

1: Distributed operating system - Wikipedia

Our work ranges from finding new paradigms for expressing high-performance computing applications in high-level languages (Strout), to understanding aspects of the application when it executes, via visualization (Isaacs), to building models and systems that ensure the efficient execution of these programs (Lowenthal).

The single-instruction-multiple-data SIMD classification is analogous to doing the same operation repeatedly over a large data set. This is commonly done in signal processing applications. Multiple-instruction-single-data MISD is a rarely used classification. While computer architectures to deal with this were devised such as systolic arrays, few applications that fit this class materialized. Multiple-instruction-multiple-data MIMD programs are by far the most common type of parallel programs. According to David A. Patterson and John L. Hennessy, "Some machines are hybrids of these categories, of course, but this classic model has survived because it is simple, easy to understand, and gives a good first approximation. It is alsoâ€”perhaps because of its understandabilityâ€”the most widely used scheme. Bit-level parallelism From the advent of very-large-scale integration VLSI computer-chip fabrication technology in the s until about , speed-up in computer architecture was driven by doubling computer word size â€”the amount of information the processor can manipulate per cycle. Historically, 4-bit microprocessors were replaced with 8-bit, then bit, then bit microprocessors. This trend generally came to an end with the introduction of bit processors, which has been a standard in general-purpose computing for two decades. Not until the early s, with the advent of x architectures, did bit processors become commonplace. Instruction-level parallelism A canonical processor without pipeline. A canonical five-stage pipelined processor. A computer program is, in essence, a stream of instructions executed by a processor. These processors are known as subscalar processors. These instructions can be re-ordered and combined into groups which are then executed in parallel without changing the result of the program. This is known as instruction-level parallelism. Advances in instruction-level parallelism dominated computer architecture from the mids until the mids. These processors are known as scalar processors. The canonical example of a pipelined processor is a RISC processor, with five stages: The Pentium 4 processor had a stage pipeline. Most modern processors also have multiple execution units. These processors are known as superscalar processors. Instructions can be grouped together only if there is no data dependency between them. Scoreboarding and the Tomasulo algorithm which is similar to scoreboarding but makes use of register renaming are two of the most common techniques for implementing out-of-order execution and instruction-level parallelism. Task parallelism Task parallelisms is the characteristic of a parallel program that "entirely different calculations can be performed on either the same or different sets of data". Task parallelism involves the decomposition of a task into sub-tasks and then allocating each sub-task to a processor for execution. The processors would then execute these sub-tasks concurrently and often cooperatively. Task parallelism does not usually scale with the size of a problem. Distributed shared memory and memory virtualization combine the two approaches, where the processing element has its own local memory and access to the memory on non-local processors. Accesses to local memory are typically faster than accesses to non-local memory. A logical view of a non-uniform memory access NUMA architecture. Computer architectures in which each element of main memory can be accessed with equal latency and bandwidth are known as uniform memory access UMA systems. Typically, that can be achieved only by a shared memory system, in which the memory is not physically distributed. A system that does not have this property is known as a non-uniform memory access NUMA architecture. Distributed memory systems have non-uniform memory access. Computer systems make use of caches â€”small and fast memories located close to the processor which store temporary copies of memory values nearby in both the physical and logical sense. Parallel computer systems have difficulties with caches that may store the same value in more than one location, with the possibility of incorrect program execution. These computers require a cache coherency system, which keeps track of cached values and strategically purges them, thus ensuring correct program execution. Bus snooping is one of the most common methods for keeping track of which values are being accessed and thus should be purged. Designing large, high-performance cache coherence systems is a

very difficult problem in computer architecture. As a result, shared memory computer architectures do not scale as well as distributed memory systems do. Parallel computers based on interconnected networks need to have some kind of routing to enable the passing of messages between nodes that are not directly connected. The medium used for communication between the processors is likely to be hierarchical in large multiprocessor machines. Classes of parallel computers[edit] Parallel computers can be roughly classified according to the level at which the hardware supports parallelism. This classification is broadly analogous to the distance between basic computing nodes. These are not mutually exclusive; for example, clusters of symmetric multiprocessors are relatively common. Multi-core processor A multi-core processor is a processor that includes multiple processing units called "cores" on the same chip. This processor differs from a superscalar processor, which includes multiple execution units and can issue multiple instructions per clock cycle from one instruction stream thread ; in contrast, a multi-core processor can issue multiple instructions per clock cycle from multiple instruction streams. Each core in a multi-core processor can potentially be superscalar as well—that is, on every clock cycle, each core can issue multiple instructions from one thread. A processor capable of concurrent multithreading includes multiple execution units in the same processing unit—that is it has a superscalar architecture—and can issue multiple instructions per clock cycle from multiple threads. Temporal multithreading on the other hand includes a single execution unit in the same processing unit and can issue one instruction at a time from multiple threads. Symmetric multiprocessing A symmetric multiprocessor SMP is a computer system with multiple identical processors that share memory and connect via a bus. Distributed computing A distributed computer also known as a distributed memory multiprocessor is a distributed memory computer system in which the processing elements are connected by a network. Distributed computers are highly scalable. The terms " concurrent computing ", "parallel computing", and "distributed computing" have a lot of overlap, and no clear distinction exists between them.

2: Distributed computing - Wikipedia

In this book, the authors review important developments and discuss new strategies for the performance evaluation of parallel, distributed and emergent systems.

Description[edit] Structure of monolithic kernel, microkernel and hybrid kernel-based operating systems A distributed OS provides the essential services and functionality required of an OS but adds attributes and particular configurations to allow it to support additional requirements such as increased scale and availability. To a user, a distributed OS works in a manner similar to a single-node, monolithic operating system. That is, although it consists of multiple nodes, it appears to users and applications as a single-node. This separation increases flexibility and scalability. The kernel[edit] This article needs attention from an expert in Computing. The specific problem is: See questions asked as comments in the "kernel" section. WikiProject Computing may be able to help recruit an expert. A kernel of this design is referred to as a microkernel. These components are the part of the OS outside the kernel. These components provide higher-level communication, process and resource management, reliability, performance and security. The components match the functions of a single-entity system, adding the transparency required in a distributed environment. In addition, the system management components accept the "defensive" responsibilities of reliability, availability, and persistence. These responsibilities can conflict with each other. A consistent approach, balanced perspective, and a deep understanding of the overall system can assist in identifying diminishing returns. Separation of policy and mechanism mitigates such conflicts. Architecture and design must be approached in a manner consistent with separating policy and mechanism. In doing so, a distributed operating system attempts to provide an efficient and reliable distributed computing framework allowing for an absolute minimal user awareness of the underlying command and control efforts. This is the point in the system that must maintain a perfect harmony of purpose, and simultaneously maintain a complete disconnect of intent from implementation. However, this opportunity comes at a very high cost in complexity. The price of complexity[edit] In a distributed operating system, the exceptional degree of inherent complexity could easily render the entire system an anathema to any user. As such, the logical price of realizing a distributed operation system must be calculated in terms of overcoming vast amounts of complexity in many areas, and on many levels. This calculation includes the depth, breadth, and range of design investment and architectural planning required in achieving even the most modest implementation. Each of these design considerations can potentially affect many of the others to a significant degree. This leads to a massive effort in balanced approach, in terms of the individual design considerations, and many of their permutations. As an aid in this effort, most rely on documented experience and research in distributed computing. History[edit] Research and experimentation efforts began in earnest in the s and continued through s, with focused interest peaking in the late s. A number of distributed operating systems were introduced during this period; however, very few of these implementations achieved even modest commercial success. Fundamental and pioneering implementations of primitive distributed operating system component concepts date to the early s. These pioneering efforts laid important groundwork, and inspired continued research in areas related to distributed computing. These breakthroughs provided a solid, stable foundation for efforts that continued through the s. The accelerating proliferation of multi-processor and multi-core processor systems research led to a resurgence of the distributed OS concept. The introduction focused upon the requirements of the intended applications, including flexible communications, but also mentioned other computers: Finally, the external devices could even include other full-scale computers employing the same digital language as the DYSEAC. For example, the SEAC or other computers similar to it could be harnessed to the DYSEAC and by use of coordinated programs could be made to work together in mutual cooperation on a common taskâ€¦ Consequently[,] the computer can be used to coordinate the diverse activities of all the external devices into an effective ensemble operation. Each member of such an interconnected group of separate computers is free at any time to initiate and dispatch special control orders to any of its partners in the system. As a consequence, the supervisory control over the common task may initially be loosely distributed throughout the system and

then temporarily concentrated in one computer, or even passed rapidly from one machine to the other as the need arises. It was completed and delivered on time, in May This was a " portable computer ", housed in a tractor-trailer , with 2 attendant vehicles and 6 tons of refrigeration capacity. Lincoln TX-2[edit] Described as an experimental input-output system, the Lincoln TX-2 emphasized flexible, simultaneously operational input-output devices, i. The design of the TX-2 was modular, supporting a high degree of modification and expansion. This technique allowed multiple program counters to each associate with one of 32 possible sequences of program code. These explicitly prioritized sequences could be interleaved and executed concurrently, affecting not only the computation in process, but also the control flow of sequences and switching of devices as well. Much discussion related to device sequencing. The full power of the central unit was available to any device. The TX-2 was another example of a system exhibiting distributed control, its central unit not having dedicated control. Intercommunicating Cells[edit] One early effort at abstracting memory access was Intercommunicating Cells, where a cell was composed of a collection of memory elements. A memory element was basically a binary electronic flip-flop or relay. Within a cell there were two types of elements, symbol and cell. Each cell structure stores data in a string of symbols, consisting of a name and a set of parameters. Information is linked through cell associations. Information was accessed in two ways, direct and cross-retrieval. Direct retrieval accepts a name and returns a parameter set. Cross-retrieval projects through parameter sets and returns a set of names containing the given subset of parameters. This was similar to a modified hash table data structure that allowed multiple values parameters for each key name. Cellular memory would have many advantages: The constant-time projection through memory for storing and retrieval was inherently atomic and exclusive. The authors were considering distributed systems, stating: We wanted to present here the basic ideas of a distributed logic system with We must, at all cost, free ourselves from the burdens of detailed local problems which only befit a machine low on the evolutionary scale of machines.

3: CSE - Parallel and Distributed File Systems

Extreme! Computing. High Performance Distributed and Parallel Systems Research Department of Computer Science Indiana University.

It is now possible to simulate to obtain a high performance distributed system. The late various physical phenomena before making any real question is to determine a suitable programming model prototype. Such a situation was made possible thanks to that provides transparency, interoperability, reliability, high performance computers that become cost-effective scalability and performance. Since such systems appear when based on standard technologies, like PC clusters. Another choice is to combine the two different worlds. Such multiphysics applications require huge computing power. Hopefully, with the availability of high bandwidth wide area networks, it is nowadays feasible to couple several computing resources together to obtain a high performance distributed system. It exists to be run on such a computing infrastructure. Several different design issues have been tackled. For example, scalability is achieved between frameworks which do not put too much burden on the two parallel CORBA objects by involving all members programmers. Most of these tools [3, 20, 2] are targeted of both collections in the communication: Such a performance is obtained while preserving Interface [18] and provide their own abstractions the semantics of CORBA and in particular inter-component or object models or more often an ad-hoc ability with standard CORBA objects. API without adhering to existing standards. Moreover, MPI is being increasingly used by a simulation code for parallel processing and between simulation codes for distributed processing. Multiphysics applications based on such an approach are limited, monolithic and complex to maintain. Concentration Most of the problems stem from the utilization of a parallel oriented middleware MPI for inter-code communication, whereas such interactions should be handled in a distributed object model from the Object Management Group CORBA and to let it be used for the design of media. Figure 1 shows a simple example with only code coupling tools. As such our solution is not completely new. Each physic is simulated by a dedicated parallel code; the codes are developed by independent teams, the basic mechanisms to facilitate their design and their in different languages: A scheduler coordinates the control flows of the codes. In order to efficiently handle inter-code communications, we extend the CORBA Common Object Request Broker Architecture specification with the notion of parallel object. Hence, the encapsulation of a parallel code into a CORBA object is easy and efficient. Though latency-tolerant algorithms are under investigation, a large amount of data will still need to be transferred. In Section 3, the concept of parallel object is defined. The presentation of previous related works is followed by a discussion of their visualization and steering application limitations. Its goal is to analyze, to design and to develop a software environment. Section 5. Section 6 concludes the paper. It adds another level of motivation and analysis constraints which stem from the human interaction loop: As a user can connect and disconnect to a running application from a priori conceived, we focus on multiphysics applications. Such any machine, a distributed-oriented middleware appears well-suited to handle such interactions. Throughout this paper, the term code denotes a set of functions in any program. Most of the time, simulation codes are parallel in order to handle the huge Multiphysics applications should be

able to run in amount of computation or data required by a simulation. The general case is an het- Let us introduce two representative examples. Computational multiphysics application The networks interconnecting the machines used by an application can be of various types, ranging from Inter- An example of such application is given by the Hydro- net to high bandwidth wide area networks. Contrary to a widespread belief, we have shown [5] Another issue is the need for interoperability. For ex- that C ORBA communications can transparently and ef- ample, two organizations may temporally cooperate in a ficiently utilize all kinds of networks and in particu- project without wanting to share their code. Hence, in- lar high performance networks: In all these situations, the programming model should let the application run at the maximum efficiency al- Scalability issue lowed by the resources allocated to it. The second issue, which is the topic of this paper, comes 2. In our propo- sition, C ORBA is in charge of handling communications It is appealing to extend the approach used to pro- between parallel codes. We assumed these codes have gram parallel machines to grids: In the ent codes into a single executable, generally based on general case, the data to be sent resp. Though, it is meaningful in some situations, this ap- receiver. A solution is to serialize the data and to use proach appears limited and too much restrictive for a standard C ORBA invocation to transfer them. Besides not supporting the two organization sce- lution is not satisfactory because the data are transferred nario presented hereinbefore, it obliterates the notion of in only one communication. A mechanism being able to code maintenance and evolution: To handle the complexity of multiphysics applica- tions, it appears important to use an adequate paradigm Parallelism It consists in several control flows that act at each level. As it is obvious that computational code together to realize an operation. Many parallel codes are should be written according to a parallel paradigm " written accordingly to a Single Program Multiple Data brought by a message passing interface or by a parallel S PMD execution model: The most used interface the upper level. A distributed object oriented approach to handle communication between such control flows is appears better suited to manage inter-code interactions. Interoperability and dynamic connections are re- This paper assumes that each control flow is located quired because they enable different organizations to into its own process and that processes are deployed over set up a multiphysics application without sharing their different nodes. A node is a pair of processor-memory codes. The localization transparency is also important that needs to use a network to communicate with other as it eases the deployment. C ORBA [14] appear very well suited as they are Parallel applications usually deal with some global well established and widely available. However, with data such as matrices or graphs that are generally scat- respect to our objective, communication performance tered over the different processes accordingly to some and scalability need to be taken into account. The question is to define an object oriented model Communication performance issue that allows a parallel code to be embedded in a kind of As the localization is determined at runtime, two object and that supports efficient and scalable commu- C ORBA objects can be interconnected through any kinds nications between two such objects. A inout diffusion array myarray ; solution is brought with the concept of parallel object. Definition A distributed parallel object is a dis- tributed object whose execution model is parallel. Consequences A parallel object needs to be associ- ated to several control flows. This association can be static, for ex- ability to perform non-blocking invocations. Such asyn- ample defined at the object creation, or it can be dy- chronous invocation mechanism is based on the return- namic, i. Data generated by an Extended-I DL compiler use M PI oper- distribution issues are managed by a modification of the ations to handle data redistributions. Binding to a S PMD object is carried out through a spe- More recently, the O MG has adopted a specification [12] cific method spmd bind which is a collective form of that defines the architecture for data parallel program- the usual bind method. Thus, the distributions associated with a parallel object at runtime. This approach requires a tion model. For example, a standard C ORBA client The three previous approaches add support for paral- can invoke operations of a parallel object without notic- lel processing to C ORBA at the expense at modifications ing it is parallel. These approaches are unsatis- interoperability. Therefore, such approach is not portable. Sec- Modification of the O RB is even worst because there are tion 4. Ideally, any data distribution should be pos- choice leads first to the inability to dynamically change sible and should be dynamically exchangeable. Sec- the data distribution, but more important, it is almost im- tion 4. The O MG approach supports also a limited explained in Section 4. Moreover, it is a low chronizations. However, it supersedes

PACO in sev- objects. This choice stems also from the object communications but it uses C ORBA to communi- consideration that the parallelism of an object appears cate with other objects. The user invokes an operation of Figure 3. This model is in concordance with our constraints since the structure of parallel codes are not changed. Only some C ORBA code is added to commu- the context information to perform parallel invocations nicate with other objects. It is then possible to associate an actual object independently of parallel operations. For each parallel opera- tion, an operation context is created. With this context, 4. This file contains a description of the parallel operations of the object, the arguments 4. An example is given in Section 4. This is illustrated in Figure 4. The distributed data are handled by the distribu- in Figure 5. The Interface Manager acts as a proxy that tion libraries. The role of this operation is to to its Interface Manager. ORB Step 1 connection: An example of connection and communication between two parallel ob- The Output Manager allows a parallel code client or jects A and B. It intercepts the invocations of all in- stances of the parallel code to have only one invocation to the standard C ORBA object. To simplify its utilization, it is 4. This flexibility is achieved with the help of a fabric This section deals with the creation and the configu- design pattern. A needs to be cast down into the right type with a narrow data distribution library provides two kinds of A PI.

4: index © IEEE Computer Society

Secondly, software development on parallel systems is supported by parallel programming languages which can provide further transparency. In a distributed DBMS, data and the applications that access that data can be localized at the same site.

Models[edit] Many tasks that we would like to automate by using a computer are of question-answer type: In theoretical computer science, such tasks are called computational problems. Formally, a computational problem consists of instances together with a solution for each instance. Instances are questions that we can ask, and solutions are desired answers to these questions. Theoretical computer science seeks to understand which computational problems can be solved by using a computer computability theory and how efficiently computational complexity theory. Traditionally, it is said that a problem can be solved by using a computer if we can design an algorithm that produces a correct solution for any given instance. Such an algorithm can be implemented as a computer program that runs on a general-purpose computer: Formalisms such as random access machines or universal Turing machines can be used as abstract models of a sequential general-purpose computer executing such an algorithm. However, it is not at all obvious what is meant by "solving a problem" in the case of a concurrent or distributed system: Three viewpoints are commonly used: Parallel algorithms in shared-memory model All processors have access to a shared memory. The algorithm designer chooses the program executed by each processor. One theoretical model is the parallel random access machines PRAM that are used. Shared-memory programs can be extended to distributed systems if the underlying operating system encapsulates the communication between nodes and virtually unifies the memory across all individual systems. A model that is closer to the behavior of real-world multiprocessor machines and takes into account the use of machine instructions, such as Compare-and-swap CAS, is that of asynchronous shared memory. There is a wide body of work on this model, a summary of which can be found in the literature. Models such as Boolean circuits and sorting networks are used. Similarly, a sorting network can be seen as a computer network: Distributed algorithms in message-passing model The algorithm designer only chooses the computer program. All computers run the same program. The system must work correctly regardless of the structure of the network. A commonly used model is a graph with one finite-state machine per node. In the case of distributed algorithms, computational problems are typically related to graphs. Often the graph that describes the structure of the computer network is the problem instance. This is illustrated in the following example. Different fields might take the following approaches: Centralized algorithms[citation needed] The graph G is encoded as a string, and the string is given as input to a computer. The computer program finds a coloring of the graph, encodes the coloring as a string, and outputs the result. Parallel algorithms Again, the graph G is encoded as a string. However, multiple computers can access the same string in parallel. Each computer might focus on one part of the graph and produce a coloring for that part. The main focus is on high-performance computation that exploits the processing power of multiple computers in parallel. Distributed algorithms The graph G is the structure of the computer network. There is one computer for each node of G and one communication link for each edge of G . Initially, each computer only knows about its immediate neighbors in the graph G ; the computers must exchange messages with each other to discover more about the structure of G . Each computer must produce its own color as output. The main focus is on coordinating the operation of an arbitrary distributed system. For example, the Cole-Vishkin algorithm for graph coloring [39] was originally presented as a parallel algorithm, but the same technique can also be used directly as a distributed algorithm. Moreover, a parallel algorithm can be implemented either in a parallel system using shared memory or in a distributed system using message passing. Complexity measures[edit] In parallel algorithms, yet another resource in addition to time and space is the number of computers. Indeed, often there is a trade-off between the running time and the number of computers: If a decision problem can be solved in polylogarithmic time by using a polynomial number of processors, then the problem is said to be in the class NC. Perhaps the simplest model of distributed computing is a synchronous system where all nodes operate in a lockstep fashion. In such systems, a central complexity measure is the number of synchronous communication

rounds required to complete the task. Let D be the diameter of the network. On the one hand, any computable problem can be solved trivially in a synchronous distributed system in approximately $2D$ communication rounds: On the other hand, if the running time of the algorithm is much smaller than D communication rounds, then the nodes in the network must produce their output without having the possibility to obtain information about distant parts of the network. In other words, the nodes must make globally consistent decisions based on information that is available in their local D -neighbourhood. Many distributed algorithms are known with the running time much smaller than D rounds, and understanding which problems can be solved by such algorithms is one of the central research questions of the field [44]. Typically an algorithm which solves a problem in polylogarithmic time in the network size is considered efficient in this model. Another commonly used measure is the total number of bits transmitted in the network cf. Other problems[edit] Traditional computational problems take the perspective that we ask a question, a computer or a distributed system processes the question for a while, and then produces an answer and stops. However, there are also problems where we do not want the system to ever stop. Examples of such problems include the dining philosophers problem and other similar mutual exclusion problems. In these problems, the distributed system is supposed to continuously coordinate the use of shared resources so that no conflicts or deadlocks occur. There are also fundamental challenges that are unique to distributed computing. The first example is challenges that are related to fault-tolerance. Examples of related problems include consensus problems , [46] Byzantine fault tolerance , [47] and self-stabilisation. Synchronizers can be used to run synchronous algorithms in asynchronous systems. Before the task is begun, all network nodes are either unaware which node will serve as the "coordinator" or leader of the task, or unable to communicate with the current coordinator. After a coordinator election algorithm has been run, however, each node throughout the network recognizes a particular, unique node as the task coordinator. For that, they need some method in order to break the symmetry among them. For example, if each node has unique and comparable identities, then the nodes can compare their identities, and decide that the node with the highest identity is the coordinator. The algorithm suggested by Gallager, Humblet, and Spira [54] for general undirected graphs has had a strong impact on the design of distributed algorithms in general, and won the Dijkstra Prize for an influential paper in distributed computing. Many other algorithms were suggested for different kind of network graphs , such as undirected rings, unidirectional rings, complete graphs, grids, directed Euler graphs, and others. A general method that decouples the issue of the graph family from the design of the coordinator election algorithm was suggested by Korach, Kutten, and Moran. The coordinator election problem is to choose a process from among a group of processes on different processors in a distributed system to act as the central coordinator. Several central coordinator election algorithms exist. A complementary research problem is studying the properties of a given distributed system. The halting problem is undecidable in the general case, and naturally understanding the behaviour of a computer network is at least as hard as understanding the behaviour of one computer. In particular, it is possible to reason about the behaviour of a network of finite-state machines. One example is telling whether a given network of interacting asynchronous and non-deterministic finite-state machines can reach a deadlock.

5: Parallel computing - Wikipedia

Parallel computing resource: At the top is the parallel computing resource, which could be distributed memory systems such as IBM SP2, network of workstations (NOW) and shared memory systems such as SGI Origin

A Discourse Concerning the Mechanical Operation of the Spirit Modifiable birth influences : surgery and trauma Complete 5 String Banjo Player (Grv (Grv) Internet jobs for the rest of us Simple Secrets to Easy Weight Loss Building scalable web applications using the cloud Open space and double locks : the Hindutva appropriation of female gender Eva Hellman Irata international code of practice The Secret Purposes Understanding NE code rules on grounding and bonding Dsm iv tr study guide Clinical endocrinology of companion animals Eucharistic Gems 127 The crystal bible judy hall What is a marine mammal? The crimson petal. Hands on Mathematics, Numbers from 1-10 Brain-Inspired IT II Decision and Behavioral Choice Organized by Natural and Artificial The future of the United States. Gandhis Seven Steps to Global Change Tax planning techniques for the closely held corporation Meditation 12 : Ellis, George, and mysteries Night Voices, Night Journeys Rights of American fisherman in British North American waters Social assistance dynamics in Europe Mountain village in Nepal Vital records of Wrentham, Massachusetts, to the year 1850 . Incredible Journey Through Life with Christ (incredible testimonials of my life) Martha and Mary of Bethany The rise of Hitler Standards-based integrated library Hacking with smartphone book Law express contract law revision guide Both sides of time Pt. 1. To the dissolution of the Abbey, 1540 A.D. A Complete and revised edition of the debate on the Jesuits Estates Act Sap qm best practices Howl moving castle 18eighteen magazine torrent Adaptation of books