

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".

Contents

Abstract	5
1 Introduction to firmware and BIOS evolution	6
1.1 Historical context of BIOS	6
1.1.1 Definition and origin	6
1.1.2 Functionalities and limitations	7
1.2 Modern BIOS and UEFI	8
1.2.1 Transition from traditional BIOS to UEFI (Unified Extensible Firmware Interface)	8
1.2.2 An other way with <i>coreboot</i>	8
1.3 Shift in firmware responsibilities	10
2 Characteristics of ASUS KGPE-D16 mainboard	11
2.1 Overview of ASUS KGPE-D16 hardware	12
2.2 Chipset	13
2.3 Processors	15
2.4 Baseboard Management Controller	16
3 Key components in modern firmware [WIP]	18
3.1 General structure of coreboot	18
3.1.1 Bootblock stage	19
3.1.2 Romstage	19
3.1.3 Ramstage	20
3.1.4 Payload	20
3.2 Advanced Configuration and Power Interface	20
3.3 System Management Mode	21
3.4 AMD Platform Security Processor and Intel Management Engine	21
4 Memory initialization and training algorithms [WIP]	23
4.1 Importance of memory initialization	23
4.2 Memory training algorithms	23
4.3 Practical examples	24
5 Firmware and hardware virtualization [WIP]	25
5.1 Introduction to hardware virtualization	25
5.2 Role of BIOS/UEFI in virtualization	25
5.3 Security and freedom considerations	25
5.4 Future trends in firmware and virtualization	25
Conclusion	26
5.5 Summary of key points	26
5.6 Call for action	26

Bibliography

27

GNU Free Documentation License

32

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, discussing how modern firmware supports and enhances virtualized environments. Security considerations and future trends in firmware development are also addressed, emphasizing the need for continued research and advocacy for free software-compatible hardware.

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 [34]. The term BIOS itself was coined by Gary Kildall, who developed the CP/M (Control Program for Microcomputers) operating system [73]. 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: 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 [36][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 [34].

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 [88]. 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 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 [24]. 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 [30][70]. 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 [30]. UEFI was designed to address the shortcomings of the traditional BIOS, providing several key improvements.



Figure 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 [24][72]. 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 [25][70].

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 [24]. 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][18].

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 [30].

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[64]. 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 [69]. 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[64]. 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 [63].



Figure 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 [54]. For example, it provides a *verified boot* function, allowing to prevent rootkits and other attacks based on firmware modifications [62]. 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 [51], since these blobs are used for certain functionalities such as memory initialization and hardware management.



Figure 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 [52].

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 [30][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 [30]. 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 [30]. 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 6: 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 [33]. Indeed, this mainboard can be operated with a fully free firmware such as GNU Boot [53].

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 [ASUS_kgpe_d16_manual]. 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 7: 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 8: 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][100].

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][103]. It is essentially an ALink bus controller and includes the hardware interrupt controller, the IOAPIC. Interrupts from peripheral always pass through the northbridge (fig. 9), since it translates ALink to HyperTransport for the CPUs and contains the IOMMU [11].

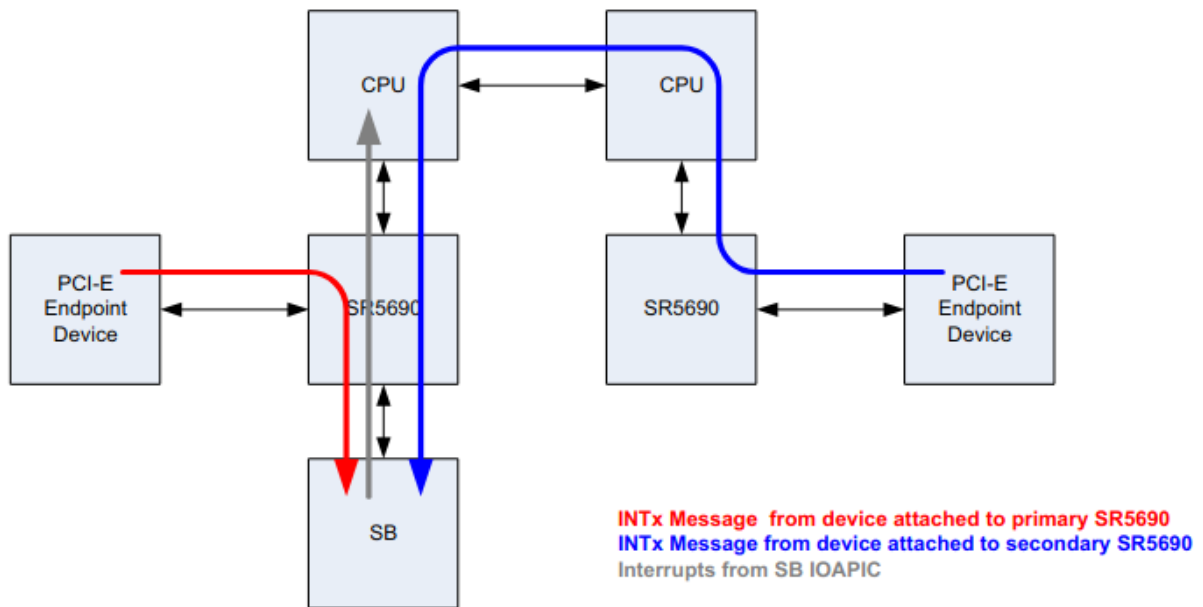


Figure 9: Functional diagram presenting the IOAPIC function of the SP5100, ASUS (2011)

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 [107]. 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 [58]. 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. 10) [11].

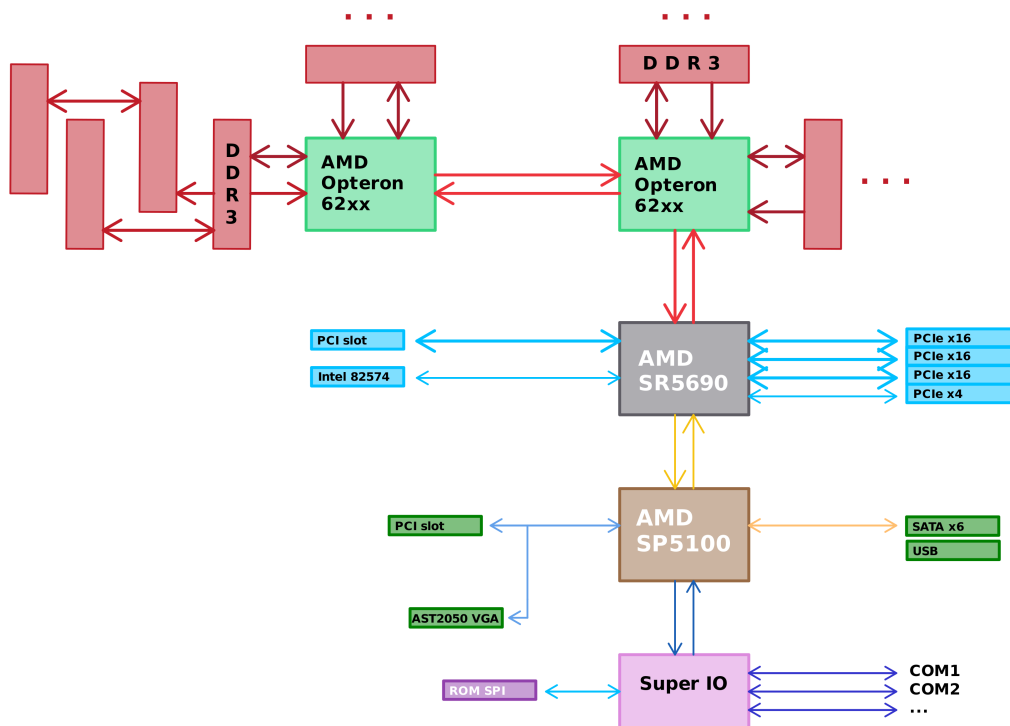


Figure 10: Functional diagram of the KGPE-D16 chipset (CC BY-SA 4.0, 2024)

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 [83].




Figure 11: 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. 11, 12) [7][71]. 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][40].

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 12: 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 [71]. 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 [76]. 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 [74]. 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 [26]. 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 [77]. 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 [78]. 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 [101], a libre firmware that can be run on the AST2050 thanks to Raptor Engineering [65], 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 [43].

Chapter 3

Key components in modern firmware [WIP]

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. 13) [62].




Figure 13: *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 14).

Each stage is compiled as a separate binary and inserted into the CBFS with custom compression. The **bootblock** 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).

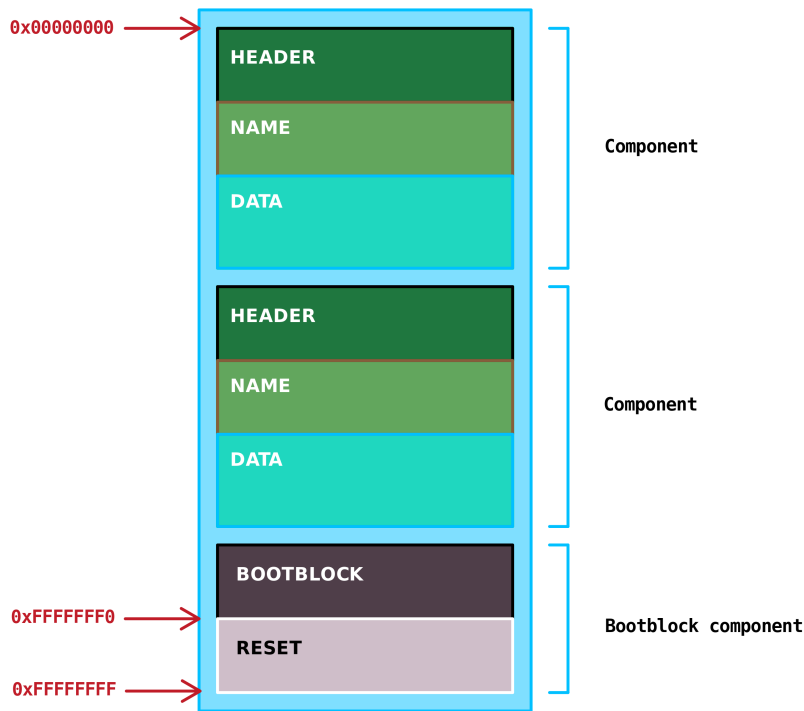


Figure 14: *coreboot* ROM architecture (CC BY-SA 4.0, 2024)

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.

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.

The **bootblock** occupies the last 20k 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 in the ROM. This block is a mandatory component of the ROM as it also contains the entry point of the firmware.

Upon startup, this stage is responsible for the initial hardware setup, which involves identifying and configuring the CPU. This process is particularly significant for AMD Family 10h/15h processors, where the firmware sets up the Bootstrap Processor (BSP) and executes the necessary initialization routines. Using the BSP, it enables the processor's cache, a small but fast type of memory that stores frequently accessed data, improving overall system speed by reducing data access time. However, the cache is there used as memory since DDR DIMMs are not available yet. It is done by programming the **Memory Type Range Registers (MTRRs)**, which define how different ranges of system memory are accessed, such as whether they are cacheable or non-cacheable, used to optimize memory performance on normal operation. The firmware will then set the stack pointer, allocate memory for the BSS, and decompress and load the next stage. On x86 platforms, this process also includes updating the CPU microcode, initializing the timer, and transitioning from 16-bit real mode to 32-bit protected mode. The **bootblock** is responsible for loading the romstage, or the verstage if verified boot is enabled.

3.1.2 Romstage

The firmware also configures the Advanced Programmable Interrupt Controller (APIC), which manages how the processor handles interrupts to ensure the system can respond efficiently to hardware and software events. Lastly, it sets up the routing for HyperTransport (HT) technology, a high-speed communication protocol used

for data exchange between the processor and the northbridge, ensuring that data flows smoothly between these components. Memory training and optimization are key functions of the firmware. During this process, the firmware adjusts 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 for a large amount of CPU cores, as supported by the KGPE-D16.

3.1.3 Ramstage

Effective resource allocation is essential for system stability, particularly in complex configurations involving multiple CPUs and peripherals. The firmware manages resource allocation, resolving any conflicts between hardware components to prevent resource contention and ensure smooth operation.

Power management is another critical aspect managed by the firmware. Through the Advanced Configuration and Power Interface (ACPI), the firmware can adjust the power states of various components, optimizing energy consumption across the system. This not only reduces power usage but also minimizes heat generation, which is vital for maintaining system longevity and efficiency.

Security is a major concern in modern systems, and the KGPE-D16s firmware includes features designed to protect the system from unauthorized access and maintain the integrity of its operations. This includes support for IOMMU, which is crucial for preventing unauthorized direct memory access (DMA) attacks, particularly in virtualized environments.

3.1.4 Payload

TODO

3.2 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[21]. 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[37]. 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[61]. 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[21]. 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[37]. 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.

ASUS KGPE-D16 Example: The ASUS KGPE-D16 mainboard, which is designed for server and high-performance computing environments, utilizes 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[ASUS_kgpe_d16_manual].

3.3 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 must be isolated from the OS and applications. Introduced by Intel, SMM operates in an environment separate from the main operating system, offering a secure and controlled space for executing sensitive operations[46].

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[39]. 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[44]. 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 securely apply firmware updates and respond to security threats without being affected by potentially compromised system software[27]. 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[48].

ASUS KGPE-D16 Example: The ASUS KGPE-D16 mainboard utilizes SMM to perform critical management tasks that need to be isolated 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*, an open-source firmware project, supports SMM, but its implementation is typically minimal compared to traditional BIOS/UEFI systems. 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[17]. 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 [66]. 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[57].

3.4 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[45]. 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[45]. 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 microcontroller 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) [22]. 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.

The Intel ME, however, 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 [35]. 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.

ASUS KGPE-D16 Example: The ASUS KGPE-D16 mainboard does not include the AMD Platform Security Processor (PSP) or the Intel ME.

Chapter 4

Memory initialization and training algorithms [WIP]

4.1 Importance of memory initialization

- Steps involved in initializing the memory controller
- Critical role in system stability and performance
- **ASUS KGPE-D16 Example:** Memory initialization process on the KGPE-D16 mainboard

Memory training involves several steps:

1. ****Detection and Initialization**:** The BIOS detects the installed memory modules, determining their size, speed, and type.
2. ****Configuration and Timing Setup**:** The BIOS configures the memory controller settings, including timings for memory access such as CAS latency, RAS to CAS delay, and other parameters[24].
3. ****Training and Calibration**:** The BIOS performs tests and adjustments to calibrate the memory system, ensuring stable operation at optimal speeds by adjusting signal voltages and testing data integrity[108].

These steps are crucial for modern systems, where improper memory configuration can lead to instability, data corruption, or suboptimal performance.

Memory timings, such as CAS latency, RAS to CAS delay, and others, must be finely tuned to ensure optimal performance. The BIOS uses a combination of predefined profiles and dynamic adjustments to achieve the best balance between speed and stability. Advanced timing optimization involves setting these parameters to ensure that memory operations are performed with minimal latency and maximum throughput[70].

4.2 Memory training algorithms

- Techniques used for training memory
- Optimization of timings and voltage settings
- Challenges in multi-core CPU environments
- **ASUS KGPE-D16 Example:** Specific algorithms used for memory training in the mainboard and their performance outcomes

To optimize memory performance, the BIOS employs various training algorithms and calibration techniques. These methods test the memory under different conditions and make necessary adjustments to improve stability and efficiency. Key techniques include voltage adjustments, data integrity testing, and signal timing calibration[72].

Voltage adjustments involve tweaking the power supplied to the memory modules to ensure reliable operation. Data integrity testing checks that data can be accurately read and written, while signal timing calibration fine-tunes the delays between different memory operations to minimize latency.

4.3 Practical examples

- Real-world scenarios where firmware played a crucial role in system performance
- Analysis of firmware updates and their impact on the KGPE-D16 mainboard
- User experiences and testimonials highlighting the importance of firmware
- **ASUS KGPE-D16 Example:** Specific case studies and firmware updates for the mainboard

Chapter 5

Firmware and hardware virtualization [WIP]

5.1 Introduction to hardware virtualization

- Definition and purpose of virtualization
- How firmware interacts with virtualized environments
- **ASUS KGPE-D16 Example:** Virtualization capabilities and performance on the mainboard

5.2 Role of BIOS/UEFI in virtualization

- Initialization and configuration for virtual machines
- Resource allocation and management
- **ASUS KGPE-D16 Example:** BIOS/UEFI settings and their impact on virtualization efficiency on the KGPE-D16

5.3 Security and freedom considerations

- Security risks associated with virtualization
- Measures taken by firmware to mitigate risks
- **ASUS KGPE-D16 Example:** Security measures implemented in the mainboard's firmware to support secure virtualization

5.4 Future trends in firmware and virtualization

- Emerging advancements and their impact on firmware
- Predictions for the evolution of BIOS/UEFI in virtualization
- **ASUS KGPE-D16 Example:** Potential future firmware updates and their expected impact on the mainboard's virtualization capabilities

Conclusion [WIP]

5.5 Summary of key points

- Recap of the evolution and current state of firmware
- Importance of understanding modern BIOS functionalities
- **ASUS KGPE-D16 Example:** Summary of the mainboard's features and firmware contributions

5.6 Call for action

- Advocacy for free software-compatible hardware
- Encouraging research and development in libre firmware solutions
- A libre BIOS is very important[81].

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] R. E. Brown et al. "LinuxBIOS as an Open-Source Firmware Alternative". In: *Proceedings of the 2003 Linux Symposium*. 2003.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] Intel Corporation. *Advanced Configuration and Power Interface (ACPI) Specification*. Intel Corporation, 1996. URL: <https://uefi.org/specifications>.

- [22] Intel Corporation. *Intel Converged Security and Management Engine (CSME) Security White Paper*. Tech. rep. 2020. URL: <https://www.intel.com/content/dam/www/public/us/en/security-advisory/documents/intel-csme-security-white-paper.pdf>.
- [23] Intel Corporation. *System Management Mode*. Tech. rep. 2016. URL: <https://www.intel.com/content/www/us/en/developer/articles/technical/system-management-mode.html>.
- [24] Intel Corporation. *Unified Extensible Firmware Interface (UEFI)*. <https://www.intel.com/content/www/us/en/architecture-and-technology/unified-extensible-firmware-interface.html>. 2020.
- [25] Microsoft Corporation. *UEFI Firmware*. <https://docs.microsoft.com/en-us/windows-hardware/drivers/bringup/uefi-firmware>. 2019.
- [26] Jane Doe. "DDR2 Memory Controller in the ASpeed AST2050". In: *Memory Systems Review* 22.2 (2015), pp. 33–40.
- [27] 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>.
- [28] 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.
- [29] 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 informatičnaja tekhnologija* 29.4 (2022), pp. 27–41. ISSN: 2074-7128.
- [30] UEFI Forum. *UEFI Specification*. <https://uefi.org/specifications>. 2021.
- [31] UEFI Forum. *Unified Extensible Firmware Interface*. <https://uefi.org/>. 2024.
- [32] 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>.
- [33] Free Software Foundation. *Respects Your Freedom (RYF) Certification*. Accessed: 2024-08-17. 2017. URL: <https://ryf.fsf.org/products/VikingsD16>.
- [34] Paul Freiberger and Michael Swaine. *Fire in the Valley: The Birth and Death of the Personal Computer*. McGraw-Hill, 2000.
- [35] 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>.
- [36] Jimmy Grewal. *The Creation of the IBM PC*. Armonk Institute. 2024.
- [37] Michael Gschwind. "Advanced Configuration and Power Interface: The Operating System Perspective". In: *IEEE Micro* 20 (2000), pp. 82–89. DOI: 10.1109/40.888702.
- [38] 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.
- [39] 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>.
- [40] 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.
- [41] Micron Technology Inc. *Technical Note: DDR3 ZQ Calibration*. TN-41-02. 2008.
- [42] 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.
- [43] Michael Jones. "Customizing OpenBMC for ASpeed AST2050". In: *Open Source Firmware Journal* 5.1 (2017), pp. 12–18.

- [44] 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.
- [45] 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.
- [46] Ronald D. Krebs, Vincent Zimmer, and Suresh Marisetty. *Beyond BIOS: Developing with the Unified Extensible Firmware Interface*. 3rd. Intel Press, 2017. ISBN: 978-0974364906.
- [47] Christoph Lameter. “NUMA (Non-Uniform Memory Access): An Overview”. In: *Queue* 11 (July 2013). DOI: 10.1145/2508834.2513149.
- [48] 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.
- [49] 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.
- [50] Samsung Electronics Co. Ltd. *DDR3 SDRAM Specification Rev 1.4*. TN-41-02. Nov. 2011.
- [51] GNU Boot project maintainers. *Frequently Asked Questions*. <https://www.gnu.org/software/gnuboot/web/faq.html>. Accessed: 2024-07-23.
- [52] GNU Boot project maintainers. *GNU Boot — Free your BIOS today!* [Online; accessed 7-May-2024]. 2024. URL: <https://www.gnu.org/software/gnuboot/>.
- [53] GNU Boot project maintainers. *GNU Boot — Status*. [Online; accessed 7-May-2024]. 2024. URL: <https://www.gnu.org/software/gnuboot/web/status.html>.
- [54] 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.
- [55] Ron Minnich. “coreboot: Status and some history”. In: 2006.
- [56] Ron Minnich, Stefan Reinauer, and Patrick Georgi. “coreboot: Open-Source Firmware Platform”. In: *Google Research* (2017). URL: <https://research.google/pubs/pub45424/>.
- [57] Benjamin Mohr. *A Comparative Analysis of Bootloaders*. Tech. rep. University of Freiburg, 2012.
- [58] Nuvoton Technology Corporation. *Nuvoton W83795G/ADG Hardware Monitor Datasheet*. Accessed: 2024-08-17. URL: <https://www.nuvoton.com/>.
- [59] 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.
- [60] Timothy Pearson. “The World Beyond x86”. In: 2014.
- [61] ACPICA Project. *ACPI Component Architecture Programmer Reference*. Accessed: 2024-08-03. 2017. URL: <https://acpica.org/documentation>.
- [62] coreboot Project. *coreboot Documentation*. 2023. URL: <https://doc.coreboot.org/>.
- [63] coreboot project. *coreboot Payloads*. <https://www.coreboot.org/Payloads>. Accessed: 2024-07-23.
- [64] coreboot project. *coreboot: Open Source Firmware*. <https://www.coreboot.org/>. Accessed: 2024-07-23.
- [65] Raptor Engineering LLC. *Raptor Engineering website*. [Online; accessed 8-May-2024]. 2009-2024. URL: <https://raptorengineering.com/>.
- [66] Stefan Reinauer et al. “The coreboot Open Source BIOS - A Review”. In: *Usenix Annual Technical Conference*. 2008.
- [67] 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.
- [68] Ronald H Rosenberg. *Open architecture computer systems*. IEEE Computer Society Press, 1994.
- [69] 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>.
- [70] M. E. Russinovich, D. A. Solomon, and A. Ionescu. *Windows Internals, Part 1*. 6th. Microsoft Press, 2012.

- [71] 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>.
- [72] 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.
- [73] 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/>.
- [74] John Smith. "Remote KVM-over-IP on the ASpeed AST2050". In: *Journal of Embedded Computing* 14.3 (2014), pp. 45–49.
- [75] 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.
- [76] 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>.
- [77] 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>.
- [78] 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>.
- [79] Sorbonne Université/CNRS. *Annuaire LIP6*. [Online; accessed 7-May-2024]. 2024. URL: <https://www.lip6.fr/recherche/resultat.php?keyword=&find=Rechercher+au+LIP6>.
- [80] Sorbonne Université/CNRS. *Laboratoire d'Informatique de Paris 6*. [Online; accessed 7-May-2024]. 2024. URL: <https://www.lip6.fr/>.
- [81] Ward Vandewege. "Coreboot: the view from the FSF". In: 2008.
- [82] M. Versen and W. Ernst. "Row hammer avoidance analysis of DDR3 SDRAM". eng. In: *Microelectronics and reliability* 114 (2020), pp. 113744–. ISSN: 0026-2714.
- [83] Vikings GmbH. *Vikings Hardware Recommendations for KGPE-D16*. Accessed: 2024-08-17. URL: <https://wiki.vikings.net/KGPE-D16>.
- [84] 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.
- [85] 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.
- [86] 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>.
- [87] 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.
- [88] Wikipedia contributors. *BIOS — Wikipedia, The Free Encyclopedia*. <https://en.wikipedia.org/w/index.php?title=BIOS&oldid=1240397019>. [Online; accessed 16-August-2024]. 2024.
- [89] 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.
- [90] 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.
- [91] 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.
- [92] 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.

- [93] 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.
- [94] 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.
- [95] 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.
- [96] 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.
- [97] 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.
- [98] 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.
- [99] 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.
- [100] 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.
- [101] 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>.
- [102] 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>.
- [103] 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.
- [104] 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.
- [105] 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.
- [106] 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>.
- [107] Winbond Electronics Corporation. *WINBOND W83667HG-A Datasheet*. Accessed: 2024-08-17. URL: <https://www.winbond.com/>.
- [108] 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.
- [109] 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/>](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.