UV Distributed Information Management
Summer term 2023

Supplementary Material

This document provides supplementary material regarding terms, topics, and questions that have been discussed during the lecture, and is supposed to support the students.

1 Data Management

This section covers supplementary material of part 1 on *Data Management*.

1.1 Computer Science Terminology

This section contains descriptions of all terms that have been discussed during the preliminary talk (Vorbesprechung, in german). Figure 1 summarizes the most important parts of a computer system (CPU, caches, main memory, and SSD/HDD) for our course.

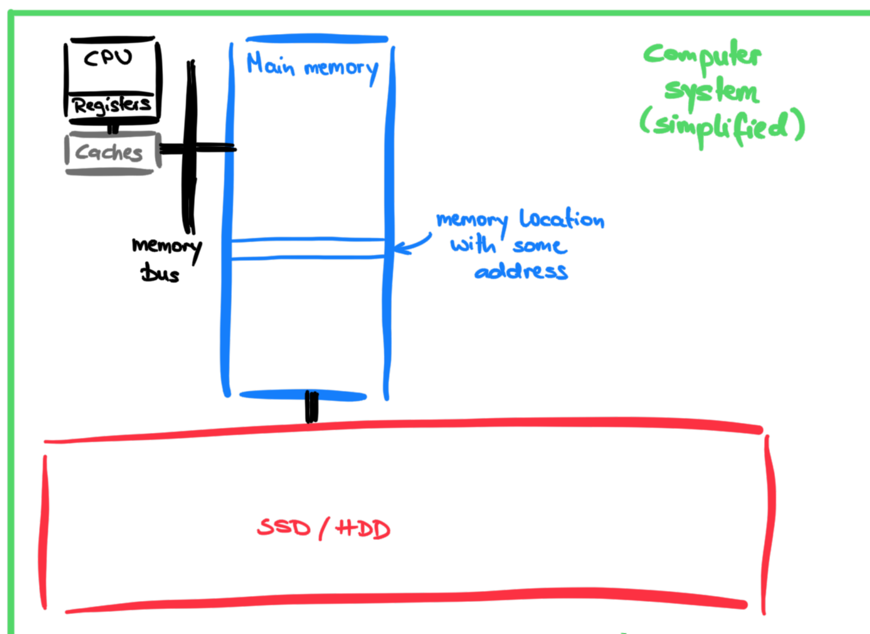

Figure 1: A computer system as geographical map (simplified).

Processor / CPU (Prozessor, in german) CPU[1] stands for *Central Processing Unit* and this is (as the name suggests) the main computational unit in a computer system. A CPU has a defined set of so-called *instructions* that basically denotes the language it understands (in binary format[2], i.e., a sequence of bits each of which can be 0 or 1). Here, it is important to realize that a CPU understands *exactly* and *only* this set of instructions and nothing else. When we switch on the power of our computer system, the CPU starts to run in a so-called *cycle* that naively executes one instruction after another, i.e., the CPU does not have any notion of intelligence. It simply executes the instructions we tell it to execute, e.g., by writing code in some programming language that is translated into the set of instructions our CPU understands. We will use the terms CPU and processor synonymously.

Main memory (RAM; Hauptspeicher, in german) Main memory[3] is a so-called *volatile* memory in a computer system that stores code and data during execution of an application. Typical main memory sizes range from $1 - 2$GB to sometimes > 100GB (nowadays; 1GB = 1Gigabyte = 2^{30}Bytes[4] = about a billion Bytes). Importantly, information that is stored on volatile memory is lost as soon as the power is switched off, i.e., *all* bits are set to 0. However, main memory is quite fast as it does not have to care about durability. Every single bit of code and data during the execution of an application is stored in main memory (at some point). When we execute an application/our code, the corresponding bits are transferred from main memory to the CPU instruction by instruction, each of which is then executed by the CPU.

HDD (Festplatte, in german) HDD[5] stands for *Hard Disk Drive* and denotes a durable, non-volatile memory that is used to store information permanently (even when the power is switched off). Typical HDD sizes range from tens of GB to sometimes tens or even hundreds of TB (nowadays; 1TB = 1Terabyte = 2^{40}Bytes = about a trillion Bytes). Although HDDs provide much more storage, they are orders of magnitude slower compared to main memory. Therefore, computer systems want to minimize accesses to information that is stored on the HDD if possible (and rather retrieve it from main memory). Nonetheless, sometimes it is simply not possible to avoid the HDD access because information that must not be lost is stored there. Once our computer system is running, information (i.e., bytes) is regularly transferred from the slower to the faster memory (e.g., from HDD to main memory) in order to speed up the overall system performance.

SSD SSD[6] stands for *Solid-State Drive* and denotes the de-facto successor technology to HDDs. SSDs provide properties that are similar to HDDs (i.e., durable, non-volatile memory and information is stored permanently) but the underlying technology is completely different (HDD: magnentic storage, SSD: flash storage). As a result, SSDs are faster than HDDs but still slower than main memory. Despite being faster than HDDs, computer systems still try to minimize access to information that is stored on SSDs (as it is still slow compared to main memory access). Typical SSD sizes range from tens of GB to sometimes a few TB (nowadays).

Cache (Zwischenspeicher, in german) A cache[7] is essentially a small but very fast piece of memory in a computer system, which is used to temporarily store data that has been

[1] Central Processing Unit: `https://en.wikipedia.org/wiki/Central_processing_unit`
[2] Binary format: `https://en.wikipedia.org/wiki/Binary_file`
[3] Main memory: `https://en.wikipedia.org/wiki/Computer_data_storage#Primary_storage`
[4] Byte: `https://en.wikipedia.org/wiki/Byte`
[5] HDD: `https://en.wikipedia.org/wiki/Hard_disk_drive`
[6] SSD: `https://en.wikipedia.org/wiki/Solid-state_drive`
[7] Cache: `https://en.wikipedia.org/wiki/Cache_(computing)`

transferred from main memory to the CPU. The objective of a cache is to speed up future accesses to the data/code that is *cached* (i.e., currently stored in the cache) and to avoid the long way to main memory. Modern computer systems typically have multiple cache levels (L1, L2, L3): For $x < y$, Lx is smaller and faster than Ly since Lx is physically (or geographically if we consider the computer systems as a map) closer to the CPU than Ly. Typical cache sizes range from tens of KB (L1; 1KB = 1Kilobyte = 2^{10}Bytes = about a thousand Bytes) to tens or hundreds of MB (L3; 1MB = 1Megabyte = 2^{20}Bytes = about one million Bytes).

Memory hierarchy (Speicherhierarchie, in german) This term refers to the hierarchy of memory[8] units that are present in basically every (modern) computer system. We have already learned about main memory, caches, SSDs, and HDDs, but there are even more memory units like CPU registers. The memory hierarchy is a pyramid that tells us about the *size* (i.e., storage capacity), the *speed* (to access some code/data that is stored there), and the volatility (i.e., is the content of the memory unit lost or not when we switch off our system) of a particular memory unit. The smallest, fastest, volatile, and (typically) most expensive memory units are located at the top of this hierarchy, whereas the largest, slowest, non-volatile, and (typically) cheapest memory units are at the bottom. Figure 2 visualizes the hierarchy: At the top, we find the CPU registers, which are tiny memory units (only a few Bytes) that are *inside* the CPU. Subsequently, we find the different levels of caches, main memory, SSDs, and HDDs (in this order). Further below, we find very large but very slow memory units like magnetic bands that are used for long-term backups (over decades). There exist more types of memory units in a modern computer system but for our purpose the memory units shown in Figure 2 suffice. With this hierarchy, we can directly see why systems typically want to minimize the accesses to HDD and even SSD: The performance gap between main memory and HDD is about 5 orders of magnitude, i.e., an access to main memory is 10^5 = about $100,000$ times faster than the equivalent access to HDD (and about 10^2 = about 100 times faster than the equivalent access to SSD). On the one hand, this is because main memory is based on a faster technology that SSDs and HDDs, respectively. On the other hand, this is also due to the physical (or geographical) location *within* our computer system, i.e., memory units that are physically (or geographically) located nearby the CPU are faster because every signal in a computer system "*travels*" at the speed of light ($299,792,458$m/s; we cannot overcome this physical limitation). Consequently, data/code that resides in, say, the L1 cache can be accessed much faster than from memory units that are farther away from the CPU. Figure 1 illustrates this principle in a computer system. As a real-world analogy, we can consider traveling by car at a constant speed of 100km/h from Salzburg to (a) Linz (about 130km) and (b) Vienna (about 300km). Our journey to Linz will take us about 1.3 hours, whereas our journey to Vienna will take us about 3 hours. A similar reasoning can be applied for the other levels in our hierarchy.

Physical level (Physische Ebene, in german) In a computer system, this level (or layer) refers to the physical components, e.g., main memory, the CPU, and HDDs, and denotes the lowest level where *everything* is encoded in bits and Bytes (i.e., a sequence of 8 consecutive bits). At this level, a computer system is just an electrical circuit (or many interconnected electrical circuits) that modifies the bits. The fact that we can use a computer system as we do nowadays is based on additional levels that hide the complexity of this level from us.

Instruction(s) (Instruktion(en), in german) We already know that a CPU has a defined

[8] Memory hierarchy: `https://en.wikipedia.org/wiki/Memory_hierarchy`

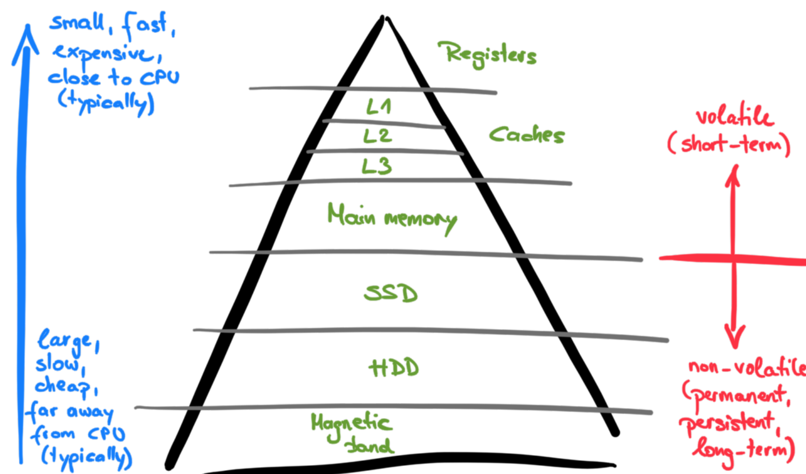

Figure 2: The portion of the memory hierarchy that is relevant for our course.

set of so-called *instructions* that basically denotes the language it understands in binary format. An instruction[9] has a defined behavior and represents an operation *encoded* as sequence of bits (binary format). (Simplified) example instructions include `ADD two numbers and store the result in a register`, or `LOAD the content of the location in memory that refers to a particular variable`, or `JUMP/BRANCH to some other instruction` (to represent control-flow, i.e., encode `if-else`, loops, and function calls in high-level programming languages). Every code that we write in a (high-level) programming language is (at some point) translated into a sequence of such instructions, and instructions effectively determine the behavior our application and, ultimately, our computer system. Once press the power button, it starts to execute the very first instruction until the startup process is finished. Then, everything we do on our machine (e.g., navigate with the mouse pointer, write some document, browse through the internet, or execute our own code) results in many instructions ($> 10^9$ instructions per second) that are execute by the CPU in the background and we observe the result of these instructions (e.g., the mouse pointer moves, a character appears in our document, a website is loading, or our code makes a particular step).

Operation(s) (Operation(en), in german) An operation[10] is the result of an instruction that is executed by the CPU. For example, an instruction may encode addition of two numbers (i.e., the `ADD` instruction) with the result being stored in some register. The corresponding operation is "Addition" and the execution of this instruction changes the value of a register. In this sense, an operation can be seen as the behavior of our computer system when executing the corresponding instruction. Furthermore, instructions are typically "small", i.e., they serve a very restricted purpose. If we want to add 5 numbers and store the result in a register, then we need to execute multiple `ADD` instructions consecutively because one `ADD` instruction can only add two numbers at a time. From this point of view, operations may also be considered to contain multiple instructions (depending on the context).

Address (Adresse, in german) An address[11] in a computer system uniquely identifies a particular location in memory. Abstractly, memory can be seen as a big box with many tiny

[9]Instruction(s): `https://en.wikipedia.org/wiki/Instruction_set_architecture#Instructions`
[10]Operation(s) in computer systems: `https://en.wikipedia.org/wiki/Operator_(computer_programming)`
[11]Memory address: `https://en.wikipedia.org/wiki/Memory_address`

slots each of which can, for example, store 64 bits. In order to refer unambiguously to a particular slot, the slots are numbered consecutively from 0 to $n - 1$ (for n slots), just like addresses in the real world. If an application uses a variable, then each variable is associated with one or multiple such slots in memory. Whenever the variable is read or modified, the computer system then automatically accesses the corresponding slot(s) in memory and retrieves or overwrites the content of the slot(s). The storage capacity of our memory is then the sum of all these tiny slots in this box.

Parallelism (Parallelismus, in german) (Hardware) Parallelism[12] refers to the *physical* ability of a computer system to perform multiple actions truly *simultaneously*. This is only possible if there exist multiple physical resources (typically CPUs) that can independently work on different tasks. As a real-world analogy, we can consider k students in class (the CPUs) that can perform different tasks independently (from start to finish, without interruption). Now consider $10 \cdot k$ tasks that we need to work on during class. A single student has to work on every single task one after another. Contrarily, k students can work on k tasks *in parallel*. Therefore, to overall time to finish $10 \cdot k$ tasks is 10 times lower for k students, i.e., we get a 10× speed up in our system.

Concurrency (Nebenläufigkeit, in german) This term refers to the *illusion* that multiple actions seemingly happen simultaneously, but they actually do *not* happen in parallel because there are not enough physical resources for parallelism. Consider a single CPU that can only execute a single instruction at a time and k tasks that are supposed to be executed simultaneously (e.g., we use multiple applications like a browser, an editor, an email client, and so on). In order to create the illusion of seemingly parallel execution of all k tasks, the instructions of each task are divided into blocks, e.g., one block may contain 10 instructions. The CPU starts to execute the first 10 instructions of task_1. Then, the CPU *switches* to task_2 and executes the first 10 instructions of task_2. The same principle is applied until the first 10 instructions have been executed for every task, and then the CPU starts to continue in this manner with the next block of instructions for task_1 (i.e., instructions 11-20 of task_1). In computer systems, a CPU can typically execute $> 10^9$ instruction per second, hence a human is not able to properly observe the fact that the tasks are not really execute in parallel but only *concurrently*[13]. Reusing the real-world analogy from our discussion on parallelism, this means that we only have a single student that works on k tasks *concurrently*. Each task may be to write down a particular Wikipedia article on a sheet of paper and a block of 10 instructions means that the student can write down 10 characters from the corresponding article. The student starts writing down the first 10 characters of article_1, then the first 10 characters of article_2 (on a different sheet of paper), and so on. After writing down the first 10 characters of each article, the student starts over and writes down the next 10 characters of each article until every article is fully replicated. In the end, the result is identical as if the student would have written down each article completely (one article after another).

Compiler (Übersetzer, in german) A compiler[14] refers to a tool in computer science that allows us to translate code that is written in a (high-level) programming language (e.g., Java) into a representation that can be executed on our CPU. That is, a compiler translated a given code into a sequence of bits (binary format), which encodes a sequence of instructions that is semantically equivalent to our code written in some programming language. The CPU is then able to execute one instruction of this sequence of instructions

[12] Parallelism: `https://en.wikipedia.org/wiki/Parallel_computing`
[13] Concurrency: `https://en.wikipedia.org/wiki/Concurrent_computing`
[14] Compiler: `https://en.wikipedia.org/wiki/Compiler`

after another, and this effectively reflects the behavior of our code. Typically, a portion of our code (e.g., `if-else`) is translated into multiple instructions.

Virtual machine (VM; virtuelle Maschine, in german) Sometimes, we want to execute code that is based on instructions that our native CPU does *not* understand, and it is no option to use a different hardware (i.e., CPU). For example, ARM processors[15] that is often used in smartphones is very different from x86 processors[16], which are common in Desktop computers. In this case, a virtual machine[17] (VM) may help: We can use a so-called *hypervisor* software[18] like VirtualBox[19] to create a virtual machine that provides the functionality of a "real" physical computer system (CPU, main memory, and the likes). However, all this functionality is implemented in software (rather than in hardware). For example, we can create a virtual machine that runs ARM instructions and *host* it using, e.g., VirtualBox, on our Desktop computer that is based on the x86 architecture. VirtualBox then acts as link between our x86 and our ARM architecture. Moreover, virtual machines are also quite useful for reproducibility, i.e., to reproduce a particular behavior on a different physical machine. The latter is also the reason why we use virtual machines in this course: It is easier for the instructor to reproduce an error that you encountered when you and the instructor use the same virtual machine.

Time (Zeit, in german) This term is mostly used in combination with *complexity*, i.e., computer scientist often argue about *time complexity*[20]. It refers to the amount of *runtime* it takes on a computer system to execute an algorithm[21] that serves a specific task. The time complexity is then the overall amount of runtime to finish the algorithm. Assuming that an operation in a computer system takes a constant amount of time, the time complexity is estimated by "counting" the number of operations the system needs to execute. We will use this term in a simplified manner (no formal definition) to argue whether something runs fast or rather slow in a computer system.

Space (Raum, in german) This term is also mostly used in combination with *complexity* and is orthogonal to time complexity. However, *space complexity*[22] refers to the amount of *memory* it takes on a computer system to execute an algorithm that serves a specific task. The space complexity is then the overall amount of memory that is consumed to finish the algorithm. Similar to time complexity, we will use t his term in an informal manner to argue whether something needs a large or rather small amount of memory in a computer system.

Performance (Performanz, in german) Performance[23] is a rather general term that refers to the amount of *meaningful* work in a computer system, i.e., how *effectively* a system can achieve its goal or task. In our course, performance is mostly related to runtime and memory, but also to availability and data transmission when discussing distributed systems.

Efficiency (Effizienz, in german) Efficiency[24] is often confused with performance and in some sense they are related. However, efficiency refers to how "*good*" a particular system

---

[15] The ARM architecture family: `https://en.wikipedia.org/wiki/ARM_architecture_family`
[16] The x86 architecture: `https://en.wikipedia.org/wiki/X86`
[17] Virtual Machine: `https://en.wikipedia.org/wiki/Virtual_machine`
[18] Hypervisor: `https://en.wikipedia.org/wiki/Hypervisor`
[19] VirtualBox: https://en.wikipedia.org/wiki/VirtualBox
[20] Time complexity: `https://en.wikipedia.org/wiki/Time_complexity`
[21] Algorithm: `https://en.wikipedia.org/wiki/Algorithm`
[22] Space complexity: `https://en.wikipedia.org/wiki/Space_complexity`
[23] Performance: `https://en.wikipedia.org/wiki/Computer_performance`
[24] Efficiency: `https://en.wikipedia.org/wiki/Efficiency`

or algorithm is able to achieve its goal or task in terms of resource utilization. In our course, efficiency is mostly related to runtime and memory. For example, a computer system is runtime- and memory-efficient if it achieves its goal or task while consuming a only a little amount of time and memory (i.e., it is not wasteful). Nowadays, systems are also often required to be *energy-efficiency*, i.e., to consume only little energy while achieving its goal or task. Of course, a system can also achieve its goal or task by following a different, less efficient strategy but the user/company may pay the price (e.g., by paying higher energy or hardware costs).

Scalability (Skalierbarkeit, in german) This term refers to the ability of a computer systems to handle larger amounts of work by possibly adding additional hardware resources such as CPUs, memory, or even full machines (in a distributed system). For example, assume our computer system is required to execute a task X_1 of size $|X_1|$, and our system is able to do this in t_1 seconds. Now, we double the amount of work, i.e., we execute a task X_2 of size $|X_2| = 2 \cdot |X_1|$. As a consequence, our system will probably not finish the task X^2 in t_1 seconds, but it will need more time $t_2 > t_1$. The factor $\frac{t_2}{t_1}$ is called *scalability factor*[25], and we want to keep this factor constant and small. Another option is to double the number of hardware resources, e.g., the number of CPUs in our system. Then, we may be able to keep our initial runtime t_1 even for task X_2, but only because we leverage additional computational power. In this case, we scale the CPUs and, again, we want to a constant and low number of additional CPUs that are required to keep t_1. As a real-world analogy, we can borrow the scenario of our discussion on parallelism: A student is able to achieve 1 task in t_1 seconds and we want to finish 10 tasks. If we stick with a single student, the overall time t_{10} will be about 10 times higher. Contrarily, if we *scale* out to 10 students each of which works on a different task independently in parallel, then we may end up with all 10 tasks being finished in t_1.

Redundancy (Redundanz, in german) In general, redundancy[26] refers to the duplication of a resource where each duplicate serves the same (or a very similar) purpose. Examples include multiple copies of the same data or multiple wires that connect a building to the internet. Depending on the specific context, redundancy can have a positive or negative connotation. On the positive side, redundancy allows us to tolerate errors in systems because we can simply use a different copy to achieve our goal or task. For example, if we want our system to be available 24 hours, 7 days a week, 365 days a year, we can have multiple different ways to connect to the internet. If our primary connection is down, we simply use one of the other connections. On the negative side, redundancy in data may also cause anomalies and inconsistencies as well as performance degradation because we need to keep all the data copies up to date.

Transparency (Transparenz, in german) This term has many different meanings even in the context of computer science. In our course, we will mainly use the adjective form of this term in combination with functionality. A functionality is considered *transparent* to the user if the complexity of this functionality is hidden from the user. For example, cloud services are typically transparent to the user because the user does not need to care about how and where the data is stored without being lost. Another example is the transparent translation from website URLs into IP addresses (i.e., addresses that uniquely identify a website in the world wide web).

---

[25] Scalability: `https://en.wikipedia.org/wiki/Scalability`

[26] Redundancy: `https://en.wikipedia.org/wiki/Data_redundancy` and `https://en.wikipedia.org/wiki/Redundancy_(engineering)`

Database system (DBS; Datenbanksystem, in german) A database system[27] (DBS) is a dedicated software system that manages (typically) large amounts of data and provides functionality to users to access and modify this data transparently.

Query (Anfrage, in german) Users typically interact with database systems using so-called *queries*. A query is formulated in a so-called *query language*, which is a well-defined computer language that allows a user to specify what she/he wants the database system to do, e.g., retrieve or modify a specific portion of the data or create new entries.

Transaction (Transaktion, in german) Transactions[28] are commonly known from electronic banking systems, where transaction are used to transfer some amount of money between two users. In database systems, this concept is generalized to encapsulate multiple operations into one *transaction*. Each transaction is typically an *atomic* and *independent* unit, that is, a transaction is executed in an all-or-nothing manner and transactions do not influence each other if they are executed concurrently. If the database system fails to execute a transaction successfully, then the transaction is reverted such that the system is in the state it has been before the transaction. Revisiting the electronic banking system, we intuitively observe that this is the desired behavior as we would not be happy if the money is lost during the transfer (i.e., neither of the two users has the money in the end).

Index An *index*[29] is a data structure in a database system that is used as shortcut to some portion of the data in order to speed up queries. There exist different types of indexes (depending on their purpose), but for our course, it suffices to view an index as black box that provides a shortcut. A real-world analogy can be found at the end of basically every technical book: The index lists words together with the page numbers where the respective term occurs.

Imperative (Imperativ, in german) In imperative programming[30] we specify *how* the computer system achieves its goal or task. In other words, we use statements to specify the commands the computer system must execute, i.e., we describe the single steps the computer must execute one after another to finally achieve its goal or task. In the most extreme case, we need to specify every single tiny step one by one. As a real-world analogy, we can tell a person to get us some milk, and we specify every single step the person needs to do (i.e., specify the supermarket, the shelve, the brand, and so on – the person has no freedom but naively follows our instructions).

Declarative (Deklarativ, in german) In contrast to imperative programming, declarative programming[31] refers to a programming paradigm where we only specify *what* we want the computer system to achieve, but do not care about how it achieves it. In the most extreme case, we do not know anything about how the system achieves its goal or task. Reusing the real-world analogy from our discussion on imperative programming, we would only tell the person to get us some milk and nothing else (i.e., we do not care about the specific supermarket, shelve, brand, and so on – the person can decide on her/his own).

Trade-off (Kompromiss, in german) This is a general term that is used in many different domains. In computer science, this term is important because computer systems often

[27] Database system: `https://en.wikipedia.org/wiki/Database`
[28] Transaction: `https://en.wikipedia.org/wiki/Database_transaction`
[29] Index: `https://en.wikipedia.org/wiki/Database_index`
[30] Imperative programming: `https://en.wikipedia.org/wiki/Imperative_programming`
[31] Declarative programming: `https://en.wikipedia.org/wiki/Declarative_programming`

have to *trade* some property in order to enable another property. A traditional trade-off[32] in computer systems is the so-called *time-space trade-off*: If we allow the system to consume more memory, then this may speed up to computation. Contrarily, if we allow the system to consume more time (e.g., be slower), then we may end up consuming less (or even no) memory at all. As a real-world analogy, we consider a person that wants to know the individual results of 10 calculations where the second calculation depends on the first calculation, the third depends on the second one, and so on. If we allow the person to use a sheet of paper to write down the results of previous calculations, she/he can simply reuse these results to calculation the remaining results. Contrarily, if we do *not* allow the person to use a sheet of paper , she/he must recalculate every previous result over and over again. Obviously, the usage of a sheet of paper (i.e., more memory) speeds up the process, i.e., we traded space for time. During our course, we will learn about different trade-offs in the design of a database system.

Distributed (Verteilt, in german) This course is about *distributed* information management, where *distributed*[33] refers to the property of a software system to be physically scattered over multiple physical machines each of which is located in different geographical locations. In order to communicate with each other, the physical machines are interconnected using a dedicated network or simply the internet. Typically, distributed systems provide their functionality to the user transparently.

1.2 Database Systems Terminology

Data integrity (Datenintegrität, in german) This term has different meanings. In the context of information security, data integrity refers to the *prevention of unauthorized modification of information*. In other words, data integrity defines the *correctness* or *validity* of the data upon modification (done by humans or machines). In our context, *consistency (or integrity) constraints* are used to describe conditions that must be satisfied for the data to be correct. The data are *consistent* (or *of integrity*) if all the constraints are satisfied.

Example: If an attribute A stores the balance of a bank account, we may want to disallow values that are smaller than EUR -5,000 (i.e., we allow a credit of at most EUR 5,000 per bank account). If a customer then tries to lend more than EUR 5,000 from our bank, this constraint is violated and the system may prevent it.

Key (Schlüssel, in german) In the relational model (record-based tables), we need a way to distinguish the rows. Let K be a subset of the attributes (columns) of a relation (table). K is called *super key* if the attributes of K suffice to uniquely identify a tuple (row) in the relation. K is a *candidate key* if K is a super key and K cannot be further reduced (i.e., no attributes can be removed) without losing the super key property. A candidate key that consists of attributes that are rarely subject to updates is typically chosen as *primary key*.

Example: The social security number, the bank account identifier, or the registration number at universities are prototypical examples of primary keys. However, also a combination of multiple attributes can serve as primary key, e.g., the combination of firstname, lastname, and birthyear.

Schema In our lecture, the schema refers to the overall design of a database, i.e., it defines the structure of the data (similar to a variable declaration in a programming language) and the relationships between the data. We distinguish between *relation* and *database schema*. The relation schema refers to the schema of a single relation (table; in the relational

[32] Trade-off: `https://en.wikipedia.org/wiki/Trade-off`

[33] Distributed systems: `https://en.wikipedia.org/wiki/Distributed_computing`

model), whereas the database schema refers to the collection of all relation schemata in the database.

Caveat: The term *schema* may have a different meaning in some database systems, e.g., a schema may subsume multiple tables. Therefore, we recommend to read the manual of the database system at hand. Nonetheless, we use the term *schema* as described above.

2 Pointers to Additional Material

In this section, we give pointers (references) to other material/videos that mostly provide more in-depth knowledge on certain topics related to database systems:

Database Courses at the University of Salzburg The Database Research Group at the University of Salzburg teaches many topics in the area of database systems in more depth. Please check the respective websites for more information on the specific topics:

- Databases 1: `https://dbresearch.uni-salzburg.at/teaching/2021ss/db1/`
- Databases 2: `https://dbresearch.uni-salzburg.at/teaching/2020ws/db2/`
- Databases Tuning: `https://dbresearch.uni-salzburg.at/teaching/2021ss/dbt/`
- Advanced Databases: `https://dbresearch.uni-salzburg.at/teaching/2020ws/adb/`
- Non-Standard Database Systems: `https://dbresearch.uni-salzburg.at/teaching/2021ss/nsdb/`
- Similarity Search: `https://dbresearch.uni-salzburg.at/teaching/2020ws/ssdb/`

Big Data Engineering on Youtube Jens Dittrich, a database systems and data science professor at the Saarland University, streams his lecture on Big Data Engineering live on Youtube (old videos are also available):
`https://www.youtube.com/user/jensdit/videos`

Youtube Channel of the CMU Database Group The Carnegie Mellon University is one of the top universities in Computer Science. The Database Group has many videos in their Youtube channel that cover most of the topics we covered in class in more detail:
`https://www.youtube.com/c/CMUDatabaseGroup/videos`

- Intro to Database Systems: `https://www.youtube.com/playlist?list=PLSE8ODhjZXjbohkNBWQs_otTrBTrjyohi`
- Advanced Database Systems: `https://www.youtube.com/playlist?list=PLSE8ODhjZXjasmrEd2_Yi1deeE360zv5O`