



HAL
open science

60 Years of Mastering Concurrent Computing through Sequential Thinking

Sergio Rajsbaum, Michel Raynal

► **To cite this version:**

Sergio Rajsbaum, Michel Raynal. 60 Years of Mastering Concurrent Computing through Sequential Thinking. ACM SIGACT News, 2020, 51 (2), pp.59-88. 10.1145/3406678.3406690 . hal-03162635

HAL Id: hal-03162635

<https://hal.science/hal-03162635>

Submitted on 17 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

60 Years of Mastering Concurrent Computing through Sequential Thinking *

Sergio Rajsbaum[†], Michel Raynal^{*,°}

[†]Instituto de Matemáticas, UNAM, Mexico

^{*}Univ Rennes IRISA, 35042 Rennes, France

[°]Department of Computing, Hong Kong Polytechnic University

rajsbaum@im.unam.mx

raynal@irisa.fr

Abstract

Modern computing systems are highly concurrent. Threads run concurrently in shared-memory multi-core systems, and programs run in different servers communicating by sending messages to each other. Concurrent programming is hard because it requires to cope with many possible, unpredictable behaviors of the processes, and the communication media. The article argues that right from the start in 1960's, the main way of dealing with concurrency has been by reduction to sequential reasoning. It traces this history, and illustrates it through several examples, from early ideas based on mutual exclusion (which was initially introduced to access shared physical resources), passing through consensus and concurrent objects (which are immaterial data), until today distributed ledgers. A discussion is also presented, which addresses the limits that this approach encounters, related to fault-tolerance, performance, and inherently concurrent problems.

Keywords: Agreement, Asynchrony, Atomicity, Concurrent object, Consensus, Consistency condition, Crash failure, Fault-tolerance, Ledger, Linearizability, Message-passing, Mutual exclusion, Progress condition, Read/write register, Sequential thinking, Sequential specification, State machine replication, Synchronization, Total order broadcast, Universal construction.

*I must appeal to the patience of the wondering reader,
suffering as I am from the sequential nature of human communication.*

Edsger W. Dijkstra (1903–1987) [26]

*Le véritable lieu de naissance est celui
où l'on a porté pour la première fois un coup d'œil intelligent sur soi-même.¹*

Marguerite Yourcenar (1903-1987) [97]

L'unicité de la pensée claire.²
Jean Fourastié (1907-1990) [37]

*This article is an extended version of the article [76], which appeared in the *Communications of the ACM*.

¹The real place of birth is the one where for the first time we took an intelligent look at ourselves.

²The uniqueness of clear thinking.

1 Introduction

Sequential reasoning is natural and easier. The human brain behaves as a multiprocessor computer, which performs many tasks simultaneously, naturally and frequently. However, despite the fact we are good at processing parallel information, it is difficult to be aware of the activities we perform concurrently, and when we try to raise awareness, we end up distracting ourselves and reducing the quality of what we are doing. Only after intense training can we, like a musician, conduct several activities simultaneously.

It is much easier to reason sequentially, doing only one thing at a time, than to understand situations where many things occur simultaneously. Furthermore, we are limited by our main communication mechanisms with others, spoken or written language, which are inherently sequential. These convey information in parallel through the voice tone, facial and body expressions, writing style, etc., but we are often unaware of it. Thus, while we are “parallel processors”, and we live in a world where multiple things happen at the same time, we usually reason by reduction to a sequential world.

The same happens in computing. It is much easier to reason about a sequential program, than about one in which operations are executed concurrently.

The grand challenge. For more than fifty years, one of the most daunting challenges in information science and technology lies in mastering concurrency. Concurrency, once a specialized discipline for experts, is forcing itself onto the entire IT community because of two disruptive phenomena: the development of networking communications, and the end of the ability to increase processors speed at an exponential rate. Increases in performance can only come through concurrency, as in multicore architectures. Concurrency is also critical to achieve fault-tolerant, distributed available services, as in distributed data bases and cloud computing. Yet, software support for these advances lags, mired in concepts from the 1960s such as semaphores. The problem is compounded by the inherently non-deterministic nature of concurrent programs: even minor timing variations may generate completely different behavior. Sometimes tricky forms of concurrency faults can appear, such as *data races* (where even though two concurrent threads handle a shared data item in a way that is correct from each thread’s perspective, a particular run-time interleaving produces inconsistent results), and others, such as *deadlocks* (situation where some threads still have computations to execute, while none of them can proceed as each is waiting for another one to proceed), *improper scheduling* of threads or processes, *priority inversions* (where some processes do not proceed because the resources they need are unduly taken away by others) and various kinds of *failures* of the processes and communication mechanisms. The result is that it is very difficult to develop concurrent applications. Concurrency is also the means to achieve high performance computing, but in this paper we are not concerned with such applications.

A simple illustration of synchronization difficulty. Let us consider the following classical illustration of the concurrency difficulties in software engineering. A bank account is shared by a group of people. The rule is that if the balance drops below a certain threshold, L , some high interest will be charged. Thus, each time a member of the group wants to withdraw some amount of money, x , she first needs to send a message to the bank to make sure the balance is greater than or equal to $L + x$. Only then she will send a message asking to withdraw x from the account. Without any synchronization, it is impossible to maintain the invariant that the balance of the account is always at least L , unless of course no withdrawals are ever done. Even assuming that the participants can directly access the account, synchronization is needed. Namely, suppose members of the group can issue `read()` operations, that directly return the current balance in the account, and execute `withdraw(x)` operations that reduce the balance by x . If Alice asks for the balance and gets back a value $z > L + x$, she cannot then issue a `withdraw(x)` operation, because she might be a bit slower than Bob, who could very fast issue `read()` and then `withdraw(x)`, just after of Alice invoked `read()` but before she invokes `withdraw(y)`.

What does concurrent computing through sequential thinking mean? Instead of trying to reason directly about concurrent computations, the idea is to transform problems in the concurrent domain into simpler problems in the sequential domain, yielding benefits for specifying, implementing, and verifying concurrent programs. It is a two-sided strategy, together with a bridge connecting them:

- Sequential specifications for concurrent programs.
- Concurrent implementations.
- Consistency conditions relating concurrent implementations to sequential specifications.

Although a program is concurrent, the specification of the object (or service) that is implementing is usually through a *sequential specification*, stating the desired behavior only in executions where the processes execute one after the other, and often through familiar paradigms from sequential computing (such as queues, stacks and lists). In the previous example, when we state the rule that if the balance drops below a certain threshold, L , some high interest will be charged, we are thinking of an account which is always in an *atomic* state, i.e., a state where the balance is well defined. This makes it easy to understand the object being implemented, as opposed to a truly concurrent specification which would be hard or unnatural. Thus, instead of trying to modify the well-understood notion of say, a queue, we stay with the usual sequential specification, and move to another level of the system the meaning of a concurrent implementation of a queue.

The second part of the strategy is to provide *implementation techniques* for efficient, scalable, and fault-tolerant concurrent objects. We will discuss only a few of the many techniques that have been developed, such as mutual exclusion algorithms, simulations of an easier to program model on top of a different system, state machine replication, quorums, broadcast abstractions.

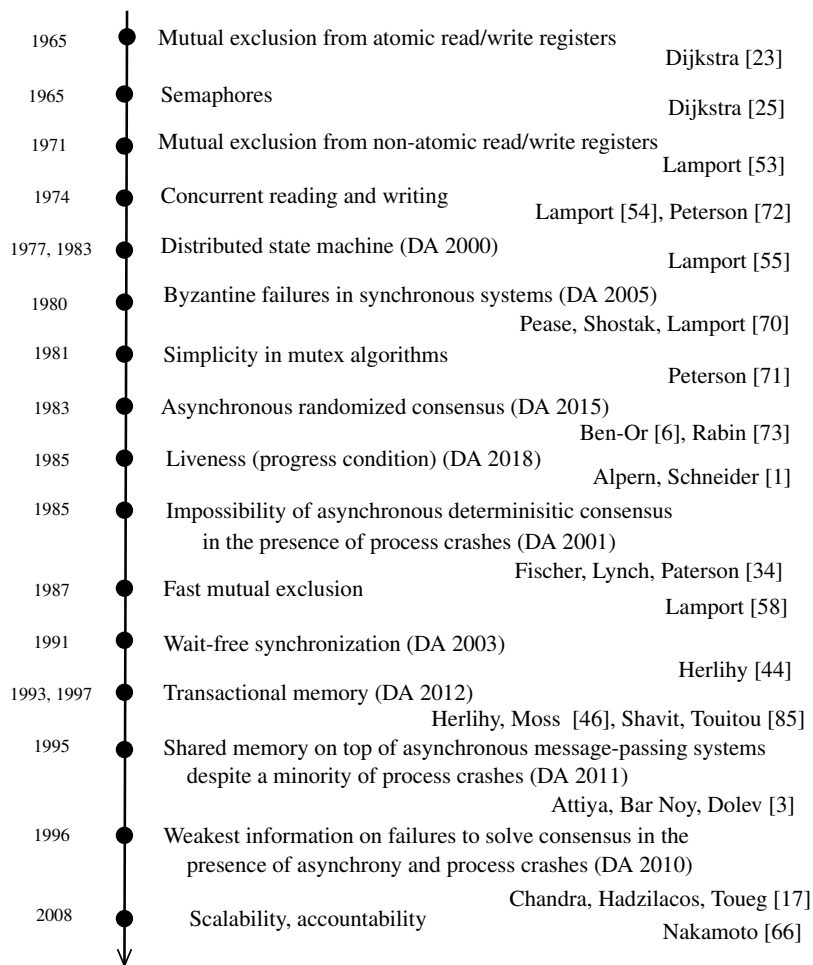
Bridging techniques show how to obtain concurrent executions that appear to the processes as if the operations invoked on the object were executed atomically, in some sequential interleaving. This is captured by the notion of a *consistency* condition, which defines the way concurrent invocations to the operations of an object correspond to a sequential interleaving, which can then be tested against the its sequential specification. We will concentrate on linearizability, among the many consistency conditions that have been proposed.

Everything is NOT reducible to sequential thinking. While a core branch of development in concurrent computing has been by reduction to sequential thinking, it is important to stress that not all problems are of this type. Some applications are not required to provide the illusion of sequential executions. Among them there are distributed graph algorithms, resource allocation problems, clock synchronization, and applications that involve large amounts of data that are continually updated, to cite a few.

Organization of the article. This article concentrates on a few significant examples of sequential reasoning used to master concurrency, highlighting fundamental notions of this approach. More specifically, it focuses on sequential specifications, linearizability, progress conditions, universal constructions, the need to solve consensus for fault-tolerance, strong shared objects as a way of solving consensus, and distributed ledgers.

To have a concrete illustration of the ideas, the article describes several algorithms. It starts by exploring a brief history of the approach, trying to understand in more detail its meaning and significance (Section 2). Then it tells the story through concrete examples, starting with mutual exclusion algorithms (Section 3) and shows a *natural* evolution from synchronization suited the sharing of physical resources to the sharing of digital (immaterial) resources/objects (Section 4).

Then, the article discusses the construction read/write registers on top of crash-prone asynchronous message-passing systems (Section 5), and more generally, considers then construction of the family of objects defined by a sequential specification (Section 6). A central technique is the *state machine*



Some of the previous papers were awarded the ACM-EATCS Edsger W. Dijkstra Prize in Distributed Computing. Created in 2000, this award is given to outstanding papers on the principles of distributed computing, whose significance and impact on the theory and/or practice of distributed computing have been evident for at least a decade). In the history line, “[aa] DA bcde” means “paper(s) referenced [aa] received the Dijkstra Award in the year bcde”.

Sidebar 1: History of synchronization: a few important dates

replication (Section 7), used to implement a replicated fault-tolerant system that seems to behave as a single sequential state machine. The article then describes how it is made possible to implement such a sequential machine in a fault-tolerant way through powerful synchronization objects (Section 8). It also discusses the *distributed ledger* trend of doing so in a highly scalable, tempered-proof way (Section 9). The approach has some limitations, both in terms of efficiency and applicability (Section 10). Finally a conclusion presents the evolution line, which motivated this article (Section 10).

2 Elements of the Approach: an Historical Perspective

The history of concurrency is long and there is an enormous amount of development, on many perspectives. As far as concurrency in distributed computing is concerned, the interested reader will find many

results (both in shared-memory and message passing) in textbooks such as [5, 14, 40, 50, 54, 65, 78, 81, 84, 91] to cite a few. Many more books exist on other approaches, such as systems, formal methods and security. The focus here is the history about the theoretical ideas of sequential reasoning to master concurrency. A few milestones are in the Sidebar 1.

When the story started. To understand what does concurrent computing through sequential thinking means, let us go back to its origins in the 1960's and 1970's, where a variety of systems were being built, both in hardware and software, moving away from the original paradigm of a single user interacting with a single sequential computer. As soon as communication and interaction among users and sequential processes was becoming possible, the desire to take advantage of these possibilities emerged: to share resources, to replicate components for fault-tolerance, to compute in a collaborative way, to distribute resources and data geographically for accessibility. All such motivations are relevant to the approach that concerns us, both to understand its meaning, and to study the tools to implement it. And indeed, the approach surveyed in the present article is relevant in different areas, still today: databases, operating systems, networking, distributed computing, and concurrent programming.

The early days were in the 1960's with the seminal work of Dijkstra and others in multi-programming (see [11] where are collected the very first articles which gave rise to concurrent programming), then in the 1970's with database and fault-tolerance communities. At the end of the 1970's, distributed computing had become a possibility. Soon after, the Internet was born and it was being used to send email and do file transfer. Additionally, it was hoped that the Internet could be used to run applications distributed over many machines, but there was little understanding of how that could be accomplished. It was motivated by a lack of understanding basic theory, that theoretical distributed computing was born, the first ACM conference on *Principles of Distributed Computing* in 1982, and the first impossibility results such as the first version the famous FLP impossibility result in 1983 [35], and the highlighting of lower bounds (e.g., [34, 72]).

Already in 1967 there were debates about multi-processor computing and Amdahl's Law predicting speedup when using multiple processors. In the late 1970s a move occurred from multiprocessors with shared memory to multicomputers with distributed memory communicating by sending messages. In the 1990s, the importance of shared memory with multicores returns, as it meets the barriers of energy expenditure, and the limits of making processors increasingly faster, emphasizing that the exponential growth prophesied by Moore's Law refers to packaging more and more components on the same chip, that is more and more parallel computing. And a new era of distributed systems is entered, which is motivated by new distributed services such as ledger-based applications, which must be open, scalable, and tolerate arbitrarily malicious faults. Nevertheless, while both parallel computing and distributed computing involve concurrency, it is important to stress the fact they address different computing worlds (see Sidebar 2).

At the origin: mutual-exclusion to share physical resources. Concurrent computing began in 1961 with what was called *multiprogramming* in the Atlas computer, where concurrency was simulated – as we do when telling stories where things happen concurrently – interlacing the execution of sequential programs. Concurrency was born in order to make efficient use of a sequential computer, which can execute only one instruction at a time, giving users the illusion that their programs are all running simultaneously, through the operating system. A collection of early foundational articles on concurrent programming appears in [11].

As soon as the programs being run concurrently began to interact with each other, it was realized how difficult it is to think concurrently. By the end of the 1960s there was already talk of a crisis: programming was done without any conceptual foundation and lots of programs were riddled with subtle race-related errors causing erratic behaviors. In 1965 Dijkstra [25] discovered that the *mutual exclusion*

As far as terminology is concerned we consider the following definitions (from [81]).

- *Parallel computing.* Parallel computing addresses concepts, methods, and strategies which allow us to benefit from parallelism (simultaneous execution of distinct threads or processes) when one has to implement a computation. The essence of parallel computing lies in the decomposition of the computation in independent computation units and exploit their independence to execute as many of them as possible in parallel (simultaneously) so that the resulting execution is time-efficient.
- *Distributed computing.* Distributed computing arises when one has to solve a problem involving geographically distributed entities (processors, nodes, sensors, peers, agents, etc.), such that each entity only has a partial knowledge of the many input parameters involved in the problem to be solved. Because their knowledge is partial, these computing entities must cooperate to solve the problem. They also must cope with their environment, which can be modeled as adversaries, such as asynchrony, failures, mobility, etc. These adversaries create an uncertainty on the state of the system, uncertainty that has to be understood and mastered if one wants to produce correct distributed software.

As we can see, parallel and distributed computing are in some sense dual: one consists in decomposing a computation into independent entities, while the other consists in allowing pre-existing entities – whose distribution is not under the control of the programmer – to cooperate in the presence of adversaries such as the net effect of asynchrony and process failures.

Sidebar 2: Distributed computing versus parallel computing: two different faces of concurrency

of parts of code is a fundamental concept of programming, and opened the way for the first books of principles on concurrent programming which appeared at the beginning of the 1970s.

Interestingly, a very early impossibility result, showing a lower bound on the size of memory needed for two-process mutual exclusion [23], was one of the motivations for the systematic study of distributed computing by pioneers of the field (see for example [66]).

Fault-tolerance. Some of the earliest sources of the approach followed in this article come from fault-tolerance concerns, dating back to the 1960's, and in fact going back to von Neumann, as described by Avižienis in [6], where he describes the sources of fault-tolerant techniques and some techniques used up to day: "Considerable effort has been continuously directed toward practical use of massive triple modular redundancy (TMR) in which logic signals are handled in three identical channels and faults are masked by vote-taking elements distributed throughout the system." The idea begins to emerge, that "The fact that a task is executed by several processors is invisible to the application software."

In 1973 the design of systems such as the SIFT [95] project at SRI began to emerge, for highly reliable computer systems for aircraft control. This project was notable for producing the Byzantine generals problem and its solutions [72], and indeed another source of motivation for the birth of theoretical research on distributed computing. Among other points, the study of Byzantine processes showed that triple modular redundancy does not work when processes can exhibit Byzantine failures (in the presence of such failures, many problems cannot be solved only if one third or more of the processes can be Byzantine).

The primary backup approach to fault-tolerance is an elementary early way of producing the illusion of a single computer, originating in 1976 or earlier [89]. A survey of about 30 years on replication collects many techniques, some from the very practical perspective [21]. It is about replicating data or a service that will behave in a manner indistinguishable from the behavior of some non-replicated reference system running on a single non-faulty node. More modern techniques such as *virtual syn-*

chrony, *viewstamped replication*, *quorums*, etc, have their origins in Lamport's seminal 1978 paper [57] proposing the fundamental *state machine replication* paradigm.

Networking. In addition to providing the motivation to build distributed systems and resource sharing, the distributed system community started early to design broadcast algorithms that are central to implementation of distributed systems that behave as if they were sequential. The proceedings of the Sixth Symposium on Data Communications (SIGCOMM) in 1979 already includes papers on integration of a DBMS into a network environment, concurrency control in distributed database systems, and the reliable update of replicated databases.

An early system was the Stony Brook System, discussed in [3]. This paper is remarkable for noticing what we call now the *Coordinated Attack Problem* where two generals need to coordinate by sending messages that may be lost, and furthermore, present one of the earliest impossibility results in the field "To show that no amount of user protocol can solve the problem in a manner to dissipate the anxiety of both parties as to the outcome of a transaction." Remarkably, this is the first computer communication problem proved to be unsolvable. Still today, introductory classes about computer networking teach this result, teach that TCP cannot guarantee state consistency between endpoints and why this is the case, explained in a similar manner to this early paper. It is not surprising that the short documentary film from 1972 about the ARPANET is titled *Computer Networks: The Heralds of Resource Sharing*.

Resource sharing. In 1965, Robert Taylor at ARPA started to fund large programs in advanced research in computing at major universities and corporate research centers throughout the United States, especially motivated by resource sharing. Among the computer projects that ARPA supported was time-sharing, in which many users could work at terminals to share a single large computer. The thing that drove Bob Taylor at ARPA was resource sharing.

Time-sharing was developed in the late 1950's out of the realization that a single expensive computer could be efficiently utilized if a multitasking, multiprogramming operating system allowed multiple users simultaneous interactive access. The DEC's PDP-10 is the machine that made time-sharing common, and this and other features made it a common fixture in many university computing facilities and research labs during the 1970s. Time-sharing emergence as the prominent model of computing in the 1970's represented a major technological shift in the history of computing.

The first commercially successful time-sharing system, the Dartmouth Time Sharing System's creators wrote in 1968 that "any response time which averages more than 10 seconds destroys the illusion of having one's own computer" [83]. Timesharing users thought that their terminal was the computer.

The architecture of this system was influenced by the Multics operating system. Multics supported multiple CPUs; it was one of the earliest multiprocessor systems. An influential early time-sharing operating system, started in 1964. The 1950s and 1960s was the era of timesharing and it was clear that the fundamental principles of timesharing would also apply to data networks. With the appearance of computer networks in the late 1960's and early 1970's, both homogeneous multi-machine configurations (TSS-IBM) consisting of identical systems and at the other extreme is the ARPA network, work on distributed operating systems was fostered, focusing on interprocess communication. A main concern was resource sharing over the network, both physical and files (which are immaterial data).

By the spring of 1969, a well-documented, reliable version of the RC 4000 multiprogramming system was running. Brinch Hansen [12] describes it, saying that "Dijkstra [26] has demonstrated that indivisible lock and unlock operations operating on binary semaphores are sufficient primitives from a logical point of view. We have been forced to conclude, however, that the semaphore concept alone does not fulfill for our requirements of safety and efficiency in a dynamic environment in which some processes may turn out to be black sheep and break the rules of the game. Instead we have introduced message buffering within the system nucleus as the basic means of process communication".

Databases and consistency conditions. This community was the first to think about consistency in replication platforms, and to struggle with fundamental limits. They had their own version of the FLP result: replication algorithms in which applications updated replicas using simple timeout mechanisms for fault-tolerance, was shown to result in non-serializable executions [8].

While the previous approaches are some of the important sources for the sequential approach to distributed computing, they were more focused on operational approaches, like consensus, mutual exclusion and fault-tolerance. Indeed it was from the semantics side, explicitly defining the *meaning* of sequential thinking for concurrency, with the idea of *serializability*, as early as 1976, “based on suitable consistency and concurrency control operation, we show that, even though the individual application programs are running concurrently and generating requests interleaved in an apparently random fashion, the overall effect produced by the concurrency control must be the same as if the programs that write had been run sequentially in some linear order” [88]. Or also in 1976 Eswaran et al. [31], “Nonserial schedules run the risk of giving a transaction an inconsistent view of the state. So we are particularly interested in those schedules which are "equivalent" to serial schedules”.

But still in 1981 a need for a theory of concurrency control was in need. Stearns et al. [10] say "Concurrency control has been actively investigated for the past several years, and the problem for nondistributed DBMSs is well understood. A broad mathematical theory has been developed to analyze the problem, and one approach, called two-phase locking, has been accepted as a standard solution. Current research on non-distributed concurrency control is focused on evolutionary improvements to two-phase locking, " "Distributed concurrency control, by contrast, is in a state of extreme turbulence. More than 20 concurrency control algorithms have been proposed for DDBMSs, and several have been, or are being, implemented. These algorithms are usually complex, hard to understand, and difficult to prove correct (indeed, many are incorrect). Because they are described in different terminologies and make different assumptions”.

In addition to the semantics contribution, from databases come some of the earliest techniques to enforce a sequential view, under failures, such as two-phase commit dating back to ideas of Lamson and Sturgis in 1976 [63].

From the seminal book of Bernstein, Hadzilacos and Goodman [9], it is clear that it was in database transaction processing that the roots of sequential thinking are most clear, “Whether by its native capabilities or the way we educate it, the human mind seems better suited for reasoning about sequential activities than concurrent ones. This is indeed unfortunate for the study of concurrency control algorithms. Inherent to the study of such algorithms is the need to reason about concurrent executions.”

3 Mutual Exclusion

In the mutual exclusion problem described and solved by Dijkstra [25] there is a collection of asynchronous processes, each alternately executing a critical and a noncritical section, that must be synchronized so that no two processes ever execute their critical sections concurrently.

Mutual exclusion. A mutual exclusion algorithm consists of the code for two operations, `acquire()` and `release()`, that a process invokes to bracket a section of code called a *critical section*. The usual environment in which a mutual exclusion algorithm is executed is *asynchronous*, where process speeds are arbitrary, independent from each other. The mutual exclusion algorithm should guarantee two conditions.

- Mutual exclusion. No two processes are simultaneously executing their critical section.
- Deadlock-freedom. if one or several processes invoke concurrently `acquire()`, eventually one of them terminates its invocation, and consequently executes its critical section.

Progress conditions. When Dijkstra introduced mutual exclusion [25], he also introduced the previous progress condition, called *deadlock-freedom*. As observed by D.E. Knuth in [53], Deadlock-freedom does not prevent specific timing scenarios from occurring in which some processes can never enter their critical section. Hence, he proposed The stronger *starvation-freedom* progress condition, states that any process that invokes `acquire()` will terminate its invocation (and will consequently execute its critical section).

On mutual exclusion algorithms from atomic read/write registers. The first mutual exclusion algorithms were abstruse, difficult to understand and prove correct (some of them are collected in [77]). We describe here an elegant algorithm by Peterson [73]. The version presented in Algorithm 1 is for two processes, but can be easily generalized to n processes.

The two processes p_1 and p_2 share three read/write atomic registers, $FLAG[1]$, $FLAG[2]$, and $LAST$. Initially $FLAG[1]$, $FLAG[2]$, are down, while $LAST$ does not need to be initialized. Both processes can read all registers. Moreover, while $LAST$ can be written by both processes, only p_i , $i \in \{1, 2\}$, writes to $FLAG[i]$. *Atomic* means that the read and write operations on the registers appear as if they have been executed sequentially (hence, the notion of “last writer” associated with $LAST$ is well defined).

<pre> operation acquire() is % invoked by p_i, $i \in \{1, 2\}$ $FLAG[i] \leftarrow \text{up}$; $LAST \leftarrow i$; let $j = 3 - i$; wait $((FLAG[j] = \text{down}) \vee (LAST \neq i))$; return() end operation. operation release() is $FLAG[i] \leftarrow \text{down}$; return() end operation. </pre>
--

Algorithm 1: Peterson’s algorithm for two processes

When process p_i invokes `acquire()`, it first raises its flag, thereby indicating it is competing, and then writes its name in $LAST$ indicating it is the last writer of this register. Next process p_i repeatedly reads $FLAG[j]$ and $LAST$ until it sees $FLAG[j] = \text{down}$ or it is no longer the last writer of $LAST$. When this occurs, p_i terminates its invocation. The operation `release()` consists in a simple lowering of the flag of the invoking process. The read and write operations on $FLAG[1]$, $FLAG[2]$, and $LAST$ are totally ordered (atomicity), which facilitates the proof of the mutual exclusion and starvation-freedom properties.

Mutual exclusion was the first mechanism for mastering concurrent programming through sequential thinking, and lead to the identification of notions that began to give a scientific foundation to the approach, such as the concepts of *progress condition* and *atomicity*.

Fast mutual exclusion and adaptive algorithms. The previous algorithm can be easily generalized to solve mutual exclusion in a set of $n \geq 2$ processes. Many n -process mutual exclusion algorithms have been proposed, in which each process must solve $(n - 1)$ conflicts to access the critical section. An algorithm in which the number of read and write accesses to shared registers is constant in contention-free scenarios appears in Lamport [60]. This article is the origin of research on *adaptive algorithms*, whose complexity depends on the concurrency pattern in which operations are invoked.

Atomicity from non-atomic read/write registers. The previous algorithms implements mutual exclusion using underlying atomic read/write registers. In fact, this hardware atomicity is not required, Lamport [55], showed that mutual exclusion can be achieved using only *safe* registers [59]. Several

algorithms building atomic read/write registers from non-atomic read/write registers are described in e.g., [78, 91].

On the database side. The concept of a *transaction* was introduced in database as a computation unit (usually, an operation-based translation of a query expressed in a specific query language) [43]. The management of transactions introduced the notion of *concurrency control*, which gave rise to several approaches to ensure that transactions appear as if they had been executed sequentially [9, 71].

Transactional memory. The concept of *transactional memory* (TM) was introduced by M. Herlihy and J. Moss in 1993 [48], and then investigated from a pure software point of view (STM) by N. Shavit and D. Touitou in 1997 [87].

The aim of a TM/STM system is to discharge the programmers from the management of synchronization in multiprocess programs that access concurrent objects. To that end, an TM/STM system provides the programmer with the concept of a transaction. Basically, the job of the programmer is to design each process of the application as a sequence of transactional code and non-transactional code, where a transaction is any piece of code that accesses concurrent objects, but contains no explicit synchronization statement, and non-transactional code does not access concurrent objects. It is then the job of the underlying TM/STM system to provide the illusion that each transaction appears as being executed atomically (see Sidebar 3, where each read or write operation is replaced by a transaction). Executing each transaction in a critical section would solve the problem, but this would be inefficient. So, for efficiency, a TM/STM system must allow transactions to execute concurrently. The major parts of a TM/STM systems execute transaction in a speculative mode at the end of with a transaction is committed or aborted. According to the TM/STM system, the recovery of a transaction can be under the control of either the system or the invoking process. Examples of STM systems based on different underlying principles can be found in [15, 24].

As we can see, a TM/STM system allows the programmer to concentrate on the problem it has to solve and not on the way the required synchronization must be implemented. In this sense it provides the programmer with a higher abstraction level. It is important to see that a transaction can be any piece of code (and not a code obtained from a specific query language as in databases). TM/STM provides programmers with a tool from which they can see executions as sequences of transactional codes.

The important point here is that both concurrency control in database and transactional memory aim at providing an abstraction level at which the users see an execution as if it was produced by a sequential processor.

4 From (Physical) Resources to (Immaterial) Objects

From physical resources to services. At the beginning, a critical section was encapsulating the use of a physical resource, which by its own nature, is sequentially specified (e.g., disk, printer, processor). Conceptually not very different, it was then used to protect concurrent accesses to preserve consistency of simple data (such as a file in the readers/writers problem [22]). However, when critical sections began to be used to encapsulate more general *shared objects*, new ideas were needed.

Data are not physical resources. A shared object is different from a physical object, in that it does not a priori require exclusive access; a process can read the data of a file while another process concurrently modifies it. The *mutex-free* (also called *lock-free*) approach (introduced by Lamport in [56]), makes possible to envisage implementations of purely digital objects in which operation executions are free

from mutual exclusion and can overlap in time, none of them depending of the others to terminate [46] (see progress conditions defined below).

Consistency conditions. Wherever concurrent accesses to share data take place, a consistency condition is needed to define what does it mean to correctly execute concurrently operations, especially in the presence of buffers and memory caches (that are defined only in sequential executions, such as read/write operations). Instead of transforming a concurrent execution into sequential execution (as in mutual exclusion), the idea appears to enforce only *virtual sequentiality*, namely, from an external observer point of view, everything must appear as if the operations were executed sequentially, thereby reducing –at a higher abstraction layer– concurrent computing to sequential computing. When the total order on the operations is required to respect the order on non-overlapping operations, this virtual sequentiality is called *atomicity* [59] or *linearizability* [51] (these two terms are synonyms). This is illustrated in Sidebar 3, which describes an execution in which three processes access an atomic read/write register R .

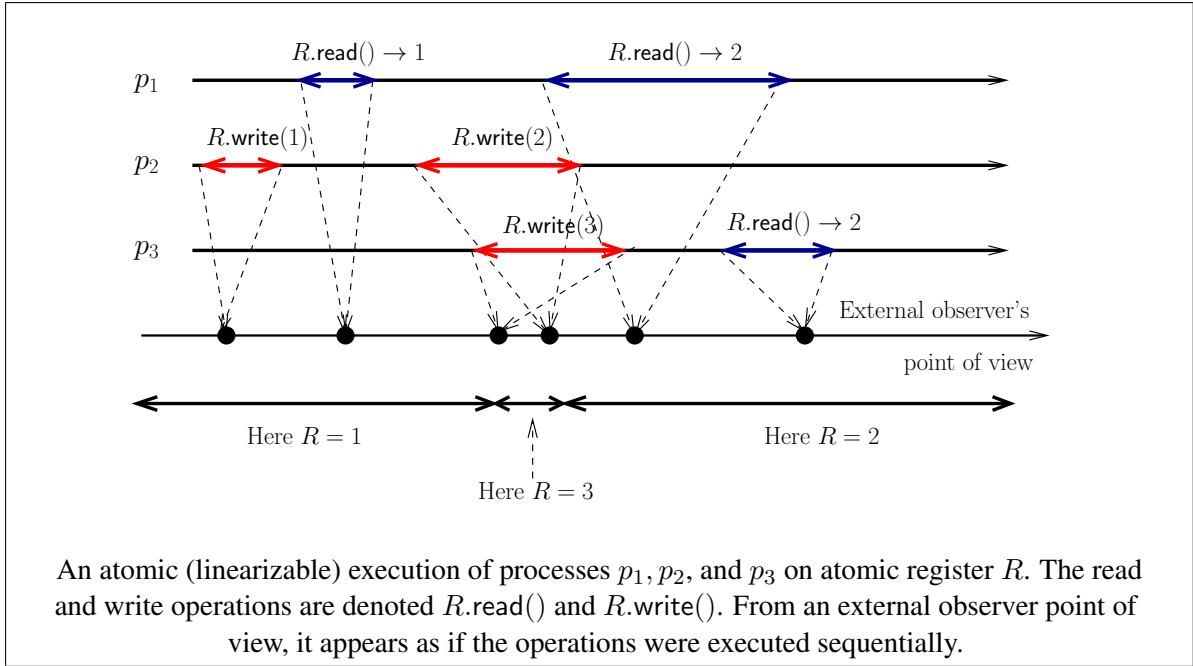
From serializability to linearizability. Since early on in 1976, in the database context, *serializability* [71, 88] of transactions that aggregate many operations without locking and unlocking entities was generally accepted as the right notion of correctness: to require that transactions appear to have executed atomically. In the concurrent programming the equivalent notion of *sequential consistency* was used, but for individual operations [58]. Later on, realizing that this type of condition is not composable, *linearizability* [51] required additionally that this sequential order must also preserve the global ordering of non-overlapping operations. Linearizability has the advantage of being composable, over sequential consistency: a system made of linearizable implementations is linearizable.

The environment has an impact on computations: crash failures. Let us remark that mutual exclusion cannot work when one has to implement an object in the presence of asynchrony and process *crashes* (premature halting). If a process crashes inside its critical section, mutual exclusion will never be released, and no other process will be able to access the object. It follows that the use of mutual exclusion (locks) is limited in the presence of asynchrony and process crashes.

On progress conditions in the presence of crash failures. Three progress conditions have been proposed for the implementation of the operations of data objects in an environment where processes are asynchronous and may crash. They are the following ones, going from the stronger to the weaker (see Table 1).

- The *wait-freedom* progress condition states that if a process invokes an object operation, and does not crash, it terminates its invocation [46]. This means that it terminates whatever the behavior of the other processes (e.g., some of them being crashed, and others being concurrently involved in object operations).
- The *non-blocking* progress condition states that if several processes concurrently invoke operations on the object, at least one of them terminates [51].
- The *obstruction-freedom* progress condition states that if a process invokes an operation, does not crash during this invocation, and all other processes stop accessing the internal representation of the object during a long enough period, then the process terminates its operation [47].

Let us remark that the wait-freedom and non-blocking progress conditions are independent of both the failure pattern and the concurrency pattern. They can be seen as the “corresponding” of starvation-freedom and deadlock-freedom in asynchronous crash-prone system. Differently, obstruction-freedom is dependent on the concurrency pattern.



Sidebar 3: An atomic execution of a read/write register

Lock-based implementations	Mutex-free implementations
	Obstruction-freedom [46]
Deadlock-freedom [25]	Non-blocking [51]
Starvation-freedom [53]	Wait-freedom [47]

Table 1: Progress conditions for the implementation of concurrent objects

Other progress conditions are investigated in [52, 93, 96]. Progress in the context of specific contention and failures patterns is addressed in [29, 78, 92].

5 Read/Write Registers on Top of Message-Passing Systems

The read/write shared register abstraction provides several advantages over message passing: a more natural transition from uniprocessors, and simplifies programming tasks. For this reason, concurrent systems that support shared memory have wide acceptance in both research and commercial computing.

It is relatively easy to build atomic read/write registers on top of a reliable asynchronous message-passing system (e.g. [79]), but if processes may crash, more involved algorithms are needed. Two important results are presented by Attiya, Bar-Noy and Dolev in [4]:

- An algorithm that implements an atomic read/write register on top of a system of n asynchronous message-passing processes, where at most $t < n/2$ of them may crash.
- A proof of the impossibility of building an atomic read/write register when $t \geq n/2$.

The section presents the algorithm, referred to as the *ABD Algorithm*, which illustrates the importance of the ideas of reducing concurrent thinking to sequential reasoning. A more detailed proof can be found in [4, 5, 81], as well as other algorithms.

Design principles of ABD: each written value has an identity. Each process is both a client and a server. Let REG be the multi-writer multi-reader (MWMR) register that is built (hence any process is allowed to read and write the register). On its client side a process p_i can invoke the operations $REG.write(v)$ (to write a value v in REG , and $REG.read()$ to obtain its current value. On its server side, a process p_i manages two local variables: reg_i which locally implement REG , and $timestamp_i$ which contains a timestamp made up of a sequence number (which can be considered as a date) and a process identity j . The timestamp $timestamp_i$ constitutes the “identity” of the value v saved in reg_i (namely, this value was written by this process at this time). Any two timestamps $\langle sn_i, i \rangle$ and $\langle sn_j, j \rangle$ are totally ordered by their lexicographical order; namely, $\langle sn_i, i \rangle < \langle sn_j, j \rangle$ means $(sn_i < sn_j) \vee (sn_i = sn_j \wedge i < j)$.

Design principles of ABD: intersecting quorums. The basic mechanism on which ABD relies on a query/response message exchange pattern. A process p_i broadcasts a query to all the processes and waits for acknowledgments from a majority of them. Such a majority *quorum* set, has the following properties. As $t < n/2$, waiting for acknowledgments from a majority of processes can never block forever the invoking process. Moreover, the fact that any two quorums have a non-empty intersection implies the atomicity property of the read/write register REG .

```

operation  $REG.write(v)$  issued by process  $p_i$  is
  build a new tag  $tag$  identifying this write operation;
  % Phase 1: acquire information on the system state %
  broadcast WRITE_REQ( $tag$ );
  wait acknowledgments from a majority of processes,
  each carrying  $tag$  and a sequence number;
  % Phase 2 : update system state %
   $ts \leftarrow \langle msn + 1, i \rangle$  where  $msn$  is
  the greatest sequence number previously received;
  broadcast WRITE( $tag, v, ts$ );
  wait acknowledgments carrying  $tag$  from a majority of proc.;
  return().

when WRITE_REQ( $tag$ ) is received from  $p_j, j \in \{1, \dots, n\}$  do
  send to  $p_j$  an acknowledgment carrying  $tag$ , and
  the sequence number contained in  $timestamp_i$ .

when WRITE( $tag, v, ts$ ) is received from  $p_j, j \in \{1, \dots, n\}$  do
  if ( $timestamp_i < ts$ ) then
     $timestamp_i \leftarrow ts; reg_i \leftarrow v$  end if;
  send to  $p_j$  an acknowledgment carrying  $tag$ .

```

Algorithm 2: Operation $REG.write(v)$: client and server behavior for a process p_i

The operation $REG.write(v)$. This operation is implemented by Algorithm 2. When a process p_i invokes $REG.write(v)$, it first creates a tag denoted (tag) which will identify the query/response messages generated by this write invocation. Then (phase 1), it executes a first instance of the query/response exchange pattern to learn the highest sequence number saved in the local variables $timestamp_j$ of a majority of processes p_j . When this is done, p_i computes the timestamp ts which will be associated with the value v it wants to write in REG . Finally (phase 2), p_i starts a second query/response pattern in which it broadcasts the pair (v, ts) to all the processes. When, it has received the associated acknowledgments from a quorum, p_i terminates the write operation.

On its server side, a process p_i that receives a WRITE_REQ message sent by a process p_j during phase 1 of a write operation, sends it back an acknowledgment carrying the sequence number associated

with the last value it saved in reg_i . When it receives `WRITE_REQ` message sent by a process p_j during phase 2 of a write operation, it updates its local data reg_i implementing *REG* if the received timestamp is more recent (with respect to the total order on timestamps) than the one saved in $timestamp_i$, and, in all cases, it sends back to p_j and acknowledgment (so p_j terminates its write).

It is easy to see that, due to the intersection property of quorums, the timestamp associated with a value v by the invoking process p_i is greater than the ones of the write operations that terminated before p_i issued its own write operation. Moreover, while concurrent write operations can associate the same sequence number with their values, these values have different (and ordered) timestamps.

The operation *REG.read* (). Algorithm 3 implements operation *REG.read* (), with a similar structure as the implementation of operation *REG.write* (). Namely, it is made up of two phases, each one being an instance of the query/response communication pattern. In the first phase, the invoking process obtains a pair (value, associated timestamp) from a minority of processes, from which – thanks to the total order on timestamps – it can extract the most recent value, that it will return as the result of the read operation.

```

operation REG.read () is
  build a new tag tag identifying this read operation;
  % Phase 1: acquire information on the system state %
  broadcast READ_REQ (tag);
  wait acknowledgments from a majority of processes,
  each carrying tag and a pair (value,timestamp);
  let ts be the greatest timestamp received,
  and v the value associated with this timestamp;
  % Phase 2 : update system state %
  broadcast WRITE (tag, v, ts);
  wait ACK_WRITE (tag) from a majority of proc.;
  return (v).

when READ_REQ (tag) is received from  $p_j, j \in \{1, \dots, n\}$  do
  send to  $p_j$  an ack. carrying tag,  $reg_i$  and  $timestamp_i$ .
  
```

Algorithm 3: Operation *REG.read* (): client and server behavior for a process p_i

Notice that the following scenario can occur, which involves two read operations *read1* and *read2* on a register *REG* by the processes p_1 and p_2 , respectively, and a concurrent write operation *REG.write*(v) issued by a process p_3 (Fig. 1). Let $ts(v)$ be the timestamp associated with v by p_3 .

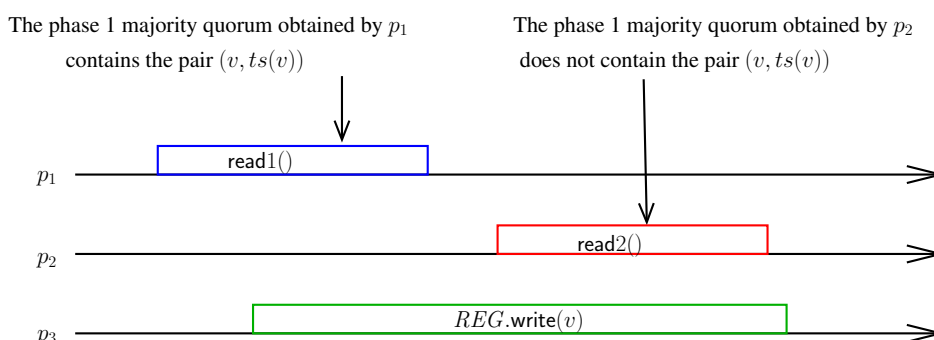


Figure 1: New/old inversion scenario

It is possible that the phase 1 majority quorum obtained by p_1 includes the pair $(v, ts(v))$, while the

one obtained by p_2 does not. If this occurs, the first read operation `read1` obtains a value more recent than the one obtained by the second `read2`, which violates atomicity. This can be easily solved by directing each read operation to write the value it is about to return as a result. In this way, when `read1` terminates and returns v , this value is known by a majority of processes despite asynchrony, concurrency, and a minority of process crashes. This phenomenon (called *new/old inversion*) is prevented by the phase 2 of a read operation.

The combination of intersecting quorums and timestamps allows for the implementation of atomic read/write registers in asynchronous message-passing systems where a minority of process may crash. Hence, sequential thinking on shared registers can be used at the upper abstraction level.

6 The World of Concurrent Objects

Objects defined by a sequential specification. A read/write register is a special case of an immaterial object. In general, an *object* is defined by the set of operations that processes can invoke, and by an automaton, which specifies the behavior of the object when these operations are invoked sequentially. The automaton specifies, for each state, and each possible operation invocation, a response to that invocation, and a transition to a new state. A stack for example, is easily specified in this way. The operations are `push(v)`, to add v at the top of the stack; if the stack is full, it returns the control value `full`. Similarly, if the stack is not empty, the operation `pop()` returns the value at the top of the stack and suppresses it from the stack; and it returns the control value `empty` if the stack is empty.

A concurrent stack can be implemented by executing the operations `pop()` and `push()` using mutual exclusion. As already indicated, this strategy to create a total order does not work if processes may crash. The *state machine replication* mechanism [57] is a general way of implementing an object by asynchronous crash-prone processes, that invoke operations on the object concurrently.

Implementing a state machine is easy if no process crash. This is no longer the case in crash-prone asynchronous systems, where the implementation of a state machine relies on the consensus object.

Consensus. At the core of many sequential reasoning for concurrent programming situations (including state machine replication) are agreement problems. A common underlying abstraction is the *consensus* object. It has a single operation denoted `propose()`, that a process can invoke once. If a process invokes `propose(v)`, the invocation eventually returns a value v' . This sequential specification is defined by the following properties.

- Validity. If an invocation returns v then there is a `propose(v)`.
- Agreement. No two different values are returned.
- Termination. If a process invokes `propose()` and does not crash, it returns a value.

Consensus objects are universal in the sense that (together with read/write registers), they can be used to implement, despite asynchrony and process crashes, any object defined by a sequential specification. The consensus-based state machine replication technique provides an illustration of this claim, as discussed below.

All objects are not equal in a crash-prone environment. It turns out that an object as simple as a concurrent stack cannot be implemented by asynchronous processes, which communicate using read/write registers only, if any operation invoked by a process that does not crash must return (independently of the speed or crashes of the other processes). Such an implementation of an object is said to be *wait-free* [46].

A way of measuring the synchronization power of an object in the presence of asynchrony and process crashes is by its *consensus number* [46]. The consensus number of an object O is the greatest integer n , such that it is possible to wait-free implement a consensus object for n processes from any

number of objects O and atomic read/write registers. The consensus number of O is ∞ if there is no such greatest integer. As an example, the consensus number of a Test&Set object or a stack object is 2, while consensus number of a Compare&Swap or LL/SC object is ∞ . The power and limits of shared memory systems is addressed in [49].

7 State Machine Replication

The *state machine replication* mechanism [57] is the main approach to implement an object in a concurrent system, with asynchrony and process crash failures in message-passing systems [57, 85], and in multiprocessors where each processor has a local memory [78]. The idea is for the processes to agree on a sequential order of the concurrent invocations, and then each one to simulate the sequential specification automaton locally. We illustrate here the approach with a mechanism for reaching the required agreement: a total order broadcast abstraction.

Total order broadcast. The TO-broadcast abstraction [45, 81] is an important primitive in distributed computing, that ensures that all correct processes receive messages in the same order (we do not define them more formally here). It is used through two operations, `TO_broadcast()` and `TO_deliver()`. A process invokes `TO_broadcast(m)`, to send a message m to all other processes. As a result, processes execute `TO_deliver()` when they receive a (totally ordered) message. The TO-broadcast abstraction is defined by the following properties (the first three are safety, while the last two are liveness properties). It is assumed without loss of generality that all messages are different.

- TO-validity. If a process executes `TO_deliver(m)` (i.e., to-delivers the a message m), then a process executes `TO_broadcast(m)`.
- TO-integrity. If a process executes `TO_deliver(m)` and `TO_deliver(m')`, then $m \neq m'$.
- TO-order. If a process executes `TO_deliver(m)` and `TO_deliver(m')` in this order, then no process executes these operations in the reverse order.
- TO-termination-1. If a process executes `TO_broadcast(m)` and does not crash, it eventually executes `TO_deliver(m)`.
- TO-termination-2. If a process executes `TO_deliver(m)`, then every process that does not crash executes `TO_deliver(m)`.

TO-broadcast illustrates one more general idea within the theory of mastering concurrent programming through sequential thinking: the identification of communication abstractions that facilitate building concurrent objects defined by a sequential specification.

State machine replication based on TO-broadcast. A concurrent implementation of object O is described in Algorithm 4. It is a universal construction, as it works for any object O defined by a sequential specification. The object has operations `op x ()`, and a transition function $\delta()$ (assuming δ is deterministic), where $\delta(state, op_x(param_x))$ returns the pair $\langle state', r \rangle$, where $state'$ is the new state of the object and res the result of the operation.

Let p_1, \dots, p_n be the set of asynchronous crash-prone processes. Each process p_i is both client (it can invoke operations on O) and server (it participates in the implementation of O). The idea of the construction is simple. Each process p_i has a copy $state_i$ of the object, and the TO-broadcast abstraction is used to ensure that all the processes p_i apply the same sequence of operations to their local representation $state_i$ of the object O . When a process p_i invokes an operation it builds a message `sent_msg` composed of two fields: `sent_msg.op` which contains the operation and `sent_msg.proc` which contains the identity of the invoking process. Then p_i to-broadcasts `sent_msg` and waits until its operation has been executed on its local copy of O . On its server side, a process p_i executes an infinite

```

when operation  $op_x(param_x)$  is invoked by the client  $p_i$  do
   $result_i \leftarrow \perp$ ; let  $sent\_msg = \langle op_x(param_x), i \rangle$ ;
  TO_broadcast( $sent\_msg$ );
  wait ( $result_i \neq \perp$ ); return ( $result_i$ ).

background task  $T$  is
  repeat forever
     $rec\_msg \leftarrow TO\_deliver()$ ;
     $\langle state_i, res \rangle \leftarrow \delta(state_i, rec\_msg.op)$ ;
    if ( $rec\_msg.proc = i$ ) then  $result_i \leftarrow res$  end if
  end repeat.

```

Algorithm 4: TO-broadcast-based construction

loop in which it first waits for the next message to-delivery. Then, it computes the next state of the object O , and, if it is the process that invoked the operation, it writes its result into its local variable $result_i$ to allow the operation to terminate. The correction of this simple universal construction follows directly from the properties of the to-broadcast abstraction [45, 81].

Implementing TO-broadcast from consensus. Algorithm 5 is a simple construction of TO-broadcast on top of an asynchronous system enriched with consensus objects [45].

Each process p_i manages four local variables: a sequence number sn_i initialized to 0, a set of message $delivered_i$ initialized to \emptyset , a queue $to_deliverable_i$ initialized to the empty sequence ϵ , and an auxiliary variable res_i . Let $broadcast(m)$ stand for “**for each** $j \in \{1, \dots, n\}$ **do** send(m) to p_j **end for**”. If the invoking process does not crash during its invocation, all processes receive m ; if it crashes an arbitrary subset of processes receive m . To simplify the presentation, it is assumed that a process can send a message to itself.

```

when  $p_i$  invokes TO_broadcast( $m$ ) do send( $m$ ) to itself.

when  $m$  is received for the first time do
  broadcast( $m$ );  $delivered_i \leftarrow delivered_i \cup \{m\}$ .

when ( $to\_deliverable_i$  contains messages not yet to-delivered) do
  let  $m =$  first message  $\in to\_deliverable_i$  not yet to-delivered;
  TO_deliver( $m$ ).

background task  $T$  is
  repeat forever
    wait( $delivered_i \setminus to\_deliverable_i \neq \emptyset$ );
    let  $seq = (delivered_i \setminus to\_deliverable_i)$ ;
    order the messages in  $seq$ ;
     $sn_i \leftarrow sn_i + 1$ ;  $res_i \leftarrow CS[sn_i].propose(seq)$ ;
    add  $res_i$  at the end of  $to\_deliverable_i$ 
  end repeat.

```

Algorithm 5: TO-broadcast from consensus

When a process p_i invokes TO_broadcast(m) it sends the message to itself, which entails its broadcast, and only then p_i adds m to its local set $delivered_i$. When a process receives a message m from another process for the first time, it does the same. It follows that when a process does not crash during its broadcast of a message m , all processes receive it. Hence, if a process p_j adds m to $delivered_j$, so do at least all the processes that do not crash.

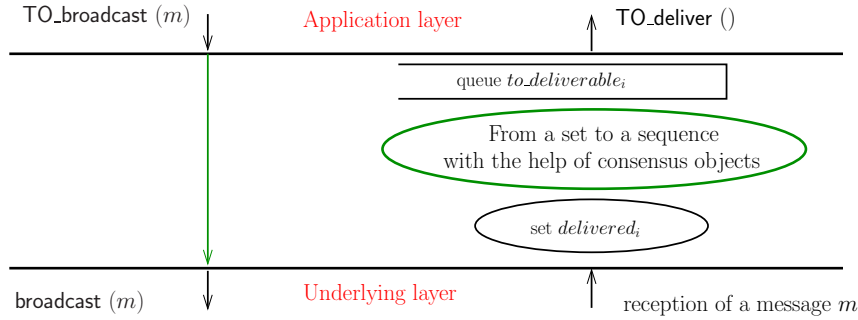


Figure 2: Structure of the consensus-based implementation of TO-broadcast

When, the queue $to_deliverable_i$ of a process p_i contains messages not yet locally to-delivered, p_i to-delivers them in the order in which they appear in $to_deliverable_i$.

The core of the algorithm is the background task T . A consensus object $SC[k]$ is associated with the iteration number k . A process p_i waits until there are messages in the set $delivered_i$ and not yet in the queue $to_deliverable_i$. When this occurs, process p_i computes this set of messages (seq) and order them. Then it proposes seq to the consensus instance $SC[k]$. This instance returns a sequence saved in res_i , which is added by p_i at the end of its local queue $to_deliverable_i$. The correctness of this algorithm relies on the properties of the consensus object. For any $k \geq 1$, the consensus instance $CS[k]$ returns the same sequence of messages to all the processes that invoke it. As processes execute instances in the same order, their queue $to_deliverable_i$ eventually contain the same sequence of messages. Formal proofs of this algorithms can be found in [20, 81].

While their styles are different, these two citations capture the universality issues encountered in asynchronous fault-tolerant distributed computing.

- *In sequential systems, computability is understood through the Church-Turing Thesis: anything that can be computed, can be computed by a Turing Machine. In distributed systems, where computations require coordination among multiple participants, computability questions have a different flavor. Here, too, there are many problems which are not computable, but these limits to computability reflect the difficulty of making decisions in the face of ambiguity, and have little to do with the inherent computational power of individual participants [49].*
- *A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable.*

L. Lamport, email Message-Id: <8705281923.AA09105@jumbo.dec.com>.

Sidebar 4: Two citations on universality

8 When Are Universal Constructions Possible?

An impossibility. A fundamental result in distributed computing is the impossibility to design a (deterministic) algorithm that solves consensus in the presence of asynchrony, even if only one process may crash, either in message-passing [36] or read/write shared memory systems [64]. Given that consen-

sus and TO-broadcast are equivalent, the state machine replication algorithm presented above cannot be implemented in asynchronous systems where processes can crash.

Thus, sequential thinking for concurrent computing has studied properties about the underlying system that enable the approach to go through. There are several ways of considering computationally stronger (read/write pr message-passing) models (see, e.g. [78, 81]), where state machine replication can be implemented. Some ways, mainly suited to message-passing systems, are presented in Sidebar 5. We discuss next a different way, through powerful communication hardware.

The case of enriched read/write systems. Nearly all read/write systems usually provide processes with synchronization-oriented atomic operations such as Test&Set, Compare&Swap, or the pair of operations Load Link/Store Conditional (LL/SC in short). These operations have a consensus number greater than 1. More specifically, the consensus number of Test&Set is 2, while the consensus number of both Compare&Swap and the pair LL/SC, is $+\infty$. Namely, 2-process (but not a 3-process) consensus can be implemented from Test&Set, despite crash failures. Compare&Swap (or LL/SC) can implement consensus for any number of processes. Hence, for any n , any object can be implemented in an asynchronous n -process read/write system enriched with Compare&Swap (or LL/SC), despite up to $n - 1$ process crashes. Furthermore, that are implementations that tolerate arbitrary, malicious failures [13, 81].

Ways of circumventing the impossibility of solving consensus:

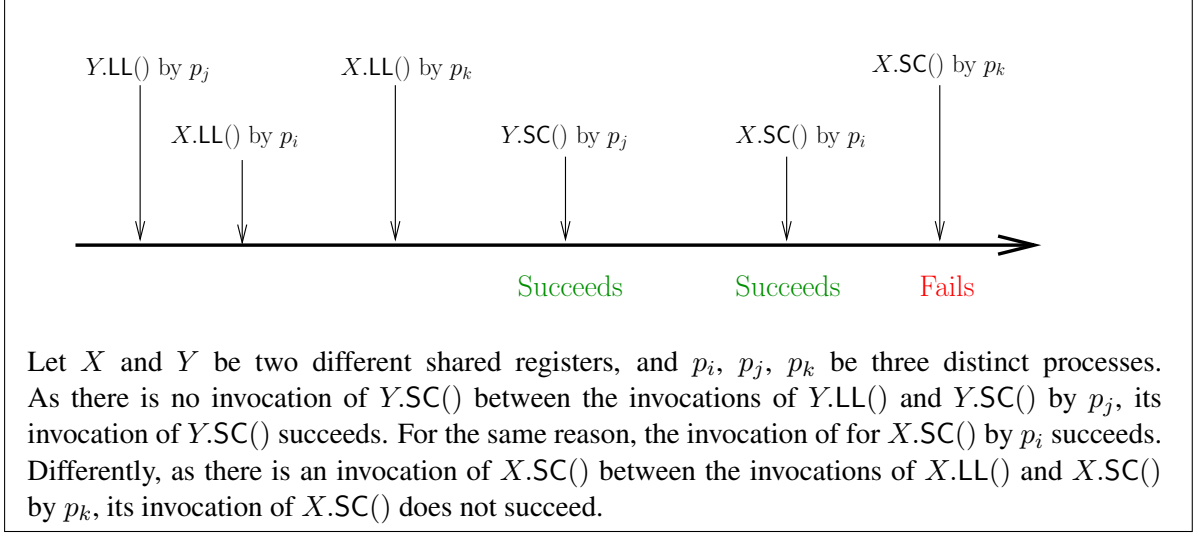
- The failure detector approach [19] abstracts away synchrony assumptions sufficient to distinguish between slow processes and dead processes.
- In *eventually synchronous systems* [28, 30] there is a time after which the processes run synchronously. The celebrated Paxos algorithm is an example [61].
- Using *random coins* [7, 70, 75] consensus is solvable with high probability.
- Using synchronization operations with consensus number n to solve consensus among n processes, in a fully asynchronous system.
- Consensus is solvable in some situations where not all combinations of input values are possible [69].

Sidebar 5: Circumventing consensus impossibility

Consensus from the pair LL/SC. The intuition of how the LL/SC operations work is as follows. Consider a memory location M accessed only by the operations LL/SC. Assumed that if a process invokes $M.SC(v)$ it has previously invoked $M.LL()$. The operation $M.LL()$ is a simple read of M which returns the current value of M . When a process p_i invokes $M.SC(v)$ the value v is written into M if and only if no other process invoked $M.SC()$ since its (p_i) last invocation of $M.LL()$. If the write succeeds $M.SC()$ returns `true`, otherwise it returns `false` (see Sidebar 6).

Algorithm 6 is a simple implementation of consensus object from the pair of operations LL/SC, which tolerates any number of process crashes. The consensus object is represented by the memory location M initialized to the default value \perp , which cannot be proposed). Each process manages a local variable val_i and a Boolean b_i .

When a process p_i invokes the operation `propose(v)` it first reads the value of M (first invocation of $M.LL()$) from which it obtains a value val_i . If $val_i \neq \perp$, it is the value decided by the consensus object and p_i returns it. If $val_i = \perp$, no value has yet been decided and possibly several processes are competing to impose their proposal as the decided value. Each of them invokes $M.SC()$. Due the semantics of the pair LL/SC one and only one of them succeeds. The winner returns its value, and the



Sidebar 6: An execution of LL/SC operations

```

operation propose( $v$ ) is
   $val_i \leftarrow M.LL();$ 
  if ( $val_i \neq \perp$ ) then return( $val_i$ )
    else  $b_i \leftarrow M.SC(v);$ 
      if  $b_i$  then return( $v$ )
      else  $val_i \leftarrow M.LL();$  return( $val_i$ )
    end if
  end if.

```

Algorithm 6: Consensus from the operations LL/SC

other competing processes read again the value of M (second invocation of $M.LL()$) and return the value proposed by the winner.

A simple stacking-based universal construction. As consensus objects can be built from the pair of operations LL/SC (Algorithm 6) and TO-broadcast communication abstraction can be built on top of consensus objects (Algorithm 5), their stacking allows us to use the universal construction Algorithm 4 to obtain an implementation of any sequentially-defined object, which copes with the net effect of asynchrony and process failures. This construction can give the reader a feeling for the distributed ledgers discussed in the next section.

A direct universal construction Algorithm 7 (based on an algorithm introduced in [32], simplified in [80]) is a direct universal construction (does not use an intermediate layer of TO-broadcast) of an object O with transition function δ , for n processes.

The shared memory is composed of the two following data structures.

- An array on n atomic single-writer multi-reader registers, $BOARD[1..n]$. While any process can read $BOARD[i]$, only process p_i can write it. Each register $BOARD[i]$ is composed of two fields: $BOARD[i].op$ which contains the last object operation invoked by p_i and $BOARD[i].sn$ which contains the associated local sequence number.
- An atomic register, $STATE$, accessed with the operations $LL()$ and $SC()$. It is made of three fields: $STATE.value$ contains the current state of the object under construction, $STATE.sn[1..n]$ is an array of local sequence number, and $STATE.res[1..n]$ which is an array of results. More pre-

cisely, $STATE.res[i]$ contains the result of the last object operation issued by p_j , and $STATE.sn[j]$ contains its sequence number.

```

when the operation  $op_x(param_x)$  is locally invoked do
   $sn_i \leftarrow sn_i + 1$ ;  $BOARD[i] \leftarrow \langle op_x(param_x), sn_i \rangle$ ;
  apply();
   $state_i \leftarrow STATE.LL()$ ; return( $state_i.res[i]$ ).

internal procedure apply() is
   $state_i \leftarrow STATE.LL()$ ;
   $board_i \leftarrow [BOARD[1], BOARD[2], \dots, BOARD[n]]$ ;
  for  $\ell \in \{1, \dots, n\}$  do
    if ( $board_i[\ell].sn = state_i.sn[\ell] + 1$ )
      then ( $state_i.value, state_i.res[\ell] \leftarrow \delta(state_i.value, pairs_i[\ell].op)$ ; % line A
             $state_i.sn[\ell] \leftarrow state_i.sn[\ell] + 1$  % line B
      end if
    end for;
   $success \leftarrow STATE.SC(state_i)$ ;
  if ( $\neg success$ ) then
     $state_i \leftarrow STATE.LL()$ ;
    if ( $sn_i = state_i.sn[i] + 1$ )
      then same as lines A and B with  $\ell = i$ ;
       $STATE.SC(state_i)$ 
    end if
  end if.

```

Algorithm 7: Universal construction for LL/SC-enriched shared memory systems (code for p_i)

Each process p_i manages a local sequence number sn_i and two local variables, denoted $board_i$ and $state_i$, which will contain local copies of $BOARD$ and $STATE$, respectively.

When a process p_i invokes an operation $op_x(param_x)$ on O , it informs all the processes of it by storing the pair $\langle op_x(param_x), sn_i \rangle$ in $BOARD[i]$. It executes then the internal procedure $apply()$ (which is the core of the construction). When it returns from $apply()$, it returns the result that has been deposited in $STATE.res[i]$. As there is no waiting statement in $apply()$, if the invoking process does not crash, it terminates its operation on O . Hence, the progress condition for object O is wait-freedom.

When p_i executes $apply()$, it first atomically reads the register $STATE$ (invocation of $STATE.LL()$), whose value is saved in its local variable $state_i$, reads the content of the array $BOARD$ and saves it in its local variable $board_i$. Let us remark that, while the the reading of each register $BOARD[j]$ is atomic, the array $BOARD$ is read asynchronously and consequently the reading of the whole array $BOARD$ is not at all atomic. When this is done, p_i starts a speculative execution, which consists in a "for" loop, with one iteration per process p_ℓ . If the last operation announced by p_ℓ is the next to be applied (according to its view of p_ℓ 's local sequence numbers), p_i applies p_ℓ 's operation to its local view of the current state of O , namely $state_i$. When this has been done for each process p_ℓ , p_i tries to write the new resulting state in $STATE$ (this is done by the invocation of $STATE.SC(state_i)$). If $STATE.SC(state_i)$ returns `true`, the speculative execution succeeded: p_i 's operation has been executed, as have also been operations from other processes, and consequently p_i 's invocation of $apply()$ terminates. Otherwise, the speculative execution failed. In this case, process p_i reads again $STATE$ (second invocation of $STATE.LL()$). If its operation has not been executed, p_i speculatively executes it on $state_i$, and tries to commit it by invoking $STATE.SC(state_i)$. If this invocation returns `true` its operation is taken into account. If it returns `false`, another process p_k invoked successfully $STATE.SC(state_k)$ between the invocations of $STATE.LL()$ and $STATE.SC()$ by p_i . But in this case, due to the fact that LL/SC are atomic operations, necessarily when p_k read $BOARD[i]$ it was informed of p_i 's operation and consequently executed it. Hence, the result obtained by p_i from $STATE.res[i]$ is the one associated with its last operation.

9 Distributed Ledgers

Since ancient times, ledgers have been at the heart of commerce, to represent concurrent transactions by a permanent list of individual records sequentialized by date (Fig. 3). Today we are beginning to see algorithms that enable the collaborative creation of digital distributed ledgers with properties and capabilities that go far beyond traditional physical ledgers. All participants within a network can have their own copy of the ledger. Any of them can append a record to the ledger, which is then reflected in all copies in minutes or even seconds. The records stored in the ledger can stay temper-proof, using cryptographic techniques.




Figure 3: Ledger object: a crucial issue for the processes is to agree on the next block to add

Ledgers as universal constructions. Mostly known because of their use in cryptocurrencies, and due to its *blockchain* incarnation [68], from the perspective of this paper a *distributed ledger* is a byzantine fault-tolerant replicated implementation of a specific *ledger* object. The ledger object has two operations, `read()` and `append()`. Its sequential specification defines it as a list of blocks. A block X can be added at the end of the list with the operation `append(X)`, while a `read()` returns the whole list. In the case of a cryptocurrency, X may contain a set of transactions. (See [33, 81] for a formalization of the ledger object.)

Thus, a ledger object, as any other object, can be implemented in a distributed, fault-tolerant way, using the state machine replication technique. Furthermore, it can then be used as a universal construction of an object O defined by a state machine with a transition function δ . To do so, when a process invokes `append(X)`, X consists of a transition to be applied to the state machine. The state of the object is obtained through a `read()` invocation, which returns the sequence of operations which have been sequentially appended to the ledger, and then locally applying them starting from the initial state of the object (see [81] for more details).

Three remarkable properties. The apparently innocent idea of a `read()` operation that returns the list of commands that have been applied to the state machine, opens the discussion of one of the remarkable points of distributed ledgers that has brought them to such wide attention. The possibility of guaranteeing a *temper-proof* list of commands. The blockchain implementation is by using cryptographic hashes that link each record to the previous one (although it actually has been known in cryptography community for years [67]).

The ledger implementation used in Bitcoin showed that it is possible to have a state machine replication tolerating Byzantine failures that *scales* to hundreds of thousands of processes. The cost is temporality sacrificing consistency— forks can happen at the end of the blockchain, which means that the last few records in the blockchain may have to be withdrawn.

The third remarkable property brought to the public attention by distributed ledgers is the issue of who the participants can be. As opposed to classic algorithms for mastering concurrency through

sequential thinking, the participants do not have to be a priori-known, can vary with time, and may even be anonymous. Anyone can append a block, and read the ledger (although there are also permissioned versions where participants have to be registered, and even hybrid models). In a sense, a distributed ledger is an open distributed database, with no central authority, where the data itself is distributed among the participants.

Agreement in dynamic systems. Bitcoin’s distributed ledger implementation is relatively simple to explain in the framework of state machine replication. Conceptually it builds on randomized consensus (something that had already been carefully studied in traditional approaches, e.g. Sidebar 5), through the following ingenious technique to implement it. Whenever several processes want to concurrently append a block, they participate in a lottery. Each process selects a random number (by solving cryptographic puzzles) between 0 and some large integer K , and the one that gets a number smaller than $k \ll K$, wins, and has the right to append its desired block. The implementation details of the lottery (by a procedure called *proof of work*) are not important for this paper; what is important here, is that processes cannot cheat by biasing the random number they get. Thus, with high probability only one wins. However, from time to time, more than one process wins and a *fork* happens, with more than one block being appended at the end of the ledger. Again, for the purpose of this paper, it suffices to say that only one branch eventually prevails (in Bitcoin this is achieved by always appending to the longest branch). This introduces a new interesting idea into the paradigm of mastering concurrency through sequential thinking: a tradeoff between faster state machine replication, and temporary loss of consistency. In other words, the x operations at the very end of the ledger, for some constant x (which depends on the assumptions about the environment) cannot yet be considered committed. To be sure (with high probability) that an operation has permanently been applied to the ledger, a process has to wait until it is at a depth greater than x in the list of blocks.

10 On the Limits of the Approach

It is intuitively clear, and it has been formally proved since a long time that linearizability is an expensive requirement. Recent papers in the context of shared memory programming, argue that it is often possible to improve performance of concurrent data structures by relaxing their semantics (see, e.g. [16, 44, 86, 90, 94]). This approach has been studied even in the case of randomized algorithms: *distributional linearizability* is a correctness condition for randomized relaxed data structures [1], defined in terms of a sequential specification and a cost function measuring the deviation from the sequential specification. Relaxations of queues and stacks that are *not* sequentially specified and can be implemented using only read/write operations are presented in [18]. Concurrent specifications are discussed further at the end of this section.

In the context of distributed systems, *eventual consistency* is widely deployed to achieve high availability by guaranteeing that if no new updates are made to a given data item, eventually all accesses to that item will return the last updated value [96]. Eventual consistency (also called optimistic replication), which is deployed in some distributed systems, has origins in early mobile computing. A system that has achieved eventual consistency is often said to have converged.

In the case of distributed ledgers, we have seen the benefit that can be gained by relaxing the sequential approach to mastering concurrency: branches at the end of the blockchain (such as Bitcoin) temporarily violate a consistent view of the ledger. Still, ledger implementations based on the blockchain technology suffer from a performance bottleneck due to the requirement of ordering all transactions in a single list, which has prompted the exploration of partially ordered ledgers, based on directed acyclic graphs such as those based on Iota, Tangle, or Hedera Hashgraph systems. The benefit is scalability to thousands of processes, that instead of communicating with each other to decide on a single leader that will append a block, they avoid communication altogether, using random numbers.

The *CAP Theorem* formalizes a fundamental limitation of the approach of mastering concurrency through sequential reasoning: at most two of the following three properties are achievable, Consistency (linearizability), Availability, Partition tolerance [41, 42]. This may give an intuition of why distributed ledgers implementations have temporary forks. An alternative is a cost in availability, and postpone the property that every non-failing participant returns a response for all operations in a reasonable amount of time. We have already seen in the ABD algorithm that the system continues to function and upholds its consistency guarantees, provided that only a minority of processes may fail.

Finally, another fundamental limitation to the approach of mastering concurrency through sequential reasoning is that not all concurrent problems of interest have sequential specifications. Many examples are discussed in [17]. Thus, the need for a formalism that extends the usual way of specifying a concurrent object through an automaton, to one that specifies the output of the object also in concurrent invocations of operations. An *interval-sequential* automaton generalizes the usual sequential automaton, and *interval-linearizability* is the corresponding notion, that associates concurrent executions to the specification defined by the interval-automaton. This is the formalism described in [17], that generalizes linearizability to arbitrary concurrent specifications.




Figure 4: An evolution line

11 Conclusion

The aim of this article was to show how does the theme of reducing concurrent programming to sequential reasoning weaves through history since the early days and along different domains (operating systems, databases, shared memory and distributed systems, cryptocurrencies, etc), to build complex concurrent systems. The thread brings in a scientific foundation through common conceptual tools, such as sequential specifications, progress and consistency conditions, synchronization abstractions like consensus, communication mechanisms such as broadcast and gossiping, fault-tolerance techniques, etc. It evolves from concrete resource-oriented mutual exclusion in a failure free-context, through immaterial objects and failures, to universal constructions of replicated state machines, to current trends on dynamic, temper-proof distributed ledgers. The associated evolution line is depicted in Fig. 4. The deep continuity lasting more than sixty years, is now exploring its frontiers, looking for roundabouts to the inherent limitations of the approach.

Acknowledgments

This work has been supported by the following projects: French ANR 16-CE40-0023-03 DESCARTES, devoted to modular structures in distributed computing, INRIA-UNAM *Équipe Associée* LiDiCo (at the Limits of Distributed Computing), and UNAM PAPIIT IN109917 and IN106520.

References

- [1] Alistarh D., Brown T., Kopinsky J., Li J. and Nadiradze G., Distributionally linearizable data structures. *Proc. 30th on Symposium on Parallelism in Algorithms and Architectures (SPAA'15)*, ACM Press, pp. 133–142 (2018)
- [2] Alpern B. and Schneider F.B., Defining liveness. *Information Processing Letters*, 21(4):181-185 (1985)

- [3] Akkoyunlu E.A., Ekanadham K., and Huber R.V., Some constraints and tradeoffs in the design of network communications. *Proc. 5th ACM Symposium on Operating Systems Principles (SOSP'75)*, ACM Press, pp. 67-74 (1975)
- [4] Attiya H., Bar-Noy A., and Dolev D., Sharing memory robustly in message-passing systems. *Journal of the ACM*, 42(1):121-132 (1995)
- [5] Attiya H. and Welch J., *Distributed computing: fundamentals, simulations and advanced topics*, (2d Edition), Wiley, 414 pages (2004)
- [6] Avizienis A., Design of fault-tolerant computers. *Proc. Fall Joint Computer Conference AFIPS*, pp. 733-743 (1967)
- [7] Ben-Or M., Another advantage of free choice: completely asynchronous agreement protocols. *Proc. 2nd ACM Symposium on Principles of Distributed Computing (PODC'83)*, pp. 27-30 (1983)
- [8] Bernstein Ph. A. and Goodman N., An algorithm for concurrency control and recovery in replicated distributed databases. *ACM Transactions on Database Systems* 9(4):596-615 (1984)
- [9] Bernstein Ph. A., Hadzilacos V., and Goodman N., *Concurrency control and recovery in database systems*. Addison-Wesley, 370 pages, ISBN 978-0201107159 (1987)
- [10] Bernstein, Ph. A. and Goodman, N., Concurrency control in distributed database systems. *ACM Computing Surveys*, 13(2):85-221. (1981)
- [11] Brinch Hansen P. (Editor), *The origin of concurrent programming*. Springer, 534 pages (2002)
- [12] Brinch Hansen P., The nucleus of a multiprogramming system. *Communications of the ACM*, 13(4):238-241. (1970)
- [13] Cachin Ch., State machine replication with Byzantine faults. In *Replication*, Springer LNCS 5959, pp. 169-184 (2011)
- [14] Cachin Ch., Guerraoui R., and Rodrigues L., *Reliable and secure distributed programming*, Springer, 367 pages, ISBN 978-3-642-15259-7 (2011)
- [15] Cachopo J. and Rito-Silva A., Versioned boxes as the basis for transactional memory. *Science of Computer Programming*, 63(2):172-175 (2006)
- [16] Calciu I., Sen S., Balakrishnan M., and Aguilera M., How to implement any concurrent data structure for modern servers. *ACM Operating Systems Review*, 51(1):24-32 (2017)
- [17] Castañeda A., Rajsbaum S., and Raynal M., Unifying concurrent objects and distributed tasks: interval-linearizability. *Journal of the ACM*, 65(6):45:1-45:42 (2018)
- [18] Castañeda A., Rajsbaum S., and Raynal M., What can be done with consensus number one: relaxed queues and stacks. *arXiv:2005.05427v1*, 32 pages, submitted to publication (2020)
- [19] Chandra T.D., Hadzilacos V., and Toueg S., The weakest failure detector for solving consensus. *Journal of the ACM*, 43(4):685-722 (1996)
- [20] Chandra T.D. and Toueg S., Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225-267 (1996)
- [21] Charron-Bost N., Pedone F., and Schiper A. (Editors), *Replication: theory and practice*, Springer LNCS 5959 (2010)
- [22] Courtois P.-J., Heymans F., and Parnas D.L., Concurrent control with readers and writers. *Communications of the ACM*, 14(10):667-668 (1971)

- [23] Cremers A. and Hibbard T.N., An algebraic approach to concurrent programming control and related complexity problems, *Symposium on Algorithms and Complexity: New Directions and Recent Results*, Pittsburgh, Pennsylvania (1976)
- [24] Dice D., Shalev O., and Shavit N., Transactional locking II. *Proc. 20th Int'l Symposium on Distributed Computing (DISC'06)*, Springer, LNCS 4167, pp. 194-208 (2006)
- [25] Dijkstra E.W., Solution of a problem in concurrent programming control. *Communications of the ACM*, 8(9):569 (1965)
- [26] Dijkstra E.W., Cooperating sequential processes. In *Programming Languages*, (F. Genuys Ed.), Academic Press, pp. 43-112 (1968)
- [27] Dijkstra E.W., Hierarchical ordering of sequential processes. *Acta Informatica*, 1(1):115-138 (1971)
- [28] Dolev D., Dwork C., and Stockmeyer L., On the minimal synchronism needed for distributed consensus. *Journal of the ACM*, 34(1):77-97 (1987)
- [29] Durand A., Raynal M., and Taubenfeld G., Set agreement and renaming in the presence of contention-related crash failures. *Proc. 20th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'18)*, Springer LNCS 11201, pp. 269-283 (2018)
- [30] Dwork C., Lynch N. and Stockmeyer L., Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2), 288-323 (1988)
- [31] Eswaran, K.P., Gray J.N., Lorie R.A., and Traiger I.L., The notions of consistency and predicate locks in a database system. *Commun. ACM* 19(11) (November 1976), 624–633.
- [32] Fatourou P. and Kallimanis N.D., Highly-efficient wait-free synchronization. *Theory of Computing Systems*, 55:475-520 (2014)
- [33] Fernández Anta A., Konwar K., Georgiou Ch., and Nicolaou N., Formalizing and implementing distributed objects layers. *ACM SIGACT News*, 49(2):58-76 (2018)
- [34] Fischer M.J. and Lynch N.A., A lower bound for the time to assure interactive consistency. *Information Processing Letters*, 14(4):183-186 (1982)
- [35] Fischer M.J., Lynch N.A., and Paterson M.S., Impossibility of distributed consensus with one faulty process. *Proc. 2d ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*, ACM Press, pp. 1-7 (1983)
- [36] Fischer M.J., Lynch N.A., and Paterson M.S., Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374-382 (1985)
- [37] Fourastié J., *Comment mon cerveau s'informe*, *Journal d'une recherche*. Robert Laffont, 262 pages (1974)
- [38] Gafni E. and Guerraoui R., Generalizing universality. *Proc. 22nd Int'l Conference on Concurrency Theory (CONCUR'11)*, Springer LNCS 6901, pp. 17-27 (2011)
- [39] Gafni E. and Rajsbaum S., Recursion in distributed computing. *Proc. 12th Int'l Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'10)*, Springer LNCS 6366, pp. 362-376 (2010)
- [40] Garg V.K., *Elements of Distributed Computing*. Wiley-Interscience, 423 pages (2002)
- [41] Gilbert S. and Lynch N., Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web-services. *ACM SIGACT News*, 33(2):51-59 (2002)
- [42] Gilbert S. and Lynch N., Perspectives on the CAP theorem. *IEEE Computer* 45(2):30-36 (2012)
- [43] Gray J. and Reuter A., *Transactions processing: concepts and techniques*, Morgan Kaufmann, 1070 pages, ISBN 978-1558601901 (1992)

- [44] Haas A., Henzinger T.A., Holzer A., Kirsch Ch.M, Lippautz M., Payer H., Sezgin A., Sokolova A., and Veith H., Local linearizability for concurrent container-type data structures. *Proc. 27th Int'l Conference on Concurrency Theory, (CONCUR'16)*, LIPIcs Vol. 59(6)1-:15 (2016)
- [45] Hadzilacos V. and Toueg S., A modular approach to fault-tolerant broadcasts and related problems. *Tech Report 94-1425*, 83 pages, Cornell University (1994) Extended version of "Fault-Tolerant Broadcasts and Related Problems", in *Distributed systems, 2nd Edition*, Addison-Wesley/ACM, pp. 97-145 (1993)
- [46] Herlihy M.P., Wait-free synchronization. *ACM Transactions on Programming Languages and Systems*, 13(1):124-149 (1991)
- [47] Herlihy M.P., Luchangco V., and Moir M., Obstruction-free synchronization: double-ended queues as an example. *Proc. 23th Int'l IEEE Conference on Distributed Computing Systems (ICDCS'03)*, IEEE Press, pp. 522-529 (2003)
- [48] Herlihy M.P. and Moss J.E.B., Transactional memory: architectural support for lock-free data structures. *Proc. 20th ACM Int'l Symposium on Computer Architecture (ISCA'93)*, ACM Press, pp. 289-300 (1993)
- [49] Herlihy M., Rajsbaum S., and Raynal M., Power and limits of distributed computing shared memory models. *Theoretical Computer Science*, 509:3-24 (2013)
- [50] Herlihy M. and Shavit N., *The art of multiprocessor programming*. Morgan Kaufmann, 508 pages, ISBN 978-0-12-370591-4 (2008)
- [51] Herlihy M.P. and Wing J.M, Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463-492 (1990)
- [52] Imbs D., Raynal M., and Taubenfeld G., On asymmetric progress conditions. *Proc. 29th ACM Symposium on Principles of Distributed Computing (PODC'10)*, ACM Press, pp. 55-64 (2010)
- [53] Knuth D.E., Additional comments on a problem in concurrent programming control. *CACM*, 9(5):321-322 (1966)
- [54] Kshemkalyani A.D. and Singhal M., *Distributed computing: principles, algorithms and systems*. Cambridge University Press, 736 pages (2008)
- [55] Lamport L., A new solution of Dijkstra's concurrent programming problem. *Communications of the ACM*, 17(8):453-455 (1974)
- [56] Lamport L., Concurrent reading and writing. *Communications of the ACM*, 20(11):806-811 (1977)
- [57] Lamport L., Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558-565 (1978)
- [58] Lamport L., How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Transactions on Computers*, 28(9):690-691 (1979)
- [59] Lamport L., On interprocess communication, Part I: basic formalism. *Distributed Computing*, 1(2):77-85 (1986)
- [60] Lamport L., A fast mutual exclusion algorithm. *ACM Transactions on Computer Systems*, 5(1):1-11 (1987)
- [61] Lamport L., The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133-169 (1998)
- [62] Lamport L., Shostack R. and Pease M., The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382-401 (1982)
- [63] Lamson B., Sturgis H., Crash Recovery in a Distributed Data Storage System. *Technical Report, Computer Science Laboratory, Xerox*, Palo Alto Research Center (1976)
- [64] Loui M. and Abu-Amara H., Memory requirements for agreement among unreliable asynchronous processes. *Advances in Computing Research*, 4:163-183, JAI Press (1987)

- [65] Lynch N. A., *Distributed algorithms*. Morgan Kaufmann Pub., 872 pages, ISBN 1-55860-384-4 (1996)
- [66] Lynch N. A., Some perspectives on PODC. *Distributed Computing*, 16(2):71–74 (2003)
- [67] Merkle, R.F., A digital signature based on a conventional encryption function. *Proc. Int’l Conference on Advances in Cryptography*, Springer LNCS 293, pp. 369-378 (1987)
- [68] Nakamoto S., Bitcoin: a peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf> (2008)
- [69] Mostéfaoui A., Rajsbaum S. and Raynal M., Conditions on input vectors for consensus solvability in asynchronous distributed systems. *Journal of the ACM* 50(6):922–954 (2003)
- [70] Mostéfaoui A., Moumen H., and Raynal M., Signature-free asynchronous binary Byzantine consensus with $t < n/3$, $O(n^2)$ messages, and $O(1)$ expected time. *Journal of the ACM*, 62(4), Article 31, 21 pages (2015)
- [71] Papadimitriou C.H., The serializability of concurrent database updates. *Journal of the ACM*, 26(4): 631–653, (1979)
- [72] Pease M., R. Shostak R. and Lamport L., Reaching agreement in the presence of faults. *Journal of the ACM*, 27:228-234 (1980)
- [73] Peterson G.L., Myths about the mutual exclusion problem. *Information Processing Letters*, 12(3):115-116 (1981)
- [74] Peterson G.L., Concurrent reading while writing. *ACM Transactions on Programming Languages and Systems*, 5:46-55 (1983)
- [75] Rabin M., Randomized Byzantine generals. *Proc. 24th IEEE Symposium on Foundations of Computer Science (FOCS’83)*, IEEE Computer Society Press, pp. 116-124 (1983)
- [76] Rajsbaum S. and Raynal M., Mastering concurrent computing through sequential thinking: a half-century evolution. *Communications of the ACM*, 63(1):52-61 (2020)
- [77] Raynal M., *Algorithms for mutual exclusion*. The MIT Press, 107 pages, ISBN 0-262-18119-3 (1986)
- [78] Raynal M., *Concurrent programming: algorithms, principles and foundations*. Springer, 515 pages, ISBN 978-3-642-32026-2 (2013)
- [79] Raynal M., *Distributed algorithms for message-passing systems*. Springer, 510 pages, ISBN 978-3-642-38122-5 (2013)
- [80] Raynal M., Distributed universal constructions: a guided tour. *Bulletin of the European Association of Theoretical Computer Science (EATCS)*, 121(1):64-96 (2017)
- [81] Raynal M., *Fault-tolerant message-passing distributed systems: an algorithmic approach*, Springer, 459 pages, ISBN 978-3-319-94140-0 (2018)
- [82] Raynal M., Stainer J., and Taubenfeld G., Distributed universality. *Algorithmica*. 76(2):502-535 (2016)
- [83] Ritchie D., The evolution of the Unix time-sharing system. *Proc. Symposium on Language Design and Programming Methodology*, Springer-Verlag, pp. 25–36 (1979)
- [84] Santoro N., *Design and analysis of distributed algorithms*, Wiley-Interscience, 589 pages, ISBN 0-471-71997-8 (2007)
- [85] Schneider F.B., Implementing fault-tolerant services using the state machine approach. *ACM Computing Surveys*, 22(4):299–319 (1990)
- [86] Shavit N., Data structures in the multicore age. *Communications of the ACM*, 54(3):76–84 (2011)
- [87] Shavit N. and Touitou D., Software transactional memory. *Distributed Computing*, 10(2):99-116 (1997)

- [88] Stearns, R.C., Lewis P.M., and Rosenkrantz D.J., Concurrency control for database systems. *Proc. 17th Symposium on Foundations of Computer Science (FOCS'76)*, IEEE Press, pp. 19–32 (1976)
- [89] Stonebraker M. and Neuhold E.J., A distributed database version of INGRES. *Proc. 2nd Berkeley Workshop on Distributed Data Management and Computer Networks*, pp. 19-36 (1977)
- [90] Talmage, E. and Welch J. L., Relaxed data types as consistency conditions. *Algorithms*, 11(5):61 (2018)
- [91] Taubenfeld G., *Synchronization algorithms and concurrent programming*. Pearson Prentice-Hall, 423 pages, ISBN 0-131-97259-6 (2006)
- [92] Taubenfeld G., Contention-sensitive data structure and algorithms. *Proc. 23rd Int'l Symposium on Distributed Computing (DISC'09)*, Springer, LNCS 5805, pp. 157-171 (2009)
- [93] Taubenfeld G., The Computational Structure of Progress Conditions. *Proc. 24th Int'l Symposium on Distributed Computing (DISC'10)*, Springer LNCS 6343, pp. 221-235 (2010)
- [94] Terry D.B., Prabhakaran V., Kotla R., Balakrishnan M., Aguilera M.K., and Abu-Libdeh A., Consistency-based service level agreements for cloud storage. *Proc. 24th ACM Symposium on Operating Systems Principles*, ACM Press, pp. 309-324 (2013)
- [95] J. H. Wensley J.H., Lamport L., Goldberg J. Green M.W., Levitt K.N., Melliar-Smith M., and Weinstock C.B., SIFT: Design and analysis of a fault-tolerant computer for aircraft control. *Proceedings of the IEEE*, 66(10):1240-1255 (1978)
- [96] Vogels W., Eventually consistent. *Communications of the ACM*, 52(1):40–44 (2009)
- [97] Yourcenar M., *Mémoires d'Hadrien*. Gallimard Ed., 347 pages (1951)