

A Survey of Real-Time Operating Systems

Anthony Serino, *Member, IEEE*, Liang Cheng *Senior Member, IEEE*

Abstract—The paper is a literature survey of RTOS (Real-Time Operating System) development. GPOS (General Purpose Operating System) existed before RTOS but did not meet performance requirements for time sensitive systems. Many GPOS have put forward adaptations to meet the requirements of real-time performance, and this paper compares RTOS and GPOS and shows their pros and cons. Furthermore, comparisons among select RTOS such as VxWorks, RTLinux, and FreeRTOS have been conducted in terms of scheduling, kernel, and priority inversion. Various tools for WCET (Worst-Case Execution Time) estimation are discussed. This paper also presents a use case of RTOS, i.e. JetOS, and future advancements such as multi-core RTOS, new RTOS architecture and RTOS security.

Index Terms—Real-time operating system, General purpose operating system, RTLinux, FreeRTOS, VxWorks, multi-core RTOS.



1 INTRODUCTION

RTOS (Real-Time Operating System) is used in a wide range of industrial systems, such as process control systems, avionics, and nuclear power plants. Most real-time operating systems run on embedded systems consisting of pieces of hardware that work as controllers with dedicated functions within mechanical or electronic systems. RTOS is critical for those mechanical or electronic systems with real-time requirements because systems could not operate safely without it. In many real-time systems a missed deadline can lead to disastrous consequences.

1.1 Hard Real Time vs. Soft Real Time

Hard real time systems and soft real time systems are both used in industry for different tasks [15]. The primary difference between hard real time systems and soft real time systems is that their consequences of missing a deadline differ from each other. For instance, performance (e.g. stability) of a hard real time system such as an avionic control system or a nuclear power plant, is dependent on the timeliness of the operation results and the correctness of the results. However, for soft real time systems such as a multimedia on-demand system, their performance is solely dependent on the results. A hard real time system is used in systems that are time sensitive and must meet their deadlines in order to avoid disaster. However, deadlines in soft real time system are less strict so that if they miss their deadlines it does not result in disaster. Part of this is because hard real time systems do not have an easy way to recover from a failure; where as a soft real time system can be time-elastic. Some of the advantages of a hard real time system is that it has a faster reaction time than soft real time systems. The hard real time system also has predefined deadlines and predictable performance, where the soft real time systems may experience degraded performance.

This survey of RTOS covers related technologies and gives an introduction of them. It also includes select topics including new adaptations on solving priority inversion, a newly purposed platform for WCET (worst-case execution time) tool development, and a RTOS that supports multi-core processing.

This paper starts with a comparison between GPOS and RTOS, and their advantages and disadvantages for general computing and for real-time systems. Section 2 discusses different techniques used to handle task scheduling and solve priority inversion. Section 3 provides a comparison between three RTOS implementations, namely VxWorks, RTLinux, and FreeRTOS. This comparison covers kernels of the operating systems, schedulers used, and how they solve priority inversion. The next section discusses WCET analysis tools used today in industry along with the development on a new platform to build and compare these tools. Section 5 provides a use case of a recently developed RTOS, i.e. JetOS in avionics. The final section of this paper discusses future developments of RTOS systems ranging from a multi-core RTOS (HIPPEROS) to a new RTOS platform (HERCULES).

2 GPOS vs RTOS

A general comparison between RTOS and GPOS is provided below to show the advantages and weaknesses of these operating systems when used for real time system applications.

2.1 Task Scheduling

The first part of differences is the way where GPOS and RTOS perform their task scheduling. Generally speaking, a GPOS utilizes a fairness policy that allows all processes to share the processor. This hinders the GPOS ability to handle time sensitive tasks because it cannot guarantee task dispatch latencies. The RTOS uses a priority based preemptive scheduling system. This enables a high priority task to take the processor from a lower priority task, and allows the high priority task to run without interruption.

- A. Serino was with a REU program hosted by the Department of Computer science and engineering at Lehigh University. E-mail: serinot@yahoo.com
- L. Cheng is with the Department of Computer science and engineering at Lehigh University. E-mail: cheng@lehigh.edu

The RTOS also utilizes every resource at its disposal to get peak performance.

2.1.1 Scheduling techniques

Round robin scheduling. This form of scheduling uses processor time sharing and gives every process with the same priority a set of time slices [4] [5], each of which corresponds to a fixed amount of processor or CPU time. When a process uses up a CPU time slice the scheduler forces it out of the CPU so that it takes turns to use the CPU resource with other processes. This scheduling technique may lead to relatively large overhead of context switching.

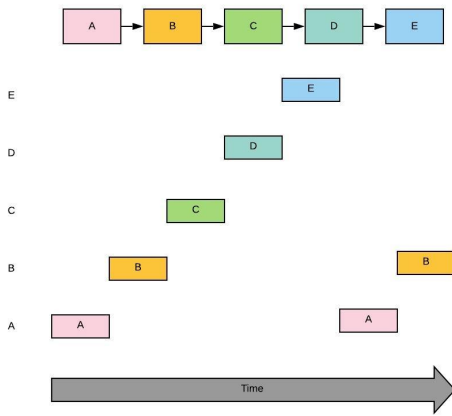


Figure 1. Round-robin Scheduling

First-in-first-out (FIFO) scheduling. The FIFO scheduler will try to execute the task with the highest priority first, but when the processes have the same priority they are scheduled in the order of their arrivals. The first task there will run to completion before starting the next one, or until a higher priority task takes the processor [4] [5].




Figure 2. FIFO Scheduling

Rate-monotonic scheduling. This scheduling method is a static-priority algorithm that sets the priority level for each task in the order of their period information. Short period tasks execute frequently, and a long period refers to infrequent execution. Short period tasks are given higher priorities while long period ones are given low priorities. This lets high priority tasks run first, and it is best used when there are well defined periodic tasks with the same CPU run time length [4] [5].




Figure 3. Rate-monotonic Scheduling

Earliest-deadline-first (EDF) scheduling. This scheduling method computes the priority of processes dynamically based on their arrival time and the execution requirements and deadlines so that it schedules the task with the earliest deadline first [4] [5]. EDF scheduler is more capable of making all deadlines to be met when system load is high comparing to rate-monotonic scheduling [4] [5].




Figure 4. EDF Scheduling

2.1.2 Kernel

Generally the GPOS does not support preemption. Preemption is when processes and threads with higher priorities can take the processor from lower priority tasks. By not allowing preemption the GPOS suffers when trying to complete a task that is time sensitive as it can make the task wait and thus miss its deadline. The GPOS is unable to cancel system calls even if they are from a lower priority task and create unpredictable delays. The advantages of its kernel are in its support for widely used application programming interfaces (APIs) and customizable operating system components for application-specific demands.

The RTOS fully supports preemption where it imposes an upper bound on how long the preemption has its interrupts disabled. This allows a high priority task to run to completion without interruption and thus let time sensitive tasks meet their deadlines. The kernel of a RTOS try to use the least amount of resources possible. To keep it simple only services with short execution paths are allowed in the kernel, and its process loading happens outside along with its file systems. This architecture makes it so that if one of these systems fails it does not corrupt other services or the kernel. The benefit of the RTOS kernel is that there is only a

small core of fundamental operating system services in the kernel. These fundamental services are signal, timers, and the scheduler [14].

2.2 Priority Inversion

Priority inversion occurs when two tasks of different priorities share a resource and the higher priority task cannot get the resource from the lower priority task. This is an issue for time sensitive tasks as it can prevent it from meeting its deadline by forcing it to wait for a resource. This needs to be fixed as it can result at best in blocking, but can also result in chain blocking or worse a full deadlock. There are a couple of protocols to solve priority inversion.

2.2.1 Priority inheritance

The priority inheritance protocol is where if a higher priority process is waiting for a lower priority process for a resource the process scheduling algorithm gives a higher/highest priority to the lower priority task on the processor so that it cannot be preempted by a different task until it completes the execution of its critical section. For example, there is a process A on the processor of priority 2 and there is a higher priority process B that could not preempt it at priority 7. The process scheduling algorithm would assign process A priority 9 temporarily so that it finishes its critical section on the processor without being interrupted by other processes which would delay process B further. This allows process B to get access to the resource as fast as possible, and once the resource is freed by process A the process scheduling algorithm reverts process A back to its original priority of 2 [2].

2.2.2 Priority ceiling

The priority ceiling is where each resource is assigned a priority ceiling where the priority is equal to the highest priority of any task which may lock the resource. This works by temporarily raising the priority of tasks in certain situations. Basically if process A tries to preempt the critical section of another process to execute in its own critical section Z; then the priority of the new section should be higher than the priority of inherited properties of all the preempted sections. If this fails then process A is denied entry into the critical section and is suspended [2].

2.2.3 Other priority inversion solutions

Priority remapping. As an improvement to the priority inheritance protocol, the priority remapping method expands the highest priority from 64 to 128 without change its interface. This means that users still only see 64 priorities, but its internal task creation is multiplied by 2. Its internal priority is extended to (0,2,4,...,126,128) even priority, where the odd priorities are left for changing when priority inversion happens.

Priority exchange. This method is also an improvement on priority inheritance protocol. It swaps the tasks priorities when a higher priority task is blocked by a lower priority task. The priorities will be swapped back after the critical section finishes running the lower priority task [3].

2.3 Modified GPOS

A modified GPOS is the result produced by people who have tried to adapt and change GPOS to have the same capabilities of both GPOS and RTOS. It has a major advantage in that if it works like a GPOS it can be used more widely for a greater amount of tasks, and if it has the capabilities of an RTOS then it supports tasks for time critical systems. For example, Linux 2.6 added preemption [11]. Generally when a GPOS is directly modified to support RTOS functionality high-resolution timers will be used. This modification enables the process scheduler to make the system more reactive and event-driven. However, Linux 2.6 is not as fast as other RTOS and the low latency patches for the GPOS timers do not solve the priority inversion issue [14].

Another way people have tried to improve GPOS is by introducing a new architecture called dueling kernels where the GPOS is ran on top of an RTOS [14]. This architecture sends real time tasks to run on the RTOS, with a higher priority than other tasks running on the GPOS. The RTOS gives these tasks the ability to preempt the tasks on GPOS, and then gives the CPU back to the GPOS when it finishes running the high priority tasks. This system has an issue where the tasks running on the RTOS have limited use of the GPOS services due to preemption issues. This causes RTOS to recreate services that exist in the GPOS. RTOS tasks also cannot use the memory management unit which is used by the GPOS for non-realtime processes. GPOS services that are ported often have different vendor extensions that do not work with other vendor's extensions.

3 VxWORKS, FREERTOS, RTLinux

The RTOS chosen for comparisons in this paper are some of the industries' top competing RTOS. VxWorks and RTLinux have been extensively compared to each other through research due to the continuing development on both the commercially available VxWorks and the free RTLinux. FreeRTOS being a more recent RTOS compared to VxWorks and RTLinux has seen little comparisons with either VxWorks or RTLinux directly. Important aspects to compare them include their kernels, schedulers, and how they handle priority inversion.

3.1 Kernel

3.1.1 RTLinux

RTLinux has a special design for its kernel because it has two kernels. RTLinux uses a specialized real time kernel called the RTCore [4]. The second kernel is the standard Linux kernel which is used for regular applications that do not have time constraints. Both interrupt handling and thread handling are controlled by RTCore. The RTCore will send these interrupts to the appropriate interrupt handler. This RTCore also restricts the Linux kernel by making it unable to disable interrupts to make sure it does not interfere with process scheduling. Thus the Linux kernel can only run when there is a task that is not real time. Real time applications can communicate with Linux kernels through first-in-first-out pipes. This enables the RTCore API threading that helps programmers learn to program

on this system. The dual kernels gives RTLinux the full functionality of Linux, while adding real time capacity to it [6].

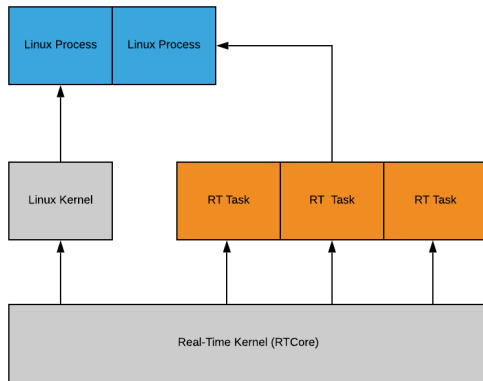


Figure 5. RTLinux Architecture Overview

3.1.2 VxWorks

VxWorks uses a single microkernel to handle basic kernel functions [4]. Additional functions like file sharing and networking have to be loaded from provided libraries. This system provides flexibility to fit its functionality without loosening its constraints on available memory and resources [8] [9].

3.1.3 FreeRTOS

FreeRTOS also utilizes a single microkernel to handle real time tasks. This kernel supports dynamic scheduling or a priority based scheduler, blocking and deadlock avoidance, and scheduler suspension. It can utilize fully featured API, or a lightweight API [8] [9].

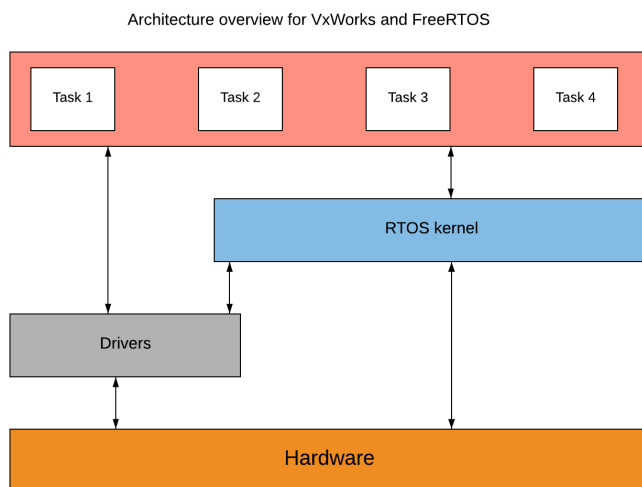


Figure 6. Architecture Overview for VxWorks and FreeRTOS

These kernels are similar in how they all have a means to handle real-time tasks. Some of the major contrasts for RTLinux is that it supports a dual kernel which allows it to handle a wide variety of tasks at the cost of being a larger kernel by having both the standard Linux kernel and the RTCore. While, the structure of VxWorks and FreeRTOS are similar in both using micro kernels.

3.2 Scheduler

The scheduler of the RTOS is an important part of how an RTOS decides the next task to be put on the processor, and to make sure that all tasks meet their deadlines. This section of the paper will discuss similarities and differences between the schedulers used by RTLinux, VWorks, and FreeRTOS.

3.2.1 RTLinux

RTLinux has a flexible scheduler by allowing different scheduling techniques to be used based off the programmers needs. The RTLinux scheduler supports FIFO scheduling, EDF scheduling, and rate-monotonic scheduling. This flexibility allows for better suited schedulers for different systems [5].

3.2.2 VxWorks

VxWorks uses a preemptive round-robin scheduling algorithm. Task priorities can range from 0 to 255 where 0 is the highest priority [11]. If a task with a higher priority than the one on the processor is ready to run, then the lower priority task will be suspended so that the higher priority task can be ran. If the two tasks have the same priority then they go into round-robin scheduling. If a resource is unavailable then the processor swaps back to a lower priority task until the resource is available. VxWorks supports POSIX API which makes the system FIFO. This makes the system more flexible and easier to meet different industrial needs [5].

3.2.3 FreeRTOS

FreeRTOS uses a dynamic preemptive priority based scheduling algorithm [9]. This scheduler can allow the user to choose to run processes in a co-operative manner or using a preemptive policy. The difference is that the preemptive policy always runs the highest priority task, and when two tasks have the same priority they share CPU time. The co-operative manner only allows context switches to occur by calling a function or when a task gets blocked.

Both RTLinux and VxWorks use the priority inheritance protocol. However, FreeRTOS does not use one of the typical ways of dealing with priority inversion; it deals with deadlocks formed by priority inversion by enforcing non-blocking tasks and by blocking tasks for fixed amounts of time.

4 WCET TOOLS

WCET analysis tools give an estimated worst case execution time of a task. Modern processor components like caches and pipelines complicate the task of estimating the WCET [18]. If a tool does not take pipelines or cache behavior into consideration it may over-estimate the WCET by multiple orders of magnitude.

4.1 AbsInt aiT

AbsInt aiT is a worst-case execution time analysis tool. This tool solves the cache and pipeline issue by statically analyzing the task's cache and pipeline behavior. This enables it to get the correct upper bounds of WCET of the task. This tool uses the technique of abstract interpretation, and offers a graphical user interface to visualize the WCET path, and allows for an interactive way to inspect pipelines and caches [18]. The abstract interpretation is a method uses a semantic based method for safe and static program analysis. The overall method applied by AbsInt aiT breaks it into multiple steps to find the WCET.

1. Reconstruction of the control flow
2. Value analysis: computation of address ranges for instructions access memory
3. Cache analysis: classifies memory references as cache hits or misses
4. Pipeline analysis: predicts the behavior of the program on the processor
5. Path analysis: determine the worst case execution path
6. Analysis of loops and recursive procedures

Steps 2, 3, and 4 are done with abstract interpretation; path analysis is done using integer linear programming [18].

4.2 Bound-T Tool

The Tidorum Ltds Bound-T tool is software used for static analysis of code to estimate its WCET and stack usage for embedded systems. This tool in its current state is not going through further development, and is currently open source. The Bound-T tool uses the same static analysis approach as that used by AbsInt aiT. The Bound-T tool was developed for local and safe analysis of simple control flows. This system used a PA (Presburger Arithmetic)-based analysis where an approximate model of computer arithmetic is used, which assumes that integer variables never overflow or wrap around. However this approximation leads to an issue of not being able to find a feasible execution path. In order for Bound-T to start getting back on track to finding WCET for embedded systems it either has to drop the PA-based analysis, or find a new form of preliminary analysis to ensure that the PA analysis can be applied [16].

4.3 OTAWA Toolbox

The OTAWA toolbox is a proposed WCET analysis tool framework that allows it to host researchers WCET algorithms and includes an abstraction layer that separates the hardware analysis from the instruction set architecture. This framework should also allow for the comparison between tools to be easier to do, and support the development of new tools. The OTAWA toolbox was used in the MERASA project to help find the most optimal WCET analysis tool for a multicore processor running mixed-critical workloads [17].

In summary, the aiT tool and the Bound-T tool both use static analysis to find the WCET. The difference is that Bound-T has run into trouble in that it uses PA-based analysis that can no longer reliably find upper bounds.

The aiT tool is a popular tool for specific processors. The OTAWA framework should allow for newly developed WCET analysis tools to be compared with existing WCET analysis tools, and help the further development of WCET tools by providing a standard framework and C++ library [17].

5 A USE CASE OF RTOS

Avionics is a good example where real time systems are used. There has been recent development of a specific type of RTOS, i.e. JetOS, to fully meet the ARINC653 international standard for aircraft usage. JetOS originates from POK RTOS, an open source project, which partially meets the ARINC653 standards. The critical parts that have been reworked or added to meet the standard include pok's scheduler, network stack, memory manager, added separate memory, and a reduction of the kernel size for less errors. The system uses ordinary partitions which separates memory, and system partitions are used to utilize services outside of the ARINC653 standard. Both of these types of partitions from the kernel point of view are the same. Currently there is a working prototype of JetOS and it is in its testing stages [13].

The kernel of JetOS drops POK's AADL (Architecture Analysis and Design Language) configuration tools for XML based configuration files. Furthermore it dropped the SPARC platform in favor of building JetOS on top of a platform that other avionic systems use. The platform chosen to replace it was the x86 and powerPC. The kernel is built to support multiple schedulers because different partitions can utilize different schedulers, and is configured statically where the number of partitions, partition memory size, port, names, etc. cannot be changed [13].

Each partition may have one or more processes. Partitions are scheduled based on a round-robin algorithm. Intra-partition schedulers implement lock-wait-unlock and priority scheduling. The resources are pre-allocated to ensure reliability. The memory is pre-allocated to every partition. These partitions are scheduled differently than kernel modules where partitions are run in user mode with time and space constraints [13].

6 FUTURE DIRECTIONS

With the rising need for both faster computing speed and better resource usage, multi-core computing has existed in general-use computers for a number of years. However, RTOS have been using a mono-core system that is becoming obsolete in terms of speed and resource usage. A multi-core RTOS as new architecture for embedded systems aims at providing a major increase in both speed and resource usage including energy efficiency [10].

6.1 Multicore Real-Time Operating System

A challenge in developing multicore RTOS is how to evaluate it according to the industry standards. Uni-core systems meet the requirements of standards such as ARINC653 and AUTOSAR. Researchers have put forward

the use of general purpose operating systems, such as Linux, in order to provide an environment where they could evaluate a multicore real-time system. Although this approach has the advantage of being able to reuse code that already exists, it runs into issues where Linux was not built to support hard real-time systems or to meet the constraints for safety-critical applications. HIPPEROS, a multicore RTOS project launched in 2010, is built from scratch so that it can implement hard real-time techniques with multicore design principles. Its kernel is built to scale with an increasing amount of cores [10].

6.1.1 Kernel

HIPPEROS' kernel is able to run on multiple different architectures and platforms with an arbitrary number of cores. It uses distributed asymmetric micro-kernel architecture that allows each core to execute a local part of the kernel. This enables a dedicated core to execute the kernel's system calls, scheduler, and resource handling and frees up parts of the kernel to be able to be executed in parallel. The kernel is configurable in that the developer or system designer can select the scheduling policy or resource allocation protocol at its build time. To manage hard real-time tasks it uses a popular process model that gives executable and timing information like deadline, period, and worst-case execution time.

a. Asymmetric Kernel Architecture

The problem with symmetric kernel architecture design is that the cores are executed with the same kernel code and protected data structures with fine-grained lock mechanisms. The asymmetric design allows for one core to fully dedicate itself to running the scheduler and dispatching the processes to other cores. For the HIPPEROS project the dedicated core is called the master core and is responsible for managing global resources, scheduler, system calls, and message passing to allow for the kernel to be executed in parallel. This works as follows whenever there is a scheduling decision the master core must be woken up to notify the slave core (any core other than the master core) to perform a context switch. The slave cores can only perform context switch after receiving an inter-processor interrupt (IPI) for the master core. This system of master and slave cores does not require locking mechanisms, and it is expected to be able to handle up to 8 cores before overloading the master core without using clustering for enhancement [1] [10].

6.1.2 Scheduler

The scheduler's API is preemptive and priority based. When a task switches state (blocked to ready), a scheduler module is called and it decides if context switching must occur due to tasks priorities. If a task misses its deadline, then a configurability policy enacts a range of responses that can terminate the process, ignore the event, or change the priority of the process [10].

6.2 HERCULES Framework

HERCULES is a project for the development of a high-performance real-time architecture for low-power embedded systems. Estimated effects of the HERCULES project include a reduction in energy consumption, productivity improvement in programming and maintaining advanced computing systems, increased concurrency and parallelism in applications, and enhanced trust of embedded systems. The objective of the project is to introduce predictability into embedded high-performance computing. There are two use cases for the HERCULES project: avionic and automobiles.

The avionic use case considers future airplanes that will have more complex needs in regards to image processing or computer vision, which may be used during landing, surveillance activities, and navigation. Moreover, the number of cameras on board an airplane is expected to increase, possible direct video streams for the pilot and crew, and possible automation to extract meaningful data from the video streams. Machine learning techniques can be applied and have been proven to do image processing at the cost of higher algorithmic complexity and computational requirements. A visual object tracking application based on Airbus high-speed machine learning techniques was used to test the HERCULES framework with various programming models and GPU-based platforms. [7].

The automobile use case is involved with autonomous driving for valet parking. The HERCULES framework was chosen to be tested for valet parking for three reasons. First it has subset functionalities required for self-driving cars, and testing is affordable within the time frame of the HERCULES project. Second, the project team has simulators and test environments to create the scenarios. Third, the agents drive at low speeds, which offers clear safety advantages. This use case study shows four major areas where algorithms will be applied are the perception area (sensor data processing), the data fusion area (creates environmental model), the decision area (action and path planning), and localization area (GPS and MAP data management) [7].

These use case studies show the flexibility of the HERCULES framework, which can be adapted to different industries.

6.3 Security Research related to RTOS

As RTOS has been used in devices and SCADA (Supervisory Control And Data Acquisition) systems for time-sensitive and mission-critical tasks, its security is an important area of research. Several security vulnerabilities of RTOS have been discussed in a report on RTOS security [?], such as lack of authentication enforcement, inefficiency of encryption, code injection, exploiting shared memory, priority inversion, denial of service attacks, and inter-process communication attacks.

Adversaries may attack the devices or embedded systems through accessing the remote debugging tools of RTOS. For example, US-Cert Vulnerability Note VU362332 [19] stated that VxWorks debug service was enabled by default so that an attacker might be able to fully compromise the embedded systems and create severe damages to the physical processes controlled by the SCADA systems through remote memory dump and remote function calls.

Thus security should be a top priority when designing, implementing, and configuring RTOS for remote debugging tools.

RTOS security may also be enhanced by detecting attacks based on systems call sequences, network activities, and power consumption patterns. For example, a class of attacks, known as payload attacks and successfully launched by Stuxnet, modifies PLC control programs (i.e., the payload for PLC firmware which may be a RTOS) and causes damages to the physical system. In [21] the authors studied firmware-level detection of PLC payload attacks that alter the timing behavior of the payload. WCET analysis and monitoring have also been used to enhance RTOS security [22].

RTOS also needs to address the memory fragmentation issue to improve its reliability and avoid program stalls as field devices such as PLCs (Programmable Logic Controllers) using RTOS may continuously run for many years without rebooting [20].

7 CONCLUSION

This survey covered the basics of how RTOS have been developed from GPOS, and how they have different advantages and disadvantages for general computing and dealing with real-time embedded systems. The discussion on the comparison among RTOS implementations that are used in industry, namely VxWorks, RTLinux, and FreeRTOS, showed some of their differences and similarities in their kernels, schedulers, and how they handle priority inversion. Then we described some of the tools used in industry for WCET analysis, and a new platform developed to help develop new tools and compare them with existing tools. The discussion of JetOS showed how existing RTOS can be adapted and changed to meet industry specific requirements. This evolves into how future RTOS will need to be able to utilize multiple cores instead of monocoresh based systems that are in place today in industry. Finally the security aspect of RTOS was discussed.

ACKNOWLEDGMENTS

This work is supported by the U.S. National Science Foundation (NSF) under Award Numbers CNS-1757787 and CNS/CPS-1646458, and by the U.S. Department of Energy under Award Number DE-OE0000779. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsors of the research.

REFERENCES

- [1] Juan Rivas, Joel Goossens, Xavier Poczekajlo, Antonio Paolillo. Implementation of memory centric scheduling for COTS multi-core real-time systems. In Proceedings of the 31st Euromicro Conference on Real-time Systems. 2019.
- [2] Lui Sha, Ragunathan Rakumar, John Lehoczky. Priority inheritance protocols: an approach to real-time synchronization. IEEE Transactions on Computers, Vol. 39, No. 9, pp. 1175-1185. September 1990.
- [3] Silambarasan, Ramanatha Venkatesan. Handling of priority inversion problem in RTLinux using priority ceiling protocol. In Proceedings of the International Journal of Advanced Engineering Research and Science (IJAERS). Vol. 3, Issue 6. June 2016.
- [4] Daniel Forsberg and Magnus Nilsson. Comparison between scheduling algorithms in RTLinux and VxWorks. Technical Report, Computer Science and Engineering at the University of Linkping, November 2006.
- [5] Stefan Holmer, Osker Hermansson. A comparison between the scheduling algorithms used in RTLinux and in VxWorks - both from a theoretical and a contextual view, Technical Report, Computer Science and Engineering at the University of Linkping, 2006.
- [6] Federico Reghenzani, Glueppe Massari, William Fornaclari. The real-time Linux kernel: A survey on Preempt RT. ACM Computing Surveys. Volume 52, Issue 1, February 2019.
- [7] Marko Bertogna. High-Performance Real-time Architectures for Low-Power Embedded Systems. H2020-EU.2.1.1. Project Website: <https://cordis.europa.eu/project/rcn/199161/factsheet/en>. 2016-2018.
- [8] Zhu, Ming-Yuan. Understanding FreeRTOS: A Requirement Analysis. CoreTek Systems, Inc., Beijing, China, Technical Report, 2011.
- [9] Rich Goyette. An analysis and description of the inner Workings of the FreeRTOS kernel." Course Report for SYSC5701: Operating System Methods for Real-Time Applications, Department of Systems and Computer Engineering, Carleton University. April 2007.
- [10] Antonio Paoillo, Oliver Dersenfans, Vladimir Svoboda, Joel Goossens, Ben Rodriguez. A New congurable and parallel embedded real-time micro-kernel for multi-core platforms. In Proceedings of the ECRTS Workshop on Operating Systems Platforms for Embedded Real-Time applications (ECRTS-OSPERT'15), July 2015.
- [11] Sukhyun Seo, Junsu Kim, Su Min Kim. An analysis of embedded operating systems: Windows CE, Linux, VxWorks, uC/OS-II, and OSEK/VDX. Journal of Applied Engineering Research, Vol. 12, No. 18, pp. 7976-7981, 2017.
- [12] C. M. Krishna and Kang G. Shin. Real-Time Systems (McGraw-Hill Series in Computer Science). McGraw-Hill College. December 1996.
- [13] Mallachiev, K.M. Pakulin, Nikolay Khoroshilov, Alexey. Design and architecture of real-time operating system. Proceedings of the Institute for System Programming of the RAS. Vol. 28, pp. 181-192. 2016.
- [14] Paul Leroux. RTOS versus GPOS: What is best for embedded development? Embedded Computing Design. January 2005.
- [15] Nisar Ahmed. Understanding real time OS concepts. International Journal of Scientific Engineering Research. Volume 6, Issue 9, pp. 67-70, September 2015.
- [16] Niklas Holsti. Status of the Bound-T time and stack analyser. <http://www.bound-t.com/>.
- [17] C. Ballabriga, H. Cass, C. Rochange, and P. Sainrat. OTAWA: An Open Toolbox for Adaptive WCET Analysis. In: Min S.L., Pettit R., Puschner P., Ungerer T. (eds) Software Technologies for Embedded and Ubiquitous Systems. SEUS 2010. Lecture Notes in Computer Science, vol 6399. Springer, Berlin, Heidelberg, 2010.
- [18] AbsInt. The Industry Standard for Static Timing Analysis. <https://www.absint.com/ait/index.htm>
- [19] US-Cert, Vulnerability Note VU362332, Wind River Systems VxWorks debug service enabled by default,

<https://www.kb.cert.org/vuls/id/362332/>

- [20] B. Zhu, A. Joseph and S. Sastry, A taxonomy of cyber attacks on SCADA systems. 2011 International Conference on Internet of Things and 4th International Conference on Cyber, Physical and Social Computing, Dalian, pp. 380-388, 2011.
- [21] Huan Yang, Liang Cheng, and Mooi Choo Chuah, Detecting payload attacks on programmable logic controllers (PLCs), IEEE Conference on Communications and Network Security (CNS), Beijing, China, May 2018.
- [22] Mohammad Hamad, Zain A. H. Hammadeh, Selma Saidi, Vassilis Prevelakis, and Rolf Ernst. Prediction of abnormal temporal behavior in real-time systems. In Proceedings of the 33rd Annual ACM Symposium on Applied Computing (SAC'18). ACM, New York, NY, USA, pp. 359-367. 2018.
Weider Yu (Dipti Baheti, and Jeremy Wai), Real-Time Operating System Security, http://people.cs.ksu.edu/~danielwang/Investigation/RTOS_Security/RTOS_Security.pdf

Anthony Serino III Anthony Serino III is a undergraduate at Misericordia University. He is majoring in Computer Science with minors in IT Security and Mathematics. He gained his first research experience as a REU student funded by the NSF as a part of the the REU program at Lehigh University.

Liang Cheng Dr. Liang Cheng received his PhD from Rutgers, The State University of New Jersey in 2002. He focuses his research on enabling intelligent infrastructure based on real-time sensing, model-driven data analytics, and machine learning. He is an expert in cyber-physical systems and ad hoc networks. His research has been funded by U.S. federal funding agencies such as NSF, DOE, DOT, and DARPA, Pennsylvania government, and companies including ABB, Agere Systems, East Penn Manufacturing, and PPL. He has advised 7 Ph.D. students to their graduation, supervised 2 postdocs, advised 23 Master's degree theses, and co-authored more than 100 peer-reviewed papers. More information is available at <http://liangcheng.info>.