# Microkernel

In computer science, a **microkernel** (often abbreviated as **µ-kernel**) is the near-minimum amount of software that can provide the mechanisms needed to implement an operating system (OS). These mechanisms include low-level address space management, thread management, and inter-process communication (IPC).

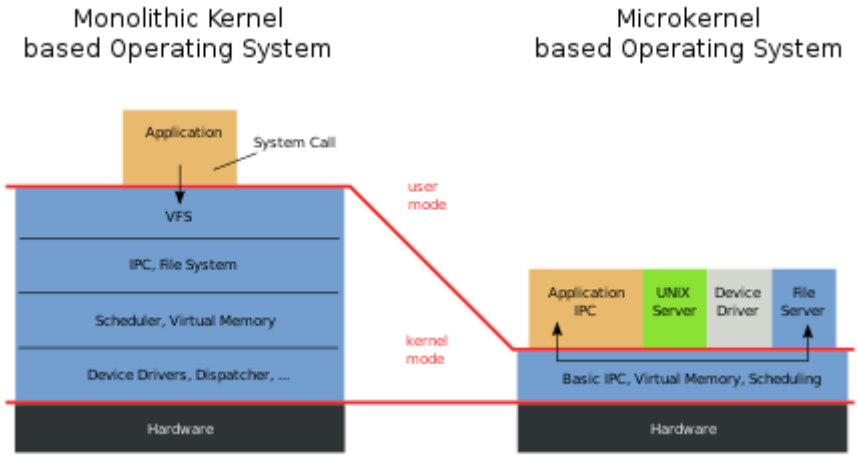


Structure of monolithic and microkernel-based operating systems, respectively

If the hardware provides multiple rings or CPU modes, the microkernel may be the only software executing at the most privileged level, which is generally referred to as supervisor or kernel mode. Traditional operating system functions, such as device drivers, protocol stacks and file systems, are typically removed from the microkernel itself and are instead run in user space.[1]

In terms of the source code size, microkernels are often smaller than monolithic kernels. The MINIX 3 microkernel, for example, has only approximately 12,000 lines of code.[2]

## Contents

# History

Microkernels trace their roots back to Danish computer pioneer Per Brinch Hansen and his tenure in Danish computer company Regnecentralen where he led software development efforts for the RC 4000 computer.[3] In 1967, Regnecentralen was installing a RC 4000 prototype in a Polish fertilizer plant in Puławy. The computer used a small real-time operating system tailored for the needs of the plant. Brinch Hansen and his team became concerned with the lack of generality and reusability of the RC 4000 system. They feared that each installation would require a different operating system so they started to investigate novel and more general ways of creating software for the RC 4000.[4] In 1969, their effort resulted in the completion of the RC 4000 Multiprogramming System. Its nucleus provided inter-process communication based on message-passing for up to 23 unprivileged processes, out of which 8 at a time were protected from one another. It further implemented scheduling of time slices of programs executed in parallel, initiation and control of program execution at the request of other running programs, and initiation of data transfers to or from peripherals. Besides these elementary mechanisms, it had no built-in strategy for program execution and resource allocation. This strategy was to be implemented by a hierarchy of running programs in which parent processes had complete control over child processes and acted as their operating systems.[5][6]

Following Brinch Hansen's work, microkernels have been developed since the 1970s.[7] The term microkernel itself first appeared no later than 1981.[8] Microkernels were meant as a response to changes in the computer world, and to several challenges adapting existing "mono-kernels" to these new systems. New device drivers, protocol stacks, file systems and other low-level systems were being developed all the time. This code was normally located in the monolithic kernel, and thus required considerable work and careful code management to work on. Microkernels were developed with the idea that all of these services would be implemented as user-space programs, like any other, allowing them to be worked on monolithically and started and stopped like any other program. This would not only allow these services to be more easily worked on, but also separated the kernel code to allow it to be finely tuned without worrying about unintended side effects. Moreover, it would allow entirely new operating systems to be "built up" on a common core, aiding OS research.

Microkernels were a very hot topic in the 1980s when the first usable local area networks were being introduced.. The AmigaOS Exec kernel was an early example, introduced in 1986 and used in a PC with relative commercial success. The lack of memory protection, considered in other respects a flaw, allowed this kernel to have very high message-passing performance because it did not need to copy data while exchanging messages between user-space programs.[9]

The same mechanisms that allowed the kernel to be distributed into user space also allowed the system to be distributed across network links. The first microkernels, notably Mach created by Richard Rashid, proved to have disappointing performance, but the inherent advantages appeared so great that it was a major line of research into the late 1990s. However, during this time the speed of computers grew greatly in relation to networking systems, and the disadvantages in performance came to overwhelm the advantages in development terms.

Many attempts were made to adapt the existing systems to have better performance, but the overhead was always considerable and most of these efforts required the user-space programs to be moved back into the kernel. By 2000, most large-scale Mach kernel efforts had ended, although Apple's macOS, released in 2001, still uses a hybrid kernel called XNU, which combines a heavily modified (hybrid) OSF/1's Mach kernel (OSFMK 7.3 kernel) with code from BSD UNIX,[10][11] and this kernel is also used in iOS, tvOS, and watchOS. Windows NT, starting with NT 3.1 and continuing with Windows 10, uses a hybrid kernel design. As of 2012, the Mach-based GNU Hurd is also functional and included in testing versions of Arch Linux and Debian.

Although major work on microkernels had largely ended, experimenters continued development. It has since been shown that many of the performance problems of earlier designs were not a fundamental limitation of the concept, but instead due to the designer's desire to use single-purpose systems to implement as many of these services as possible. Using a more pragmatic approach to the problem, including assembly code and relying on the processor to enforce concepts normally supported in software led to a new series of microkernels with dramatically improved performance.

Microkernels are closely related to exokernels.[12] They also have much in common with hypervisors,[13] but the latter make no claim to minimality and are specialized to supporting virtual machines; the L4 microkernel frequently finds use in a hypervisor capacity.

# Introduction

Early operating system kernels were rather small, partly because computer memory was limited. As the capability of computers grew, the number of devices the kernel had to control also grew. Throughout the early history of Unix, kernels were generally small, even though they contained various device drivers and file system implementations. When address spaces increased from 16 to 32 bits, kernel design was no longer constrained by the hardware architecture, and kernels began to grow larger.

The Berkeley Software Distribution (BSD) of Unix began the era of larger kernels. In addition to operating a basic system consisting of the CPU, disks and printers, BSD added a complete TCP/IP networking system and a number of "virtual" devices that allowed the existing programs to work 'invisibly' over the network. This growth continued for many years, resulting in kernels with millions of lines of source code. As a result of this growth, kernels were prone to bugs and became increasingly difficult to maintain.

The microkernel was intended to address this growth of kernels and the difficulties that resulted. In theory, the microkernel design allows for easier management of code due to its division into user space services. This also allows for increased security and stability resulting from the reduced amount of code running in kernel mode. For example, if a networking service crashed due to buffer overflow, only the networking service's memory would be corrupted, leaving the rest of the system still functional.

# Inter-process communication

Inter-process communication (IPC) is any mechanism which allows separate processes to communicate with each other, usually by sending messages. Shared memory is, strictly defined, also an inter-process communication mechanism, but the abbreviation IPC usually refers to message passing only, and it is the latter that is particularly relevant to microkernels. IPC allows the operating system to be built from a number of smaller programs called servers, which are used by other programs on the system, invoked via IPC. Most or all support for peripheral hardware is handled in this fashion, with servers for device drivers, network protocol stacks, file systems, graphics, etc.

IPC can be synchronous or asynchronous. Asynchronous IPC is analogous to network communication: the sender dispatches a message and continues executing. The receiver checks (polls) for the availability of the message, or is alerted to it via some notification mechanism. Asynchronous IPC requires that the kernel maintains buffers and queues for messages, and deals with buffer overflows; it also requires double copying of messages (sender to kernel and kernel to receiver). In synchronous IPC, the first party (sender or receiver) blocks until the other party is ready to perform the IPC. It does not require buffering or multiple copies, but the implicit rendezvous can make programming tricky. Most programmers prefer asynchronous send and synchronous receive.

First-generation microkernels typically supported synchronous as well as asynchronous IPC, and suffered from poor IPC performance. Jochen Liedtke assumed the design and implementation of the IPC mechanisms to be the underlying reason for this poor performance. In his L4 microkernel he pioneered methods that lowered IPC costs by an order of magnitude.[14] These include an IPC system call that supports a send as well as a receive operation, making all IPC synchronous, and passing as much data as possible in registers. Furthermore, Liedtke introduced the concept of the *direct process switch*, where during an IPC execution an (incomplete) context switch is performed from the sender directly to the receiver. If, as in L4, part or all of the message is passed in registers, this transfers the in-register part of the message without any copying at all. Furthermore, the overhead of invoking the scheduler is avoided; this is especially beneficial in the common case where IPC is used in an remote procedure call (RPC) type fashion by a client invoking a server. Another optimization, called *lazy scheduling*, avoids traversing scheduling queues during IPC by leaving threads that block during IPC in the ready queue. Once the scheduler is invoked, it moves such threads to the appropriate waiting queue. As in many cases a thread gets unblocked before the next scheduler invocation, this approach saves significant work. Similar approaches have since been adopted by QNX and MINIX 3.

In a series of experiments, Chen and Bershad compared memory cycles per instruction (MCPI) of monolithic Ultrix with those of microkernel Mach combined with a 4.3BSD Unix server running in user space. Their results explained Mach's poorer performance by higher MCPI and demonstrated that IPC alone is not responsible for much of the system overhead, suggesting that optimizations focused exclusively on IPC will have limited impact.[15] Liedtke later refined Chen and Bershad's results by making an observation that the bulk of the difference between Ultrix and Mach MCPI was caused by capacity cache-misses and concluding that drastically reducing the cache working set of a microkernel will solve the problem.[16]

In a client-server system, most communication is essentially synchronous, even if using asynchronous primitives, as the typical operation is a client invoking a server and then waiting for a reply. As it also lends itself to more efficient implementation, most microkernels generally followed L4's lead and only provided a synchronous IPC primitive. Asynchronous IPC could be implemented on top by using helper threads. However, experience has shown that the utility of synchronous IPC is dubious: synchronous IPC forces a multi-threaded design onto otherwise simple systems, with the resulting synchronization complexities. Moreover, an RPC-like server invocation sequentializes client and server, which should be avoided if they are running on separate cores. Versions of L4 deployed in commercial products have therefore found it necessary to add an asynchronous notification mechanism to better support asynchronous communication. This signal-like mechanism does not carry data and therefore does not require buffering by the kernel. By having two forms of IPC, they have nonetheless violated the principle of minimality. Other versions of L4 have switched to asynchronous IPC completely.[17]

As synchronous IPC blocks the first party until the other is ready, unrestricted use could easily lead to deadlocks. Furthermore, a client could easily mount a denial-of-service attack on a server by sending a request and never attempting to receive the reply. Therefore, synchronous IPC must provide a means to prevent indefinite blocking. Many microkernels provide timeouts on IPC calls, which limit the blocking time. In practice, choosing sensible timeout values is difficult, and systems almost inevitably use infinite timeouts for clients and zero timeouts for servers. As a consequence, the trend is towards not providing arbitrary timeouts, but only a flag which indicates that the IPC should fail immediately if the partner is not ready. This approach effectively provides a choice of the two timeout values of zero and infinity. Recent versions of L4 and MINIX have gone down this path (older versions of L4 used timeouts). QNX avoids the problem by requiring the client to specify the reply buffer as part of the message send call. When the server replies the kernel copies the data to the client's buffer, without having to wait for the client to receive the response explicitly.[18]

## Servers

Microkernel servers are essentially daemon programs like any others, except that the kernel grants some of them privileges to interact with parts of physical memory that are otherwise off limits to most programs. This allows some servers, particularly device drivers, to interact directly with hardware.

A basic set of servers for a general-purpose microkernel includes file system servers, device driver servers, networking servers, display servers, and user interface device servers. This set of servers (drawn from QNX) provides roughly the set of services offered by a Unix monolithic kernel. The necessary servers are started at system startup and provide services, such as file, network, and device access, to ordinary application programs. With such servers running in the environment of a user application, server development is similar to ordinary application development, rather than the build-and-boot process needed for kernel development.

Additionally, many "crashes" can be corrected by simply stopping and restarting the server. However, part of the system state is lost with the failing server, hence this approach requires applications to cope with failure. A good example is a server responsible for TCP/IP connections: If this server is restarted, applications will experience a "lost" connection, a normal occurrence in a networked system. For other services, failure is less expected and may require changes to application code. For QNX, restart capability is offered as the QNX High Availability Toolkit.[19]

# Device drivers

Device drivers frequently perform direct memory access (DMA), and therefore can write to arbitrary locations of physical memory, including various kernel data structures. Such drivers must therefore be trusted. It is a common misconception that this means that they must be part of the kernel. In fact, a driver is not inherently more or less trustworthy by being part of the kernel.

While running a device driver in user space does not necessarily reduce the damage a misbehaving driver can cause, in practice it is beneficial for system stability in the presence of buggy (rather than malicious) drivers: memory-access violations by the driver code itself (as opposed to the device) may still be caught by the memory-management hardware. Furthermore, many devices are not DMA-capable, their drivers can be made untrusted by running them in user space. Recently, an increasing number of computers feature IOMMUs, many of which can be used to restrict a device's access to physical memory.[20] This also allows user-mode drivers to become untrusted.

User-mode drivers actually predate microkernels. The Michigan Terminal System (MTS), in 1967, supported user space drivers (including its file system support), the first operating system to be designed with that capability.[21] Historically, drivers were less of a problem, as the number of devices was small and trusted anyway, so having them in the kernel simplified the design and avoided potential performance problems. This led to the traditional driver-in-the-kernel style of Unix,[22] Linux, and Windows NT. With the proliferation of various kinds of peripherals, the amount of driver code escalated and in modern operating systems dominates the kernel in code size.

# Essential components and minimality

As a microkernel must allow building arbitrary operating system services on top, it must provide some core functionality. At a minimum, this includes:

- Some mechanisms for dealing with address spaces, required for managing memory protection
- Some execution abstraction to manage CPU allocation, typically threads or scheduler activations
- Inter-process communication, required to invoke servers running in their own address spaces

This minimal design was pioneered by Brinch Hansen's Nucleus and the hypervisor of IBM's VM. It has since been formalised in Liedtke's *minimality principle*:

> A concept is tolerated inside the microkernel only if moving it outside the kernel, i.e., permitting competing implementations, would prevent the implementation of the system's required functionality.[16]

Everything else can be done in a usermode program, although device drivers implemented as user programs may on some processor architectures require special privileges to access I/O hardware.

Related to the minimality principle, and equally important for microkernel design, is the separation of mechanism and policy, it is what enables the construction of arbitrary systems on top of a minimal kernel. Any policy built into the kernel cannot be overwritten at user level and therefore limits the generality of the microkernel.[12] Policy implemented in user-level servers can be changed by replacing the servers (or letting the application choose between competing servers offering similar services).

For efficiency, most microkernels contain schedulers and manage timers, in violation of the minimality principle and the principle of policy-mechanism separation.

Start up (booting) of a microkernel-based system requires device drivers, which are not part of the kernel. Typically this means that they are packaged with the kernel in the boot image, and the kernel supports a bootstrap protocol that defines how the drivers are located and started; this is the traditional bootstrap procedure of L4 microkernels. Some microkernels simplify this by placing some key drivers inside the kernel (in violation of the minimality principle), LynxOS and the original Minix are examples. Some even include a file system in the kernel to simplify booting. A microkernel-based system may boot via multiboot compatible boot loader. Such systems usually load statically-linked servers to make an initial bootstrap or mount an OS image to continue bootstrapping.

A key component of a microkernel is a good IPC system and virtual-memory-manager design that allows implementing page-fault handling and swapping in usermode servers in a safe way. Since all services are performed by usermode programs, efficient means of communication between programs are essential, far more so than in monolithic kernels. The design of the IPC system makes or breaks a microkernel. To be effective, the IPC system must not only have low overhead, but also interact well with CPU scheduling.

# Performance

On most mainstream processors, obtaining a service is inherently more expensive in a microkernel-based system than a monolithic system.[12] In the monolithic system, the service is obtained by a single system call, which requires two *mode switches* (changes of the processor's ring or CPU mode). In the microkernel-based system, the service is obtained by sending an IPC message to a server, and obtaining the result in another IPC message from the server. This requires a context switch if the drivers are implemented as processes, or a function call if they are implemented as procedures. In addition, passing actual data to the server and back may incur extra copying overhead, while in a monolithic system the kernel can directly access the data in the client's buffers.

Performance is therefore a potential issue in microkernel systems. The experience of first-generation microkernels such as Mach and ChorusOS showed that systems based on them performed very poorly.[15] However, Jochen Liedtke showed that Mach's performance problems were the result of poor design and implementation, specifically Mach's excessive cache footprint.[16] Liedtke demonstrated with his own L4

microkernel that through careful design and implementation, and especially by following the minimality principle, IPC costs could be reduced by more than an order of magnitude compared to Mach. L4's IPC performance is still unbeaten across a range of architectures.[23][24][25]

While these results demonstrate that the poor performance of systems based on first-generation microkernels is not representative for second-generation kernels such as L4, this constitutes no proof that microkernel-based systems can be built with good performance. It has been shown that a monolithic Linux server ported to L4 exhibits only a few percent overhead over native Linux.[26] However, such a single-server system exhibits few, if any, of the advantages microkernels are supposed to provide by structuring operating system functionality into separate servers.

A number of commercial multi-server systems exist, in particular the real-time systems QNX and Integrity. No comprehensive comparison of performance relative to monolithic systems has been published for those multiserver systems. Furthermore, performance does not seem to be the overriding concern for those commercial systems, which instead emphasize reliably quick interrupt handling response times (QNX) and simplicity for the sake of robustness. An attempt to build a high-performance multiserver operating system was the IBM Sawmill Linux project.[27] However, this project was never completed.

It has been shown in the meantime that user-level device drivers can come close to the performance of in-kernel drivers even for such high-throughput, high-interrupt devices as Gigabit Ethernet.[28] This seems to imply that high-performance multi-server systems are possible.

## Security

The security benefits of microkernels have been frequently discussed.[29][30] In the context of security the minimality principle of microkernels is, some have argued, a direct consequence of the principle of least privilege, according to which all code should have only the privileges needed to provide required functionality. Minimality requires that a system's trusted computing base (TCB) should be kept minimal. As the kernel (the code that executes in the privileged mode of the hardware) has unvetted access to any data and can thus violate its integrity or confidentiality, the kernel is always part of the TCB. Minimizing it is natural in a security-driven design.

Consequently, microkernel designs have been used for systems designed for high-security applications, including KeyKOS, EROS and military systems. In fact common criteria (CC) at the highest assurance level (Evaluation Assurance Level (EAL) 7) has an explicit requirement that the target of evaluation be "simple", an acknowledgment of the practical impossibility of establishing true trustworthiness for a complex system. Again, the term "simple" is misleading and ill-defined. At least the Department of Defense Trusted Computer System Evaluation Criteria introduced somewhat more precise verbiage at the B3/A1 classes:

> "The TCB shall [implement] complete, conceptually simple protection mechanisms with precisely defined semantics. Significant system engineering shall be directed toward minimizing the complexity of the TCB, as well as excluding from the TCB those modules that are not protection-critical."
>
> — Department of Defense Trusted Computer System Evaluation Criteria

In 2018, a paper presented at the Asia-Pacific Systems Conference claimed that microkernels were demonstrably safer than monolithic kernels by investigating all published critical CVEs for the Linux kernel at the time. The study concluded that 40% of the issues could not occur at all in a formally verified microkernel, and only 4% of the issues would remain entirely unmitigated in such a system.[31]

# Third generation

More recent work on microkernels has been focusing on formal specifications of the kernel API, and formal proofs of the API's security properties and implementation correctness. The first example of this is a mathematical proof of the confinement mechanisms in EROS, based on a simplified model of the EROS API.[32] More recently (in 2007) a comprehensive set of machine-checked proofs was performed of the properties of the protection model of seL4, a version of L4.[33]

This has led to what is referred to as *third-generation microkernels*,[34] characterised by a security-oriented API with resource access controlled by capabilities, virtualization as a first-class concern, novel approaches to kernel resource management,[35] and a design goal of suitability for formal analysis, besides the usual goal of high performance. Examples are Coyotos, seL4, Nova,[36][37] Redox and Fiasco.OC.[36][38]

In the case of seL4, complete formal verification of the implementation has been achieved,[34] i.e. a mathematical proof that the kernel's implementation is consistent with its formal specification. This provides a guarantee that the properties proved about the API actually hold for the real kernel, a degree of assurance which goes beyond even CC EAL7. It was followed by proofs of security-enforcement properties of the API, and a proof demonstrating that the executable binary code is a correct translation of the C implementation, taking the compiler out of the TCB. Taken together, these proofs establish an end-to-end proof of security properties of the kernel.[39]

# Examples

Some examples of microkernels are:

- The L4 microkernel family
- Redox
- Zircon

# Nanokernel

The term *nanokernel* or *picokernel* historically referred to:

- A kernel where the total amount of kernel code, i.e. code executing in the privileged mode of the hardware, is very small. The term *picokernel* was sometimes used to further emphasize small size. The term *nanokernel* was coined by Jonathan S. Shapiro in the paper *The KeyKOS NanoKernel Architecture* (https://web.archive.org/web/20110621235229/http://www.cis.upenn.edu/~KeyKOS/NanoKernel/NanoKernel.html). It was a sardonic response to Mach, which claimed to be a microkernel while Shapiro considered it monolithic, essentially unstructured, and slower than the systems it sought to replace. Subsequent reuse of and response to the term, including the picokernel coinage, suggest that the point was largely missed. Both *nanokernel* and *picokernel* have subsequently come to have the same meaning expressed by the term microkernel.
- A virtualization layer underneath an operating system, which is more correctly referred to as a hypervisor.
- A hardware abstraction layer that forms the lowest-level part of a kernel, sometimes used to provide real-time functionality to normal operating systems, like Adeos.

There is also at least one case where the term nanokernel is used to refer not to a small kernel, but one that supports a nanosecond clock resolution.[40]

## See also

- Kernel (computer science)
  - Exokernel
  - Hybrid kernel
  - Loadable kernel module
  - Monolithic kernel
- Microservices
- Tanenbaum–Torvalds debate
- Trusted computing base
- Unikernel
- Multi-Environment Real-Time

## References


1. Herder, Jorrit N. (23 February 2005). "Toward a True Microkernel Operating System" (http://www.minix3.org/theses/herder-true-microkernel.pdf) (PDF). *minix3.org*. Retrieved 22 June 2015.
2. "read-more" (http://wiki.minix3.org/doku.php?id=www:documentation:read-more). Retrieved 20 December 2016.
3. "2002 Computer Pioneer Award Recipient" (https://www.computer.org/web/awards/pioneer-per-hansen). IEEE Computer Society. Retrieved 13 September 2016.
4. Brinch Hansen, Per (2004). *A Programmer's Story: The Life of a Computer Pioneer* (http://brinch-hansen.net/memoirs/contents.html). Retrieved 13 September 2016.
5. Brinch Hansen, Per (April 1969). *RC 4000 Software: Multiprogramming System* (http://brinch-hansen.net/papers/1969a.pdf) (PDF) (Technical report). Regnecentralen. Retrieved 13 September 2016.
6. Brinch Hansen, Per (1970). "The Nucleus of a Multiprogramming Operating System" (http://www.brinch-hansen.net/papers/1970a.pdf) (PDF). *Communications of the ACM*. **13** (4): 238–250. CiteSeerX 10.1.1.105.4204 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.105.4204). doi:10.1145/362258.362278 (https://doi.org/10.1145%2F362258.362278). S2CID 9414037 (https://api.semanticscholar.org/CorpusID:9414037).
7. .Wulf, William; Cohen, Ellis; Corwin, William; Jones, Anita; Levin, Roy; Pierson, C.; Pollack, Fred (June 1974). "HYDRA: The Kernel of a Multiprocessor Operating System". *Communications of the ACM*. **17** (6): 337–345. doi:10.1145/355616.364017 (https://doi.org/10.1145%2F355616.364017). S2CID 8011765 (https://api.semanticscholar.org/CorpusID:8011765).
8. Rashid, Richard; Robertson, George (December 1981). *Accent: A communication oriented network operating system kernel*. SOSP '81 Proceedings of the eighth ACM symposium on Operating systems principles. Pacific Grove, California, USA. pp. 64–75. doi:10.1145/800216.806593 (https://doi.org/10.1145%2F800216.806593).
9. Sassenrath, Carl (1986). *Amiga ROM Kernel Reference Manual*. Exec.
10. Jim Magee. *WWDC 2000 Session 106 - Mac OS X: Kernel* (https://www.youtube.com/watch?v=ggnFoDqzGMU). 14 minutes in.
11. "Porting UNIX/Linux Applications to Mac OS X" (https://developer.apple.com/library/mac/#documentation/Porting/Conceptual/PortingUnix/glossary/glossary.html#//apple_ref/doc/uid/TP40002859-TPXREF101). Apple. Retrieved 26 April 2011.
12. Liedtke, Jochen (September 1996). "Towards Real Microkernels". *Communications of the ACM*. **39** (9): 70–77. doi:10.1145/234215.234473 (https://doi.org/10.1145%2F234215.234473). S2CID 2867357 (https://api.semanticscholar.org/CorpusID:2867357).

13. Heiser, Gernot; Uhlig, Volkmar; LeVasseur, Joshua (January 2006). "Are Virtual-Machine Monitors Microkernels Done Right?" (http://os.ibds.kit.edu/65_747.php). *ACM SIGOPS Operating Systems Review*. ACM. **40** (1): 95–99. doi:10.1145/1113361.1113363 (https://doi.org/10.1145%2F1113361.1113363). S2CID 7414062 (https://api.semanticscholar.org/CorpusID:7414062).
14. Liedtke, Jochen (December 1993). *Improving IPC by kernel design*. 14th ACM Symposium on Operating System Principles. Asheville, NC, USA. pp. 175–88. CiteSeerX 10.1.1.40.1293 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.40.1293).
15. Chen, J. Bradley; Bershad, Brian N. (December 1993). *The Impact of Operating System Structure on Memory System Performance* (https://people.eecs.berkeley.edu/~prabal/resources/osprelim/CB93.pdf) (PDF). SOSP '93 Proceedings of the fourteenth ACM symposium on Operating systems principles. Asheville, NC, USA. pp. 120–133. doi:10.1145/168619.168629 (https://doi.org/10.1145%2F168619.168629).
16. Liedtke, Jochen (December 1995). *On µ-Kernel Construction*. SOSP '95 Proceedings of the fifteenth ACM symposium on Operating systems principles. Copper Mountain Resort, CO, USA. pp. 237–250. doi:10.1145/224056.224075 (https://doi.org/10.1145%2F224056.224075).
17. Elphinstone, Kevin; Heiser, Gernot (November 2013). *From L3 to seL4: What Have We Learnt in 20 Years of L4 Microkernels?*. SOSP '13 Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles. Farmington, PA, USA. pp. 133–150. doi:10.1145/2517349.2522720 (https://doi.org/10.1145%2F2517349.2522720).
18. "Synchronous Message Passing" (http://www.qnx.com/developers/docs/7.0.0/#com.qnx.doc.neutrino.sys_arch/topic/ipc_Sync_messaging.html). Retrieved 14 July 2019.
19. "The QNX High Availability Toolkit" (https://web.archive.org/web/20050824015422/http://www.qnx.com/download/download/8107/QNX_High_Availability_Toolkit.pdf) (PDF). Archived from the original (http://www.qnx.com/download/download/8107/QNX_High_Availability_Toolkit.pdf) (PDF) on 24 August 2005.
20. Wong, William (27 April 2007). "I/O, I/O, It's Off to Virtual Work We Go" (http://www.electronicdesign.com/embedded/io-io-its-virtual-work-we-go). *Electronic Design*. Retrieved 8 June 2009.
21. Alexander, Michael T. (1971). "Organization and Features of the Michigan Terminal System". *Proceedings of the November 16–18, 1971, Fall Joint Computer Conference*. **40**: 589–591. doi:10.1145/1478873.1478951 (https://doi.org/10.1145%2F1478873.1478951). S2CID 14614148 (https://api.semanticscholar.org/CorpusID:14614148).
22. Lions, John (1 August 1977). *Lions' Commentary on UNIX 6th Edition, with Source Code*. Peer-To-Peer Communications. ISBN 978-1-57398-013-5.
23. Liedtke, Jochen; Elphinstone, Kevin; Schönberg, Sebastian; Härtig, Hermann; Heiser, Gernot; Islam, Nayeem; Jaeger, Trent (May 1997). *Achieved IPC performance (still the foundation for extensibility)* (http://ieeexplore.ieee.org/xpl/RecentCon.jsp?punumber=4643). 6th Workshop on Hot Topics in Operating Systems. Cape Cod, MA, USA: IEEE. pp. 28–31.
24. Gray, Charles; Chapman, Matthew; Chubb, Peter; Mosberger-Tang, David; Heiser, Gernot (April 2005). *Itanium—a system implementor's tale* (http://www.usenix.org/publications/library/proceedings/usenix05/tech/general/gray.html). USENIX Annual Technical Conference. Annaheim, CA, USA. pp. 264–278.
25. van Schaik, Carl; Heiser, Gernot (January 2007). *High-performance microkernels and virtualisation on ARM and segmented architectures* (https://web.archive.org/web/20070426092901/http://www.ertos.nicta.com.au/publications/). 1st International Workshop on Microkernels for Embedded Systems. Sydney, Australia: NICTA. pp. 11–21. Archived from the original (http://ertos.nicta.com.au/publications) on 26 April 2007. Retrieved 1 April 2007.

26. Härtig, Hermann; Hohmuth, Michael; Liedtke, Jochen; Schönberg, Sebastian (October 1997). "The performance of µ-kernel-based systems" (http://portal.acm.org/citation.cfm?id=266660&dl=ACM&coll=&CFID=15151515&CFTOKEN=6184618). *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*: 66–77. doi:10.1145/268998.266660 (https://doi.org/10.1145%2F268998.266660). ISBN 0-89791-916-5. S2CID 1706253 (https://api.semanticscholar.org/CorpusID:1706253).
27. Gefflaut, Alain; Jaeger, Trent; Park, Yoonho; Liedtke, Jochen; Elphinstone, Kevin J.; Uhlig, Volkmar; Tidswell, Jonathon E.; Deller, Luke; et al. (2000). *The Sawmill multiserver approach*. 9th ACM SIGOPS European Workshop. Kolding, Denmark. pp. 109–114. CiteSeerX 10.1.1.25.8376 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.25.8376).
28. Leslie, Ben; Chubb, Peter; FitzRoy-Dale, Nicholas; Götz, Stefan; Gray, Charles; Macpherson, Luke; Potts, Daniel; Shen, Yueting; Elphinstone, Kevin; Heiser, Gernot (September 2005). "User-level device drivers: achieved performance". *Journal of Computer Science and Technology*. **20** (5): 654–664. doi:10.1007/s11390-005-0654-4 (https://doi.org/10.1007%2Fs11390-005-0654-4). S2CID 1121537 (https://api.semanticscholar.org/CorpusID:1121537).
29. Tanenbaum, Andrew S. "Tanenbaum-Torvalds debate, part II" (http://www.cs.vu.nl/~ast/reliable-os/).
30. Tanenbaum, A., Herder, J. and Bos, H. (May 2006).
31. Biggs, Simon; Lee, Damon; Heiser, Gernot (2018). "The Jury Is In: Monolithic OS Design Is Flawed: Microkernel-based Designs Improve Security" (https://dl.acm.org/doi/abs/10.1145/3265723.3265733). *Proceedings of the 9th Asia-Pacific Workshop on Systems*. Jeju Island, Republic of Korea: Association for Computing Machinery. pp. 1–7. doi:10.1145/3265723.3265733 (https://doi.org/10.1145%2F3265723.3265733).
32. Shapiro, Jonathan S.; Weber, Samuel. *Verifying the EROS Confinement Mechanism* (https://web.archive.org/web/20160303174121/http://www.eros-os.org/papers/oakland2000.ps). IEEE Conference on Security and Privacy. Archived from the original (http://www.eros-os.org/papers/oakland2000.ps) on 3 March 2016.
33. Elkaduwe, Dhammika; Klein, Gerwin; Elphinstone, Kevin (2007). *Verified Protection Model of the seL4 Microkernel* (http://ertos.org/publications/papers/Elkaduwe_GE_07.abstract). submitted for publication.
34. Klein, Gerwin; Elphinstone, Kevin; Heiser, Gernot; Andronick, June; Cock, David; Derrin, Philip; Elkaduwe, Dhammika; Engelhardt, Kai; Kolanski, Rafal; Norrish, Michael; Sewell, Thomas; Tuch, Harvey; Winwood, Simon (October 2009). *seL4: Formal verification of an OS kernel* (http://www.sigops.org/sosp/sosp09/papers/klein-sosp09.pdf) (PDF). 22nd ACM Symposium on Operating System Principles. Big Sky, MT, USA.
35. Elkaduwe, Dhammika; Derrin, Philip; Elphinstone, Kevin (April 2008). *Kernel design for isolation and assurance of physical memory* (https://web.archive.org/web/20100424035229/http://www.ertos.nicta.com.au/publications/papers/Elkaduwe_DE_08.abstract). 1st Workshop on Isolation and Integration in Embedded Systems. Glasgow, UK. doi:10.1145/1435458 (https://doi.org/10.1145%2F1435458). Archived from the original (http://ertos.nicta.com.au/publications/papers/Elkaduwe_DE_08.abstract) on 24 April 2010. Retrieved 17 August 2009.
36. "TUD Home: Operating Systems: Research: Microkernel & Hypervisor" (http://www.inf.tu-dresden.de/index.php?node_id=2697). *Faculty of Computer Science*. Technische Universität Dresden. 12 August 2010. Retrieved 5 November 2011.
37. Steinberg, Udo; Kauer, Bernhard (April 2010). *NOVA: A Microhypervisor-Based Secure Virtualization Architecture*. Eurosys 2010. Paris, France. pp. 209–222. doi:10.1145/1755913.1755935 (https://doi.org/10.1145%2F1755913.1755935).
38. Lackorzynski, Adam; Warg, Alexander (March 2009). *Taming Subsystems – Capabilities as Universal Resource Access Control in L4* (http://portal.acm.org/citation.cfm?id=1519135&dl=ACM). IIES'09: Second Workshop on Isolation and Integration in Embedded Systems. Nuremberg, Germany. CiteSeerX 10.1.1.629.9845 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.629.9845).

39. Klein, Gerwin; Andronick, June; Elphinstone, Kevin; Murray, Toby; Sewell, Thomas; Kolanski, Rafal; Heiser, Gernot (February 2014). "Comprehensive Formal Verification of an OS Microkernel". *ACM Transactions on Computer Systems*. **32** (1): 2:1–2:70. doi:10.1145/2560537 (https://doi.org/10.1145%2F2560537). S2CID 4474342 (https://api.semanticscholar.org/CorpusID:4474342).
40. David L. Mills and Poul-Henning Kamp (28 November 2000). "The Nanokernel" (http://www.eecis.udel.edu/~mills/database/papers/nano/nano2.pdf) (PDF). Retrieved 28 August 2017.

# Further reading

- Scientific articles about microkernels (http://citeseerx.ist.psu.edu/search?q=microkernel) (on CiteSeerX), including:
  - Dan Hildebrand (1992). "An Architectural Overview of QNX". *Proceedings of the Workshop on Micro-kernels and Other Kernel Architectures*: 113–126. CiteSeerX 10.1.1.459.4481 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.459.4481). ISBN 1-880446-42-1. – the basic QNX reference.
  - Tanenbaum, A., Herder, J. and Bos, H. (May 2006). "Can We Make Operating Systems Reliable and Secure?" (https://web.archive.org/web/20170621194406/https://www.computer.org/portal/site/computer/menuitem.eb7d70008ce52e4b0ef1bd108bcd45f3/index.jsp?&pName=computer_level1&path=computer%2Fhomepage%2F0506&file=cover1.xml&xsl=article.xsl). *Computer*. **39** (5): 44–51. doi:10.1109/MC.2006.156 (https://doi.org/10.1109%2FMC.2006.156). S2CID 99779 (https://api.semanticscholar.org/CorpusID:99779). Archived from the original (http://www.computer.org/portal/site/computer/menuitem.eb7d70008ce52e4b0ef1bd108bcd45f3/index.jsp?&pName=computer_level1&path=computer/homepage/0506&file=cover1.xml&xsl=article.xsl) on 21 June 2017. Retrieved 3 April 2020. -the basic reliable reference.
  - Black, D.L., Golub, D.B., Julin, D.P., Rashid, R.F., Draves, R.P., Dean, R.W., Forin, A., Barrera, J., Tokuda, H., Malan, G., and Bohman, D. (March 1992). "Microkernel Operating System Architecture and Mach". *Journal of Information Processing*. **14** (4). – the basic Mach reference.
- *Varhol, Peter D. (January 1994). "Small Kernels Hit It Big" (https://web.archive.org/web/20060307075538/http://www.byte.com/art/9401/sec8/art2.htm). *Byte*. Archived from the original (http://www.byte.com:80/art/9401/sec8/art2.htm) on 7 March 2006. Retrieved 20 September 2017. An assessment of the present and future state of microkernel based OSes as of January 1994
- MicroKernel page (http://c2.com/cgi/wiki?MicroKernel) from the Portland Pattern Repository
- The Tanenbaum–Torvalds debate
  - The Tanenbaum-Torvalds Debate, 1992.01.29 (http://www.oreilly.com/catalog/opensources/book/appa.html)
  - Tanenbaum, A. S. "Can We Make Operating Systems Reliable and Secure? (http://www.computer.org/csdl/mags/co/2006/05/r5044-abs.html)".
  - Torvalds, L. Linus Torvalds about the microkernels again, 2006.05.09 (http://www.realworldtech.com/forums/index.cfm?action=detail&id=66630&threadid=66595&roomid=11)
  - Shapiro, J. "Debunking Linus's Latest (https://web.archive.org/web/20160922022726/http://www.coyotos.org/docs/misc/linus-rebuttal.html)".
  - Tanenbaum, A. S. "Tanenbaum-Torvalds Debate: Part II (http://www.cs.vu.nl/~ast/reliable-os/)".

---