ensure they are functioning correctly. The BSP also determines the type of reset that has occurred whether it's a cold reset, which happens when the system is powered on from an off state, or a warm reset, which is a restart without turning off the power. Identifying the reset type is crucial for deciding which initialization procedures need to be executed [4][9].

2.4 Baseboard Management Controller

The Baseboard Management Controller (BMC) on the KGPE-D16 motherboard, specifically the ASpeed AST2050, plays a role in the server's architecture by managing out-of-band communication and control of the hardware. The AST2050 is based on an ARM926EJ-S processor, a low-power 32-bit ARM architecture designed for embedded systems [100]. This architecture is well-suited for BMCs due to its efficiency and capability to handle multiple management tasks concurrently without significant resource demands from the main system.

The AST2050 features several key components that contribute to its functionality. It includes an integrated VGA controller, which enables remote graphical management through KVM-over-IP (Keyboard, Video, Mouse), a critical feature for administrators who need to interact with the system remotely, including BIOS updates and troubleshooting [98]. Additionally, the AST2050 integrates a dedicated memory controller, which supports up to 256MB of DDR2 RAM. This allows it to handle complex tasks and maintain responsiveness during management

operations [34]. The BMC also features a network interface controller (NIC) dedicated to management traffic, ensuring that remote management does not interfere with the primary network traffic of the server. This separation is vital for maintaining secure and uninterrupted system management, especially in environments where uptime is critical [101]. Another important architectural aspect of the AST2050 is its support for multiple I/O interfaces, including I2C, GPIO, UART, and USB, which allow it to interface with various sensors and peripherals on the motherboard [102]. This versatility enables comprehensive monitoring of hardware health, such as temperature sensors, fan speeds, and power supplies, all of which can be managed and configured through the BMC.

When combined with OpenBMC [127], a libre firmware that can be run on the AST2050 thanks to Raptor Engineering [88], the architecture of the BMC becomes even more powerful. OpenBMC takes advantage of the AST2050's architecture, providing a flexible and customizable environment that can be tailored to specific use cases. This includes adding or modifying features related to security, logging, and network management, all within the BMC's ARM architecture framework [55].

Chapter 3

Key components in modern firmware

3.1 General structure of coreboot

The firmware of the ASUS KGPE-D16 is crucial in ensuring the proper functioning and optimization of the mainboard's hardware components. For this to be done efficiently, *coreboot* is organized in different stages (fig. 3.1) [85].




Figure 3.1: *coreboot*'s stages timeline, by *coreboot* project (CC BY-SA 4.0, 2009)

Being a complex project with ambitious goals, *coreboot* decided early on to establish a file-system-based architecture for its images (also called ROMs). This special file-system is CBFS (which stands for coreboot file system). The CBFS architecture consists of a binary image that can be interpreted as a physical disk, referred to here as ROM. A number of independent components, each with a header added to the data, are located within the ROM. The components are nominally arranged sequentially, although they are aligned along a predefined boundary (fig. 3.2).

Each stage is compiled as a separate binary and inserted into the CBFS with custom compression. The bootblock stage is usually not compressed, while the ramstage and the payload are compressed with LZMA. Each stage loads the next stage at a given address (possibly decompressing it in the process).

Some stages are relocatable and can be placed anywhere in the RAM. These stages are typically cached in the CBMEM for faster loading times during wake-up. The CBMEM is a specific memory area used by the *coreboot* firmware to store important data structures and logs during the boot process. This area is typically allocated in the system's RAM and is used to store various types of runtime information that it might need to reference after the initial boot stages.

In general, *coreboot* manages main memory through a structured memory map (fig. 3.1), allocating specific address ranges for various hardware functions and system operations. The first 640KB of memory space is typically unused by coreboot due to historical reasons. Graphics-related operations use the VGA address range and the text mode address ranges. It also reserves the higher for operating system use, ensuring that critical system components like the IOAPIC and TPM registers have dedicated address spaces. This structured approach helps maintain system stability and compatibility across different platforms and allows for a reset vector fixed at an address ($0xFFFFF0$), regardless of the ROM size. Payloads are typically loaded into high memory, above the

reserved areas for hardware components and system resources. The exact memory location can vary depending on the system's configuration, but generally, payloads are placed in a region of memory that does not conflict with the firmware code or the reserved memory map areas, such as the ROM mapping ranges. This placement ensures that payloads have sufficient space to execute without interfering with other critical memory regions allocated [25].

0x00000 - 0x9FFFF	Low memory (first 640KB). Never used.
0xA0000 - 0xAFFFF	VGA graphics address range.
0xB0000 - 0xB7FFF	Monochrome text mode address range. Few motherboards use it, but the KGPE-D16 does.
0xB8000 - 0xBFFFF	Text mode address range.
0xFEC00000	IOAPIC address.
0xFED44000 - 0xFED4FFFF	Address range for TPM registers.
0xFF000000 - 0xFFFFFFFF	16 MB ROM mapping address range.
0xFF800000 - 0xFFFFFFFF	8 MB ROM mapping address range.
0xFFC00000 - 0xFFFFFFFF	4 MB ROM mapping address range.
0xFEC00000 - DEVICE MEM HIGH	Reserved area for OS use.

Table 3.1: *coreboot* memory map

3.1.1 Bootblock stage

The bootblock is the first stage executed after the CPU reset. The beginning of this stage is written in assembly language, and its main task is to set everything up for a C environment. The rest, of course, is written in C. This stage occupies the last 20k (fig. 3.2) of the image and within it is a main header containing information about the ROM, including the size, component alignment, and the offset of the start of the first CBFS component. This block is a mandatory component as it also contains the entry point of the firmware.

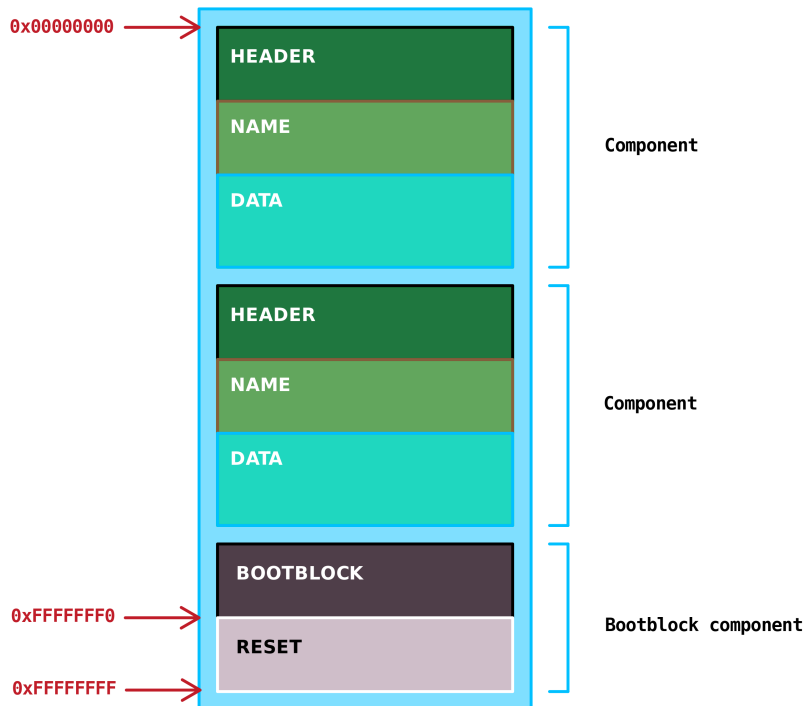


Figure 3.2: *coreboot* ROM architecture (CC BY-SA 4.0, 2024)

Upon startup, the first responsibility of the bootblock is to execute the code from the reset vector located at the conventional reset vector in 16-bit real mode. This code is specific to the processor architecture and, for our board, is stored in the architecture-specific sources for x86 within *coreboot*. The entry point into *coreboot* code

is defined in two files in the `src/cpu/x86/16bit/` directory: `reset16.inc` and `entry16.inc`. The first file serves as a jump to the `_start16bit` procedure defined in the second. Due to space constraints this function must remain below the 1MB address space because the IOMMU has not yet been configured to allow anything else.

During this early initialization, the Bootstrap Core (BSC) performs several critical tasks while the other cores remain dormant. These tasks include saving the results (and displaying them if necessary) of the Built-in Self-Test (BIST), formerly known as POST; invalidating the TLB to prevent any address translation errors; determining the type of reset (e.g., cold start or warm start); creating and loading an empty Interrupt Descriptor Table (IDT) to prevent the use of "legacy" interrupts from real mode until protected mode is reached. In practice, this means that at the slightest exception, the BSC will halt. The code then switches to 32-bit protected mode by mapping the first 4 GB of address space for code and data, and finally jumps to the 32-bit reset code labeled `_protected_start`.

Once in protected mode, which constitutes the "normal" operating mode for the processor, the next step is to set up the execution environment. To achieve this, the code contained in `src/cpu/x86/32bit/entry32.inc`, followed by `src/cpu/x86/64bit/entry64.inc`, and finally `src/arch/x86/bootblock_crt0.S`, establishes a temporary stack, transitions to long mode (64-bit addressing) with paging enabled, and sets up a proper exception vector table. The execution then jumps to chipset-specific code via the `bootblock_pre_c_entry` procedure. Once these steps are completed, the bootblock has a minimal C environment. The procedure now involves allocating memory for the BSS, and decompressing and loading the next stage.

The jump to `_bootblock_pre_entry` leads to the code files `src/soc/amd/common/block/cpu/car/cache_as_ram.S` and `src/vendorcode/amd/agesa/f15tn/gcccar.inc`, which are specific to AMD chipsets. It's worth noting that these files were developed by AMD's engineers as part of the *AGESA* project. The operations performed at this stage are related to pre-RAM memory initialization. All cores of all processors (up to a limit of 64 cores) are started. The *Cache-As-Ram* is configured using the Memory-type range registers. These registers allow the specification of a specific configuration for a given memory area [9]. In this case, the area that should correspond to physical memory is mapped to the cache, while other areas, such as PCI or other bus zones, are configured accordingly. A specific stack is set up for each core of each processor (within the arbitrary limit of 64 cores and 7 nodes, meaning 7 Core 0s). Core 0s receive 16KB, while the Bootstrap Core (BSC) gets 64KB. The other cores receive 4KB each. All cores except the BSC are halted and will restart during the romstage. Finally, the execution jumps to the entry point of the *bootblock* written in C, labeled `bootblock_c_entry`. This entry point is located in `src/soc/amd/stoneyridge/bootblock/bootblock.c` and is specific to AMD processors. It is the first C routine executed, and its role is to verify that the current processor is indeed the BSC, allowing the function `bootblock_main_with_basetime` to be called exclusively by the BSC.

We are now in the file `src/lib/bootblock.c`, written by Google's team, and entering the `bootblock_main_with_basetime` function, which immediately calls `bootblock_main_with_timestamp`. At this stage, the goal is to start the romstage, but a few more tasks need to be completed.

The `bootblock_soc_early_init` function is called to initialize the I2C bus of the southbridge. The `bootblock_fch_early_init` function is invoked to initialize the SPI buses (including the one for the ROM) and the serial and "legacy" buses of the southbridge. The CMOS clock is then initialized, followed by the pre-initialization of the serial console. The code then calls the `bootblock_mainboard_init` function, which enters, for the first time, the files specific to the ASUS KGPE-D16 motherboard: `src/mainboard/ASUS/kgpe-d16/bootblock.c`. This code performs the northbridge initialization via the `bootblock_northbridge_init` function found in `src/northbridge/amd/amdfam10/bootblock.c`. This involves locating the HyperTransport bus and enabling the discovery of devices connected to it (e.g., processors). The southbridge is initialized using the `bootblock_southbridge_init` function from `src/southbridge/amd/sb700/bootblock.c`. This function, largely programmed by Timothy Pearson from Raptor Engineering, who performed the first coreboot port for the ASUS KGPE-D16, finalizes the activation of the SPI bus and the connection to the ROM memory via SuperIO. The state of a recovery jumper is then checked (this jumper is intended to reset the CMOS content, although it is not fully functional at the moment, as indicated by the `FIXME` comment in the code). Control then returns to `bootblock_main` in `src/lib/bootblock.c`.

At this point, everything is ready to enter the romstage. *coreboot* has successfully started and can now continue its execution by calling the `run_romstage` function from `src/lib/prog_loaders.c`. This function begins by locating the corresponding segment in the ROM via the southbridge and SPI bus using `prog_locate`, which utilizes the SPI driver in `src/drivers/cbfs_spi.c`. The contents of the romstage are then copied into the

cache-as-ram by `cbfs_prog_stage_load`. Finally, the `prog_run` function transitions to the romstage after switching back to 32-bit mode.

3.1.2 Romstage

The *romstage* in *coreboot* serves the critical function of early initialization of peripherals, particularly system memory. This stage is crucial for setting up the necessary components for the platform's operation, ensuring that everything is in place for subsequent stages of the boot process. During this phase, *coreboot* configures the Advanced Programmable Interrupt Controller (APIC), which is responsible for correctly handling interrupts across multiple CPUs, especially in systems using Symmetric Multiprocessing (SMP). This includes setting up the Local APIC on each processor and the IOAPIC, part of the southbridge, to ensure that interrupts from peripherals are routed to the appropriate CPUs. Additionally, the firmware configures the HyperTransport (HT) technology, a high-speed communication protocol that facilitates data exchange between the processor and the northbridge, ensuring smooth data flow between these components.

The *romstage* begins with a call to the `_start` function, defined in `src/cpu/x86/32bit/entry32.inc` via `src/arch/x86/assembly_entry.S`. We then enter the `cache_as_ram_setup` procedure, written in assembly language, located in `src/cpu/amd/car/cache_as_ram.inc`. This procedure configures the cache to load the future *ramstage* and initialize memory based on the number of processors and cores present. Once this is completed, the code calls `cache_as_ram_main` in `src/mainboard/asus/kgpe-d16/romstage.c`, which serves as the main function of the *romstage*. In the `cache_as_ram_main` function, after reducing the speed of the HyperTransport bus, only the Bootstrap Core (BSC) initializes the spinlocks for the serial console, the CMOS storage memory (used for saving parameters), and the ROM. At this point, the HyperTransport bus is enumerated, and the PCI bridges are temporarily disabled. The port 0x80 of the southbridge, used for motherboard debugging with *Post Codes*, is also initialized. These codes indicate the status of the boot process and can be displayed using special PCI cards connected to the system. The SuperIO is then initialized to activate the serial port, allowing the serial console to follow *coreboot's* progress in real-time. If everything proceeds as expected, the code 0x30 is sent, and the boot process continues.

If the result of the Built-in Self-Test (BIST), saved during the *bootblock*, shows no anomalies, all cores of all nodes are configured, and they are placed back into sleep mode (except for the Core 0s). If everything goes well, the code 0x32 is sent, and the process continues. Using the `enable_sr5650_dev8` function, the southbridges P2P bridge is activated. Additionally, a check is performed to ensure that the number of physical processors detected does not exceed the number of sockets available on the board. If any issues were detected during the BIST, the machine will halt, and the error will be displayed on the console. Otherwise, the process continues, and the default hardware information table is constructed, and the microcode of the physical processors is updated if necessary. If everything proceeds correctly, the code 0x33 and then 0x34 is sent, and the process continues. The information about the physical processors is retrieved using `amd_ht_init`, and communication between the two sockets is configured via `amd_ht_fixup`. This process includes disabling any defective HT links (one per socket in this AMD Family 15h chipset). If everything is working as expected, the code 0x35 is sent, and the boot process continues. With the `finalize_node_setup` function, the PCI bus is initialized, and a mapping is created (`setup_mb_resource_map`). If all goes well, the code 0x36 is sent. This is done in parallel across all Core 0s, so the system waits for all cores to finish using the `wait_all_core0_started` function. The communication between the northbridge and southbridge is prepared using `sr5650_early_setup` and `sb7xx_51xx_early_setup`, followed by the activation of all cores on all nodes, with the system waiting for all cores to be fully initialized. If everything is successful, the code 0x38 is sent.

At this point, the timer is activated, and a warm reset is performed via the `soft_reset` function to validate all configuration changes to the HT, PCI buses, and voltage/power settings of the processors and buses. This results in a system reboot, passing again through the *bootblock*, but much faster this time since the system recognizes the warm reset condition. Once this reboot is complete, the HyperTransport bus is reconfigured into isochronous mode (switching from asynchronous mode), finalizing the configuration process.

Memory training and optimization are also key functions of the firmware during the *romstage*. This process involves adjusting memory settings, such as timings, frequencies, and voltages, to ensure that the installed memory modules operate efficiently and stably. This step is crucial for achieving optimal performance, especially when

dealing with large amounts of RAM and many CPU cores, as supported by the KGPE-D16. We'll see that in detail during the next chapter.

After memory initialization, the process returns to the `cache_as_ram_main` function, where a memory test is performed. This involves writing predefined values to specific memory locations and then verifying that the values can be read back correctly. If everything passes successfully, the CBMEM is initialized and one sends code `0x41`. At this point, the configuration of the PCI bus is prepared, which will be completed during the ramstage by configuring the PCI bridges. The system then exits `cache_as_ram_main` and returns to `cache_as_ram_setup` to finalize the process.

`coreboot` then transitions to the next stage, known as the postcar stage, where it exits the cache-as-RAM mode and begins using physical RAM.

3.1.3 Ramstage

The ramstage performs the general initialization of all peripherals, including the initialization of PCI devices, on-chip devices, the TPM (if not done by verstage), graphics (optional), and the CPU (setting up the System Management Mode). After this initialization, tables are written to inform the payload or operating system about the existence and current state of the hardware. These tables include ACPI tables (specific to x86), SMBIOS tables (specific to x86), coreboot tables, and updates to the device tree (specific to ARM). Additionally, the ramstage locks down the hardware and firmware by applying write protection to boot media, locking security-related registers, and locking SMM (specific to x86) [85]. Effective resource allocation is essential for system stability, particularly in complex configurations involving multiple CPUs and peripherals. This stage manages initial resource allocation, resolving any conflicts between hardware components to prevent resource contention and ensure smooth operation and security, which is a major concern in modern systems. This includes support for IOMMU, which is crucial for preventing unauthorized direct memory access (DMA) attacks, particularly in virtualized environments (however there are still vulnerabilities that can be exploited, such as sub-page or IOTLB-based attacks or even configuration weaknesses [72][70]).

3.1.3.1 Advanced Configuration and Power Interface

The Advanced Configuration and Power Interface (ACPI) is a critical component of modern computing systems, providing an open standard for device configuration and power management by the operating system (OS). Developed in 1996 by Intel, Microsoft, and Toshiba, ACPI replaced the older Advanced Power Management (APM) standard with more advanced and flexible power management capabilities [29]. At its core, ACPI is implemented through a series of data structures and executable code known as ACPI tables, which are provided by the system firmware and interpreted by the OS. These tables describe various aspects of the system, including hardware resources, device power states, and thermal zones. The ACPI Specification outlines these structures and provides the necessary standardization for interoperability across different platforms and operating systems [47]. These tables are used by the OS to perform low-level tasks, including managing power states of the CPU, controlling the voltage and frequency scaling (also known as Dynamic Voltage and Frequency Scaling, or DVFS), and coordinating power delivery to peripherals.

The ACPI Component Architecture (ACPIA) is the reference implementation of ACPI, providing a common codebase that can be used by OS developers to integrate ACPI support. ACPIA includes tools and libraries that allow for the parsing and execution of ACPI Machine Language (AML) code, which is embedded within the ACPI tables [84]. One of the key tools in ACPIA is the Intel ACPI Source Language (IASL) compiler, which converts ACPI Source Language (ASL) code into AML bytecode, allowing firmware developers to write custom ACPI methods [27]. The triggering of ACPI events is managed through a combination of hardware signals and software routines. For example, when a user presses the power button on a system, an ACPI event is generated, which is then handled by the OS. This event might trigger the system to enter a low-power state, such as sleep or hibernation, depending on the configuration provided by the ACPI tables [47]. These power states are defined in the ACPI specification, with global states (G0 to G3) representing different levels of system power consumption, and device states (D0 to D3) representing individual device power levels.

The ASUS KGPE-D16 mainboard, which is designed for server and high-performance computing environments, needs ACPI for managing its power distribution across multiple CPUs and attached peripherals. ACPI is integral in

controlling the power states of various components, thereby optimizing performance and energy use. Additionally, the firmware on the KGPE-D16 uses ACPI tables to manage system temperature and fan speed, ensuring reliable operation under heavy workloads [15].

3.1.3.2 System Management Mode

System Management Mode (SMM) is a highly privileged operating mode provided by x86 processors for handling system-level functions such as power management, hardware control, and other critical tasks that are to be isolated from the OS and applications. Introduced by Intel, SMM operates in an environment separate from the main operating system, offering a controlled space for executing sensitive operations [59].

SMM is triggered by a System Management Interrupt (SMI), which is a non-maskable interrupt that causes the CPU to save its current state and switch to executing code stored in a protected area of memory called System Management RAM (SMRAM). SMRAM is a specialized memory region that is isolated from the rest of the system, making it inaccessible to the OS and preventing tampering or interference from other software [49]. Within SMM, the firmware can execute various low-level functions that require direct hardware control or need to be protected from the OS. This includes tasks such as thermal management, where the system monitors CPU temperature and adjusts performance or power levels to prevent overheating, as well as power management routines that enable efficient energy usage by adjusting power states based on system activity [56]. One of the critical security features of SMM is its role in managing firmware updates and handling system-level security events. Because SMM operates in a privileged mode that is isolated from the OS, it can apply firmware updates and could respond to security threats without being affected by potentially compromised system software [35]. However, the high privilege level and isolation of SMM also present significant security challenges. If an attacker can compromise SMM, they gain full control over the system, bypassing all security measures implemented by the OS [64]. Also, with a proprietary firmware, it means that this code with a very high privilege level cannot be audited at all, nor even replaced.

The ASUS KGPE-D16 mainboard needs SMM to perform critical management tasks that need to be done in parallel from the operating system. For example, SMM is used to monitor and manage system health by responding to thermal events and adjusting power levels to maintain system stability. SMM operates independently of the main operating system, allowing it to perform sensitive tasks securely. *coreboot* supports SMM, but its implementation is typically minimal compared to traditional proprietary firmware. In *coreboot*, SMM initialization involves setting up the System Management Interrupt (SMI) handler and configuring System Management RAM (SMRAM), the memory region where SMM code executes [19]. The extent of SMM support in *coreboot* can vary significantly depending on the hardware platform and the specific requirements of the system. *coreboot*'s design philosophy emphasizes a lightweight and fast boot process, delegating more complex management tasks to payloads or the operating system itself [89].

One of the key challenges with implementing SMM in *coreboot* is ensuring that SMI handlers are configured correctly to manage necessary system tasks without compromising security or performance. *coreboot*'s approach to SMM is consistent with its overall goal of providing a streamlined and efficient firmware solution, leaving more intricate functionalities to be handled by subsequent software layers [76].

3.1.4 Payload

The payload is the software that executes after *coreboot* has completed its initialization tasks. It resides in the CBFS and is predetermined at compile time, with no option to choose it at runtime. The primary role of the payload is to load and hand control over to the operating system. In some cases, the payload itself can be a component of the operating system [85]. Examples of payloads are *GNU GRUB*, *SeaBIOS*, *memtest86+* or even sometimes the *Linux kernel* itself.

TianoCore, a free implementation of the UEFI (Unified Extensible Firmware Interface) specification is often used as a payload [103]. It provides a UEFI environment after *coreboot* has completed its initial hardware initialization. This allows the system to benefit from the advanced features of UEFI, such as a more flexible boot manager, enhanced features, and support for modern hardware. Indeed, UEFI, and by extension *TianoCore*, includes a driver model that allows hardware manufacturers to provide UEFI-compatible drivers. These drivers can be loaded at boot time, allowing the firmware to support a wide range of modern devices that *coreboot*, with its more minimalistic and custom-tailored approach, might not support out of the box. For example, GOP drivers are responsible for setting up the graphics hardware in UEFI environments. They replace the older VGA BIOS routines used in

legacy BIOS systems. With GOP drivers, the system can initialize the GPU and display a graphical interface even before the operating system loads [81]. Hardware manufacturers can distribute proprietary UEFI drivers as part of firmware updates, making it straightforward for end-users to install and use them. This is especially useful for specialized hardware that requires specific drivers not included in the free software community. It also gives hardware vendors more control over how their devices are initialized and used, which can be an advantage for vendors but is a freedom and user control limitation.

Payloads are then definitely important parts of the firmware.

3.2 AMD Platform Security Processor and Intel Management Engine

The AMD Platform Security Processor (PSP) and Intel Management Engine (ME) are embedded subsystems within AMD and Intel processors, respectively, that handle a range of security-related tasks independent of the main CPU. These subsystems are fundamental to the security architecture of modern computing platforms, providing functions such as secure boot, cryptographic key management, and remote system management [57]. The AMD PSP is based on an ARM Cortex-A5 processor and is responsible for several security functions, including the validation of firmware during boot (secure boot), management of Trusted Platform Module (TPM) functions, and handling cryptographic operations such as key generation and storage. The PSP operates independently of the main x86 cores, which allows it to execute security functions even when the main system is powered off or compromised by malware [57]. The PSP's isolated environment ensures that sensitive operations are protected from threats that could affect the main OS.

Similarly, the Intel Management Engine (ME) is a dedicated processor embedded within Intel chipsets that operates independently of the main CPU. The ME is a comprehensive subsystem that provides a variety of functions, including out-of-band system management, security enforcement, and support for Digital Rights Management (DRM) [28]. The ME's firmware runs on an isolated environment that allows it to perform these tasks securely, even when the system is powered off. This capability is crucial for enterprise environments where administrators need to perform remote diagnostics, updates, and security checks without relying on the main OS. Intel ME enforces Digital Rights Management (DRM) through a multifaceted approach leveraging its deeply embedded, hardware-based capabilities. At the core is the Protected Execution Environment (PEE), which operates independently from the main CPU and operating system. This isolation allows to privately manage cryptographic keys, certificates, and other sensitive data critical for DRM, which can be very problematic from a user freedom perspective [42]. By handling encryption and decryption processes within this protected environment, Intel ME ensures that DRM-protected content, such as video streams, remains secure and unreachable by the user, raising concerns about the control users have over their own devices [78]. Intel ME also plays a significant role in maintaining platform integrity through the secure boot process. During secure boot, Intel ME ensures that only digitally signed and authorized operating systems and applications are loaded, which can prevent users from installing alternative or modified software on their own hardware, further restricting their freedom [104]. This is further reinforced by Intel ME's remote attestation capabilities, where the systems state is reported to a remote server. This process verifies that only systems meeting specific security standards dictated by third parties are allowed to access DRM-protected content, potentially limiting users' control over their own devices [18]. Moreover, Intel ME supports High-bandwidth Digital Content Protection (HDCP), a technology that restricts how digital content is transmitted over interfaces like HDMI or DisplayPort. By enforcing HDCP, Intel ME ensures that protected digital content, such as high-definition video, is only transmitted to and displayed on authorized devices, effectively preventing users from freely using the content they have legally acquired [62][80]. Together, these features enable Intel ME to provide a comprehensive and robust DRM enforcement mechanism. However, this also means that users have less control over their own hardware and digital content, raising serious concerns about privacy, user autonomy, and the broader implications for freedom in computing [42][54].

Added to that, Intel ME has been a source of controversy due to its deep integration into the hardware and its potential to be exploited if vulnerabilities are discovered. Researchers have demonstrated ways to hack into the ME, potentially gaining control over a system even when it is powered off [45]. These concerns have led to calls for greater transparency and security measures around the ME and similar subsystems. When comparing Intel ME and AMD PSP, the primary difference lies in their scope and functionality. Intel ME offers more extensive remote management capabilities, making it a more comprehensive tool for enterprise environments, while AMD

PSP focuses more narrowly on core security tasks. Nonetheless, both play critical roles in ensuring the security and integrity of modern computing systems.

The ASUS KGPE-D16 mainboard does not include AMD PSP nor Intel ME.

Chapter 4

Memory initialization and training

4.1 Importance of DDR3 Memory Initialization

DDR3 (Double Data Rate Type 3) is a widely used type of SDRAM (Synchronous Dynamic Random-Access Memory) that offers significant performance improvements over its predecessors, DDR and DDR2. Key features of DDR3 include higher data rates, lower power consumption, and increased memory capacity, making it essential for high-performance computing environments [115]. One of the critical aspects of DDR3 is its internal architecture, which supports data rates ranging from 800 to 1600 Mbps and operates at a lower voltage of 1.5V. This enables faster data processing and more efficient power usage, crucial for modern applications that require high-speed memory access [66]. Additionally, DDR3 memory modules are available in larger capacities, allowing systems to handle larger datasets and more complex computing tasks [5]. However, the advanced features of DDR3 come with increased complexity in its initialization and operation. For example, DDR3 uses a fly-by topology (fig. 4.1) for routing the address, command, and clock signals. In this topology, signals are routed sequentially from the memory controller to each DRAM chip, reducing signal reflections and improving overall signal integrity. This design is essential for maintaining stability at the high speeds DDR3 operates at, but it also introduces timing challenges, such as timing skew, that must be carefully managed [52].

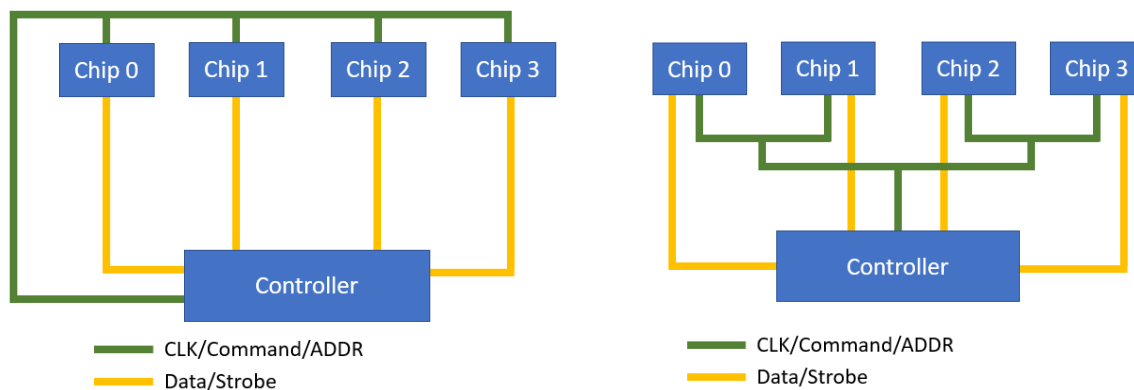


Figure 4.1: DDR3 fly-by *versus* T-topology (CC BY-SA 4.0, 2021)

Proper memory initialization ensures that the memory controller and the memory modules are correctly configured to work together, allowing for efficient data transfer and reliable operation. The initialization process involves setting various parameters, such as memory timings, voltages, and frequencies, which are critical for ensuring that the memory operates within its optimal range [66]. Failure to initialize DDR3 memory correctly can lead to several serious consequences, including system instability, data corruption, and reduced performance [99]. In the worst-case scenario, improper memory initialization can prevent the system from booting entirely, as the memory subsystem fails to function correctly. In the context of the ASUS KGPE-D16, a server motherboard designed for high-performance applications, proper DDR3 memory initialization is particularly important. The KGPE-D16 supports up to 256GB of DDR3 memory across 16 DIMM slots, and any issues during memory initialization, if non-fatal, could severely impact the system's ability to handle large datasets or maintain stable operation under heavy workloads [15]. Given the critical role that memory plays in the overall performance of the KGPE-D16, ensuring that DDR3 memory is correctly initialized is essential for achieving the desired balance of performance,

reliability, and stability in demanding server environments.

4.1.1 General steps for DDR3 configuration

DDR3 memory initialization is a detailed and essential process that ensures both the stability and performance of the system. The process involves several critical steps: detection and identification of memory modules, initial configuration of the memory controller, adjustment of timing and voltage settings, and the execution of training and calibration procedures.

The initialization begins with the detection and identification of the installed memory modules. During the BIST, the firmware reads the Serial Presence Detect (SPD) data stored on each memory module. SPD data contains crucial information about the memory module's specifications, including size, speed, CAS latency (CL), RAS to CAS delay (tRCD), row precharge time (tRP), and row cycle time (tRC). This data allows to configure the memory controller for optimal compatibility and performance.

Indeed, once the memory modules have been identified, the firmware proceeds to the initial configuration of the memory controller. This controller is governed by a state machine that manages the sequence of operations required to initialize, maintain, and control memory access. This state machine consists of multiple states that represent various phases of memory operation, such as reset, initialization, calibration, and data transfer. The transitions between these states are either automatic or command-driven, depending on the specific requirements of each phase [66][52]. This state machine is presented in the fig. 4.2. Automatic transitions, depicted by thick arrows in the automaton, occur without external intervention. These typically include transitions that ensure the memory enters a stable state, such as the transition from power-on to initialization, or from calibration to idle states. These transitions are crucial for maintaining the integrity and stability of the memory system, as they ensure that the controller progresses through necessary stages like ZQ calibration and write leveling, which are essential for proper signal timing and impedance matching [66][52][21].

On the other hand, command-driven transitions, represented by normal arrows in the automaton, require specific commands issued by the memory controller or the CPU to advance to the next state. For instance, the transition from the idle state to the data transfer state requires explicit read or write commands. Similarly, transitioning from the initialization state to the calibration state involves issuing mode register set (MRS) commands that configure the memory's operating parameters. These command-driven transitions are integral to the dynamic operation of the memory system, allowing the controller to respond to the system's operational needs and ensuring that memory accesses are performed efficiently and accurately [66][52].

The memory controller configuration involves setting up fundamental parameters such as the memory clock (MEMCLK) frequency and the memory channel configuration. The MEMCLK frequency is derived from the SPD data, while the memory channels are configured to operate in single, dual, or quad-channel modes, depending on the system architecture and the installed modules [21]. Proper configuration of the memory controller is vital to ensure synchronization with the memory modules, establishing a stable foundation for subsequent operations.

The first critical step, during the INIT phase involves the adjustment of timing and voltage settings. These settings are essential for ensuring that DDR3 memory operates efficiently and reliably. Key timing parameters include CAS Latency (CL), RAS to CAS Delay (tRCD), Row Precharge Time (tRP), and Row Cycle Time (tRC). These parameters are finely tuned to balance speed and stability [66]. The BIOS uses the SPD data to set these parameters and may also adjust them dynamically to achieve the best possible performance. Voltage settings, such as DRAM voltage (typically 1.5V for DDR3) and termination voltage (VTT), are also configured to maintain stable operation, especially under varying conditions such as temperature fluctuations [52]. Training and calibration are among the most complex and crucial stages of DDR3 memory initialization. The fly-by topology used for address, command, and clock signals in DDR3 modules enhances signal integrity by reducing the number of stubs and their lengths, but it also introduces skew between the clock (CK) and data strobe (DQS) signals [52]. This skew must be compensated to ensure that data is written and read correctly. The BIOS performs write leveling, which adjusts the timing of DQS relative to CK for each memory module. This process ensures that the memory controller can write data accurately across all modules, even when they exhibit slight variations in signal timing due to the physical layout [66].




Figure 4.2: DDR3 controller state machine

ZQ calibration is another vital procedure that adjusts the output driver impedance and on-die termination (ODT) to match the systems characteristic impedance [52]. This calibration is critical for maintaining signal integrity under different operating conditions, such as voltage and temperature changes. During initialization, the memory controller issues a ZQCL command, triggering the calibration sequence that optimizes impedance settings. This ensures that the memory system can operate with tight timing tolerances, which is crucial for systems requiring high reliability. Read training is also essential to ensure that data read from the memory modules is interpreted correctly by the memory controller. This process involves adjusting the timing of the read data strobe (DQS) to align perfectly with the data being received. Proper read training is necessary for reliable data retrieval, which directly impacts system performance and stability.

In summary, the DDR3 memory initialization process in systems like the ASUS KGPE-D16 involves a series of detailed and interdependent steps that are critical for ensuring system stability and performance. These include the detection and identification of memory modules, the initial configuration of the memory controller, precise adjustments of timing and voltage settings, and rigorous training and calibration procedures.

4.2 Memory initialization techniques

4.2.1 Memory training algorithms

Memory training algorithms are designed to fine-tune the operational parameters of memory modules, such as timing, voltage, and impedance. These algorithms play a crucial role in achieving the optimal performance of DDR3 memory systems, particularly in complex multi-core environments where synchronization and timing are challenging. The primary algorithms used in memory training include ZQ calibration and write leveling. Optimizing timing and voltage settings is a critical aspect of memory training. The memory controller adjusts parameters such as CAS latency, RAS to CAS delay, and other timing characteristics to ensure that data is read and written with minimal delay and maximum accuracy. Voltage adjustments are also crucial, as they help stabilize the operation of memory modules by ensuring that the power supplied is within the optimal range, compensating for any variations due to temperature or other factors [52][21][44].

ZQ calibration is a critical step in DDR3 memory initialization that ensures the proper impedance matching of the output driver and on-die termination (ODT) resistance. Impedance matching is crucial for maintaining signal integrity by minimizing reflections and ensuring reliable data transmission between the memory controller and the DRAM modules. It is initiated by sending ZQCL (ZQ Calibration Long) commands to the DDR3 DIMMs. Each ZQCL command triggers a long calibration cycle within the DRAM module. The purpose of this calibration is to adjust the output driver impedance and the ODT resistance to match the specified target impedance. This adjustment compensates for process variations, voltage fluctuations, and temperature changes that can affect the impedance characteristics of the DRAM module [44].

A bit in the DRAM Controller Timing register is set to 1 to send the ZQCL command, and an address bit is also set to 1 to indicate that the ZQCL command should be directed to the memory module. Upon receiving the ZQCL command, the DRAM module begins the calibration process. This involves a series of internal adjustments where the DRAM module measures its current impedance and compares it against the target impedance. The module then modifies its internal settings to reduce the difference between the current and target impedance values [44][66]. This process is iterative, meaning that it may require multiple adjustments to converge on the optimal impedance settings. The calibration is designed to ensure that the DRAM module's impedance remains within a tight tolerance, which is critical for high-speed data communication. The ZQ calibration process is time-sensitive. After issuing the ZQCL command, the system must wait for 512 memory clock cycles (MEMCLKs) to allow the calibration to complete. This delay is necessary because the calibration involves both measurement and adjustment phases, which require precise timing to ensure accuracy [44]. If the system does not wait the full 512 MEMCLKs, the calibration may be incomplete, leading to suboptimal impedance matching and potential signal integrity issues, such as reflections or noise on the data lines.

During the ZQ calibration, the DRAM module adjusts its output driver impedance, which controls the strength of the signals it sends out. The stronger the signal, the less susceptible it is to noise, but if the impedance is too high or too low, it can cause signal distortion or reflections. The ODT resistance is also calibrated to properly terminate signals that reach the end of a data line. Proper termination is essential to prevent signal reflections that could interfere with the integrity of the data being transmitted. The ZQCL command adjusts these settings by fine-tuning the resistance values based on the modules feedback, ensuring that the signal paths are optimized for both transmission and termination. Once the ZQ calibration is complete, the DCT register bit is reset to 0, indicating that the calibration command has been processed. The memory controller then verifies that the DRAM module has correctly adjusted its impedance settings. This verification process may involve additional test signals sent across the memory bus to confirm that signal integrity meets the required standards. If the calibration is successful, the memory subsystem is now properly calibrated and ready for normal operation. In systems with LRDIMMs or RDIMMs, additional steps may be necessary to ensure that all ranks and channels are calibrated correctly, particularly in multi-rank configurations where impedance matching can be more complex. However, in systems with complex memory configurations, such as those using multiple DIMMs per channel or operating at higher memory frequencies, the ZQ calibration process becomes even more critical. The calibration may need to be repeated at different operating points to ensure that the memory subsystem remains stable across all conditions. This could involve performing multiple ZQCL calibrations at different memory frequencies, or under different thermal conditions, to account for the dynamic nature of memory operation in modern systems.

In seed-based algorithms, an initial "seed" value is used as a reference point for the calibration process. The memory controller iteratively adjusts the impedance based on feedback from the memory module, refining the calibration with each iteration. This method provides a more precise calibration, particularly in systems where fine-tuned impedance matching is critical for high-frequency operations [58]. Also, while seed-based methods can accelerate the convergence of calibration, they require careful selection of initial seed values to avoid suboptimal or even faulty impedance settings [44].

Write leveling is another critical aspect of memory training, particularly in DDR3 systems that use a fly-by topology. It involves using the physical layer (PHY) to detect the edge of the Data Strobe (DQS) signal in synchronization with the clock (CK) signal on the DIMM (Dual In-line Memory Module) during write access. The DQS signal is a timing signal generated by the memory controller that accompanies data (DQ) during read and write operations. For write operations, the DQS signal must be perfectly aligned with the CK signal to ensure that data is correctly written to memory cells. Indeed, in systems using a fly-by topology, the DQS signal might arrive at different times for different memory devices on the same module due to the signal traveling through different lengths of trace. Write leveling compensates for this skew by adjusting the timing of the DQS signal relative to the CK signal for each lane (a group of data lines) [21]. This training is performed on a per-channel and per-DIMM basis, ensuring that each memory module is correctly synchronized with the memory controller, minimizing timing mismatches that could lead to data corruption.

Using seed-based algorithms, the memory controller sets an initial delay value and then iteratively adjusts it based on the feedback received from the memory module. This process ensures that the DQS signal is correctly aligned with the CK signal at the memory module's pins, minimizing the risk of data corruption and ensuring reliable write operations [66][44]. Seed-based write leveling offers improved precision but must be finely tuned to account for the specific characteristics of the memory module and the overall system architecture [44].

In contrast to seed-based algorithms, seedless methods do not rely on an initial reference value. Instead, they dynamically adjust the impedance and timing parameters during the calibration process. Seedless ZQ calibration continuously monitors the impedance of the memory module and makes real-time adjustments to maintain optimal matching. This approach can be beneficial in environments where the operating conditions are highly variable, as it allows for more flexible and adaptive calibration [58]. Similarly, seedless write leveling dynamically adjusts the DQS timing based on real-time feedback from the memory module. This method is particularly useful in systems where the memory configuration is frequently changed or where the operating conditions vary significantly [52][44]. The traditional ZQ calibration methods, while effective, often struggle with matching impedance perfectly across all conditions. A master thesis by Gopikrishna [44] builds upon these traditional methods by proposing enhancements that involve more sophisticated calibration approaches, leading to better impedance matching and overall memory performance [44].

4.2.2 BIOS and Kernel Developer Guide (BKDG) recommendations

The BIOS and Kernel Developer Guide (BKDG from AMD [9]) is a technical manual aimed at BIOS developers and operating system kernel programmers. It provides in-depth documentation on the AMD processor architecture, system initialization processes, and configuration guidelines. The document is essential for understanding the proper initialization sequences, including those for DDR3 memory, to ensure system stability and performance, particularly for AMD Family 15h processors.

The initialization of DDR3 memory begins with configuring the DDR supply voltage regulator, which ensures that the memory modules receive the correct power levels. Following this, the Northbridge (NB) P-state is forced to NBP0, a state that guarantees stable operation during the initial configuration phases. Once these preliminary steps are completed, the initialization of the DDR physical layer (PHY) begins, which is critical for setting up the communication interface between the memory controller and the DDR3 modules. PHY fence training deals with overall signal alignment at the physical interface, while ZQ calibration focuses on impedance matching, and write leveling addresses timing alignment during write operations. Each process involves different methods as PHY fence training uses iterative timing adjustments, ZQ calibration uses impedance adjustments via the ZQ pin, and write leveling adjusts DQS timing relative to CK during writes. These processes are critical for configuring DDR3 DIMMs and ensuring stable and reliable operation, especially when booting from an unpowered state such as ACPI S4 (hibernation), S5 (soft off), or G3 (mechanical off).

4.2.2.1 DDR3 initialization procedure

DDR3 initialization is a multi-step process that prepares both the memory controllers and the DIMMs for operation. This initialization is essential to set up the memory configuration and to ensure that the memory subsystem operates correctly under various conditions.

- **Enable DRAM initialization:** The process begins by enabling DRAM initialization. This is done by setting the `EnDramInit` bit in the `D18F2x7C_dct` register to 1. The `D18F2x7C_dct` register is a specific configuration register within the memory controller that controls various aspects of the DRAM initialization process. Enabling this bit initiates the sequence of operations required to prepare the memory for use. After setting this bit, the system waits for 200 microseconds to allow the initialization command to propagate and stabilize.
- **Deassert memory reset:** Next, the memory reset signal, known as `MemRstX`, is deasserted by setting the `DeassertMemRstX` bit in the `D18F2x7C_dct` register to 1. Deasserting `MemRstX` effectively takes the memory components out of their reset state, allowing them to begin normal operation. The system then waits for an additional 500 microseconds to ensure that the memory reset is fully deasserted and the memory components are stable.
- **Assert clock enable (CKE):** The next step involves asserting the clock enable signal, known as 'CKE', by setting the `AssertCke` bit in the `D18F2x7C_dct` register to 1. The CKE signal is critical because it enables the clocking of the DRAM modules, allowing them to synchronize with the memory controller. The system must wait for 360 nanoseconds after asserting CKE to ensure that the clocking is correctly established.
- **Registered DIMMs and LRDIMMs initialization:** For systems using registered DIMMs (RDIMMs) or Load Reduced DIMMs (LRDIMMs), additional initialization steps are necessary. RDIMMs and LRDIMMs have buffering mechanisms that reduce electrical loading and improve signal integrity, especially in systems with multiple memory modules. During initialization, the BIOS programs the `ParEn` bit in the `D18F2x90_dct` register based on whether the DIMM is buffered or unbuffered. For RDIMMs, specific Register Control (RC) commands, such as RC0 through RC7, are sent to initialize the memory module's control registers. Similarly, LRDIMMs require a series of Flexible Register Control (FRC) commands, such as F0RC and F1RC, to initialize their internal registers according to the manufacturers specifications.
- **Mode Register Set (MRS):** The initialization process also involves sending Mode Register Set (MRS) commands. These commands are used to configure various operational parameters of the DDR3 memory modules, such as burst length, latency timings, and operating modes. Each MRS command targets a specific mode register within the memory module, and the exact sequence of commands is crucial for setting up the DIMMs according to the systems requirements and the DIMM manufacturers guidelines.

4.2.2.2 ZQ calibration process

ZQ calibration is a key step in DDR3 initialization, responsible for calibrating the output driver impedance and on-die termination (ODT) resistance of the DDR3 modules. Proper impedance matching is essential for maintaining signal integrity, reducing signal reflections, and ensuring reliable data communication between the memory controller and the memory modules.

- **Sending ZQCL commands:** The BIOS initiates ZQ calibration by sending two ZQCL (ZQ Calibration Long) commands to each DDR3 DIMM. ZQCL commands instruct the memory module to perform a long calibration cycle, during which the module adjusts its output driver impedance and ODT resistance to match the desired target impedance. This process compensates for variations due to manufacturing differences, voltage fluctuations, and temperature changes. To send a ZQCL command, the BIOS programs the `SendZQCmd` bit in the `D18F2x7C_dct` register to 1 and sets the `MrsAddress[10]` bit to 1, indicating that the ZQCL command should be sent to the memory module.
- **Calibration timing:** After sending the ZQCL command, the system must wait for 512 memory clock cycles (MEMCLKs) to allow the calibration process to complete. During this time, the memory module adjusts its internal impedance to ensure it matches the specified target impedance. This timing is critical, as inadequate wait times could result in incomplete or inaccurate calibration, leading to signal integrity issues and potential data errors.

- **Finalization of initialization:** Once the ZQ calibration is complete, the BIOS deactivates the DRAM initialization process by setting the `EnDramInit` bit in the `D18F2x7C_dct` register to 0. For LRDIMMs, additional configuration steps are required to finalize the initialization process. These steps include programming the DCT registers to monitor for errors and ensure that the LRDIMMs are operating correctly.

4.2.2.3 Write leveling process

The BIOS and Kernel Developer Guide (BKDG) provides a comprehensive approach to the write leveling process, which is essential for ensuring correct data alignment during write operations in DDR3 memory systems. Write leveling is particularly crucial in systems utilizing a fly-by topology, where timing skew between the clock and data signals can introduce significant challenges. This kind of algorithms were not necessary for DDR2, for example. If the target operating frequency is higher than the lowest supported MEMCLK frequency, the BIOS must perform multiple passes to achieve proper write leveling. The MEMCLK is the clock signal that synchronizes the communication between the memory controller and the memory modules.

During each pass, the memory subsystem is configured for a progressively higher operating frequency:

- **Pass 1:** The memory subsystem is configured for the lowest supported MEMCLK, ensuring that initial timing adjustments are made under the most stable conditions.
- **Pass 2:** The subsystem is then adjusted for the second-lowest MEMCLK, gradually increasing the operating frequency while fine-tuning the alignment of the DQS and CK signals.
- **Pass N:** This process continues until the highest MEMCLK supported by the system is reached, ensuring that the memory operates reliably at its maximum speed.

This step-wise configuration ensures that the memory system is stable across all supported operating frequencies, minimizing the risk of timing errors during write operations, especially as frequencies increase and timing margins become tighter. The configuration process varies depending on whether the DIMM is a Registered DIMM (RDIMM) or an Unregistered DIMM (UDIMM). RDIMMs include an additional buffer to improve signal integrity, which is particularly important in systems with multiple DIMMs. The steps common to both types include a preparation with the DDR3 Mode Register Commands (see fig. 4.2). Mode registers in DDR3 memory are used to configure various operational parameters such as latency settings, burst length, and write leveling. One of the key mode registers is `MR1_dct`, which is specific to DDR3 and controls certain features of the memory module, including write leveling. `MR1_dct` is used to enable or disable specific functions such as write leveling and output driver settings. The `dct` suffix refers to the Data Control Timing that is specific to this register's function in managing the timing and control of data operations within the memory module. For RDIMMs, a 4-rank module is treated as two separate DIMMs, where each rank is essentially a separate memory module within the same DIMM. The first two ranks are the primary target for the initial configuration. The remaining two ranks are treated as non-target and are configured separately.

Then, these steps are followed, still common to both RDIMMs and UDIMMs:

- **Step 1A: Output Driver and ODT configuration for target DIMM:**
 - For the first rank (target):
 - * Set `MR1_dct[1:0][Level] = 1` to enable write leveling.
 - * Set `MR1_dct[1:0][Qoff] = 0` to ensure the output drivers are active.
 - For other ranks:
 - * Set `MR1_dct[1:0][Level] = 1` to prepare for write leveling.
 - * Set `MR1_dct[1:0][Qoff] = 1` to deactivate the output drivers for ranks that are not currently being leveled.
 - If there are two or more DIMMs per channel, or if there is one DIMM per three channels:
 - * Program the target ranks `RttNom` (nominal termination resistance value) for `RttWr` termination, which helps in managing signal integrity during the write process by ensuring the correct impedance matching.

- **Step 1B: Configure non-target RttNom to normal operation:**
 - After the initial configuration, the RttNom values for the non-target ranks are set to their normal operating states.
 - A wait time of 40 MEMCLKs is observed to ensure the configuration settings are stable before proceeding.
- **Step 3: PHY configuration:**
 - The PHY is then configured to measure and adjust the timing delays accurately for each data lane. The PHY layer is responsible for converting the signals from the memory controller into a form that can be transmitted over the physical connections to the memory modules.
- **Step 4: Perform write leveling:**
 - The actual write leveling process is executed, where the DQS signal timing is adjusted to ensure it aligns perfectly with the CK signal at the memory modules pins, ensuring that data is written accurately.
- **Step 5: Disable PHY configuration post-measurement:**
 - After completing the write leveling process, the PHY configuration is disabled to stop further timing measurements and adjustments, locking in the calibrated settings.
- **Step 6: Program the DIMM to normal operation:**
 - Finally, the DIMM is reprogrammed to its normal operational state, resetting Qoff and Level to 0 to conclude the write leveling process and return to standard operation.

For each DIMM, the BIOS must calculate the coarse and fine delays for each lane in the DQS Write Timing register:

- **Coarse Delay Calculation:** This involves setting a basic delay based on a seed value specific to the platform. The seed value is determined during initial system configuration and serves as a starting point for further delay adjustments.
- **Critical Delay Determination:** The minimum of the coarse delays for each lane and DIMM is considered the critical delay. This delay is crucial for ensuring that all data lanes are correctly synchronized.
- **Platform-Specific Seed:** The seed ranges between -1.20ns and +1.20ns, providing a small adjustment range to fine-tune the timing based on the specific characteristics of the platform. This seed value can differ for the first pass compared to subsequent passes, allowing for incremental adjustments as the system stabilizes.

4.3 Current implementation and potential improvements [WIP]

4.3.1 Current implementation in coreboot on the KGPE-D16 [WIP]

- Overview of the current DDR3 initialization process in *coreboot* on the KGPE-D16
- Analysis of strengths and weaknesses of the current implementation

The process starts by calling the `fill_mem_ctrl` function from `src/northbridge/amd/amdfam10/raminit_sysinfo_in_ram.c`. At this current step, only the BSC is running the firmware code. This function iterates over all memory controllers (one per node) and initializes their corresponding structures with the system information needed for the RAM to function. This includes the addresses of PCI nodes (important for DMA operations) and SPD addresses, which are internal ROMs in each memory slot containing crucial information for detecting and initializing memory modules. If successful, the system posts codes 0x3D and then 0x40. The `raminit_amdmct` function from `src/northbridge/amd/amdfam10/raminit_amdmct.c` is then called. This function, in turn, calls `mctAutoInitMCT_D` from `src/northbridge/amd/amdmct/mct_ddr3/mct_d.c`, which is responsible for the initial memory initialization, predominantly written by Raptor Engineering.

At this stage, it is assumed that memory has been pre-mapped contiguously from address 0 to 4GB and that the previous code has correctly mapped non-cacheable I/O areas below 4GB for the PCI bus and Local APIC access for processor cores.

The following prerequisites must be in place from the previous steps:

- The HyperTransport bus configured, and its speed is correctly set.
- The SMBus controller is configured.
- The BSP is in unreal mode.
- A stack is set up for all cores.
- All cores are initialized at a frequency of 2GHz.
- If we were using saved values, the NVRAM would have been verified with checksums.

The memory controller for the BSP is queried to check if it can manage ECC memory, which is a type of memory that includes error-correcting code to detect and correct common types of data corruption.

For each node available in the system, the memory controllers are identified and initialized using a `DCTStatStruc` structure defined in `src/northbridge/amd/amdmct/mct_ddr3/mct_d.h`. This structure contains all necessary fields for managing a memory module. The process includes:

- Retrieving the corresponding field in the `sysinfo` structure for the node.
- Clearing fields with zero.
- Initializing basic fields.
- Initializing the controller linked to the current node.
- Verifying the presence of the node (checking if the processor associated with this controller is present). If yes, the SMBus is informed.
- Pre-initializing the memory module controller for this node using `mct_preInitDCT`.

The `coreboot` code compensates for the delay between DQS and DQ signals, as well as between CMD and DQ. This is handled in the `DQSTiming_D` function.

Finally, if the RAM is of the ECC type, error-correcting codes are enabled, and the function ends by activating power-saving features if requested by the user.

TODO (continue notes from PROJET)

```
1 void main(int a, int b)
2 {
3     return 0;
4 }
```

Listing 4.1: *Example C code*

We saw that in (lst. 4.1).

4.3.2 Potential enhancements [WIP]

- Identifying areas for improvement in the current implementation
- Potential enhancements to memory training algorithms and configuration settings
- Broader applicability of these improvements to other systems using `coreboot`

Chapter 5

Virtualization of the operating system through firmware abstraction

In contemporary computing systems, the operating system (OS) no longer interacts directly with hardware in the same way it did in earlier computing architectures. Instead, the OS operates within a highly abstracted environment, where critical functions are managed by various firmware components such as ACPI, SMM, UEFI, Intel Management Engine (ME), and AMD Platform Security Processor (PSP). This layered abstraction has led to the argument that the OS is effectively running in a virtualized environment, akin to a virtual machine (VM).

5.1 ACPI and abstraction of hardware control

The Advanced Configuration and Power Interface (ACPI) provides a standardized method for the OS to manage hardware configuration and power states, effectively abstracting the underlying hardware complexities. ACPI abstracts hardware details, allowing the OS to interact with hardware components without needing direct control over them. This abstraction is similar to how a hypervisor abstracts physical hardware for VMs, enabling a consistent interface regardless of the underlying hardware specifics.

According to Bellosa [17], the abstraction provided by ACPI not only simplifies the OS's interaction with hardware but also limits the OS's ability to fully control the hardware, which is instead managed by ACPI-compliant firmware. This layer of abstraction contributes to the virtualization-like environment in which the OS operates.

More importantly, the ACPI Component Architecture (ACPIA) is a critical component integrated into the Linux kernel, serving as the foundation for the system's ACPI implementation [26]. ACPIA provides the core ACPI functionalities, such as hardware configuration, power management, and thermal management, which are essential for modern computing platforms. However, its integration into the Linux kernel has brought significant complexity and code overhead, making Linux heavily dependent on ACPIA for managing ACPI-related tasks.

ACPIA is a large and complex project, with its codebase encompassing a wide range of functionalities required to implement ACPI standards. The integration of ACPIA into the Linux kernel significantly increases the kernel's overall code size. An example of that can easily be reproduced with a small experiment (lst. 5.2).

```
1 $ git clone https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git
2 $ cd linux/drivers/acpi
3 $ find . -name "*.c" -o -name "*.h" | xargs wc -l
4 [...]
5 168970 total
```

Listing 5.2: *How to estimate the impact of ACPIA in Linux*

As of recent statistics, ACPIA comprises between 100,000 to 200,000 lines of code, making it one of the larger subsystems within the Linux kernel. This size is indicative of the extensive range of features and capabilities ACPIA must support, including but not limited to the ACPI interpreter, AML (ACPI Machine Language) parser, and various hardware-specific drivers. The ACPIA codebase is not monolithic; it is highly modular and consists of various components, each responsible for specific ACPI functions. For instance, ACPIA includes components

for managing ACPI tables, interpreting AML bytecode, handling events, and interacting with hardware. This modularity, while beneficial for isolating different functionalities, also contributes to the overall complexity of the system. The separation of ACPICA into multiple modules necessitates careful coordination and integration with the rest of the Linux kernel, adding to the kernel's complexity.

ACPICA's integration into the Linux kernel is designed to maintain a clear separation between the core ACPI functionalities and the kernel's other subsystems [26]. This separation is achieved through well-defined interfaces and abstraction layers, allowing the Linux kernel to interact with ACPICA without being tightly coupled to its internal implementation details. For example, ACPICA provides an API that the Linux kernel can use to interact with ACPI tables, execute ACPI methods, and manage power states. This API abstracts the underlying complexity of the ACPI implementation, making it easier for kernel developers to incorporate ACPI support without delving into the intricacies of ACPICA's internals. Moreover, ACPICA's role in interpreting AML bytecode, which is essentially a form of low-level programming language embedded in ACPI tables, adds a layer of abstraction. The Linux kernel relies on ACPICA to execute AML methods and manage hardware resources according to the ACPI specifications. This reliance further underscores the idea that ACPI acts as a virtualizing environment, shielding the kernel from the complexities of directly interfacing with hardware components.

5.2 SMM as a hidden execution layer

System Management Mode (SMM) is a special-purpose operating mode provided by x86 processors, designed to handle system-wide functions such as power management, thermal monitoring, and hardware control, independent of the OS. SMM operates transparently to the OS, executing code that the OS cannot detect or control, similar to how a hypervisor controls the execution environment of VMs.

Research by Huang and Smith [51] argues that SMM introduces a hidden layer of execution that diminishes the OS's control over the hardware, creating a virtualized environment where the OS is unaware of and unable to influence certain system-level operations. This hidden execution layer reinforces the idea that the OS runs in an environment similar to a VM, with the firmware acting as a hypervisor.

5.3 UEFI and persistence

The Unified Extensible Firmware Interface (UEFI) has largely replaced the traditional BIOS in modern systems, providing a sophisticated environment that includes a kernel-like structure capable of running drivers and applications independently of the OS. UEFI remains active even after the OS has booted, continuing to manage certain hardware functions, which abstracts these functions away from the OS.

McClellan [71] discusses how UEFI creates a persistent execution environment that overlaps with the OS's operation, effectively placing the OS in a position where it runs on top of another controlling layer, much like a guest OS in a VM. This persistence and the ability of UEFI to manage hardware resources independently further blur the lines between traditional OS operation and virtualized environments. Indeed, as we studied in a precedent chapter, UEFI is designed as a modular and extensible firmware interface that sits between the computer's hardware and the operating system. Unlike the monolithic BIOS, UEFI is composed of several layers and components, each responsible for different aspects of the system's boot and runtime processes. The core components of UEFI include the Pre-EFI Initialization (PEI), Driver Execution Environment (DXE), Boot Device Selection (BDS), and Runtime Services. Each of these components plays a critical role in initializing the hardware, managing drivers, selecting boot devices, and providing runtime services to the OS.

The PEI (Pre-EFI Initialization) phase is responsible for initializing the CPU, memory, and other essential hardware components. It ensures that the system is in a stable state before handing control to the DXE phase. In the DXE phase, the system loads and initializes various drivers required for the OS to interact with the hardware. The DXE phase also constructs the UEFI Boot Services, which provide the OS with interfaces to the hardware during the boot process. The BDS (Boot Device Selection) phase is responsible for selecting the device from which the OS will boot. It interacts with the UEFI Boot Manager to determine the correct boot path and load the OS. After the OS has booted, UEFI provides Runtime Services that remain accessible to the OS. These services include

interfaces for managing system variables, time, and hardware. UEFI also supports the execution of standalone applications, which can be used for system diagnostics, firmware updates, or other tasks. These applications operate independently of the OS, highlighting UEFI's capabilities as a minimalistic OS.

UEFI abstracts the underlying hardware from the OS, providing a standardized interface for the OS to interact with different hardware components. This abstraction simplifies the development of OSes and drivers, as they do not need to be tailored for specific hardware configurations. UEFI's hardware abstraction is one of the key features that enable it to act as a virtualizing environment for the OS [71].

5.3.1 Memory Management

UEFI provides a detailed memory map to the OS during the boot process, which includes information about available, reserved, and used memory regions. The OS uses this memory map to manage its own memory allocation and paging mechanisms. The overlap in memory management functions highlights UEFI's role in preparing the system for OS operation. This memory map includes all the memory regions in the system, categorized into different types, such as usable memory, reserved memory, and memory-mapped I/O. The OS relies on this map to understand the system's memory layout and avoid conflicts [82]. The OS extends UEFI's memory management by implementing its own memory allocation, paging, and virtual memory mechanisms. However, the OS's memory management is built on the foundation provided by UEFI, demonstrating the close relationship between the two.

5.3.2 File System Management

UEFI includes its own file system management capabilities, which overlap with those of the OS. The most notable example is the EFI System Partition (ESP), a special partition formatted with the FAT file system that UEFI uses to store bootloaders, drivers, and other critical files [38]. The ESP is a mandatory partition in UEFI systems, containing the bootloaders, firmware updates, and other files necessary for system initialization. UEFI accesses the ESP independently of the OS, but the OS can also access and manage files on the ESP, creating an overlap in file system management functions [59]. UEFI natively supports the FAT file system, allowing it to read and write files on the ESP. This support overlaps with the OS's file system management, as both UEFI and the OS can manipulate files on the ESP.

5.3.3 Device Drivers

As we studied in an earlier chapter, UEFI includes its own driver model, allowing it to load and execute drivers independently of the OS. This capability overlaps with the OS's driver management functions, as both UEFI and the OS manage hardware devices through drivers. UEFI drivers are typically used during the boot process to initialize and control hardware devices. These drivers provide the necessary interfaces for the OS to interact with the hardware once it has booted [59]. After the OS has booted, it loads its own drivers for hardware devices. However, the OS often relies on the initial hardware setup performed by UEFI drivers.

5.3.4 Power Management

UEFI provides power management services that overlap with the OS's power management functions. These services allow UEFI to manage power states and transitions independently of the OS [38]. These services ensure that the system conserves power during periods of inactivity and can quickly resume operation when needed. The OS extends UEFI's power management by implementing its own power-saving mechanisms, such as CPU throttling and dynamic voltage scaling.

5.4 Intel and AMD: control beyond the OS

Intel Management Engine (ME) and AMD Platform Security Processor (PSP) are embedded microcontrollers within Intel and AMD processors, respectively. These components run their own firmware and operate independently of the main CPU, handling tasks such as security enforcement, remote management, and digital rights management (DRM).

Bulygin [20] highlights how these microcontrollers have control over the system that supersedes the OS, managing hardware and security functions without the OS's knowledge or consent. This level of control is reminiscent of a

hypervisor that manages the resources and security of VMs. The OS, in this context, operates similarly to a VM that does not have full control over the hardware it ostensibly manages.

5.5 The OS as a virtualized environment

The combined effect of these firmware components (ACPI, SMM, UEFI, Intel ME, and AMD PSP) creates an environment where the OS operates in a virtualized or highly abstracted layer. The OS does not directly manage the hardware; instead, it interfaces with these firmware components, which themselves control the hardware resources. This situation is analogous to a virtual machine, where the guest OS operates on virtualized hardware managed by a hypervisor.

Smith and Chen [97] argues that modern OS environments, influenced by these firmware components, should be considered virtualized environments. The firmware acts as an intermediary layer that abstracts and controls hardware resources, thereby limiting the OS's direct access and control.

The presence and operation of modern firmware components such as ACPI, SMM, UEFI, Intel ME, and AMD PSP contribute to a significant abstraction of hardware from the OS. This abstraction creates an environment that parallels the operation of a virtual machine, where the OS functions within a controlled, virtualized layer managed by these firmware systems. The growing body of research supports this perspective, suggesting that the traditional notion of an OS directly managing hardware is increasingly outdated in the face of these complex, autonomous firmware components.

Conclusion

This document has explored the evolution and current state of firmware, particularly focusing on the transition from traditional BIOS to more advanced firmware interfaces such as UEFI and *coreboot*. The evolution from a simple set of routines stored in ROM to complex systems like UEFI and *coreboot* highlights the growing importance of firmware in modern computing. Firmware now plays a critical role not only in hardware initialization but also in memory management, security, and system performance optimization.

The study of the ASUS KGPE-D16 mainboard illustrates how firmware, particularly *coreboot*, plays a crucial role in the efficient and secure operation of high-performance systems. The KGPE-D16, with its support for free software-compatible firmware, exemplifies the potential of libre firmware to deliver both high performance and freedom from proprietary constraints. However, it is important to acknowledge that the KGPE-D16 is not without its imperfections. The detailed analysis of firmware components, such as the bootblock, romstage, and especially the RAM initialization and training algorithms, reveals areas where the firmware can be further refined to enhance system stability and performance. These improvements are not only beneficial for the KGPE-D16 but can also be applied to other boards, extending the impact of these optimizations across a broader range of hardware.

Moreover, the discussion on modern firmware components such as ACPI, SMM, UEFI, Intel ME, and AMD PSP demonstrates how these elements abstract hardware from the operating system, creating a virtualized environment where the OS operates more like a guest in a hypervisor-controlled system. This abstraction raises important considerations about control, security, and user freedom in contemporary computing. As we continue to witness the increasing complexity and influence of firmware in computing, it becomes crucial to advocate for free software-compatible hardware. The dependence on proprietary firmware and the associated restrictions on user freedom are growing concerns that need to be addressed. The development and adoption of libre firmware solutions, such as *coreboot* and GNU Boot, are essential steps towards ensuring that users retain control over their hardware and software environments.

It is imperative that the community of developers, researchers, and users come together to support and contribute to the development of free firmware. By fostering innovation and collaboration in this field, we can advance towards a future where free software-compatible hardware becomes the norm, ensuring that computing remains open, secure, and under the control of its users. The significance of a libre BIOS cannot be overstated, it is the foundation upon which a truly free and open computing ecosystem can be built [107]. The importance of the GNU Boot project cannot be overstated. As a fully free firmware initiative, GNU Boot represents a critical step towards achieving truly libre BIOSes, ensuring that users can maintain full control over their hardware and firmware environments. The continued development and support of GNU Boot are essential for advancing the goals of free software and protecting user freedoms in the increasingly complex landscape of modern computing.

Bibliography

- [1] ACMCS. "The Evolution of Firmware: BIOS to UEFI". In: *ACM Computing Surveys* 47.4 (2015), pp. 55–61. DOI: 10.1145/2766462.
- [2] ACPI. *ACPI Specification*. <https://www.acpi.info/spec.htm>. Accessed: 2024-07-05.
- [3] Advanced Micro Devices (AMD). *AMD Embedded Chipsets: SR5690 and SP5100*. Accessed: 2024-08-17. URL: <https://www.amd.com/en/products/embedded-chipsets>.
- [4] Advanced Micro Devices (AMD). *AMD Family 15h Models 30h-3Fh Processors BIOS and Kernel Developer's Guide*. Accessed: 2024-08-17. 2014. URL: https://www.amd.com/system/files/TechDocs/48751_15h_Mod_30h-3Fh_BKDG.pdf.
- [5] Altera. "DDR3 SDRAM Memory Interface Termination and Layout Guidelines". In: AN-520-1.0. 2008.
- [6] AMD. *AMD DDR3 Memory Controller: Technical Overview*. Available at AMD Developer Central. 2011. URL: <https://developer.amd.com/>.
- [7] AMD. *AMD Opteron 6200 Series Processor*. Available at AMD Developer Central. 2011. URL: <https://developer.amd.com/>.
- [8] AMD. *AMD Platform Security Processor (PSP)*. <https://www.amd.com/en/technologies/security>. Accessed: 2024-07-05.
- [9] AMD. "BIOS and Kernel Developers Guide (BKDG) for AMD Family 15h Models 00h-0Fh Processors Rev 3.14". In: 42301. Jan. 2013.
- [10] AMD. *HyperTransport Technology: Technical Overview*. Available at AMD Developer Central. 2011. URL: <https://developer.amd.com/>.
- [11] AMD. "SR5690/5670/5650 BIOS Developers Guide 3.00". In: 43870. Nov. 2010.
- [12] AMD. "SR5690/5670/5650 Register Programming Requirements 3.05". In: 43872. Aug. 2012.
- [13] T. Anderson. *BIOS vs. UEFI: Understanding the Modern Boot Environment*. <https://www.pcworld.com/article/3171322/bios-vs-uefi-understanding-the-modern-boot-environment.html>. 2018.
- [14] IBM Archives. *IBM Personal Computer*. <https://www.ibm.com/history/personal-computer>. 2024.
- [15] Asus. *Asus KGPE-D16 Mainboard Documentation and User Manuals*. Accessed: 2024-07-05.
- [16] Vladimir Bashun et al. "Too young to be secure: Analysis of UEFI threats and vulnerabilities". eng. In: *14th Conference of Open Innovation Association FRUCT*. Vol. 232. 14. FRUCT Oy, 2013, pp. 16–24. ISBN: 1479949779.
- [17] Frank Bellosa. "Impact of ACPI on Operating System Control". In: *Journal of Embedded Systems* 12.3 (2010), pp. 134–142.
- [18] Paul Bischoff. *Intel Management Engine: The obscure chip that does a lot for your computer*. Accessed: 2024-08-17. 2020. URL: <https://proprivacy.com/privacy-news/intel-management-engine>.
- [19] R. E. Brown et al. "LinuxBIOS as an Open-Source Firmware Alternative". In: *Proceedings of the 2003 Linux Symposium*. 2003.
- [20] Maxim Bulygin. "Chipset-Level Control: Understanding Intel ME and AMD PSP". In: *Security Architecture Journal* 18.2 (2013), pp. 45–56.
- [21] Jon Burnett. *DDR3 Design Considerations for PCB Applications (AN111)*. Presentation at Freescale Technology Forum. Freescale Semiconductor, Inc. July 2009.

- [22] H. Chang and A. Smith. "UEFI Secure Boot in Modern Computing". In: *International Journal of Information Security* 12.3 (2013), pp. 231–241. DOI: 10.1007/s10207-013-0191-1.
- [23] Kaixing Cheng et al. "Two Optimization Ways of DDR3 Transmission Line Equal-Length Wiring Based on Signal Integrity". eng. In: *International Journal of Electronics and Telecommunications* 67.3 (2021), pp. 385–394. ISSN: 2081-8491.
- [24] Ronny Chevalier et al. "Co-processor-based Behavior Monitoring: Application to the Detection of Attacks Against the System Management Mode". eng. In: vol. 2017. ACM, 2017, pp. 399–411.
- [25] Coreboot Project. *Coreboot Memory Management and Payload Allocation*. Accessed: 2024-08-17. 2024. URL: <https://doc.coreboot.org/memory-map.html>.
- [26] Intel Corporation. *ACPI Programming Reference*. Accessed: 2024-08-24. Apr. 2023. URL: <https://cdrdv2.intel.com/v1/dl/getContent/772726>.
- [27] Intel Corporation. *Advanced Configuration and Power Interface (ACPI) Specification*. Intel Corporation, 1996. URL: <https://uefi.org/specifications>.
- [28] Intel Corporation. *Intel Converged Security and Management Engine (CSME) Security White Paper*. Tech. rep. 2020. URL: <https://software.intel.com/content/dam/www/public/us/en/security-advisory/documents/intel-csme-security-white-paper.pdf>.
- [29] Intel Corporation. *Introduction to ACPI*. Accessed: 2024-08-24. Apr. 2023. URL: <https://cdrdv2.intel.com/v1/dl/getContent/772721>.
- [30] Intel Corporation. *System Management Mode*. Tech. rep. 2016. URL: <https://www.intel.com/content/www/us/en/developer/articles/technical/system-management-mode.html>.
- [31] Intel Corporation. *Unified Extensible Firmware Interface (UEFI)*. <https://www.intel.com/content/www/us/en/architecture-and-technology/unified-extensible-firmware-interface.html>. 2020.
- [32] Microsoft Corporation. *UEFI Firmware*. <https://docs.microsoft.com/en-us/windows-hardware/drivers/bringup/uefi-firmware>. 2019.
- [33] Cory Doctorow. *Intel x86 processors ship with a secret backdoor*. Accessed: 2024-08-17. 2016. URL: <https://boingboing.net/2016/06/15/intel-x86-processors-ship-with.html>.
- [34] Jane Doe. "DDR2 Memory Controller in the ASpeed AST2050". In: *Memory Systems Review* 22.2 (2015), pp. 33–40.
- [35] Christopher Domas. "The Memory Sinkhole - Unleashing an x86 Design Flaw Allowing Universal Privilege Escalation". In: *Black Hat USA* (2015). URL: <https://www.blackhat.com/docs/us-15/materials/us-15-Domas-The-Memory-Sinkhole-Unleashing-An-x86-Design-Flaw-Allowing-Universal-Privilege-Escalation-wp.pdf>.
- [36] Shawn Embleton, Sherri Sparks, and Cliff C. Zou. "SMM rootkit: a new breed of OS independent malware". eng. In: *Security and communication networks* 6.12 (2013), pp. 1590–1605. ISSN: 1939-0114.
- [37] Mark M. Ermolov, Dmitry V. Sklyarov, and Maxim S. Goryachy. "Undocumented x86 instructions to control the CPU at the microarchitecture level in modern INTEL processors". eng. In: *Bezopasnost' i informat'siya ionnykh tekhnologii* 29.4 (2022), pp. 27–41. ISSN: 2074-7128.
- [38] UEFI Forum. *UEFI Specification*. <https://uefi.org/specifications>. 2021.
- [39] UEFI Forum. *Unified Extensible Firmware Interface*. <https://uefi.org/>. 2024.
- [40] Aurelien Francillon et al. "Co-processor-based Behavior Monitoring: Application to the Detection of Attacks Against the System Management Mode". In: *arXiv* (2018). URL: <https://arxiv.org/abs/1803.02700>.
- [41] Free Software Foundation. *Respects Your Freedom (RYF) Certification*. Accessed: 2024-08-17. 2017. URL: <https://ryf.fsf.org/products/VikingsD16>.
- [42] Free Software Foundation. *The Management Engine: An Attack on Computer Users' Freedom*. Accessed: 2024-08-17. 2016. URL: <https://www.fsf.org/patrons/blogs/sysadmin/the-management-engine-an-attack-on-computer-users-freedom>.
- [43] Paul Freiberger and Michael Swaine. *Fire in the Valley: The Birth and Death of the Personal Computer*. McGraw-Hill, 2000.

- [44] Siddula Gopikrishna. "A Novel Impedance Calibration Method for Low Cost Memory Applications". Accessed: 2024-08-24. Master of Science Thesis. Hyderabad, India: International Institute of Information Technology, Hyderabad, Feb. 2021. URL: https://cdn.iiit.ac.in/cdn/web2py.iiit.ac.in/research_centres/publications/download/mastersthesis.pdf.a761b452d5c4ae57.476f70696b726973686e615f4d.pdf.
- [45] Maxim Goryachy and Mark Ermolov. "How to Hack a Turned Off Computer, or Running Unsigned Code in Intel Management Engine". In: *Black Hat Europe* (2017), pp. 1–23. URL: <https://www.blackhat.com/docs/eu-17/materials/eu-17-Goryachy-How-To-Hack-A-Turned-Off-Computer-Or-Running-Unsigned-Code-In-Intel-Management-Engine-wp.pdf>.
- [46] Jimmy Grewal. *The Creation of the IBM PC*. Armonk Institute. 2024.
- [47] Michael Gschwind. "Advanced Configuration and Power Interface: The Operating System Perspective". In: *IEEE Micro* 20 (2000), pp. 82–89. DOI: 10.1109/40.888702.
- [48] Ya Hai et al. "A wide-frequency and high-precision ZQ calibration circuit for NAND Flash memory". eng. In: *Microelectronics* 143 (2024), pp. 106051–. ISSN: 1879-2391.
- [49] John Heasman. "Implementing and Detecting an ACPI BIOS Rootkit". In: *Black Hat USA* (2007). URL: <https://www.blackhat.com/presentations/bh-usa-07/Heasman/Presentation/bh-usa-07-heasman.pdf>.
- [50] M. D. Hill and M. R. Marty. "The Impact of Caching on Multicore Performance". In: *Communications of the ACM* 51.12 (2008), pp. 48–54.
- [51] Rich Huang and John Smith. "Invisible Hypervisor: An Analysis of System Management Mode". In: *Proceedings of the 16th ACM Conference on Computer and Communications Security*. ACM. 2009, pp. 25–35.
- [52] Micron Technology Inc. *Technical Note: DDR3 ZQ Calibration*. TN-41-02. 2008.
- [53] Intel Corporation. *Intel Management Engine (Intel ME)*. <https://www.intel.com/content/www/us/en/architecture-and-technology/intel-management-engine.html>. Accessed: 2024-07-05.
- [54] io.netgarage. *Intel Management Engine*. Accessed: 2024-08-17. 2024. URL: <https://io.netgarage.org/me/>.
- [55] Michael Jones. "Customizing OpenBMC for ASpeed AST2050". In: *Open Source Firmware Journal* 5.1 (2017), pp. 12–18.
- [56] Corey Kallenberg and Xeno Kovah. "BIOS and SMM Internals". In: (2014). URL: https://opensecuritytraining.info/IntroBIOS_files/Day1_07_Advanced%20x86%20-%20BIOS%20and%20SMM%20Internals%20-%20SMM.pdf.
- [57] David Kaplan, Jeremy Powell, and Tom Woller. "AMD Memory Encryption". In: *Architectural Support for Programming Languages and Operating Systems*. 2016, pp. 149–160. DOI: 10.1145/2851141.2851148.
- [58] Dong-Seok Kim, Dong-Seok Oh, and Seok-Hoon Lee. "Design DDR3 ZQ Calibration having improved impedance matching". In: *Proceedings of the 2010 Fall Conference on Semiconductor and Display Technology*. Retrieved from <https://koreascience.kr/article/CFKO200835536002505.pdf>. Seoul, South Korea: Hanyang University, 2010, pp. 191–192.
- [59] Ronald D. Krebs, Vincent Zimmer, and Suresh Marisetty. *Beyond BIOS: Developing with the Unified Extensible Firmware Interface*. 3rd. Intel Press, 2017. ISBN: 978-0974364906.
- [60] Stefan Krempel. *Intel-Fernwartung AMT bei Angriffen auf PCs genutzt*. Accessed: 2024-08-17. 2017. URL: <https://www.heise.de/news/Intel-Fernwartung-AMT-bei-Angriffen-auf-PCs-genutzt-3739441.html>.
- [61] Christoph Lameter. "NUMA (Non-Uniform Memory Access): An Overview". In: *Queue* 11 (July 2013). DOI: 10.1145/2508834.2513149.
- [62] Michael Larabel. *HDCP 2.2 Coming To The Intel i915 Linux DRM Driver*. Accessed: 2024-08-17. 2018. URL: <https://www.phoronix.com/news/HDCP-2.2-For-i915-DRM>.
- [63] Michael Larabel. *HDCP 2.2 Support Being Worked On For Intel Linux Graphics Driver*. Accessed: 2024-08-17. 2017. URL: <https://www.phoronix.com/news/HDCP-2.2-Intel-Linux-Driver>.

- [64] Olivier Levillain et al. *How to Protect the BIOS and its Secrets*. Tech. rep. ANSSI, Eurecom, 2011. URL: https://cyber.gouv.fr/sites/default/files/IMG/pdf/Cansec_final.pdf.
- [65] Huiyong Li et al. "Reflection Reduction on DDR3 High-Speed Bus by Improved PSO". eng. In: *TheScientificWorld* 2014 (2014), pp. 257972–11. ISSN: 2356-6140.
- [66] Samsung Electronics Co. Ltd. *DDR3 SDRAM Specification Rev 1.4*. TN-41-02. Nov. 2011.
- [67] GNU Boot project maintainers. *Frequently Asked Questions*. <https://www.gnu.org/software/gnuboot/web/faq.html>. Accessed: 2024-07-23.
- [68] GNU Boot project maintainers. *GNU Boot — Free your BIOS today!* [Online; accessed 7-May-2024]. 2024. URL: <https://www.gnu.org/software/gnuboot/>.
- [69] GNU Boot project maintainers. *GNU Boot — Status*. [Online; accessed 7-May-2024]. 2024. URL: <https://www.gnu.org/software/gnuboot/web/status.html>.
- [70] Alex Markuze et al. "Understanding DMA Attacks in the Presence of an IOMMU". In: *Proceedings of the Sixteenth European Conference on Computer Systems (EuroSys '21)*. ACM, 2021. URL: <https://research.vmware.com/publications/understanding-dma-attacks-in-the-presence-of-an-iommu>.
- [71] Laura McClean. *UEFI: The Definitive Guide to Modern Firmware*. O'Reilly Media, 2017.
- [72] Ivison Medeiros et al. "Towards a dynamic data distribution management framework based on Apache Spark". In: *Journal of the Brazilian Computer Society* 23.1 (2017), pp. 1–15. DOI: 10.1186/s13173-017-0066-7. URL: <https://journal-bcs.springeropen.com/articles/10.1186/s13173-017-0066-7>.
- [73] R. Minnich and E. Hendricks. "Challenges and Progress in coreboot Development". In: *Journal of Open Source Software* 3.29 (2018), pp. 1–6. DOI: 10.21105/joss.00429.
- [74] Ron Minnich. "coreboot: Status and some history". In: 2006.
- [75] Ron Minnich, Stefan Reinauer, and Patrick Georgi. "coreboot: Open-Source Firmware Platform". In: *Google Research* (2017). URL: <https://research.google/pubs/pub45424/>.
- [76] Benjamin Mohr. *A Comparative Analysis of Bootloaders*. Tech. rep. University of Freiburg, 2012.
- [77] Nuvoton Technology Corporation. *Nuvoton W83795G/ADG Hardware Monitor Datasheet*. Accessed: 2024-08-17. URL: <https://www.nuvoton.com/>.
- [78] Danny O'Brien. *Intels Management Engine is a Security Hazard, and Users Need a Way to Disable It*. Accessed: 2024-08-17. 2017. URL: <https://www.eff.org/deeplinks/2017/05/intels-management-engine-security-hazard-and-users-need-way-disable-it>.
- [79] Alexander Ogolyuk, Andrey Sheglov, and Konstantin Sheglov. "UEFI BIOS and Intel Management Engine Attack Vectors and Vulnerabilities". eng. In: *Proceedings of the XXth Conference of Open Innovations Association FRUCT 776.20* (2017), pp. 657–662. ISSN: 2305-7254.
- [80] Linux Kernel Organization. *High-bandwidth Digital Content Protection (HDCP)*. Accessed: 2024-08-17. 2020. URL: <https://www.kernel.org/doc/html/v5.8/driver-api/mei/hdcp.html>.
- [81] OSDev Wiki. *Graphics Output Protocol (GOP)*. Accessed: 2024-08-17. 2024. URL: <https://wiki.osdev.org/GOP>.
- [82] OSDev Wiki contributors. *UEFI - OSDev Wiki*. [Online; accessed 25-August-2024]. 2024. URL: <https://wiki.osdev.org/UEFI#Memory>.
- [83] Timothy Pearson. "The World Beyond x86". In: 2014.
- [84] ACPIA Project. *ACPI Component Architecture Programmer Reference*. Accessed: 2024-08-03. 2017. URL: <https://acpica.org/documentation>.
- [85] coreboot Project. *coreboot Documentation*. 2023. URL: <https://doc.coreboot.org/>.
- [86] coreboot project. *coreboot Payloads*. <https://www.coreboot.org/Payloads>. Accessed: 2024-07-23.
- [87] coreboot project. *coreboot: Open Source Firmware*. <https://www.coreboot.org/>. Accessed: 2024-07-23.
- [88] Raptor Engineering LLC. *Raptor Engineering website*. [Online; accessed 8-May-2024]. 2009-2024. URL: <https://raptorengineering.com/>.

- [89] Stefan Reinauer et al. "The coreboot Open Source BIOS - A Review". In: *Usenix Annual Technical Conference*. 2008.
- [90] Felix Richter et al. "BIOS and UEFI firmware analysis". In: *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*. 2011, pp. 7–16.
- [91] Ronald H Rosenberg. *Open architecture computer systems*. IEEE Computer Society Press, 1994.
- [92] M. Rudolph. "LinuxBIOS: Open Source Boot Firmware". In: *Proceedings of the Linux Symposium*. 2007, pp. 159–167. URL: <https://ols.fedoraproject.org/OLS/Reprints-2007/rudolph-Reprint.pdf>.
- [93] M. E. Russinovich, D. A. Solomon, and A. Ionescu. *Windows Internals, Part 1*. 6th. Microsoft Press, 2012.
- [94] Anand Lal Shimpi. "The Bulldozer Review: AMD FX-8150 Tested". In: *AnandTech* (2011). URL: <https://www.anandtech.com/show/4955/the-bulldozer-review-amd-fx8150-tested>.
- [95] M. Shin and K. Lee. "Design and Implementation of a UEFI-Compliant Firmware Platform". In: *Journal of Computer Science and Technology* 26.2 (2011), pp. 219–230. DOI: 10.1007/s11390-011-0121-8.
- [96] Leonard J. Shustek. *In His Own Words: Gary Kildall*. Computer History Museum Blog. Accessed: August 16, 2024. 2016. URL: <https://computerhistory.org/blog/in-his-own-words-gary-kildall/>.
- [97] David Smith and Alice Chen. "Firmware as the New Hypervisor: A Virtualized Perspective". In: *Computer Security Review* 27.4 (2019), pp. 210–225.
- [98] John Smith. "Remote KVM-over-IP on the ASpeed AST2050". In: *Journal of Embedded Computing* 14.3 (2014), pp. 45–49.
- [99] Vilas Sridharan et al. "Memory Errors in Modern Systems: The Good, The Bad, and The Ugly". eng. In: *Computer architecture news* 43.1 (2015), pp. 297–310. ISSN: 0163-5964.
- [100] ASpeed Technology. "ASpeed AST2050: ARM926EJ-S Based BMC Architecture". In: *ASpeed Whitepaper* (2013). Accessed: 2024-08-21. URL: <https://www.aspeedtech.com/products.php?fPath=20&rId=29>.
- [101] ASpeed Technology. *ASpeed AST2050: Network Interface Controller for BMC*. Accessed: 2024-08-21. 2013. URL: <https://www.aspeedtech.com/products.php?fPath=20&rId=29>.
- [102] ASpeed Technology. *I/O Interfaces of the ASpeed AST2050*. Accessed: 2024-08-21. 2013. URL: <https://www.aspeedtech.com/products.php?fPath=20&rId=29>.
- [103] TianoCore Project. *TianoCore as a Coreboot Payload*. Accessed: 2024-08-17. 2024. URL: <https://doc.coreboot.org/payloads/tianocore.html>.
- [104] UEFI Forum. *What is UEFI?* Accessed: 2024-08-17. 2023. URL: <https://uefi.org/sites/default/files/resources/What%20is%20UEFI-Aug31-2023-Final.pdf>.
- [105] Sorbonne Université/CNRS. *Annuaire LIP6*. [Online; accessed 7-May-2024]. 2024. URL: <https://www.lip6.fr/recherche/resultat.php?keyword=&find=Rechercher+au+LIP6>.
- [106] Sorbonne Université/CNRS. *Laboratoire d'Informatique de Paris 6*. [Online; accessed 7-May-2024]. 2024. URL: <https://www.lip6.fr/>.
- [107] Ward Vandewege. "Coreboot: the view from the FSF". In: 2008.
- [108] M. Versen and W. Ernst. "Row hammer avoidance analysis of DDR3 SDRAM". eng. In: *Microelectronics and reliability* 114 (2020), pp. 113744–. ISSN: 0026-2714.
- [109] Vikings GmbH. *Vikings Hardware Recommendations for KGPE-D16*. Accessed: 2024-08-17. URL: <https://wiki.vikings.net/KGPE-D16>.
- [110] Dong Wang and Wei Yu Dong. "Attacking Intel UEFI by Using Cache Poisoning". eng. In: *Journal of physics. Conference series* 1187.4 (2019), pp. 42072–. ISSN: 1742-6588.
- [111] Muhammad Waqar et al. "DDR4 Data Channel Failure Due to DC Offset Caused by Intermittent Solder Ball Fracture in FBGA Package". eng. In: *IEEE access* 9 (2021), pp. 63002–63011. ISSN: 2169-3536.
- [112] Wikipedia contributors. *AGESA — Wikipedia, The Free Encyclopedia*. [Online; accessed 8-May-2024]. 2023. URL: <https://en.wikipedia.org/w/index.php?title=AGESA&oldid=1166805057>.
- [113] Wikipedia contributors. *AMD Platform Security Processor — Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=AMD_Platform_Security_Processor&oldid=1216563013.

- [114] Wikipedia contributors. *BIOS* — *Wikipedia, The Free Encyclopedia*. <https://en.wikipedia.org/w/index.php?title=BIOS&oldid=1240397019>. [Online; accessed 16-August-2024]. 2024.
- [115] Wikipedia contributors. *DDR3 SDRAM* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 8-May-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=DDR3_SDRAM&oldid=1207641521.
- [116] Wikipedia contributors. *Free software* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 30-January-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=Free_software&oldid=1196006316.
- [117] Wikipedia contributors. *Free Software Foundation* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=Free_Software_Foundation&oldid=1222269091.
- [118] Wikipedia contributors. *Free software movement* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 29-January-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=Free_software_movement&oldid=1197710495.
- [119] Wikipedia contributors. *GNU Free Documentation License* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 30-January-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=GNU_Free_Documentation_License&oldid=1193649968.
- [120] Wikipedia contributors. *GNU General Public License* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 30-January-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=GNU_General_Public_License&oldid=1199241605.
- [121] Wikipedia contributors. *GNU GRUB* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=GNU_GRUB&oldid=1217643156.
- [122] Wikipedia contributors. *GNU Project* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=GNU_Project&oldid=1205139455.
- [123] Wikipedia contributors. *Intel Management Engine* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=Intel_Management_Engine&oldid=1216703991.
- [124] Wikipedia contributors. *Laboratoire d'Informatique de Paris 6* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=Laboratoire_d%27Informatique_de_Paris_6&oldid=1222525180.
- [125] Wikipedia contributors. *Non-disclosure agreement* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 8-May-2024]. 2023. URL: https://en.wikipedia.org/w/index.php?title=Non-disclosure_agreement&oldid=1183749255.
- [126] Wikipedia contributors. *Northbridge (computing)* — *Wikipedia, The Free Encyclopedia*. [https://en.wikipedia.org/w/index.php?title=Northbridge_\(computing\)&oldid=1231509957](https://en.wikipedia.org/w/index.php?title=Northbridge_(computing)&oldid=1231509957). [Online; accessed 17-August-2024]. 2024.
- [127] Wikipedia contributors. *OpenBMC* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 8-May-2024]. 2023. URL: <https://en.wikipedia.org/w/index.php?title=OpenBMC&oldid=1183698628>.
- [128] Wikipedia contributors. *SeaBIOS* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2023. URL: <https://en.wikipedia.org/w/index.php?title=SeaBIOS&oldid=1179465237>.
- [129] Wikipedia contributors. *Southbridge (computing)* — *Wikipedia, The Free Encyclopedia*. [https://en.wikipedia.org/w/index.php?title=Southbridge_\(computing\)&oldid=1239483618](https://en.wikipedia.org/w/index.php?title=Southbridge_(computing)&oldid=1239483618). [Online; accessed 17-August-2024]. 2024.
- [130] Wikipedia contributors. *The Free Software Definition* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 29-January-2024]. 2023. URL: https://en.wikipedia.org/w/index.php?title=The_Free_Software_Definition&oldid=1192713194.
- [131] Wikipedia contributors. *The Open Source Definition* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 30-January-2024]. 2023. URL: https://en.wikipedia.org/w/index.php?title=The_Open_Source_Definition&oldid=1191447775.
- [132] Wikipedia contributors. *X86* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 7-May-2024]. 2024. URL: <https://en.wikipedia.org/w/index.php?title=X86&oldid=1221800539>.

- [133] Winbond Electronics Corporation. *WINBOND W83667HG-A Datasheet*. Accessed: 2024-08-17. URL: <https://www.winbond.com/>.
- [134] K. Wolf. "Modern Boot Firmware: Moving from BIOS to UEFI". In: *IEEE Computer Society* 39.5 (2006), pp. 42–47. DOI: 10.1109/MC.2006.156.
- [135] Jinhui Yi, Mingfu Wang, and Lidong Bai. "Design of DDR3 SDRAM read-write controller based on FPGA". eng. In: *Journal of physics. Conference series* 1846.1 (2021), pp. 12046–. ISSN: 1742-6588.

GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.

<<https://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “**Document**”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “**you**”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “**Modified Version**” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “**Secondary Section**” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “**Invariant Sections**” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “**Cover Texts**” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “**Transparent**” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “**Opaque**”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “**Title Page**” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “**publisher**” means any person or entity that distributes copies of the Document to the public.

A section “**Entitled XYZ**” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “**Acknowledgements**”, “**Dedications**”, “**Endorsements**”, or “**History**”.) To “**Preserve the Title**” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document’s license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the

latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/licenses/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008. The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright © YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with . . . Texts.” line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.