



RustyHermit: A Scalable, Rust-Based Virtual Execution Environment

Stefan Lankes¹(✉), Jonathan Klimt¹, Jens Breitbart², and Simon Pickartz³

¹ Institute for Automation of Complex Power Systems, RWTH Aachen University, Aachen, Germany

{slankes, jonathan.klimt}@eonerc.rwth-aachen.de

² Bosch Chassis Systems Control, Robert Bosch GmbH, Gerlingen, Germany
jens.breitbart@de.bosch.com

³ ParTec Cluster Competence Center GmbH, Munich, Germany
pickartz@par-tec.com

Abstract. System-level development has been dominated by programming languages such as C/C++ for decades. These languages are inherently unsafe, error-prone, and a major reason for vulnerabilities. High-level programming languages with a secure memory model and strong type system are able to improve the quality of the system software. This paper explores the programming language Rust for development of a scalable, virtual execution environment and presents the integration of a Rust-based IP stack into RUSTYHERMIT. RustyHermit is part of the standard Rust toolchain and common Rust applications are able to build on top of RUSTYHERMIT.

1 Introduction

The C programming language is still dominating system-level software development as it was designed for this exact case and is known to provide high performance. However, C is also known to be error-prone and difficult to use in large scale projects as even senior developers can hardly avoid an incorrect usage of C. Dangling pointers and missing boundary checks are other typical reasons for issues within kernel code. This is not a new observation. As described in by Cutler et al. [3], the Pilot kernel [29] and the Lisp machine [8] are early examples of the usage of a high-level language (Mesa and Lisp, respectively) for Operating System (OS) development. However, the approach has not gained acceptance and is hardly used because memory safety of high-level languages often induces runtime overhead (e.g., due to garbage collection).

Furthermore, the OS requirements changed fundamentally over the last years. The basic infrastructure within OSs was established in the seventies when hardware was expensive and resource sharing was the focus. The virtualization of hardware resources has been established for a simplified resource sharing, e.g., sharing a processor in round-robin manner. However, in the era of cloud computing, complete machines are virtualized supporting server consolidations. Virtualization is implemented as another software abstraction layer in an already

highly layered software stack. Typical modern OSs still include support for old physical protocols (e.g., floppy disks), irrelevant optimizations (e.g., disk elevator algorithms on SSDs) and backward-compatible interfaces (e.g., POSIX). Anil Madhavapeddy et al. discuss these issues in [20,21] and present unikernels, i.e., specialized library OSs, as a solution. Unikernels are built by compiling high-level languages directly into specialized single-address-space machine images. In doing so, unused code is removed by static code analysis and system calls are replaced by common function calls promising a faster resource handling. Unikernels are able to run directly on a hypervisor or bare metal on the hardware. They provide a smaller footprint compared to traditional OS kernels and have more prospect to optimize the applications, e.g., the application and the kernel can be optimized by means of Link-time Optimization (LTO).

Current Unikernels relinquish backward compatibility, often rely on uncommon programming interfaces, and barely support multi-processor systems. In [16], we present a rewrite of HERMITCORE [14] in Rust called RUSTYHERMIT and demonstrate that the performance of the Rust implementation is on a par with the original C implementation. RUSTYHERMIT is integrated into the standard runtime of Rust and its compiler infrastructure. It is trivial to port pure Rust application to RUSTYHERMIT, as it just requires a configuration change. Only applications, which bypass the Rust runtime and call directly a C library, have to port also the C library to the new system. Furthermore, existing C/C++ and Fortran applications can be linked with RUSTYHERMIT and generate a bootable image. In this paper, we focus on the integration of a Rust-based IP stack enabling the building and deployment of secure and efficient cloud applications.

The rest of this paper is structured as follows: We start with a discussion of the related work in the area of unikernels and the usage of high-level programming languages for kernel development. In Sect. 3, we give a short introduction to Rust, followed by the Sect. 4 on kernel development using Rust and the integration of the IP stack. In the Sect. 5 we compare the performance of our kernel with Linux. Finally, Sect. 6 summarizes the paper and give a short outlook.

2 Related Work

High-level programming languages provide type-safety and memory-safety as well as convenient abstractions of concurrent programming reducing the susceptibility to errors. However, kernel developers are often skeptical to use new languages because they expect them to introduce additional overhead compared to C [37] and require a redevelopment of kernel components. Yet, many research projects use high-level programming languages to benefit from new features such as a safe memory handling. New system programming languages, e.g., D [4], Nim [27], Go [7], and Rust [35], have emerged in the last decade. For nearly every language there exists an OS project.

In Rust, the compiler is able to determine when memory has to be freed avoiding the need for according runtime checks. This results in far less runtime

overhead compared to other high-level programming languages, but introduces unique memory handling at the language level. Levy et al. [17, 18] show that Rust is attractive for kernel development because it promises memory-safety while providing good performance. In addition, Balasubramanian et al. [1] show that Rust offers software fault isolation (SFI) with lower overhead and Narayanan et al. in [24] steps to realize a Rust-based verified firmware. Currently, Microsoft [22] is also analyzing Rust as a system programming language. Projects such as Redox [31], Tock [36] or teaching kernels such as our eduOS-rs [6] show that Rust is usable for OS development, but all these Rust kernels were not designed for cloud environments.

Both HERMITCORE and RUSTYHERMIT belong to the class of unikernels or library OSs. *MirageOS* [20], *IncludeOS* [2], *rumprun kernels* [9], and *OSv* [10] are typical representatives. The fundamental drawback of unikernels is the porting effort that is required to adapt existing applications to the underlying minimalistic OS. This often requires both expert work and a considerable amount of time. One objective of the Unikraft [38] project is to build unikernels targeted at specific applications, without requiring the time-consuming, expert work. Unikraft is written in C, uses newlib [30] as the C library, and LwIP [5] as the network stack.

Like Unikraft, HERMITCORE relies on *LwIP* as it is easy to combine to a kernel and the list of requirements is small. However, *LwIP* was mainly designed for embedded systems and it is a challenge to get the same performance as provided by common operating systems (e.g., Linux). For instance, Kuenzer et al. [12] can improve the performance by using a low-level API (instead of a socket interface) and to integrate checksum offloading. Further improvements can be achieved by supporting the Data Plane Development Kit (DPDK) [28]. However, this is not available for all devices.

The compatibility of unikernels to common OSs (e.g., Linux) is currently still limited. HermitTux [26] has similar objectives and realizes compatibility to Linux by rewriting system calls and using a modified C library. However, the compatibility of HermitTux is limited as not all Linux system calls have been re-implemented. RUSTYHERMIT is also not compatible to common OSs, but it offers the possibility to write portable Rust applications. Changes to the source code are not required to run the application on Linux or other OSs.

3 Introduction to Rust

Rust is a new programming language originally designed by Graydon Hoare as a replacement for C/C++. Its goal is to provide the same level of performance, but to allow for more comprehensive safety checks at compile time and by default enabled runtime checks when the compile time checks are not sufficient (e.g., array access with indices not known at compile time). We discuss only the features relevant to understand this paper, a detailed overview on Rust can be found in [11].

Rust relies on *ownership* to provide safe memory handling without runtime overhead. Each resource (e.g., memory) in Rust has a variable that is called its

owner. There is exactly one owner at a time and whenever this owner goes out of scope, the resource will be dropped and the memory freed. Ownership can be forwarded to another variable invalidating the original owner, or the owner can borrow the resource to another variable. Read only access can be provided to multiple variables at a time via immutable borrows, as long as no mutable borrow is happening at the same time. In general, these rules prevent data races, the dangling pointer problem, and pointer aliasing for mutable access. For most tasks it is possible to develop code that these rules are satisfied at compile time, however it is also possible to use `std::cell::RefCell` to bypass compile time checks, but enforce runtime checks.

Similarly to these checks, Rust provides compile time checks as well ensuring the correct execution of concurrent or parallel code. Data that is shared between threads must implement the so-called `sync` trait (the rust term for an interface) or must be wrapped into a mutex providing this trait. This rule prevents data races, as long as the synchronization mechanism (e.g., the mutex) is implemented correctly. Furthermore, the Rust compiler checks the lifetime of values shared by threads and will not compile code in which a value is not guaranteed to outlive the threads borrowing a value.

All checks named before can be circumvented by using the `unsafe` keyword. Unsafe Rust code provides the same level of control as C, e.g., it provides *raw pointers* enabling direct, unchecked memory accesses and even supports the usage of inline assembly. Code in unsafe regions should be reviewed more carefully than code that is checked by the compiler and as a result it is typically frowned upon by the Rust community. Currently, it is not possible to write a kernel without unsafe code. For instance, inline assembly is important to restore the context of the FPU. However, the RUSTYHERMIT only requires 1170 lines of unsafe code corresponding to only 1.71% of total code size.

The Rust standard library is divided into an OS-independent and an OS-dependent part. The library known as *core* library is the major part of the OS-independent library and already implements basic error/panic handling, string operations, and atomic operations. Furthermore, Rust offers the possibility to redefine the global memory allocator. This allocator is used by all other Rust codes unless explicitly circumvented. In contrast, the part known as *std* condenses the OS-dependent libraries and extends them with various data structures, console output, and thread handling. It is easily possible to create a project that does not use *std* by adding `#![no_std]` to the main file.

4 A Unikernel Written in Rust

RUSTYHERMIT is a rewrite of our 64 bit unikernel HERMITCORE [14, 15] which was written in C. RUSTYHERMIT is completely written in Rust, supports the Intel 64 Architecture and comes with support for SSE4, AVX2, and AVX512. It has multi-core and single-core multiprocessing support by the means of multithreading and multiprocessing. The Kernel supports the execution of more threads than available cores. This is an important feature for dealing with concurrent applications or to integrate performance monitoring tools. Currently, the

scheduler does not support load balancing as explicit thread placement is favored over automatic strategies. Scheduling overhead is reduced to a minimum by the employment of a dynamic timer, i.e., the kernel does not interrupt computation threads which run exclusively on certain cores and do not use any timer. To improve the security behavior, RUSTYHERMIT provides a stack guard and is completely position-independent. Consequently, the loader is able to randomize the memory layout.

4.1 Integration of RustyHermit into libstd

One major goal of RUSTYHERMIT was a complete integration into the Rust toolchain to simplify the application development. Any common Rust application should be buildable with RUSTYHERMIT. To achieve this goal, the kernel provides the required interfaces to the Standard Library (*libstd*) whilst being based only on the *core* library. The operating system abstraction layer of the Rust toolchain is relative small, so only around 26 files within a total of ~3000 lines of code are required to integrate RUSTYHERMIT into the standard library of Rust.

Most operating systems are written in C and use a common C library as interface to the kernel. These functions are typically provided by a helper crate¹ in Rust realizing an interface to the C functions. For instance, the C interface for Rust is published in the crate *libc*², however C functions are by definition unsafe.

In case of RUSTYHERMIT, the complete kernel is written in Rust and theoretically, it could be directly integrated into the Rust standard library. However, the kernel uses a set of external crates to detect processor features, programming of the interrupt controller, or log messages. As the Rust community wants to reduce the dependencies of the basic runtime libraries to external crates, we cannot integrate RUSTYHERMIT into *libstd* directly. Instead, we create two helper crates *hermit-abi*³ and *hermit-sys*⁴. The former describes only the interface to the library operating system for linkage and is included in *libstd*'s dependencies, just like the *libc* crate does for the Linux interface of the Standard Library. The latter is a helper crate, with the main purpose of building the kernel as static library from source and linking it to the application.

Separating the kernel and *libstd* into separate compilation units also allows the use of different compiler settings for each of them. Hereby, we are able to compile the kernel without FPU and AVX/SSE support and to enable it for the rest of the application. This is necessary because AVX and SSE is not longer limited to floating-point operations and the compiler would use these instructions to optimize the kernel code. The usage of AVX and SSE within the kernel would trigger interrupts to save the FPU context.

¹ Crate is a tree of modules that produces a library or executables. Much like a package in other programming languages.

² <https://crates.io/crates/libc>.

³ <https://crates.io/crates/hermit-abi>.

⁴ <https://crates.io/crates/hermit-sys>.

Listing 1.1. Extension of Cargo.toml to integrate RUSTYHERMIT

```
[target.'cfg(target_os = "hermit")'.dependencies]
hermit-sys = "0.1.*"
```

A Rust based RUSTYHERMIT application can be build by adding the *hermit-sys* crate to the application dependencies as shown in Listing 1.1 and declare it as an **external crate** in the applications source. Rust’s package manager Cargo [34] will then download the kernel’s sources, compile it, and link it to the application.

4.2 Network Support

The library operating system only provides basic features such as interrupt handling, device drivers, memory management, and scheduling. One possible solution to integrate network support, is the use of real hardware drivers. The hypervisor emulates these devices by trapping every request to the device and emulating the behavior of the real hardware (*trap and emulate*). This approach comes with an important overhead.

Another solution is to use *para-virtualization* where the hypervisor provides a simpler and faster interface for the I/O devices to the guest, who is aware of running on a hypervisor. Today, *virtio* is the standard abstraction layer [25] for these para-virtualized I/O devices on KVM-accelerated hypervisors. The driver is split into two parts: the frontend and the backend. The former is provided by the guest kernel while the backend is provided by the host. This abstraction layer can be used for para-virtualization of any I/O device. In case of a network interface, there exist at least two buffers. One buffer is handling all incoming packets, while the second buffer is handling all outgoing packets. The original version of *virtio* [32] was developed by Rusty Russell for the support of his own virtualization solution. RUSTYHERMIT provides a frontend driver in the kernel that is used to access the file system with *virtio-fs* [39] and network support.

As shown in Fig. 1, RUSTYHERMIT uses *smoltcp* [19] as a dual IPv4/IPv6 stack and is provided by *hermit-sys* to the Rust runtime. *smoltcp* is an event-driven TCP/IP stack written in Rust and designed for bare-metal, real-time systems. In principle, *hermit-sys* creates a thread, which handles all incoming packets including ARP and ICMP packets with the help of *smoltcp*. The IP stack is added to *hermit-sys* and not to the kernel, so it can use the default memory allocator of the Rust runtime and to enable hardware dependent optimizations (e.g., AVX support) as explained before Sect. 4.1. We implemented a direct interface between *smoltcp* and Rust’s standard library to forward TCP streams to all common Rust applications. All outgoing messages are directly passed through the IP stack to the *virtio* interface. Incoming messages triggers an interrupt, where the interrupt handler wakes the IP thread of *hermit-sys*. Acknowledgments and retransmission of lost messages are directly handled by this thread. If the incoming message is intended for a specific thread, which is

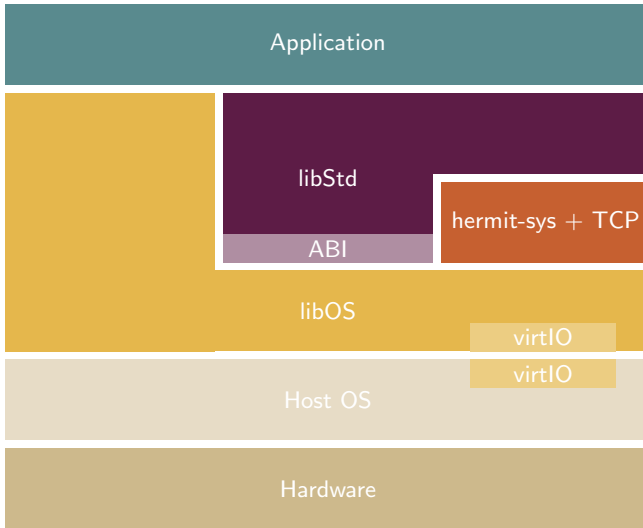


Fig. 1. Architecture overview of RUSTYHERMIT.

blocked on a socket, the IP thread wakes the thread and afterwards the thread consumes the data.

This approach works well for communication pattern in which the delay between requesting data and receiving data is high enough to hide the overhead of both the interrupts and the context switches between the threads. To get peak performance for communication pattern with a short delay, the driver switch to a polling mode. In this case, interrupts from the virtio device will be disabled. The complete communication is then realized by the application threads, which are waiting for incoming messages. If the IP stack is not used for 20ms, the driver switch back to the non-polling mode.

We implemented this behavior with *Futures* [33], which are Rust way of expressing asynchronous computation. This interface provides a polling method to check if the data is available. The usage of this standard mechanism offers the possibilities to easily check several futures asynchronously.

5 Evaluation

All benchmarks were performed on a NUMA system possessing two sockets each with 12 physical cores, exposing 24 cores in total. The CPUs are Intel Skylake CPUs (Xeon Gold 6128) clocked at 3.4 GHz, equipped with 256 GiB DDR4 RAM and 19.25 MiB L3 cache. We used a 4.18.0 Linux kernel with CentOS 8. All benchmarks are compiled with optimization level 3 and LTO.

As said before, unikernels are designed to run within a hypervisor. For the evaluation, qemu-kvm 2.12.0 is used and accelerated by KVM. All benchmarks run within virtual machines with the same setup. The network interface and

the storage is integrated by virtio to reduce the overhead. The only difference is that for Linux guests the virtual machine is configured to provide 4 GB of main memory, while RUSTYHERMIT is configured with 512 MByte main memory.

5.1 OS Micro-benchmarks

In this section we present benchmarks regarding system call overhead and scheduling. The `getpid` system call is the one with the smallest runtime and closely represents the overhead of a system call. The function `yield_now` of the Rust runtime triggers the scheduler to check if another task is ready and switches to them. In our case, the system is idle and consequently the function returns directly after the check of the ready queues. For benchmarking the system call performance, we call `getpid` and `yield_now` 1 000 000 times and measure the number of cycles the call took. Table 1 summarizes the results as average number of CPU cycles for Linux and RUSTYHERMIT. The overhead of RUSTYHERMIT is smaller as system calls are just function calls in library OSs and the runtime system is smaller compared to the Linux software stack.

Table 1. Comparison of basic system services by Linux and RUSTYHERMIT.

System activity	Linux	RustyHermit
Time to boot	≤15 s	≤1.0 s
Reserved memory	748 MByte	58 MByte
Boot image size	1.8 GByte	1.55 MByte
<code>yield_now()</code>	1439 cycles	68 cycles
<code>getpid()</code>	1147 cycles	43 cycles

Table 1 also shows memory consumption of a minimal CentOS 8 configuration, where only a secure shell server is running and compares it with the memory consumption of the smallest possible RUSTYHERMIT application. To determine these numbers, the memory consumption of the hypervisor on the host system is evaluated. The numbers show the physically allocated memory. The reserved memory in the logical address space is clearly larger because the virtual machines are configured to use up to 4 GByte memory for the Linux guest and 500 MByte for RUSTYHERMIT as guest. Both virtual machines are not fully utilized. The low memory consumption and the small image size for RUSTYHERMIT promise a better resource utilization in data centers.

To evaluate the boot time, the time between the start of the virtual machine and the first response of a ICMP-based ping request is measure. To avoid side effects from the storage device, the boot image is stored in *tmpfs*. The last step before entering the main function of the Rust application in RUSTYHERMIT is the initialization of the network stack. Therefore, the results show the minimal time to start the unikernel application within a hypervisor. While it is possible

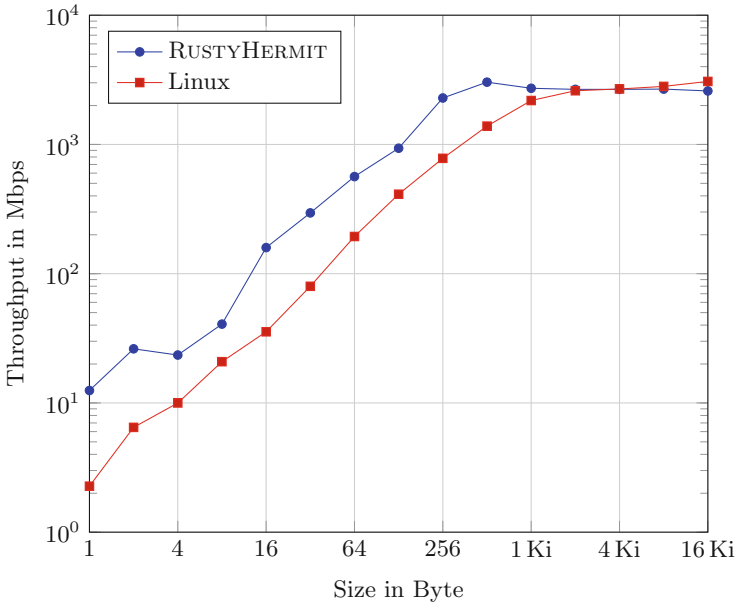


Fig. 2. Comparison of the network throughput between RUSTYHERMIT and Linux

to start applications in Linux before the network services have started, this is a rather unlikely scenario and it is more likely that other services are started between the network service start and the application start. As expected for a unikernel, RUSTYHERMIT is clearly faster in comparison to Linux which is beneficial for services requiring low latencies.

5.2 Network Performance

To determine the network performance, a benchmark is used transferring data with Rust's standard TCP stream interface. Both the server and the client are running on the same node. The sender is running in all test scenarios within a VM, while the receiver is running natively on the host system. In case of the senders, the checksums of the IP packets are built within the guest machine. All interfaces use an MTU of 1500 Bytes and the Nagle algorithm [23] is disabled.

Figure 2 compares the performance between RUSTYHERMIT and Linux. Up to a message size of the MTU, RustyHermit provides a clearly higher bandwidth in comparison to Linux. For messages with the size of 512 bytes, RustyHermit is twice as fast. Linux is more efficient at splitting messages larger than the MTU size and as a result Linux currently provides higher performance for large messages. This is something that needs to be worked on in *smoltcp* and is not directly part of RUSTYHERMIT.

6 Conclusion

In this paper, we present RUSTYHERMIT a unikernel completely written in Rust. We integrate a Rust-based IP stack not depending on C/C++. RUSTYHERMIT is published on GitHub [13] and is completely integrated into Rust’s toolchain. Consequently, common Rust applications, which do not bypass the Rust runtime and directly use OS services are able to run on RUSTYHERMIT without modifications.

We show that RUSTYHERMIT provides excellent performance in micro benchmarks and has a small memory footprint compared to a minimal CentOS 8 virtual machine image. The IP stack *smoltcp* and its integration into Rust’s standard library provide a higher bandwidth in comparison to Linux for message smaller than the MTU size. In combination with the low memory footprint is RUSTYHERMIT suitable for the development of scalable micro services.

References

1. Balasubramanian, A., Baranowski, M.S., Burtsev, A., Panda, A., Rakamari, Z., Ryzhyk, L.: System programming in rust: beyond safety. *SIGOPS Oper. Syst. Rev.* **51**(1), 94–99 (2017). <https://doi.org/10.1145/3139645.3139660>
2. Bratterud, A., Walla, A., Haugerud, H., Engelstad, P.E., Begnum, K.: IncludeOS: a resource efficient unikernel for cloud services. In: Proceedings of the 2015 IEEE 7th International Conference on Cloud Computing Technology and Science (Cloud-Com) (2015)
3. Cutler, C., Kaashoek, M.F., Morris, R.T.: The benefits and costs of writing a POSIX kernel in a high-level language. In: 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI), pp. 1–19, September 2018
4. D Language Foundation: The D Programming Language. <https://dlang.org/>. Accessed 4 Mar 2019
5. Dunkels, A.: Design and Implementation of the LwIP TCP/IP Stack. Swedish Institute of Computer Science (2001)
6. eduOS-rs: A teaching operating system written in rust. <https://rwth-os.github.io/eduOS-rs/>. Accessed 13 Feb 2019
7. Google: The Go Programming Language. <https://golang.org>. Accessed 4 Mar 2019
8. Greenblatt, R.D., Knight, T.F., Holloway, J.T., Moon, D.A.: A LISP machine. *ACM SIGIR Forum* **15**(2), 137–138 (1980)
9. Kantee, A.: Flexible operating system internals - the design and implementation of the anykernel and rump kernels. Ph.D. thesis, Department of Computer Science and Engineering, Aalto University, Aalto, Finland (2012)
10. Kivity, A., et al.: OSv - Optimizing the operating system for virtual machines. In: USENIX Annual Technical Conference (2014)
11. Klabnik, S., Nichols, C.: The Rust Programming Language (Manga Guide). No Starch Press, San Francisco (2018)
12. Kuenzer, S., et al.: Unikernels everywhere: the case for elastic CDNs. In: Proceedings of the 13th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE 2017, pp. 15–29. Association for Computing Machinery, New York (2017). <https://doi.org/10.1145/3050748.3050757>

13. Lankes, S., Breitbart, J., Pickartz, S.: Rustyhermit - a rust-based, lightweight unikernel. <https://github.com/hermitcore/libhermit-rs>. Accessed 3 Oct 2019
14. Lankes, S., Pickartz, S., Breitbart, J.: HermitCore: a unikernel for extreme scale computing. In: Proceedings of the 6th International Workshop on Runtime and Operating Systems for Supercomputers, ROSS 2016, pp. 4:1–4:8. ACM, New York (2016)
15. Lankes, S., Pickartz, S., Breitbart, J.: A low noise unikernel for extrem-scale systems. In: Knoop, J., Karl, W., Schulz, M., Inoue, K., Pionteck, T. (eds.) ARCS 2017. LNCS, vol. 10172, pp. 73–84. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-54999-6_6
16. Lankes, S., Breitbart, J., Pickartz, S.: Exploring rust for unikernel development. In: Proceedings of the 10th Workshop on Programming Languages and Operating Systems, PLOS 2019, pp. 8–15. Association for Computing Machinery, New York (2019). <https://doi.org/10.1145/3365137.3365395>
17. Levy, A., et al.: Ownership is theft: experiences building an embedded OS in Rust. ACM, New York, October 2015
18. Levy, A., Campbell, B., Ghena, B., Pannuto, P., Dutta, P., Levis, P.: The case for writing a kernel in rust. Proceedings of the 8th Asia-Pacific Workshop on Systems (APSys 2017), pp. 1–7 (2017)
19. M-Labs: uhyve - a minimal hypervisor for rustyhermit. <https://github.com/m-labs/smoltcp>. Accessed 1 Aug 2019
20. Madhavapeddy, A., et al.: Unikernels: library operating systems for the cloud. In: Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2013, pp. 461–472. ACM, New York (2013). <https://doi.org/10.1145/2451116.2451167>
21. Madhavapeddy, A., Scott, D.J.: Unikernels: rise of the virtual library operating system. ACM Queue **11**(11), 30 (2013)
22. Microsoft Security Response Center: Why rust for safe systems programming. <https://msrc-blog.microsoft.com/2019/07/22/why-rust-for-safe-systems-programming/>. Accessed 1 Aug 2019
23. Nagle, J.: Congestion Control in IP/TCP Internetworks. <https://tools.ietf.org/html/rfc896>. Accessed 10 July 2020
24. Narayanan, V., Baranowski, M.S., Ryzhyk, L., Rakamarić, Z., Burtsev, A.: Redleaf: towards an operating system for safe and verified firmware, pp. 37–44 (2019). <https://doi.org/10.1145/3317550.3321449>. <http://doi.acm.org/10.1145/3317550.3321449>
25. OASIS Virtual I/O Device (VIRTIO) TC: Virtual I/O Device (VIRTIO) Version 1.1 (2018). <https://docs.oasis-open.org/virtio/virtio/v1.1/csprd01/virtio-v1.1-csprd01.html>
26. Olivier, P., Chiba, D., Lankes, S., Min, C., Ravindran, B.: A binary-compatible unikernel. In: Proceedings of the 15th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE 2019, pp. 59–73. Association for Computing Machinery, New York (2019). <https://doi.org/10.1145/3313808.3313817>
27. Picheta, D.: Nim in action. <http://nim-lang.org/>. Accessed 4 Mar 2019
28. DPK Project: Data Plane Development Kit. <https://www.dpdk.org>. Accessed 10 July 2020
29. Redell, D.D., et al.: Pilot - an operating system for a personal computer. Commun. ACM **23**(2), 81–92 (1980)
30. RedHat: Newlib - a c library for embedded systems. <https://sourceware.org/newlib/>. Accessed 13 Feb 2019

31. Redox: A unix-like operating system written in rust. <https://www.redox-os.org>. Accessed 13 Feb 2019
32. Russell, R.: Virtio: towards a de-facto standard for virtual i/o devices. SIGOPS Oper. Syst. Rev. **42**(5), 95–103 (2008). <https://doi.org/10.1145/1400097.1400108>
33. Rust Project Developers: Future - a Representation of asynchronous computation. <https://doc.rust-lang.org/std/future/trait.Future.html>. Accessed 10 July 2020
34. Rust Project Developers: Cargo - a rust package manager. <https://doc.rust-lang.org/cargo/>. Accessed 4 Mar 2019
35. Rust Project Developers: The Rust Programming Language. <https://www.rust-lang.org>. Accessed 4 Mar 2019
36. Tock: A secure embedded operating system for cortex-m based microcontrollers. <https://www.tockos.org>. Accessed 4 Mar 2019
37. Torvalds, L.: (2004). <http://harmful.cat-v.org/software/c++/linus>
38. Unikraft: An easy way of crafting unikernels. <http://unikraft.neclab.eu>. Accessed 4 Mar 2019
39. Virtio-fs Project Developers: virtio-fs. <https://virtio-fs.gitlab.io>. Accessed 14 July 2020