

Copyright Notice

The document is provided by the contributing author(s) as a means to ensure timely dissemination of scholarly and technical work on a non-commercial basis. This is the author's version of the work. The final version can be found on the publisher's webpage.

This document is made available only for personal use and must abide to copyrights of the publisher. Permission to make digital or hard copies of part or all of these works for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. This works may not be reposted without the explicit permission of the copyright holder.

Permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the corresponding copyright holders. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each copyright holder.

IEEE papers: © IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The final publication is available at <http://ieeexplore.ieee.org>

ACM papers: © ACM. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The final publication is available at <http://dl.acm.org/>

Springer papers: © Springer. Pre-prints are provided only for personal use. The final publication is available at <link.springer.com>

Lightweight Framework for Runtime Updating of C-Based Software in Embedded Systems

Simon Holmbacka, Wictor Lund, Sébastien Lafond, Johan Lilius
Department of Information Technologies, Åbo Akademi University
Joukahaisenkatu 3-5 FIN-20520 Turku
Email: `firstname.lastname@abo.fi`

Abstract

Software updates in embedded systems are typically performed by bringing the system to stop, replacing the software and restarting the system. This process can in certain cases be very time consuming and costly, which leads to less frequent software updates. In order to establish both long uptime and up-to-date software, the software must be updated during runtime. This paper presents a runtime updating framework for embedded systems. The framework is based on FreeRTOS and mechanisms have been added to dynamically link and re-link FreeRTOS tasks to the system during runtime. Our framework enables the programmer to easily create updatable software with simple annotations to the program. Experiments demonstrate the benefits of updating software during runtime with an acceptable overhead when transferring the application state.

1 Introduction

Long uptime and up-to-date software are highly valued properties in many industrial plants, automation devices, complex machineries etc. The production rate of such devices is dependent on the uptime, which means that a long and complex reboot process highly unwanted [14] or in some cases simply not allowed [6]. Runtime updating, on the other hand, solves this problem by replacing parts of the software during runtime. Besides solving the no-reboot issue, runtime updating would also enable the possibility of transmitting new software versions on-the-fly over a communication media such as the Internet. It would reduce costs and eventual risk of having personnel physically updating software in inaccessible or remote locations.

Runtime updating of software has existed for decades and has been previously presented in many forms [11]. Its behavior and requirements for updating varies based on functionality, operating system, underlying hardware

platform etc. In this paper we present the implementation of a lightweight runtime update framework for embedded systems. The advantages of our framework are *a)* small footprint for data and program memory, which makes it suitable for even small scale embedded systems *b)* no steady state execution overhead *c)* easy annotations of application state without including explicit meta data from the programmer *d)* the framework is compatible with all FreeRTOS compatible hardware and uses the standard GCC without any modifications.

By using the framework, the programmer can create easily updatable applications as plug-in components, and on-the-fly insert them to the system. The framework is an extension to the operating system FreeRTOS [3] and is implemented using 1600 lines of c-code. Early evaluations demonstrate the benefits of updating software during runtime with regard to performance and functionality, and a low overhead for state transfer in typical applications.

2 Related Work

Runtime updating techniques exist in several forms and on different levels of abstraction. Highlevel mechanisms such as [2] and [16] update Java systems online by using component based solutions such as the OSGi. On levels closer to hardware, a component framework THINK [10] has been used to create updatable components with well defined interfaces. THINK supports re-linking of components during runtime and explicit mechanisms for starting and stopping a component. In contrast to the extensive THINK framework, we address the importance of a lightweight framework for embedded devices. Our framework is created for small to medium range embedded devices with a restricted amount of memory and limited amount of processing power. To simplify programming, the framework can be compiled directly with the standard GCC and requires no new language learning.

Wahler et al. present in [15] a component frame-

work for enabling runtime updating in real-time embedded systems. Similar to our approach, the state context in the applications is stored in a shared memory area, but instead of tagging the memory area with meta data, we define a new memory section (`.rtu`) to store data related to runtime updating.

An important question for runtime updating is the level of transparency to the programmer, and what implications the update will have on the application. In [8] the updating mechanism uses a global function pointer to dynamically point on the current function version, and change this pointer value as the software is updated. This framework requires therefore that a version tag is added to each version of the functions. The framework in [8] also requires a special compiler and certain modifications to the program in order to achieve updatable compilable applications. We solve several of these issues by introducing dynamic linking of ELF binaries. In our system, The updated functions are not referred to by their name, but their address in the linked binary.

State transfer [5, 13, 9] is one of the key issues in runtime updating. The approach for implementing state transfer is dependent on the abstraction level the mechanism is operating on. Java-based systems [9] can swap complete classes while C-based systems cannot. Hayden et al. [5] describe the development of a serialization mechanism used to trace the content of each data structure in C-based systems to determine the content of the application state. To trace the state content in a structured way, we adopted a tracing mechanism included in the DWARF library [17].

Lastly we allow the programmer to mark the state context with a special annotation. This approach reduces the transparency slightly, but in contrast to [4] requires no static analysis of the application. It neither puts constraints on the program structure or introduces any steady state overhead because no implicit monitoring is needed during runtime.

3 Implementation

This section describes the most important parts of the implementation of the runtime updating mechanism. We have implemented the mechanism as an extension for FreeRTOS [3]. The runtime updating mechanism is created using roughly 1600 SLOC. A minimum configuration with OS and runtime updating mechanism requires in total only 700 KB of program memory. To further optimize this size, a hand tailored pointer tracer would need to be implemented. Nevertheless chose the DWARF library due to its rich functionality. Sections 3.1 and 3.2 focuses on the two most important aspects of the implementation: *Dynamic binary linking* and *state transfer*.

3.1 Dynamic binary linking

The first main feature of the runtime updating mechanism is to componentize FreeRTOS tasks into separate Executable and Linkable Format (ELF) binaries. This componentization enables the possibility of re-linking software during runtime. The feature is not available in the default FreeRTOS package, and has been added to dynamically add and remove software binaries to the system. Figure 1 illustrates FreeRTOS and the runtime




Figure 1: Structure of the ELF binary-based composition

updating environment using ELF binaries for two parts: The FreeRTOS base system and the tasks schedulable by FreeRTOS. At startup, the system binary is linked with all task binaries. After this, the tasks can use kernel resources and the kernel is able to schedule FreeRTOS tasks from the linked ELF binaries. With dynamic linking, tasks can be externally compiled and inserted into the running system from external sources such as ftp, flash cards, etc. and linked into the kernel. To replace parts of the software, we have defined a set of steps to follow stop, re-link and start components. This algorithm is further described in [15] Section 3.2 with the exception of steps 4 and 5.

3.2 State transfer

A task's state is a task's obtained properties when executing to a certain point in the program. The state is usually represented by a collection of declared variables, pointers or data structures with certain values. The value of the context defining the state must be temporarily stored in a predefined place in order for the new software version to continue with the same state as the old version.

In our framework, the programmer annotates the context with `_RTU_DATA_` e.g `unsigned int _RTU_DATA_a;` This annotation allows the variable to be saved in a special data segment called `.rtu`. It gives the program a structured way of storing selected information in an application, and gives the updated version of the program a predictable way of retrieving the information. As the context notation has been placed, the programmer declares *checkpoints* `TASK_IN_SAFE_STATE()` in the program to determine point at which the program is *safe* for update. A safe state is a state in which no external events

such as inter task communication, open file descriptors, open sockets, etc., can disturb the state of the task. At this point the updating mechanism is invoked and all annotated variables are stored in the `.rtu` segment. As the new version of the program is re-linked with the FreeRTOS system, the state context can be transferred from the temporary `.rtu` segment to the context used in the program.

The runtime update is initiated as soon as any actor sets a `rtu_requested` flag high in the application. This actor can be a timer, external event or another task in the system. Fig. 2 shows two versions of a simple program. The first version of the program contains a counter

<pre> /*Version 1 of the program*/ int _RTU_DATA_ statel = 0; int main() { rtu_requested = 0; while (1) { if (rtu_requested) { /* Go into safe state */ TASK_IN_SAFE_STATE(); } statel += 1; } } </pre>	<pre> /*Version 2 of the program*/ int _RTU_DATA_ statel = 0; int main() { rtu_requested = 0; while (1) { if (rtu_requested) { /* Go into safe state */ TASK_IN_SAFE_STATE(); } statel += 2; } } </pre>
---	---

Figure 2: Example of two versions of a program with an updatable state

which increments by one unit each loop iteration. After the `rtu_requested` flag is set, the runtime update mechanism transfers the counter value to the new version. The new version of the program continues with the stored counter value, but increments this value with two units. In case the type of the context changes (e.g. from `int` to `float`) the application must perform an explicit state transform on the context [15].

3.2.1 Pointer tracer

Our framework supports the usage of pointers as part of data structures in the state context. In order to determine what data to include in the context transfer, we introduced debug information to the compiler. For this purpose we use the DWARF debug library [17] for GCC to mark the content of memory references according to data type. The DWARF library is able to follow pointer references down to the leaf node and extract information along its path. It also marks an already searched reference in order to avoid infinite search loops. With this library our framework is able to determine and transfer data types such as atomic variables, pointers, arrays, arrays of pointers, structs etc. without explicit information from the programmer.

4 Evaluation

To demonstrate early results, we performed a set of experiments with the runtime updating mechanism on top

of FreeRTOS. All experiments used a timer to trigger the runtime updating mechanism at a given moment.

4.1 Platforms

The evaluation was run on two embedded platforms with different architectures.

The first platform was a Versatile Express board equipped with an ARM Cortex-A9 based CA9 NEC CoreTile 9x4 quad-core chip running at 400 MHz with 1 GB of DDR2 main memory. This ARM device uses the ARMv7 architecture and was compiled with the ARM instruction set. Moreover, we created a port [1] of FreeRTOS for the Versatile Express, which was used as the underlying software platform. The second platform used in the experiments was an ARM926EJ-S Processor based on the Atmel ARM9 CPU, which was using the ARMv6 architecture. The CPU was running on 200 MHz and had 64MB of external DDR2 main memory. We mapped the runtime updating mechanism on an existing FreeRTOS port designed for the ARM926EJ-S device.

4.2 Software update

We evaluated two different outcomes of using the runtime updating mechanism: *Performance* and *functionality*. Later, we also performed overhead measurements on both platforms to determine at what cost a runtime update can be achieved. All tests were benchmarked against a baseline application (v1), which later was updated (v2).

4.2.1 ARM Cortex-A9

In the first case, a simple controller use-case was created for demonstrating the effects of runtime updating. Our system shown in Figure 3 consists of a common control loop in which a PID controller regulates the frame rate of a video used for entertainment purposes. The plant is a video player which outputs the frame rate value to the controller, which in turn strives to keep a framerate of 40 fps.




Figure 3: Structure of the video player and its controller

The system is also influenced by load disturbances which may alter the frame rate of the video randomly. The task of the controller is therefore to regulate the sleep time between the frames depending on what disturbances

are present at the moment. Initially the system used a normal PID controller (v1). This controller proved to be unable to keep the frame rate stable enough to satisfy the users. Our runtime updating mechanism was implemented to on-the-fly update the PID controller with software modifications (v2). The modifications added features such as minimum variance (MV) to the PID controller, which made it more suitable for the environment in which it is used. The state of the control system containing the I-term and D-term was, as described, stored before the update and restored after the update.

Figure 4 shows the frame rate output of both controllers, and the runtime updating is taking place at 115 samples. From Figure 4, it is clear that the new controller (v2) is more suitable to regulate the frame rate.




Figure 4: Improved stability after updating the controller

4.2.2 ARM9

In the second case, we used the ARM9 platform to demonstrate increased functionality in the application by updating the software. The application in question is a display used to output measurement values from the chip.

The first software version (v1) included two measurement values: On-chip temperature and Internal chip voltage. The software also included two derived values: Max chip temperature and average chip temperature. These values are based on previous sensor values and are therefore part of the state context.

Figure 5 shows the result of the runtime update: v2 includes more sensor values such as external temperature, and auxiliary voltage levels. It also displays all temperatures in both Centigrade and Fahrenheit degrees – a feature added by the new software. Furthermore, v2 of the software uses floats instead of integers for representing certain temperatures. This property is achieved by transferring the state to the newer version and performing explicit state conversion.

4.3 Overhead

Updating software requires a certain time to store/load data and for the linking of ELF binaries. Since the user experience is degraded if this overhead is too large [7],




Figure 5: Runtime update of display software. Version 1 to the left and Version 2 to the right

we set up experiments to determine the overhead based on different program configurations. Furthermore, we run the experiments on both the Cortex-A9 platform and the ARM9 platform in order to compare the results from different performance domains.

The overhead was measured as timing values. Timing measurements in FreeRTOS is derived by measuring the OS ticks (time stamp) at two points in the program and afterwards calculate the difference. We defined the overhead as the time difference between: *entering the safe state in the old version* and *start of task execution in the new version*. In order to compare time differences in two different task versions, we stored the time stamp from the old version as part of the context and transferred it to the new task version. The time stamp in FreeRTOS is defined as an unsigned long integer. The base line for measuring the overhead was a zero size transfer including only the time stamp from the old version. After this reference measurement, we conducted experiments with two settings: 1) *one array with different sizes* and 2) *Different amount of single variables*. All data types in the arrays and variables in the experiments were double precision floats.

The first experiment with one array of different sizes is presented in Table 1. As seen in the table, the overhead from increasing the array size does not influence much since the pointer tracer does not trace every single element, but only the array itself.

ARM9							
Array size	3	100	1k	10k	20k	50k	100k
Time [ms]	56	57	57	61	65	79	99

ARM Cortex-A9							
Array size	3	100	1k	10k	20k	50k	100k
Time [ms]	11	11	11	12	13	15	18

Table 1: Overhead of updating software with 1 array of context of different sizes

The second experiment used different numbers of single variables as context and is presented in Table 2. By introducing small but many variables (such as in Table 2), the overhead increases more and reaches 5 seconds

ARM9

Nr. Var.	10	100	1k	2500	5k	10k
Time [ms]	61	104	537	1311	2579	5162

ARM Cortex-A9

Nr. Var.	10	100	1k	2500	5k	10k
Time [ms]	12	15	51	112	215	426

Table 2: Overhead of updating software with different numbers of single variables

for the ARM9 in the case of 10000 variables. While this overhead is significant, the usual defined context in a task seldom includes this huge amount of variables. Table 2 also shows that a task including up to 100 variables for defining the context introduces an overhead of only 104 ms for the ARM9, which is hardly noticed by the user.

5 Conclusions

We have presented a lightweight runtime updating framework for FreeRTOS based systems. In contrast to related frameworks, we have focused on small size, easy annotation and no steady state overhead – this makes our framework suitable for systems with limited resources and predictability as key feature. The mechanism is able to transfer annotated context between application versions, which makes the update bump-less to the user. We have evaluated the mechanism on two different platforms and according to two of the major aspects: performance and functionality. The overhead for updating a given context has also been measured, and results indicate an acceptable overhead for the typical use case. Lastly, the increase in overhead showed similar trends on both platforms even if the Cortex-A9 run with the Data cache enabled and the ARM9 did not.

6 Future Work

We intend to further improve the componentization of software by integrating an already created component framework [12] into our system. With a proper component framework [10], [15], the safe state can be more clearly defined and e.g. for explicit communication between tasks. A more rigorous component framework can also more formally define safe states since the control of communication between components are more precise handled. A proper component framework rises the level of abstraction to the programmer and is especially useful in a multi-programmer environment.

Acknowledgment

This work has been supported by the Artemis JU project RECOMP: Reduced Certification Costs Using

Trusted Multi-core Platforms (Grant Agreement number 100202).

References

- [1] ÅGREN, D. Freertos cortex-a9 mpcore port. <https://github.com/ESLab/FreeRTOS---ARM-Cortex-A9-VersatileExpress-Quad-Core-port>.
- [2] BANNO, F., MARLETTA, D., PAPPALARDO, G., AND TRAMONTANA, E. Tackling consistency issues for runtime updating distributed systems. In *IPDPSW* (april 2010), pp. 1–8.
- [3] BARRY, R. *FreeRTOS Reference Manual: API functions and Configuration Options*. Real Time Engineers Ltd, 2009.
- [4] CHEN, H., YU, J., CHEN, R., ZANG, B., AND CHUNG YEW, P. Polus: A powerful live updating system. In *Proc. of the 29th Intl Conf. on Software Engineering* (2007), pp. 271–281.
- [5] CHRISTOPHER HAYDEN, EDWARD SMITH, M. H. J. F. State transfer for clear and efficient runtime updates. In *3:rd Workshop on Hot Topics in Software Upgrades (HotSWUp11)* (2011).
- [6] KLEINER, P. Satellite provider fixes business-critical error in orbiting satellite. *Boards and Solutions, ECE* (September 2012), 20–23.
- [7] MAKRIKIS, K., AND BAZZI, R. A. Immediate multi-threaded dynamic software updates using stack reconstruction. Tech. rep., 2008.
- [8] NEAMTIU, I., HICKS, M., STOYLE, G., AND ORIOL, M. Practical dynamic software updating for c. *SIGPLAN Not.* 41, 6 (June 2006), 72–83.
- [9] NOUBISSI, A., IGUCHI-CARTIGNY, J., AND LANET, J. Hot updates for java based smart cards. In *Data Engineering Workshops (ICDEW), 2011 IEEE 27th International Conference on* (april 2011), pp. 168–173.
- [10] POLAKOVIC, J., MAZARE, S., STEFANI, J.-B., AND DAVID, P.-C. Experience with safe dynamic reconfigurations in component-based embedded systems. In *Proceedings of the 10th international conference on Component-based software engineering* (Berlin, Heidelberg, 2007), CBSE’07, Springer-Verlag, pp. 242–257.
- [11] SEGAL, M., AND FRIEDER, O. On-the-fly program modification: systems for dynamic updating. *Software, IEEE* 10, 2 (march 1993), 53–65.
- [12] SLOTTE, R. A lightweight rich-component framework for real-time embedded systems. Master’s thesis, Åbo Akademi University, Finland, 2012.
- [13] SSU, K.-F., AND JIAU, H. C. Online non-stop software update using replicated execution blocks. In *24th International Computer Software and Applications Conference* (Washington, DC, USA, 2000), COMPSAC ’00, IEEE Computer Society.
- [14] TOIVONEN, H. T., AND TAMMINEN, J. Minimax robust lq control of a thermomechanical pulping plant. *Automatica* 26, 2 (Apr. 1990), 347–351.
- [15] WAHLER, M., RICHTER, S., AND ORIOL, M. Dynamic software updates for real-time systems. In *Proceedings of the 2nd International Workshop on Hot Topics in Software Upgrades* (New York, NY, USA, 2009), HotSWUp ’09, ACM, pp. 2:1–2:6.
- [16] WANG, T., ZHOU, X., WEI, J., AND ZHANG, W. Towards runtime plug-and-play software. *Quality Software, International Conference on Q* (2010), 365–368.
- [17] WORKGROUP, D. D. I. F. Dwarf debugging information format version 3. <http://dwarfstd.org/doc/Dwarf3.pdf>, 2005.