

Design of an OSEK/VDX-compatible System API for Linux

Study Thesis

by

Johannes Bauer

born December 6th, 1983 in Lichtenfels

Department of Computer Sciences 4
Distributed Systems and Operating Systems
University of Erlangen-Nuremberg

Tutors:

Prof. Dr.-Ing. Wolfgang Schröder-Preikschat

Dipl. Inf. Michael Stilkerich

Dipl. Inf. Christian Wawersich

Begin: 1. Januar 2007

Submission: 27. September 2007

Erklärung

Ich versichere, dass ich die Arbeit ohne fremde Hilfe und ohne Benutzung anderer als der angegebenen Quellen angefertigt habe und dass die Arbeit in gleicher oder ähnlicher Form noch keiner anderen Prüfungsbehörde vorgelegen hat und von dieser als Teil einer Prüfungsleistung angenommen wurde. Alle Ausführungen, die wörtlich oder sinngemäß übernommen wurden, sind als solche gekennzeichnet.

.....
Erlangen, den 27. September 2007

Ich bin damit einverstanden, dass die Arbeit durch Dritte eingesehen und unter Wahrung urheberrechtlicher Grundsätze zitiert werden darf. Im Falle eine Aufbewahrung meiner Arbeit im Staatsarchiv erkläre ich mein Einverständnis, dass die Arbeit Benutzern zugänglich gemacht wird.

.....
Erlangen, den 27. September 2007

Abstract

Embedded real time systems often need to be optimized for high availability and deterministic runtime- and scheduling behavior. The OSEK-OS operating system standard is quite fit for this purpose: by means of the priority ceiling protocol many actions and properties of the system are already known ahead of runtime allowing for a customized generation of the actual operating system code. For testing of functional properties of an OSEK-OS-conform operating system it is useful to test on a platform which has sophisticated debugging utilities available. A Linux system is suitable as most Linux distributions already innately include versatile tools. This study thesis will evaluate the possibility of simulation of an OSEK-OS-conform operating system and it's mapping onto a UNIX-process.

After a brief explanation of how a OSEK operating system works the developed code generator called *josek* will be introduced. The method of operation and particularities of *josek* will be discussed, paying special attention to the scheduler – the integral component of any operating system. It will be explained how a specially crafted stack is used in order to perform task switching and how this stack can be protected with userland means provided by any Linux-process. Problem cases which will appear during development of such an operating system are illuminated and their solution is presented. This includes cases where special compiler optimizations might cause malfunction of the generated code. After the study thesis has shown that and how it is possible to have functional components of an OSEK operating system emulated by a UNIX-process, the study thesis will be completed by a detailed performance review. Not only will the code generated by different configurations of *josek* be compared against itself, but it will also compare against *Trampoline*, another open source implementation of an OSEK operating system.

Zusammenfassung

Eingebettete Echtzeitsysteme müssen *per definitionem* darauf optimiert werden, eine hohe Verfügbarkeit zuzusichern und oft auch deterministisches Laufzeit- und Schedulingverhalten zeigen. Hierfür ist der Betriebssystemstandard OSEK-OS wie geschaffen: In OSEK-OS sind durch das Priority-Ceiling-Protokoll viele Eigenschaften und Vorgänge bereits zur Zeit der Erzeugung des Betriebssystems bekannt und können daher individuell optimiert generiert werden. Zum Testen funktionaler Eigenschaften eines OSEK-OS-konformen Betriebssystems ist es für den Entwickler von unschätzbarem Wert, das System auf einer mächtigeren Plattform als dem Zielsystem auszuführen, um Debugging-Vorgänge effizient ausführen zu können. Hierfür eignet sich ein Linux-System, da die allermeisten Distributionen bereits vielfältige und hoch entwickelte Debugging-Werkzeuge mitbringen. Diese Arbeit beschäftigt sich mit der Simulation eines OSEK-OS-konformen Betriebssystems und dessen Abbildung auf einen UNIX-Prozess.

Nach einer grundlegenden Beschreibung, wie ein OSEK-konformes Betriebssystem funktioniert wird detailliert der im Laufe der Arbeit entwickelte Codegenerator *josek* vorgestellt. Besondere Aufmerksamkeit gilt dem Scheduler und Optimierungen an diesem, da er den integralen Bestandteil eines jeden OSEK-Systems darstellt. Es wird erklärt, wie ein Stackaufbau und -schutz mit Mitteln, die einem Linux-Prozess zur Verfügung stehen, realisiert werden kann. Probleme, die es während der Entwicklung gab, werden vorgestellt, diskutiert und eine Lösung präsentiert. Hierzu zählen beispielsweise Compileroptimierungen, die die Funktionsweise des erzeugten Codes beeinträchtigen können. Nachdem die Arbeit gezeigt hat, dass und wie ein OSEK-OS auf einen UNIX-Prozess abgebildet werden kann, wird die Arbeit durch Messungen abgerundet. Hierbei werden entstehende Performancedifferenzen zwischen verschiedenen *josek* Betriebssysteminstanzen ebenso erörtert wie ein Vergleich mit dem OpenSouce-Projekt *Trampoline*.

Danksagung

Ich bedanke mich bei Internetpoker für die vielen Stunden, die mich meine Studienarbeit vergessen haben lassen. Weiterhin möchte ich mich bei Asli Tuna und Freya Onassis, Moderatoren des Top-Fernsehsenders „Astro TV“ für die vielen amüsanten Sendungen, bei denen ich beiwohnen durfte, wie wahrlich übernatürliches vollbracht wird, danken. Ich kann euch verzeihen, dass ihr auf mich aufgelegt habt, als ich mir meine telefonische Gratisberatung abholen wollte. Danke an DMAX für hilfreiche Tipps wie man sich ein Motorrad bauen kann oder wie man eines in die Luft sprengt. Der deutschen Grammatik, die völlig unleserliche, geschachtelte aber syntaktisch völlig korrekte Nebensätze zulässt, die wohl niemand mehr verstehen kann, der sie nicht mindestens drei mal gelesen hat, möchte ich selbstredend auch danken. Die freundlichen Herren von der GEZ, die mit ihren peinlichen Drohbriefen die deutsche Bundespost unterstützen haben sich selbstverständlich auch meinen Dank verdient – nichts für Ungut, die Sache mit dem Hausverbot bitte bloß nicht persönlich nehmen. Natürlich danke ich auch dir, meinem treuen Freund: Koffein. Danke an meine Leser, die genug Humor besitzen, um die ironischen Stellen dieser völlig sinnfreien Danksagung selbstständig zu erkennen.

Selbstverständlich gilt auch mein Dank meinen freundlichen Kollegen am Lehrstuhl für Informatik 4, mit denen man wirklich viele lustige Sachen erleben kann. Und danke euch, meine Eltern, dass ihr diesen ganzen Spaß finanziert...

Contents

1	Introduction	3
1.1	Motivation	3
1.2	Notation	4
1.3	Terminology	4
2	OSEK	5
2.1	Description	5
2.2	Priority Ceiling Protocol	6
2.3	Related Projects	8
2.3.1	openOSEK	8
2.3.2	Trampoline	8
3	josek	10
3.1	Overview	10
3.2	Modular Concept	11
3.3	Scheduler	11
3.3.1	Requirements	11
3.3.2	Work Principle	12
3.3.3	Optimization of the Queue Length	14
3.3.4	$O(1)$ Scheduler	17
3.3.5	Rescheduling Points	18
3.4	Tasks	20
3.4.1	taskdesc structure	20
3.4.2	Stack Creation	20
3.4.3	Context Switching	22
3.4.4	Stack Sharing	26
3.5	Stack Protection	26
3.6	Alarms and Counters	28
3.7	Event Handling	29
3.8	Interrupt Mapping	30
3.9	Hooks	31

3.10	Further Compiler Hints	33
3.10.1	const and pure Functions	33
3.10.2	noreturn Functions	34
3.10.3	Static Value Arrays	35
3.11	OIL Parsing and Internal Representation	36
4	Running KESO on josek	37
4.1	UNIX Integration	37
4.2	OSEK Backend	38
5	Performance Analysis	39
5.1	Test Conditions	39
5.2	Scheduling Overhead	40
5.3	Measurements	41
5.3.1	x86-64 vs. x86-32	42
5.3.2	<i>josek</i> vs. <i>Trampoline</i>	43
5.3.3	Comparison between x86-32, x86-64 and <i>Trampoline</i>	44
5.3.4	KESO	45
5.3.5	Other Points of Interest	46
6	Conclusion and Future Work	49
6.1	Conclusion	49
6.2	Future Work	50

Chapter 1

Introduction

Many embedded operating systems rely on a lower layer which handles the communication with the underlying hardware. In the operating system layer terms like "stack pointer" or "instruction pointer" are unknown – only the far higher abstraction of tasks is relevant there. Tasks are basically functions which have a certain priority assigned and which can be scheduled by the operating system. The *means* of scheduling tasks is of course changing internal CPU-registers, but this is completely hidden to the top layer.

1.1 Motivation

One of the many alternatives in choosing such a low, hardware-dependent layer for an embedded operating system is OSEK. OSEK-OS has many advantages: it is a fully static, real time capable operating system with a freely available specification and also has been standardized by ISO 17356. Its small overhead in terms of memory usage and code footprint make it predestined for use in highly embedded systems. The actual problem is to find a free implementation of such an OSEK conform operating system – which is exactly the point where this study thesis comes into place. It shall provide the necessary academic background for understanding the OSEK operating system and will implement the OSEK-OS system API. Step by step the interesting points in development will be explained and reasons of challenging design decisions will be balanced. The final goal is developing a portable, open-source OSEK operating system generator, which can be used to interconnect an operating system like *KESO* with either the Linux API or any other operating system that the OSEK implementation supports.

1.2 Notation

- A monospaced font is used for the following purposes:
 - Code snippets: "[...] achieved by the `ret` assembly command."
 - Filenames: "This is described in the `RULES` file."
 - Commands: "We will use `sed` to patch the code."
 - OSEK system calls and OSEK hooks: "The `ActivateTask` system call may invoke the `ErrorHook`."
- Parts which are printed with *emphasis*:
 - Names of software packages: "*josek* is written in *Java*."
 - Operating system task states: "The task changed from the *running*-state to *ready*."
- Citations are written with square braces: "Scheduling behavior is defined in [OSE05]."
- References to a certain OSEK task or resource are printed in French quotation marks: "Task »A« requires resource »X«."
- *Technical terms* are explained in the terminology section 1.3.

1.3 Terminology

- The time at which the actual OSEK operating system is constituted by generating C-code will be referred to as *system generation time* or short *generation time*.
- The Intel x86 compatible machine will be called x86, or more specifically x86-32 throughout this document (called IA-32 by Intel). The AMD64 extensions to the x86 which are called "Intel 64" by Intel (formerly known as EM64T) will be referred to as x86-64. They are not to be confused with the completely different IA-64 architecture.
- Switches which can be defined in *josek* during system generation time in order to affect the generated code will be called *generation defines* or short *defines*, if unambiguous.
- Although this study thesis refers to an OSEK/VDX compatible operating system, it will be called OSEK for simplicity throughout this document. This is common practice and is also done by the OSEK Group itself.[OSE04]

Chapter 2

OSEK

2.1 Description

The OSEK operating system specification describes an operating system interface which was designed for use in the automotive area. As many different vehicle distributors including Daimler-Benz, BMW, Volkswagen and Siemens Automotive called for one common, slim operating system interface, OSEK was developed to fulfill these needs. An OSEK operating system is intended for use in highly embedded systems – systems which typically provide only meager resources for the developer. As the whole system is limited in terms of RAM dimensions and/or text segment size, OSEK aims towards omitting unnecessary features which would be bloating the code.

OSEK introduces four so-called conformity classes. These are four different classes of operating systems (many of which are subsets of each other) specifying the minimum requirements the system needs to provide so it can be called OSEK compliant.

	BCC1	BCC2	ECC1	ECC2
Multiple activation of tasks	No	Yes	No	Only ET ¹
Number of tasks not <i>suspended</i>	8		16	
More than one task per priority	No	Yes	No	Yes
Number of events per task	0		8	
Number of task priorities	8		16	
Resources	RES_SCHEDULER	8 (including RES_SCHEDULER)		

Table 2.1: The Four OSEK Conformance Classes in Detail

¹ET = Extended Tasks, BT = Basic Tasks

Table 2.1 shows all these conformance classes. They are separated in two subclasses, BCC and ECC, which stand for "Basic Conformity Class" and "Extended Conformity Class". The main difference is, as shown by the table, that in the basic classes no events are allowed – all tasks are therefore basic tasks which will never transition into an *waiting*-state.

These classes are provided for one sole purpose: specialization. An BCC conforming OSEK operating system can, for example, omit any references to the *waiting*-state, memory allocation for event masks or delivered events and code resembling the event API.

These highly optimized OSEK systems can most easily be built when using a code generator. Such a generator will be developed and explained in the course of this paper.

2.2 Priority Ceiling Protocol

The priority ceiling protocol is *the* integral part of an OSEK operating system, described in [SRL90] and as part of the OSEK operating system specification in [OSE05], chapter 8.5. It is a method for task synchronization, based on static priorities. When using the priority ceiling protocol, it is ensured by design that no priority inversion or deadlocking because of different order of acquirement of resources can occur. For this to be achieved the OSEK system generator has to know in advance which task can occupy which resource or resources. This mapping is one piece of information present in the OIL² file, the central configuration file of the operating system. Each single resource is then assigned a priority – its so-called *ceiling priority*. This is the same priority as the top priority task which can occupy this resource. To make it clearer, a small example:




Figure 2.1: Tasks depending on resources

²OSEK Implementation Language

Figure 2.1 shows three tasks and three resources, arrows resembling the relationship "may occupy". Each task has its default priority, which also is defined in the OIL file. The code generator analyzes this graph and determines the ceiling priority of resources »X«, »Y« and »Z«. They are 16, 12 and 16 respectively.

When any of these tasks acquires a declared resource, a priority change happens: should the ceiling priority of the resource be higher than the current priority of the task, the task's priority will temporarily be raised to the priority of the resource.

This has one important implication which is essential to the functionality of the priority ceiling protocol: the priority of the task is higher than usual during occupation of a resource. Therefore this task cannot be preempted by other tasks which might require this resource. Hence a possible deadlock-situation is avoided when two tasks sharing the same two resources acquire these in different order.




Figure 2.2: Acquisition of Resources using the Priority Ceiling Protocol

This is illustrated by figure 2.2. The task »A« is activated while the operating system is in "Idle" state. It is immediately scheduled with its default priority (7) as there are no other tasks in state "ready". It first acquires the resource »Y« and afterwards »X« which raises the ceiling priority of the task to 12 and 16, respectively. Then the system is interrupted by an interrupt service routine, which activates task »C«. Task »C« has default priority 16, but as currently task »A« runs with this same priority also, »C« cannot preempt »A«. Only after »A« releases resource »X« its ceiling priority is decreased back to 12. The scheduler immediately switches to task »C«, as it now has top priority. After »C« is finished through the `TerminateTask` system call, »A« is scheduled again which releases its remaining resource »Y« and also terminates.

The priority ceiling protocol correlates with an OSEK operating system in two dimensions. On the one hand, special optimizations are possible when it is known in advance the priority ceiling protocol is in use – it will, for example, never occur that a task is preempted by another task of same priority. On the other hand the priority ceiling protocol can only be used when certain things like the priority of each single task are already known in advance (i.e. at system generation time). It is helpful to keep these two causal connections in mind during the development of an OSEK operating system.

2.3 Related Projects

2.3.1 openOSEK

openOSEK is a free OSEK implementation, which aims towards a slim implementation of a real-time-capable operating system for automotive applications.[ope07] The project seems to have one major problem, though: obviously, a top-down approach to the problem of creating an OSEK compliant operating system was used. This means that coding standards and preparations for OSEK-OS, OSEK-CAN, Time-Triggered-OSEK, OSEK-COM and OSEK-NM have already been taken. However, most of the code consists of stubs, no stable release has been made up to today. The SVN repository reveals that important components of *openOSEK* have not been touched in a period of half a year, although malfunctioning. The project has been bogged down in details: the developers provide commercial support, project donations, a Wiki, mailing lists, project forums, IRC channels and blogs – but no usable code.

2.3.2 Trampoline

The *Trampoline* project makes a solid impression upon its reviewer. It has the advantage of caring developers, a cleanly designed concept and frequent updates. Trampoline is split up into three major parts:

- *goil*, the OIL parser. It takes an OIL file as input and outputs a C-file and two header-files which resemble the configuration of the OSEK system. The *goil* parser also needs information about the destination platform. Currently the Infineon C167, PowerPC and Linux (through use of the `libpcl`³) are supported.[BBFT06]
- *viper*, the Virtual Processor Emulator. On a UNIX system it is implemented as an own process for target abstraction purposes. *viper* takes care of timers,

³Portable Coroutine Library

interrupts and alarms. It communicates with the actual OSEK process via shared memory and asynchronous POSIX signals.

- *Trampoline* provides the OSEK-Kernel. The *Trampoline* code is not generated (in contrast to *josek*), but still uses code-specific optimizations. For this purpose many preprocessor-statements are used.

The intermediate *viper* layer is introduced as it is relatively hard to keep the actual OSEK operating system code portable among different architectures. Code which interferes with the processor (like controlling the sleep states), switches context or programs the interrupt controller is highly hardware dependent. Therefore these pieces of code are replaced by *viper* code when compiled for UNIX emulation mode. This *viper* code is then interpreted by the *viper* daemon and appropriate actions are performed. The task code can therefore be interchanged between different architectures without any adaptations.

The minor disadvantages of *Trampoline* are its rather large memory footprint (e.g. any alarm needs 17 bytes, any counter needs 14 bytes on a 32-bit PPC [BBFT06]). Code readability is not too good because of many preprocessor-statements in kernel-code, which distract the reader of the code from the integral parts. The system is relatively slow as it uses the *viper* hardware abstraction layer.

Although *Trampoline* aims towards creating a OSEK system for academic and research purposes, it does not clearly show how efficient an highly specialized OSEK operating system can become when it is boiled down to its essential requirements. This is due to the addition of the *viper* abstraction layer.

Chapter 3

josek

All the required considerations mentioned in section 1.1 taken into account it is evident that one solution to the problem of creating a highly optimized OSEK operating system is writing an OSEK/VDX code generator. Thus, *josek* (short for "Java OSEK", pronounced /'dʒoʊ-sɛk/) was developed. Its characteristics will be illustrated in the next few sections.

3.1 Overview

josek is a *Java* application which basically does only two things: parse an OIL file, which serves as input and then generate ANSI-C-code (with minor parts of assembly) which resembles the actual operating system. The resulting C-Code can then be compiled with, for example, the `gcc` compiler and then be linked against the actual application. Generating the code via a separate program has many advantages compared to a static solution:

- The code can be highly optimized to the actual problem. It is, for example, possible to have systems which work completely without using timers or alarms. Not only will in these special cases the text-segment become smaller (as the dead code is removed by the generating engine), but the variables which would be held in RAM can also be discarded, hence saving valuable resources.
- Carefully generated code is much more readable than code which has been statically written and which relies on many `#ifdef` preprocessor statements. These are easily processed by any preprocessor, but for humans they are, especially when nested, hard to understand.
- Special features which have to be implemented in many parts of the resulting operating system (like a debugging feature or memory protection facil-

ities for example) can easily be implemented when using a code generator. The resulting code is left completely unaffected when these features are deactivated. This kind of aspect-oriented approach makes the *generating* code (i.e. the *josek* source itself) more readable as coherent parts are combined in one single location.

3.2 Modular Concept

josek uses a modular concept to generate OSEK code. There are two ways of enhancing the features *josek* provides:

1. Placing code into the `arch/` directory and adding references to these added code files in the global `RULES` file. The `RULES` file will be parsed during generation time and the referenced code will be added at the appropriate locations. Parsing of code in the `RULES` file allows only for minor and relatively static adaptations or conditions (like the presence of a certain generation define).
2. Adding a code hook: Code hooks are *Java* classes which are derived from the interface `CodeHook`. They provide much more comfortable means of altering the generating code as all internal data structures at the time of generation are known to the hook and every imaginable *Java* construct can be used to influence the result.

3.3 Scheduler

The essential and integral part of an OSEK operating system is the scheduler. Therefore careful attention has to be drawn to implement the scheduler not only correctly, but also efficiently. As OSEK uses task priorities and the priority ceiling protocol (explained in section 2.2) to avoid priority inversion there also exist more than one scheduling queue with different priorities.

3.3.1 Requirements

What the priority ceiling scheduler essentially does after its invocation is:

- Determine if there are tasks in the *ready*-state (include one task which might currently be in the *running*-state).
- If there are, determine the most important of these tasks according to their current priority.

- Schedule or dispatch the task with highest priority. Change the context from the currently running task to the new task, if applicable.

Thus, the special requirements which are necessary in an OSEK compliant operating system are:

- Honor the priority ceiling protocol. This includes calculation of resource priorities (which can be done in advance, i.e. at system generation time) and change of task priorities according to the resources held by that task.
- Handle multiple activation of tasks correctly. Should a task be activated more than once, it will also be run more than once (according to the limits imposed by the used OSEK conformity class). The order in which tasks are run has to be the same as the order of activation. Imagine a system which consists of two tasks »A« and »B« of same priority. If the activation order is »A«, »B«, »A« then the tasks have to be executed in exactly that order. Rearrangement is, according to the OSEK specification, illegal.[OSE05]

3.3.2 Work Principle

To fulfill the requirements mentioned in section 3.3.1, tasks are usually stored in a so-called scheduling queue. There are multiple queues, one for each priority level present in the system. Only tasks in the *ready*-, *running*- or *waiting* state are enqueued, *suspended* tasks are not. In an operating system with no currently running tasks, all queues are empty:




Figure 3.1: Scheduling Queue with all Tasks Suspended

When in this state a task is activated, it is inserted in the queue according to its default priority (figure 3.2).

Activation order is preserved because when other tasks are activated, they are inserted in the back of the priority queue as shown in figure 3.3.




Figure 3.2: Single Task Enqueued with Priority 0




Figure 3.3: Multiple Activation

Should a task then gain priority by the acquirement of resources, it will additionally be inserted in the queue according to the ceiling priority. In figure 3.4 task »A« acquires two resources and therefore first gains priority level 1 and afterwards priority 2. Tasks which gain priority due to resource acquirement are always entered in front of the queue. Consequently, when the resources held are released, the queue entry at the front of the queue is removed.




Figure 3.4: Priority Ceiling Protocol in Effect

Figure 3.5 shows what happens when a task of high priority is activated (task »C«) while another task is on its ceiling priority (task »A«): the recently activated

task is inserted at the back of the queue and will not be scheduled until task »A« releases its resource. Through the release, the priority of »A« decreases to 1 resulting in »C« with its default priority of 2 becoming the most important task in the system. It would therefore then preempt »A«.




Figure 3.5: Activation after Priority Gain

More graphically speaking the currently running task can be determined by traversing the queues from top to bottom and in each queue from left to right. The running task is the first task encountered not in the *waiting* state.

3.3.3 Optimization of the Queue Length

An important optimization is to predict the exact maximum length of each of these queues. It is determined by two factors:

1. The number of multiple activations of a task in its base queue (the queue with the task's default priority)
2. The number of different resources the tasks can acquire to gain a higher priority

As the number of multiple activations are limited due to the OSEK conformance classes and the resources which any given task uses are known at the time of code-generation, an exactly fitting queue length can be constructed.

Consider the following example with tasks »A«, »B« and »C« and resources »X«, »Y« and »Z«:

The scheduler has its maximal load when all schedulable tasks have been activated the number of times permitted by the used conformance class and restrictions in the OIL file of the OSEK operating system. Additionally all priority gains which are possible by the acquirement of all resources have to be calculated in this worst-case scenario.

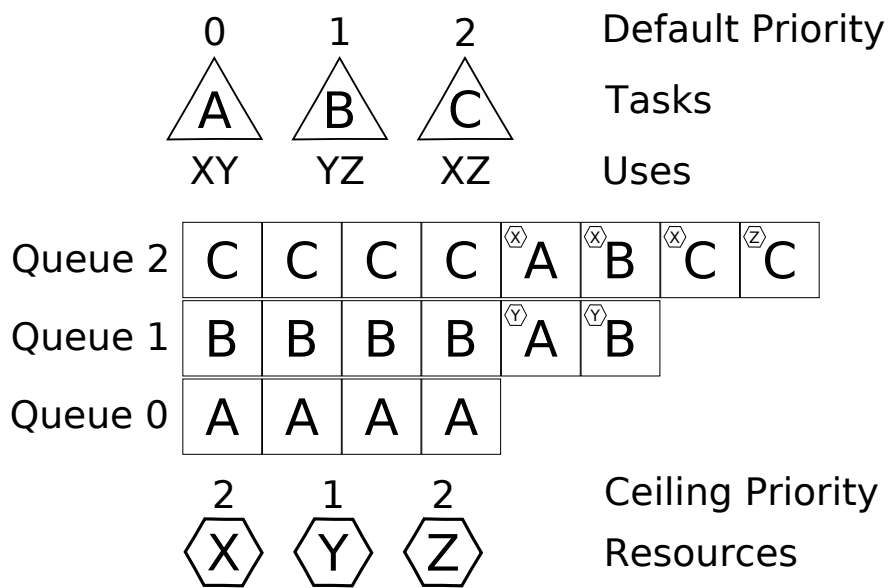


Figure 3.6: Simple Task and Resource Example

However, it soon becomes evident that there are numerous optimizations in this kind of scheme. The first, obvious optimization is possible because resources which would lead to an acquiring task's priority gain can only be occupied once:




Figure 3.7: Resources Occupied only Once

Furthermore when it occurs that there is only one task per priority level, the correct but cumbersome implementation of the scheduler can be deskilled. Multiple activation of different tasks is not possible in this scenario – the tasks need not to be explicitly scheduled by inserting an entry into the scheduling queue, but instead an activation counter can be used, as shown in figure 3.8.

Through these optimizations the scheduling queue size can be drastically reduced – in this example from 18 entries of a naive implementation down to a total of 6 entries.

Now that is known in advance how much memory the entries consume in the worst case it has to be determined how to store these entries: as a linked




Figure 3.8: Handling Multiple Activation by a Counter

list or an array. Any naive implementation would probably use the linked list as a solution, as the linked list is the standard implementation of queues in which many queue/dequeue operations occur. However there is a memory space penalty associated with this approach: taking a look at, for example, the AVR architecture it becomes clear that pointers are quite memory consuming – they’re “expensive”. Any pointer on an AVR consumes two bytes of RAM – one for each entry in the list. Considering the minimal size of 6 entries in the previous example still 12 bytes of memory would be wasted as data storage overhead. This is no acceptable implementation for an operating system which claims to primarily target highly embedded systems.

There is a possibility of reducing the problem of memory consumption. Consider the following data structure:

```

1 struct bbqueue {
2     unsigned char write, size;
3     TaskType tids[QUEUE_LENGTH];
4 };

```

In this data type resembling a queue no pointers are used. Instead there is an index `write` denoting the next element in the `tids` array which is not occupied. Additionally there is a counter `size` which yields the number of elements currently stored in the queue. Using this type of bounded buffer ensures that the overhead stays small (two bytes in total), yet it has little cost of inserting elements at the queue’s ends:

```

1 void insert_front(struct bbqueue &q, TaskType new) {
2     q->tids[(q->write - q->size - 1 + QUEUE_SIZE)
3             % QUEUE_SIZE] = new;
4     (q->size)++;
5 }

7 void insert_back(struct bbqueue &q, TaskType new) {
8     q->tids[(q->write)++] = new;
9     (q->size)++;
10 }

```

Removing elements from the queue's ends performs even better:

```
1 void remove_front(struct bbqueue &q) {
2     (q->size)--;
3 }

5 void remove_back(struct bbqueue &q) {
6     (q->write)--;
7     (q->size)--;
8 }
```

The only operation which would be costly is inserting or removing elements from the middle of the queue. This will, however, because of the nature of the scheduling queues in an OSEK conforming operating system, never happen. The data structure is therefore very well suited for use as a OSEK scheduler queue.

3.3.4 $O(1)$ Scheduler

The currently presented and optimized scheduler has a performance of $O(n)$ with n being the number of scheduling queues. This can be further reduced on some machines to have a constant cost of $O(1)$. As it is somewhat machine specific (although many architectures provide the necessary armamentarium) this optimization has not made it into the final *josek* code. As it would be a nice improvement, however, it will be discussed.

The following loop, which triggers the $O(n)$ cost, is embedded in the generated code:

```
1 for (priority = NUMBER_PRIORITIES - 1; priority >= 0;
2     priority--) {
3     /* Check if queue is nonempty */
4     [...]
5 }
```

This can be simplified in the following way: for each enqueueing process, set a bit in a global variable according to the queue number. After each dequeueing which empties the queue, clear this bit in the same global variable. The code to perform this:

```
1 /* Enqueue*/
2 active_queues |= (1 << queue_nr);

4 /* Dequeue */
5 if (!queuesize[queue_nr]) active_queues &= ~(1 <<
6     queue_nr);
```

When this is done, the queue selection algorithm can be replaced by a variant showing far greater performance:

```
1 if (!queue_nr) return;          /* All queues empty */
2 asm("bsr %1, %0\n" : "=r"(priority) : "r"(queue_nr));

4 /* Highest priority queue is now held by "priority" */
5 [...]
```

So by using the special `bsr` (Bit Search Reverse) operation [Int07b] the performance of searching all queues can be significantly improved. When the number of available priority queues exceeds 32 (or 64 on x86-64, respectively) however, an additional check is necessary and could be implemented like:

```
1 highqueue = 3;
2 for (i = 2; i >= 0; i--) {
3     if (queue_nr[i]) highqueue = i;
4     break;
5 }
6 if (highqueue == 3) return;      /* All queues empty */
7 asm("bsr %1, %0\n" : "=r"(priority) : "r"(queue_nr[
    highqueue]));
```

The presented code would allow for up to 96 different priority levels. Although it contains a tiny loop (which will almost certainly become unrolled by any good compiler) it still is a great win over the naive solution.

3.3.5 Rescheduling Points

When designing an OSEK-OS compliant operating system it is important to investigate exactly at which points in time rescheduling is necessary – and why. These are the necessary rescheduling points with a full preemptive scheduling policy, according to section 4.6.1 of [OSE05]:

- **Schedule**: A manual rescheduling point obviously occurs when the `Schedule` system call is invoked.
- **ActivateTask**: When a new task has a transition to the *ready*-state, the OSEK operating system may determine that it is more important than the currently running task. The context has to be changed to the recently activated task.
- **SetEvent**: Upon setting an event a currently blocked task might change from the *waiting*-state to *ready*. If this task is of higher priority than the currently running task, once again the context will change.

- **WaitEvent:** Analogously to `SetEvent` a task may go into the *waiting*-state after executing the `WaitEvent` system call. It won't, of course, if the event mask has already been set. This is done in accordance with section 4.6.1 of [OSE05]. But when it does, another task has to be scheduled.
- **ReleaseResource:** When releasing a resource the currently running task might lose of priority as the resource's ceiling priority might have been higher than the priority of the task prior to its acquirement. It may then happen that a switch occurs to another task.
- **TerminateTask:** After the currently running task has changed to the *suspended*-state the operating system must determine which task to schedule next.
- **ChainTask:** When the current task is exiting and a switch to another task is requested, the scheduler is invoked. However, `ChainTask` requires no ordinary rescheduling point. In fact, the OSEK-OS specification states (section 13.2.3.3 of [OSE05]) that a call of `ChainTask` will have *immediate* effect on the state of the requested task instead of resulting in multiple requests. Consider a currently running task »A« has performed an `ActivateTask` call on a task »B« having the same priority. Afterwards it calls `ChainTask(A)`. The next running task will be »A«, not »B«.
- **Initialization:** After the system initialization and entering of the idle-loop the scheduler needs to be invoked to bootstrap the whole system.

OSEK provides the possibility of configuring a system not fully preemptive (i.e. preemption may occur at any given point in time), but restrict the policy to explicit rescheduling points. This policy is then called non-preemptive scheduling and is in compliance with section 4.6.2 of [OSE05]. It has the same rescheduling points as above with the following three exceptions:

- `ActivateTask`
- `SetEvent`
- `ReleaseResource`

If rescheduling is desired after a call to one of these functions, a manual `Schedule` has to be invoked afterwards.

3.4 Tasks

3.4.1 `taskdesc` structure

For each task in a *josek* system, one array entry in the global `taskdesc` structure is reserved. This structure stores the following values:

- **const void** *stackbase: The base address of the stack, i.e. the stack pointer which is loaded at stack initialization.
- **void** *stack: The current stack pointer at the time of a context change.
- **const void** (*launch)(): The function pointer of the task.
- **enum** TASKSTATE state: The current state the according task is in (one of *suspended*, *ready*, *running* or *waiting*).
- **const PriorityType** defaultprio: An integer value indicating the task's default priority (when it does not occupy any resources).
- **unsigned char** numbertasks: The number of task activations. Used for determining if the limits of the conformance class are exceeded and for optimization as shown in section 3.3.3.
- **EventMaskType** event_mask, event_wait: The event mask of events set for that task and the event mask of events the task is waiting on.

3.4.2 Stack Creation

Before changing the stack pointer to a new memory area it has to be assured that this area has been properly initialized. The piece of memory which the stack pointer will point to after its modification is a global array. It has to be well-prepared so the assembly instructions which are executed after its change find sane values to work with. These instructions are for an x86-32 machine `pop %ebx/%esi/%edi/%ebp, popf` and `ret` in order of execution.[Int07c] Hence there have to be 16 bytes for the `pop` instructions, 4 bytes for the `popf` instruction which pops the `EFLAGS` register from the stack and 4 bytes for the return address.[Int07b] Therefore the complete context size (accessible through the variable `CTXSIZE`) is 24 bytes. Further stack analysis for architectures different from x86-32 will be shown in section 3.4.3.

Note that during this context change the contents of the floating point registers and floating point control values are *not* saved – a context change from code which uses the FPU to another context which uses the FPU is therefore not possible without a race condition.

The reasoning behind this is the following:

1. *josek* has been developed having systems with tiny resources in mind. These embedded systems often do not have any FPU at all – floating point support seems unnecessary for a real-world embedded-system scenario.
2. The context of the FPU on an x86-32 machine has a size of 108 bytes. This is almost five times the size of a "regular" context.
3. As the FPU context contains control registers which have to be set correctly the creation of an initial stack becomes more difficult. This would be a detail that attention would have to be paid to and it would therefore distract any reader of the code from the important design aspects. In particular the FPU control word would have to be set to 0x37f in order to guarantee a default FPU initialization.[Int07a]

Should an implementation of an FPU stack save be necessary after all this can easily be achieved afterwards by insertion of two assembly statement for each context saving and restoration. The x86-32 architecture provides two assembly instructions, which take pointers to allocated memory as parameters and save or restore the FPU context to or from there, respectively.[Int07c] The code to be inserted would be

```
1  sub $108, %esp          frstor (%esp)
2  fsave (%esp)          add $108, %esp
```

This code will reserve space on the stack and store the FPU registers there ("pushfpu") or load the FPU registers from the stack and free this stack space ("popfpu").

Moreover, not only the FPU registers are not saved, but none of the extended x86-operation registers are. This includes the registers provided by special CPU features like MMX, MMXExt, SSE, SSE2, 3dNow and 3dNowExt. The reasoning is obviously the same as mentioned before: these instructions are not only overkill for any embedded application, they're also absolutely machine-dependent. This is not the intended use of *josek*. All these registers (including FPU) could of course also be saved by using the `fxsave` and `fxrstor` mnemonics. This would, however, increase the size of the additional overhead from 108 bytes to 512 bytes, resulting in an extremely expensive context change.

But there are more solutions than just always saving the FPU registers:

- The compiler can be instructed to generate soft-FPU instructions. When using the `gcc` this can be achieved by the `-msoft-float` command line

switch. The generated FPU code is much slower than hardware FPU operations, but it is portable and has no problems with preemption, as it uses integer registers only.

- A resource »FPU« could be defined. Tasks which would want to use FPU instructions would have to declare this usage in the OIL file. Any FPU operations then become critical sections which have to be protected by `GetResource/ReleaseResource` calls.
- It is possible to program the CPU to deliver traps in case of FPU usage after a context change. This is done on an x86-32 system by setting the *Monitor Coprocessor* bit of the processor control register 0 (`%cr0`). Any FPU instruction will then trap when the *Task Switched* bit of the same register has been set. This is done specifically to allow for "lazy" context switching, which does only save the FPU registers on demand (i.e. when realizing that another process also wants to use the FPU). The problem is that in user space modifications of the `%cr0` processor register are usually disallowed by the operating system, rendering this approach impracticable.

3.4.3 Context Switching

Method of Operation

An interesting internal part of any operating system is the method of context switching which is used by the scheduler. There are two possible context switching methods available which differ in functionality:

- `dispatch`: Starts a task for the first time. The stack pointer set to a piece of memory which was specially crafted for stack dispatching (explained in section 3.4.2). The registers and CPU flags are popped from the stack – these need to be set to zero to avoid accidentally setting flags which might trigger traps. Afterwards a `ret` is performed which jumps to the address on very top of the stack. This is the function pointer of the task which is dispatched.
- `switchtask`: This will switch from a currently active task to another task. In order to be able to resume the currently running task appropriately, the first action is saving the processor flags and registers by pushing them onto the stack. Afterwards the current stack pointer is copied into the global `taskdesc` structure. The stack pointer of the task to be switched to is then restored from the `taskdesc` structure. CPU flags and registers are popped from the stack and a `ret` is performed, returning to the exact position the task was previously preempted from.

What registers exactly need to be saved by the callee is dependent on the architecture used and the ABI¹. Table 3.1 lists the differences in context content among architectures:

Architecture	Content of Context	Total Size (Bytes)
x86-32	%ebx, %esi, %edi	12
	%ebp	4
	EFLAGS	4
	Return Address	4
	Sum: 24	
x86-64	%rbx	8
	%r12, %r13, %r14, %r15	32
	%rbp	8
	RFLAGS	8
	Return Address	8
Sum: 64		
AVR	%r2-%r17	16
	SREG	1
	Return Address	2
Sum: 19		

Table 3.1: Content of Contexts

The calling conventions of different architectures which have been taken into account in order to create a minimal context size can be found in literature.[Fog07][Gud07]

Pitfalls

Manipulating the stack pointer directly is hazardous business: the host operating system does not condone any off-by-one-errors; should the stack pointer which is restored be off by a single byte relative to the correct position it is likely the entire system will crash. In such a case the accompanying stackframe is lost implying development tools like debuggers are rendered useless.

Then there are more subtle and compiler dependent pitfalls. Any function call which occurs at the very end of a function can be optimized away and replaced by an unconditional jump. This is called a tail-call optimization.[Sch07] Essentially the following code:

```

1 dostuff:
2     push %rbp           ; Enter

```

¹Application Binary Interface

```

3     mov %rsp, %rbp

5     [...]

7     call finalize

9     mov %rbp, %rsp    ; Leave
10    pop %rbp
11    ret

```

Will be replaced by this optimized version:

```

1  dostuff:
2     push %rbp          ; Enter
3     mov %rsp, %rbp

5     [...]

7     mov %rbp, %rsp    ; Leave
8     pop %rbp
9     jmp finalize

```

The `ret` instruction in the function `finalize` will therefore directly return into the parent function of `dostuff` instead of first returning into `dostuff` and from there on returning to the caller. The `dostuff` stackframe is bypassed this way – CPU time is saved.

It's a horrible optimization scenario for anyone directly manipulating the stack pointer, however. The reason is easy: it may be the case that `dostuff` has no parent, for example when it is executed on a specially crafted stack – as it is exactly the case in *josek*. When the optimization is activated the `ret` will jump to some undefined address and the whole operating system will in all probability crash.

This is the case when using any recent version of *gcc* and using more than `-O1` optimization. As soon as `-O2` the optimization flag `-foptimize-sibling-calls` is implied resulting in the described optimization. The naive solution would be to pass `-fno-optimize-sibling-calls` to the compiler flags of the task source file. This is extremely simple and effective, it has the disadvantage of slowing down other code, however, which is unaffected by the optimization.

josek uses a more sophisticated approach. First, all fragile functions need to be determined – by that is meant all functions which are broken by the tail-call-optimization. These are all functions which have no parent (i.e. no caller): it solely applies to tasks. Which functions tasks may call at their tail is the next target to identify. Only few are relevant:

- ShutdownOS
- TerminateTask
- ChainTask

Calling any other function at the end of a task is illegal according to the OSEK-OS system specification and therefore undefined behavior – the operating system crashing – is a perfectly legal consequence. ShutdownOS is uncritical as it will never return at all. Any optimization is fine here.

But things are different with ChainTask or TerminateTask. ChainTask also works with the optimization as long as *another* task is chained. Should a task chain itself, a return into the current context is needed and the `ret` will jump into nirvana. TerminateTask is also fine as long as the next task to be scheduled is not currently active. Should it be, maybe through multiple activation, will it crash in exactly the same way.

The problem can be avoided by a trick: first, the system calls ChainTask and TerminateTask are renamed to `__ChainTask` and `__TerminateTask`. Then two macros of the following kind are inserted into the `os.h` file:

```

1 #define TerminateTask() do { __TerminateTask(); \
2                               __asm__ volatile(""); \
3                               } while(0)

```

This will cause any call to TerminateTask to be replaced by its actual call and an (empty) volatile assembly statement afterwards. The compiler detects code after the TerminateTask call – the `asm` statement. Actually there is *no* code, but we marked the inline assembly volatile so the compiler must assume the worst case and is not allowed to move the code away either. It will therefore just do a plain call to TerminateTask and the problem described vanishes. The exact same approach applies to the ChainTask function which will also solve the described difficulties there.

The OSEK specification states, however, that TerminateTask and ChainTask shall have a return value – although they obviously *never* return.[OSE05] Therefore, applications which rely on a value being returned by those function do not compile with the above code. For such applications, a different workaround has to be used where the whole TerminateTask and ChainTask functions are programmed in assembly and essentially just call `__TerminateTask` or `__ChainTask`, respectively. This portable approach has been used in *josek* although the method described above would have greater performance – if only the OSEK specification would not require a non-returning function to provide a return value.

3.4.4 Stack Sharing

A straightforward optimization is stack sharing. It becomes possible as a direct consequence of the usage of the priority ceiling protocol. The protocol asserts, as explained in section 2.2, that the ready task with the highest priority is in the *running*-state. As a matter of fact it can safely assumed that a task will never be preempted by another task of same or lower priority. In particular it is not possible that a task is preempted by a second activation of itself. It is the most elemental optimization that each task is assigned only one piece of memory for use as a stack in contrast to reserving a stack for each instance. Moreover it is also possible for two *different* tasks to share the same piece of memory if they are mutually exclusive. This is always possible with tasks which have the same default priority in an OSEK system without extended tasks. This shared stack has to be sized to both task's needs, of course. Its size must be determined by the more greedy (in terms of memory usage) requirements.

3.5 Stack Protection

A problem commonly found in programs which use fixed stack sizes – as this is the case in any OSEK operating system – is stack overflow. Due to the fact that the stack cannot grow dynamically according to the needs of the program it will at some point hit the boundary. This can especially occur by a too high nesting level of functions: as each function call uses some stack memory (on x86-32 at least 4 bytes for the return address of the function, usually 8 bytes because the old frame pointer `%ebp` is also saved) the problem particularly arises when using recursive functions.

This can be detected when running the OSEK operating system on Linux quite comfortably using the `mprotect` system call. When generating the OSEK-OS with the generation `define stackprotector` four things are done:

1. The stack sizes of tasks are rounded up to multiples of the page size (usually 4096 bytes).
2. Between each of the tasks' stacks a separate protector memory area is inserted which has exactly the size of one page.
3. All stacks and pages are aligned at page boundaries. This is a compiler-specific parameter. Using the `gcc` compiler it can be achieved using the `__attribute__((aligned (4096)))` extension.

4. The protector-pages are marked as inaccessible through invocation of the Linux-specific `mprotect` system call. They are marked as `PROT_NONE` which means the page has no read, write or execute privileges.

Any access beyond the allocated stack memory (but within the size of one page) will then lead to an immediate synchronous program termination (`SIGSEGV`, segmentation violation). This can be easily traced to determine where the illegal access occurred.

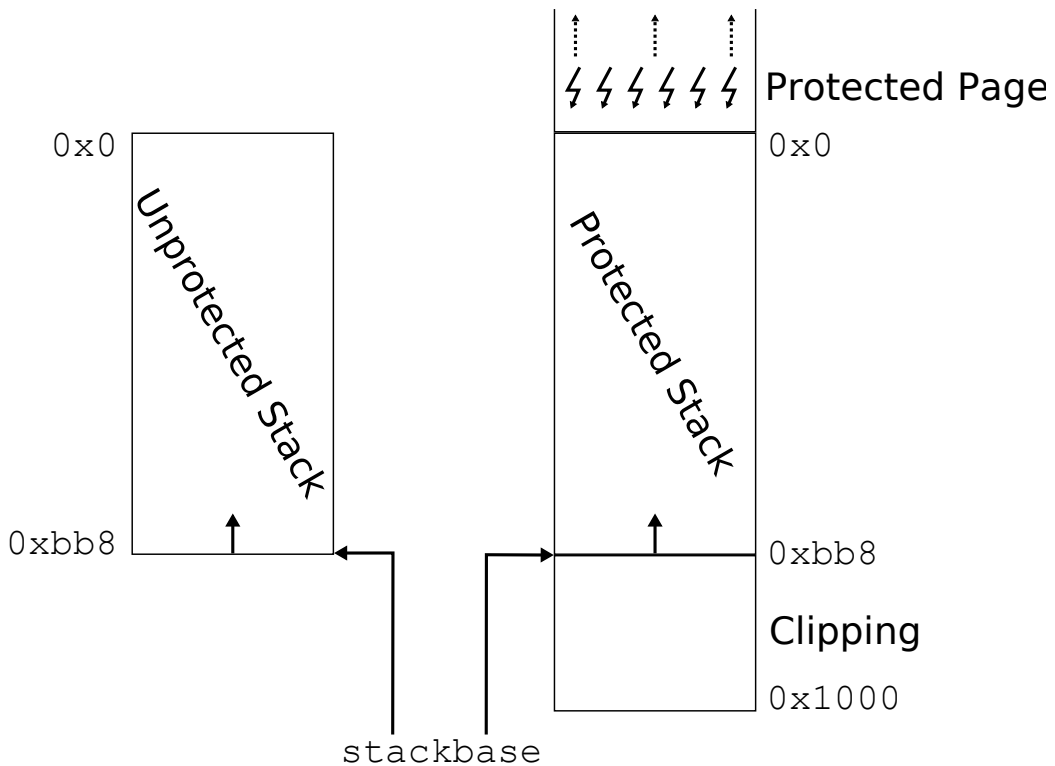


Figure 3.9: Unprotected and Protected Stack Areas

It is, of course, no exhaustive memory protection. Any task can access the whole system memory through use of arbitrary pointers. Hence, such deliberate memory corruptions are not detectable. However, it is very well suited for detection of unintentional stack overflow as is the case with a too high nesting level of functions.

The drawbacks of this type of memory protection are evident:

1. It is Linux-specific.
2. It only detects memory corruption within one page size from the top or bottom of the tasks' stack.

3. For a number of n tasks, $n + 1$ protection pages are wasted.
4. Clipping occurs, as the stack size needs to become a multiple of the page size. For example in figure 3.9 the stack size is 3000 Bytes – 1096 Bytes are clipped away and not usable.

The greatest problem probably is portability to non-Linux architectures. Therefore a "poor man's version" of the memory protection is available by use of the generation define `stackprotector_simple`. When it is in use again memory pages separate tasks' stacks, but this time they're not protected using `mprotect` but initialized with a sequence of pseudo-random numbers. As the sequence is deterministic it can easily be verified at any time through a call of the `check_stackprotector()` function. Although this method is more portable it can not detect reading of invalid memory, only writing. Another drawback is that memory corruption is only perceptible afterwards, not synchronous to the actual memory access. Debugging therefore becomes slightly more complicated.

3.6 Alarms and Counters

According to the minimum OSEK requirements (as described in section 3.3.1) any OSEK operating system needs to provide at least one alarm. Alarms are based on counters. The implementation of these counters is operating system dependent. More concretely, a process of a POSIX compliant operating system has only one possibility of mapping counters onto the superimposing operating system interface: the `SIGALRM` signal delivered by the kernel after a `alarm()` or `setitimer()` system call. `alarm()` disqualifies for the purpose of serving counter ticks as its granularity is with one second much too coarse. Therefore `setitimer()` must be used.

The man page of the `setitimer()` system call explains that any Linux process is provided three timers which vary in their functionality: one counts real time (`ITIMER_REAL`), another counts virtual time (time in which the process is scheduled in user space, `ITIMER_VIRTUAL`) and a last one which counts the time in which the process is scheduled plus the time in which the operating system performs system calls for that process (`ITIMER_PROF`).

This means that there is only one useful timer available: `ITIMER_REAL`. Therefore for any OSEK task which needs more than one alarm, a mapping of many internal alarms to this system timer has to be performed. It is realized by the *josek* counters.

Counters are handled in a simple manner: they are represented by a global variable counting every tick. The tick interval is known at system generation time and is specified in the OIL file. During generation of the OSEK operating system,

the greatest common denominator over all counter intervals is calculated. This will later on during runtime pose the interval for the `setitimer()` system call – we will refer to it as t_{min} . Each counter also has a property called `CntrTickMaxValue` which essentially is $\frac{t_{counter}}{t_{min}}$. This value will not change during runtime and can therefore be stored either in RAM or in the text-segment. Refer to section 3.10.3 for further reference on how this is done.

During the execution of the `setitimer()` handler, a tick count is increased for every counter available in the system. When the tick count reaches the `CntrTickMaxValue`, a "real" counter tick is triggered. This will cause the `CntrCurrentValue` variable to be increased, which wraps if it reaches `CntrMaxValue` – a variable also specified during generation time. Should the `CntrCurrentValue` wrap it will trigger possible alarms dependent on it.

There are three possible results in the event of a triggered alarm:

- *Callback function:* Probably the most common action after an alarm was triggered is calling a callback function. This function has to be defined in the operating system using the `ALARMCALLBACK` keyword.
- *Setting an event:* The alarm can set an event for a certain task. For this, the alarm simply calls `SetEvent`.
- *Activating a task:* If requested, the alarm will use the system call `ActivateTask` in order to activate a certain task.

3.7 Event Handling

In order to be able to deliver and receive events the `taskdesc` structure of any generated *josek* operating system can contain two values: the event *mask* and *wait* variables. *josek* automatically decides whether these values are needed and how large they need to be according to how many events are used in the system – if they are used at all.

Handling of events then is straightforward. A task clearing its event mask through the `ClearEvent` system call will directly write to the `mask` value:

```

1 StatusType ClearEvent (EventMaskType m) {
2     taskdesc[current_task].mask &= ~m;
3 }
```

When a task delivers an event, an analogous piece of code is executed, except that the bitwise `OR` is performed with the mask and a rescheduling point is introduced should the destination task have been woken up through delivery of this event (assuming a full-preemptive system has been built).

Should a task want to wait for an event through use of the `WaitEvent` system call, the *wait* value is set and the task state is set to *waiting* if this event has not yet been delivered yet, i.e. if the following condition is true:

```
1 (taskdesc[n].wait & taskdesc[n].mask) == 0
```

Should the task state be changed to the *waiting*-state a rescheduling point is also hit afterwards. The scheduler will then switch to a *ready* task (or *idle*, if none is available).

3.8 Interrupt Mapping

When mapping hardware IRQs onto a UNIX system there is basically only one choice for asynchronous userland process interruptions: POSIX signals.

This is what *josek* does: Simulating interrupts by a signal handler for `SIGUSR1` and `SIGUSR2`. As the introduction of asynchronous program interruptions always brings along concurrency close attention has to be paid that data structures which are read or modified by the ISR-handler are properly locked in critical sections. As the OSEK specification does not permit many system calls to be invoked from within an ISR this does not pose a big problem. In particular the following system calls have to be interlocked against an ISR interruption:

- `SetEvent` and `ClearEvent` have to be locked, as atomic access to the `taskdesc[n].mask` and `taskdesc[n].wait` variables has to be guaranteed.
- `Schedule` has to lock its access to the task queue as execution of the `ActivateTask` system call from within the ISR could corrupt the data structure.

Apart from these difficulties, the signal handler code simply executes the appropriate ISR handler upon decision on which signal was received.

This is the implementation of the mapping of external, arbitrary interrupt sources to a UNIX process. It might be of interest, however, to emulate device interrupts, too. It could be the case, for example, that an application opens a device driver (e.g. `/dev/can0` for a CAN-bus controller) and wishes an interrupt to be delivered as soon as the device has data available. If the device was opened in blocking mode and a `read` would be performed, the whole OSEK process would block waiting for data to arrive, effectively locking up the whole operating system. In order to solve that problem, what the code generator has to do is the following: in the OIL file file descriptors which should be monitored would be declared. At system startup, the OSEK operating system uses the `pthread_create` library call

in order to create a second thread which has the sole purpose of monitoring this file descriptor. What the monitor thread does is call the `poll` system call on the monitoring file descriptor in an infinite loop. This system call will block until data is available for reading from the file descriptor. After `poll` returns, the child uses the `kill` system call to deliver its parent thread a `SIGIO` signal. The signal handler of the parent will then call the "device interrupt" routine, which can afterwards perform the `read` on the file descriptor to read out the data which has arrived.

3.9 Hooks

The OSEK operating system specification provides a convenient means for the application developer to integrate preparation or clean-up-tasks into the environment: operating system hooks. These hooks are basically functions which are executed at special times. They are valid system wide and the times they "fire" can be one or more of the following:

- `ErrorHook`: Fires after a system call was completed unsuccessfully, but before the switch back to the task level.
- `PreTaskHook`: Fires before a task is scheduled but already in the context of the new task.
- `PostTaskHook`: Fires prior to a task being scheduled away, but still in the old task's context.
- `StartupHook`: Fires before initialization of the operating system, but before the scheduler is running.
- `ShutdownHook`: Fires when `ShutdownOS` is called to shut the operating system down.

These hooks seem pretty straightforward to implement, but the exact circumstances on how and when they are invoked make the implementing code a little bit tricky. The easiest of these calls is probably the `ErrorHook`. The only thing the generator needs to check is if the user wants an `ErrorHook` installed and insert a simple `if`-clause at the finalization points of system calls (i.e. before `return` statements) which conditionally calls the hook.

The `PreTaskHook` and `PostTaskHook` are a little bit more complicated. This is because the specification states that they both need to be executed in the context of the task to be scheduled or of the task which is currently being left, respectively. The first naive approach would probably to insert `PreTaskHook` calls at the beginning of each `TASK()` block and to insert `PostTaskHook` calls at the end of these

blocks. Thinking about this approach – parsing the C-file – will certainly yield the result that this is far more complicated than one would assume at first glance. Although insertion of the `PreTaskHook` calls is straightforward, insertion of the `PostTaskHook` calls is not – they have to be inserted *before* any `TerminateTask`, `ChainTask` or `ActivateTask` calls. And what about the user redefining some macros, which is perfectly legal? The hooks also need to be called when a task is scheduled away from or scheduled to – not just the first time at startup and at the `TerminateTask` system call.

Realizing that this approach is a dead-end road, how about calling these hooks from the scheduler? The scheduler is the part of the system which knows if a task change is necessary. This approach works well with the exception of two aspects:

- The user might call an `ActivateTask` in the `PreTaskHook`. This will again case the scheduler to be invoked. Therefore care has to be taken that the scheduler does not call the `PreTaskHook` from an inconsistent state.
- The user might want to access the task ID of the task which is scheduled to or scheduled away from, respectively. This variable has to be properly initialized in the scheduler *before* calling the hooks.

Using this method works well if attention is paid to these aspects – this is also how it is done in *josek*.

The problem with the implementation of the `StartupHook` isn't evident at first glance, either. It looks as if it would be as simple as inserting a call to the hook upon operating system initialization. And it is, almost. Consider a fully preemptive system which has two tasks called »subsidiary« having priority 1 and »important« having priority 10. Then there is the following `StartupHook`:

```
1 void StartupHook() {
2     ActivateTask(subsidiary);
3     ActivateTask(important);
4 }
```

When this piece of code is executed, the task with low priority, »subsidiary« is scheduled first, leaving »important« waiting for its turn although it has higher priority. It seems acceptable, as when you come to think of it »subsidiary« is activated *before* »important«. A closer look into the OSEK operating system specification will reveal, however, that the scheduler has not yet been activated when the `StartupHook` is called. It will only become active after *both* `ActivateTask` system calls have finished – the scheduler will then of course schedule »important« first and priority inversion is avoided.

The `ShutdownHook` is as straightforward as it seems: a simple call to this hook in the body of the `ShutdownOS` system call is fully sufficient.

3.10 Further Compiler Hints

There are function attributes which can be used to make the generated code perform better than usual. Some are compiler specific, however, and will, if so, refer to recent versions of the *gcc*.

3.10.1 `const` and `pure` Functions

The *GNU Compiler Collection* can make use of the `const` function attribute, which tells the compiler that the attributed function's return value will depend solely on the parameters of the function.[StGDC07] It will never depend on the state of the program (i.e. not read global variables). If the compiler knows this about a function, it is a safe optimization to call the function less often than actually required by the source code. In practice this means that subsequent calls to the function in question will always yield the same return value as long as they get the same input. This allows the *gcc* to optimize these subsequent calls away. A good example of a function which should definitely be declared `const` are lookup tables which are stored in the text segment as described in section 3.10.3.

To demonstrate the effect, consider the following piece of C code:

```
1 printf("%d\n", CntrMaxValue(x));
2 printf("%d\n", CntrMaxValue(x));
```

It actually requires two calls to the `CntrMaxValue` function. So without having this function declared `const`, the following code will be generated:

```
1 movzbl %spl, %ebx

3 mov %ebx, %edi ; 1st Parameter
4 callq 400eb1 <CntrMaxValueValues>

6 movzbl %al, %esi ; Value
7 mov $0x40129d, %edi ; String reference
8 mov $0x0, %eax
9 callq 4006c0 <printf@plt>

11 mov %ebx, %edi ; 1st Parameter
12 callq 400eb1 <CntrMaxValueValues>

14 movzbl %al, %esi ; Value
15 mov $0x40129d, %edi ; String reference
16 mov $0x0, %eax
17 callq 4006c0 <printf@plt>
```

It can be clearly seen that the call to `CntrMaxValueValues`, which is the function that the `CntrMaxValue` macro expands to, is called twice. It takes `%edi` as input and returns its value in register `%al`. This needs to be done twice, as the `gcc` does not the function will return the same value each call for same values of `%edi`. When the optimization is active, however, the following code is generated:

```

1  movzbl %spl, %edi                ; 1st Parameter
2  callq 400eb1 <CntrMaxValueValues>

4  movzbl %al, %ebx

6  mov   %ebx, %esi                ; Value
7  mov   $0x40129d, %edi           ; String reference
8  mov   $0x0, %eax
9  callq 4006c0 <printf@plt>

11 mov   %ebx, %esi               ; Value
12 mov   $0x40129d, %edi          ; String reference
13 mov   $0x0, %eax
14 callq 4006c0 <printf@plt>

```

The function `CntrMaxValueValues` is also called, but this time only once. Its return value `%al` is saved to register `%ebx` from where it is passed on to the two `printf` function calls.

Likewise, the function attribute `pure` can be used. It is a little weaker than the `const` attribute and essentially expresses that the function in question does not only take the arguments into account, but may also read (not modify!) the content of global variables. Functions which can be declared `pure` for the compiler to allow for optimization are `GetEvent` or `GetTaskState` for example. They are so tiny, however, that it might be even better to define them as macros to force inlining:

```

1  #define GetTaskState(id, state) ((*state)) = taskdesc[
      id].state, E_OK)

```

3.10.2 noreturn Functions

There are functions in any OSEK operating system which will never return. These are:

- ShutdownOS
- ChainTask

- `TerminateTask`

When the compiler knows a function will never return, it can optimize away any cleanup-work which would be usually be done. This means, for example, the following code:

```

1  00000000004007c0 <JOSEK_TASK_job5>:
2  4007c0: 48 83 ec 08      sub    $0x8, %rsp
3  4007c4: 31 c0           xor    %eax, %eax
4  4007c6: e8 55 07 00 00  callq 400f20 <__TerminateTask>

6                                     ; Free stack and return:
7  4007cb: 48 83 c4 08      add    $0x8, %rsp
8  4007cf: c3             retq

```

Will get optimized in a fashion that the two instructions after the call to `TerminateTask` are simply dropped.

Note that this optimization is not always usable because of a strange reason: according to the OSEK specification, `ChainTask` and `TerminateTask` shall provide a return value.[OSE05] Although it is not clearly evident why functions which never return should need a return value, the OSEK specification nevertheless states it is necessary. The method described above therefore applies only to systems which never use this imaginary return value.

3.10.3 Static Value Arrays

During development it will occur that a table of static values has to be accessible from inside *josek*. Simply using a global array of values has the advantage of easy implementation and fast lookup. However the lookup is fast at a price: the values are kept in RAM. This might pose a problem to embedded systems which have limited resources. Especially when these values are never changed during the course of the program, preserving RAM space is one big goal which needs to be achieved when writing an operating system for highly embedded devices. Although of minor importance when testing functional components of the OS on an x86 with hundreds of megabytes of RAM it is a real drawback when using a microcontroller unit which only provides 128 bytes of SRAM.

josek has a solution: during the generation of the operating system's code a generation define (`saveram`) will decide how static arrays of values are stored: without any optimization as a global array or as a lookup-table. The lookup-table basically is a lookup-function which has the sole purpose of running input data through a switch/case statement which determines the appropriate return value.

The penalty is evident: a function call is more costly in terms of CPU time than a memory lookup. Registers have to be saved and the lookup of an arbitrary

value in an equally distributed and compiler-optimized switch/case statement is of order $O(\log n)$ as opposed to $O(1)$ when using a global array. What is preferred has to be decided from case to case – it is just one opportunity *josek* provides. The last decision has to be made by the developer who generates the operating system.

3.11 OIL Parsing and Internal Representation

For parsing the input OIL configuration file, code generated by *JavaCC*² is used. It is based on a `.jj` file resembling the OIL grammar. For this purpose a modified version of the KESO grammar was used – actually the OIL grammar is a bit easier to implement as fewer recursions of objects are permitted.

The parser code generates – granted the OIL file is syntactically correct – a tree according to the input data. This tree is represented by the *Java* class `Configuration`. It can differentiate different attribute types such as boolean, integer, float or string. The OIL configuration file is bijectively mapped onto this configuration-tree. After the parser is done, the raw first pass is complete.

Afterwards, the configuration data needs to be refined and a so-called `HighLevelConfiguration` object is created. This object parses the simple configuration tree for *logical* objects. It can, in contrast to the simple configuration, also detect logical errors such as a task referring to an undefined resource. In this parsing step, all logical objects are parsed specifically according to their jobs in the OSEK operating system to be generated. These objects are:

- Alarms
- Counters
- Events
- Hooks
- Resources
- Tasks

Each single of these objects has very specific member variables. Settings which have meaning to the high level configuration (like the `CYCLETIME` of an alarm, for example) are taken from the simple configuration, values without meaning are simply discarded without emitting an error or warning.

²Java Compiler Compiler

Chapter 4

Running KESO on josek

4.1 UNIX Integration

In order to have the possibility to access UNIX device drivers from within *KESO* it is necessary to map some UNIX system calls onto a *KESO*-interface. This is why the `PosixIO` class was added to *KESO*. It provides the following functions, of which *Java* prototypes are shown:

- `int open(String path, int how)`
- `int read(int fd, char[] buf, int size)`
- `int write(int fd, char[] buf, int size)`
- `int fcntl(int fd, int cmd, long arg)`
- `void close(int fd)`

Making use of these functions which are directly mapped onto the corresponding UNIX system calls makes it possible to open a UNIX device and write or read data from it. This was tested using a Linux-compatible CAN controller. The CAN controller provides a character device (`/dev/can0`) for bus communication. When a *KESO* program wants to send data over the bus it simply writes a specially crafted structure to this character device. As soon as the write is complete, the CAN bus packet is sent over the bus. Reading works by a polling `read` system call on the driver device. If there is no data available, the read call will not block, but return unsuccessfully immediately, indicating by its return code no data was available. For simulating interrupts upon arrival of data over the bus the concept described in section 3.8 may be used.

4.2 OSEK Backend

KESO uses a far higher abstraction level for tasks than a pure OSEK operating system does. They are represented in a *Java*-style manner as classes instead of "raw" functions. Each class provides a `launch` method, which contains the actual code executed upon task activation. Management of activation and termination is performed by the `TaskService` class. It provides several methods of which the functionality in respect to an OSEK operating system is obvious:

- `Task getTaskByName(String taskName)`
- `int activate(Task taskID)`
- `int chain(Task taskID)`
- `int terminate()`

As tasks are represented by classes, an activation of a class by its name is more expensive than it would be in "raw" OSEK code. The raw OSEK implementation uses `enums` to represent task IDs, which are reduced during compile time to constant values. In *KESO*, however, the lookup of a `String` to the `Task` class has to be performed first (i.e. before calling `activateTask` or `chainTask`). An exception to this rule is, of course, a task referring to itself using the `this` class variable.

Furthermore for use of events or resources two classes named `EventService` and `ResourceService` are provided which also have APIs that are easily mappable onto an OSEK conform operating system. They behave likewise and will not be further explained.

Chapter 5

Performance Analysis

5.1 Test Conditions

For testing purposes two systems were used, a Pentium M 1400 MHz for all x86-32 measurements and an Athlon64 3700+ for x86-64 measurements. As both systems are hardly comparable, there is an additional variable which scales the timings down. This variable has been determined by a small assembly routine essentially executing a lot of no operation (`nop`) operations in a row and counting how many can be executed per second. This `BogoMIPS` value is $2.517 \cdot 10^9$ for the x86-32 and $5.973 \cdot 10^9$ for the x86-64 architecture. The reference architecture will, in all cases, be the x86-64. So when any x86-32 timings are normalized, they will be divided by the factor of $f = \frac{5.973 \cdot 10^9}{2.517 \cdot 10^9} \approx 2.373$, when the throughput (MOPs per second) of any function will be normalized, it will be multiplied by f .

To measure the throughput all functions under evaluation include a counter to a global variable. This counter is incremented every time a task unit is completed (e.g. for every `ChainTask` call or for every `ActivateTask/TerminateTask` combination). After a period of at least 100 seconds the process was interrupted via an asynchronous POSIX signal. The signal handler did an output of the counter value and terminated the whole process.

In order to also be able to measure the timings of single functions as is the case in the `GetResource/ReleaseResource` scenario the following piece of assembly code was used:

```
1  xor %eax, %eax
2  cpuid

4  rdtsc

6  mov %rax, %0
7  shl $32, %rdx
```

```

8  or %rdx, %0

10 xor %eax, %eax
11 cpuid

```

This function, which was always inlined by use of `#define` statements and guarded by `__volatile__ __asm__ inline` assembly, does the following:

- Clear the CPU instruction pipeline by a call of `cpuid`.
- Read the time stamp counter the CPU provides, store it in the given variable.
- Clear the instruction pipeline again.

The instruction pipeline has to be specifically cleared as the position of the `rdtsc` opcode may be arbitrarily rearranged within the pipeline. It therefore leads to inconsistent and strange results if not performed. The whole function – when executed twice to determine the evaluated function’s runtime – took a total of 109 CPU cycles on the reference machine. These 109 cycles were always subtracted.

Note that these measurements vary because of jitter – the 109 CPU cycles are therefore only an average. This jitter occurs because the OSEK process is not the only process in the Linux system at the time. Although it uses 99% of the CPU time, some time is spent in the Linux kernel and the OSEK process is occasionally scheduled away in favor of another process. These averaged 109 cycles are therefore the closest approximation accomplishable with user-space means on a Linux system.

5.2 Scheduling Overhead

In order to schedule between tasks, their specific contexts have to be saved and restored. The sizes of these contexts are compared in detail across different architectures in table 5.1:

Types	Machine Words			Bytes		
	x86-32	x86-64	AVR	x86-32	x86-64	AVR
General Purpose Registers	3	1	16	12	8	16
Extended Registers	0	4	0	0	32	0
Pointer Registers	1	1	0	4	8	0
Processor Flags	1	1	1	4	8	1
Return Address	1	1	2	4	8	2
Sum	6	8	19	24	64	19

Table 5.1: Scheduling Overhead in Terms of RAM Size

Note that the only architecture where the pointer size is unequal to the machine word size is the AVR, where a pointer (like the return address or stack pointer) occupies two bytes compared to a one byte machine word.

It becomes clear the AVR is not the ideal architecture for multitasking, especially since it has rather large overhead (19 bytes) compared to common SRAM sizes (usually ranging from 128 bytes up to 1024 bytes).[Atm07]

5.3 Measurements

In these following tests points it has been evaluated how much time it consumes for a task or set of tasks to perform the following actions:

- A task repeatedly calling `ChainTask` on itself.
- A task repeatedly calling `ActivateTask` on itself, afterwards calling `TerminateTask`.
- A task acquiring and releasing a resource in an infinite loop.
- Two tasks, »A« and »B« send each other events. In an infinite loop, task »A« signals »B« an event, which »B« has been waiting for. Afterwards, »A« calls `WaitEvent` to wait for an event »B« will send. Task »B« does likewise, it performs analogously.
- Task »A« activates a task of higher priority, »B«, which will immediately preempt »A«. »B« terminates immediately. Afterwards, »A« calls `ChainTask` on itself.
- A test like the one before, except that task »B« will first activate a third task »C«, which is of higher priority than »B«. »C« will immediately terminate.
- Another test like the one before, except that »C« activates a task »D« of higher priority before terminating. »D« immediately terminates.
- A task calls `Schedule` in an infinite loop.

5.3.1 x86-64 vs. x86-32

For the tests described in section 5.3, the following results were determined on the x86-32 and x86-64 platform:

	x86-32		x86-64	
	Time	n/sec	Time	n/sec
ChainTask (self)	105 ns	9.53 M	49.3 ns	20.3 M
ActivateTask (self) TerminateTask ()	139 ns	7.19 M	64.2 ns	15.6 M
GetResource () ReleaseResource ()	51.2 ns	19.5 M	23.7 ns	42.3 M
WaitEvent () ClearEvent () SetEvent ()	301 ns	3.33 M	185 ns	5.40 M
ActivateTask (HP) ChainTask (self)	275 ns	3.64 M	163 ns	6.15 M
ActivateTask (MP) ActivateTask (HP) ChainTask (self)	451 ns	2.22 M	284 ns	3.52 M
ActivateTask (MP1) ActivateTask (MP2) ActivateTask (HP) ChainTask (self)	620 ns	1.61 M	426 ns	2.35 M
Schedule ()	34.0 ns	29.4 M	14.1 ns	70.9 M

Table 5.2: *josek* x86-32 against x86-64

These values can be normalized according to the explanation in 5.1:

	Norm. x86-32		x86-64		Ratio
	Time	n/sec	Time	n/sec	
ChainTask (self)	44.2 ns	22.6 M	49.3 ns	20.3 M	0.90
ActivateTask (self) TerminateTask ()	58.6 ns	17.1 M	64.2 ns	15.6 M	0.91
GetResource () ReleaseResource ()	21.6 ns	46.4 M	23.7 ns	42.3 M	0.91
WaitEvent () ClearEvent () SetEvent ()	127 ns	7.89 M	185 ns	5.40 M	0.68

ActivateTask (HP) ChainTask (self)	116 ns	8.64 M	163 ns	6.15 M	0.71
ActivateTask (MP) ActivateTask (HP) ChainTask (self)	190 ns	5.26 M	284 ns	3.52 M	0.67
ActivateTask (MP1) ActivateTask (MP2) ActivateTask (HP) ChainTask (self)	261 ns	3.83 M	426 ns	2.35 M	0.61
Schedule ()	14.3 ns	69.8 M	14.1 ns	70.9 M	1.02

Table 5.3: *josek* x86-32 (normalized) against x86-64

Comparing these values yields an interesting result: *josek* performs better (at least relatively) on the x86-32 architecture. The difference is not to be dismissed as beside the point; it's up to 57% faster¹ when heavy scheduling occurs. This is not really surprising, as the context change is more expensive on the x86-64 than it is on the x86-32 – the reasoning behind this is explained in section 5.2.

5.3.2 *josek* vs. *Trampoline*

On the x86-64 platform these tests were also conducted against *Trampoline*:

	<i>Trampoline</i>		<i>josek</i>		Ratio
	Time	n/sec	Time	n/sec	
ChainTask (self)	911 ns	1.10 M	49.3 ns	20.3 M	18.5
ActivateTask (self) TerminateTask ()	1.15 μ s	870 k	64.2 ns	15.6 M	17.9
GetResource () ReleaseResource ()	492 ns	2.03 M	23.7 ns	42.3 M	20.8
WaitEvent () ClearEvent () SetEvent ()	2.53 μ s	395 k	185 ns	5.40 M	13.7
ActivateTask (HP) ChainTask (self)	2.05 μ s	488 k	163 ns	6.15 M	12.6
ActivateTask (MP) ActivateTask (HP) ChainTask (self)	3.17 μ s	315 k	284 ns	3.52 M	11.2

¹ $\frac{100 \cdot 3.24}{2.06} - 100 \approx 57\%$

ActivateTask (MP1)	4.36 μ s	229 k	426 ns	2.35 M	10.2
ActivateTask (MP2)					
ActivateTask (HP)					
ChainTask (self)					
Schedule ()	247 ns	4.05 M	14.1 ns	70.9 M	17.5

Table 5.4: *Trampoline* compared to *josek* on x86-64

When comparing *Trampoline* to a equal *josek* system, the *Trampoline* OSEK always seems to be at least one magnitude behind in terms of CPU time. This is primarily because of the *viper* hardware abstraction layer which *Trampoline* provides.

5.3.3 Comparison between x86-32, x86-64 and *Trampoline*

Figure 5.1 shows the results determined in the previous sections graphically.

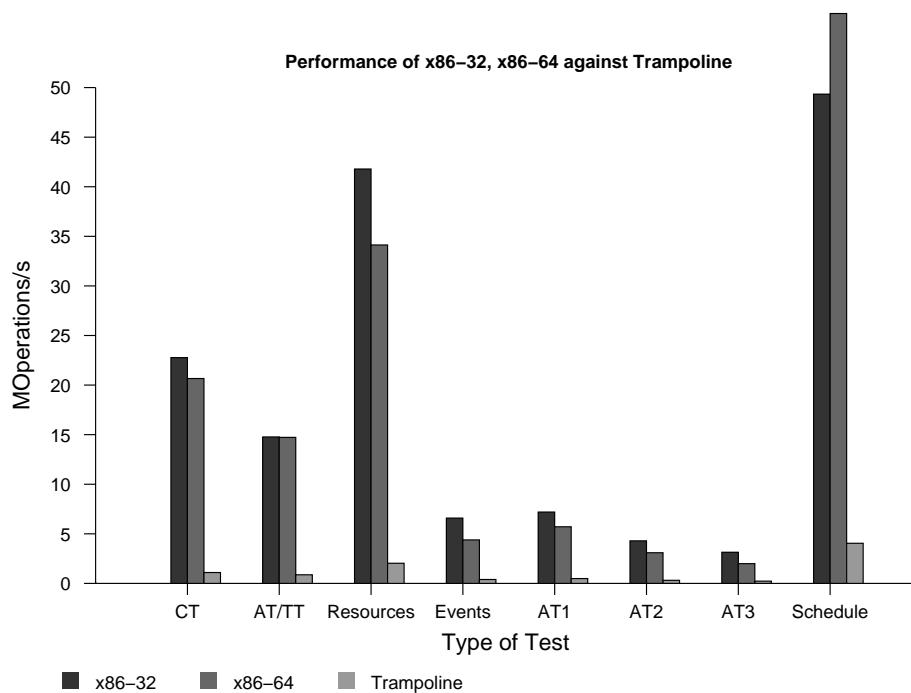


Figure 5.1: Comparison of Architectures

The particular reason the examples using resources and the one calling `Schedule` repeatedly are so fast compared to the others is because they don't need any con-

text switches. In the resource example, as there is always only one task active which repeatedly acquires and releases the resource, a context change to another task simply never becomes necessary. Although the scheduler is called every time from within `ReleaseResource`, it always decides that nothing has to be changed and the program continues right where it left of. The same thing happens in the `Schedule` example: no other task than the currently running one has been activated; the scheduling queue is looked at, but no scheduling occurs.

Calls to `ActivateTask/TerminateTask` are slower than just calling `ChainTask` because in a full-preemptive system the scheduling decision is performed twice (i.e. at each system call) instead of just once.

5.3.4 KESO

The *KESO* operating system layer which was introduced in chapter 4 can run on top of *josek*. A comparison between a "raw" *josek* system and a *KESO* system is drawn in table 5.5. The different areas of evaluation have been programmed in Java to resemble the original C-code as closely as possible. *KESO* does more than the pure *josek*, of course: null pointer checks are included before every dereferencing of a pointer, as is assured by the *Java* language specification, for example.

	<i>KESO</i>		<i>josek</i>		Ratio
	Time	n/sec	Time	n/sec	
<code>ChainTask (self)</code>	50.0 ns	20.0 M	49.3 ns	20.3 M	1.01
<code>ActivateTask (self)</code> <code>TerminateTask ()</code>	72.0 ns	13.9 M	64.2 ns	15.6 M	1.12
<code>GetResource ()</code> <code>ReleaseResource ()</code>	33.6 ns	29.8 M	23.7 ns	42.3 M	1.42
<code>WaitEvent ()</code> <code>ClearEvent ()</code> <code>SetEvent ()</code>	177 ns	5.64 M	185 ns	5.40 M	0.96
<code>ActivateTask (HP)</code> <code>ChainTask (self)</code>	175 ns	5.70 M	163 ns	6.15 M	1.08
<code>ActivateTask (MP)</code> <code>ActivateTask (HP)</code> <code>ChainTask (self)</code>	329 ns	3.04 M	284 ns	3.52 M	1.16
<code>ActivateTask (MP1)</code> <code>ActivateTask (MP2)</code> <code>ActivateTask (HP)</code> <code>ChainTask (self)</code>	455 ns	2.20 M	426 ns	2.35 M	1.07

Table 5.5: *KESO* compared to "raw" *josek* on x86-64

There is one stunning result in the above numbers: the event-example above is faster when running on the *KESO* system, although *KESO* introduces overhead compared to pure *josek*. It seems paradox at first. The result can be explained, however. A run of `cachegrind`, a tool in the *valgrind* package, shows that the example running on top of *KESO* has a 2.9% D1-cache miss rate while the "raw" *josek* has a miss rate of 8.9% – more than three times as much. To verify the results are solely an example of unfortunate cache usage, the test was also conducted on an AMD X2 4800 (2 Cores, 1 MB L2-cache per core, 2.4 GHz clock rate) and on an Intel Core2 6400 (2 Cores, 2 MB L2-cache per core, 2.13 GHz clock rate). The results prove the above explanation is true: on the X2 *KESO* runs with 8.06 MOps/s and 2.9% D1-cache miss rate while *josek* performs with 6.48 MOps/s and 9.5% miss rate (performance ratio of 0.80). In contrast on the Intel Core2 *KESO* yields 6.2 MOps/s with 0% D1-cache misses and *josek* gets 7.16 MOps/s with also 0% cache misses (performance ratio of 1.15).

The reason why the Core2 performs better than both AMD processors is most likely the higher cache associativity of the Intel processor. It has an 8-way D1-cache in contrast to the 2-way D1-caches of the AMD processors. Why cache associativity is of great importance and may impact performance in the way described above is well explained in literature.[HS89]

In any case it becomes clear that the overhead introduced by *KESO* is well within the margin of measurement inaccuracy. This is a great win for the *Java* based solution as it means the additional checks have negligible overhead in CPU-time, but nevertheless do they play an important role in saving development time.

5.3.5 Other Points of Interest

The relevance of a small context size can be nicely evaluated by artificially inflating the context size and measuring how performance degrades. Therefore in the following test on an x86-64 (i.e. one machine word equals 8 bytes), more and more machine words were added to the context, although unnecessary for functional purposes. The test program solely performing the `ChainTask` example was executed, each time with increased context size. `push` and `pop` operations were inserted at the appropriate places in `switchtask` and `dispatch`.

Figure 5.2 clearly shows that the runtime (which is almost equal to the context switch time in this example) grows linearly with the size of the context.

A particularity of the code using `GetResource/ReleaseResource` has not become clear so far, however. It's the difference in runtime of the two resource functions. For this purpose, both have been separately analyzed using the CPU cycle count method described in section 5.1.

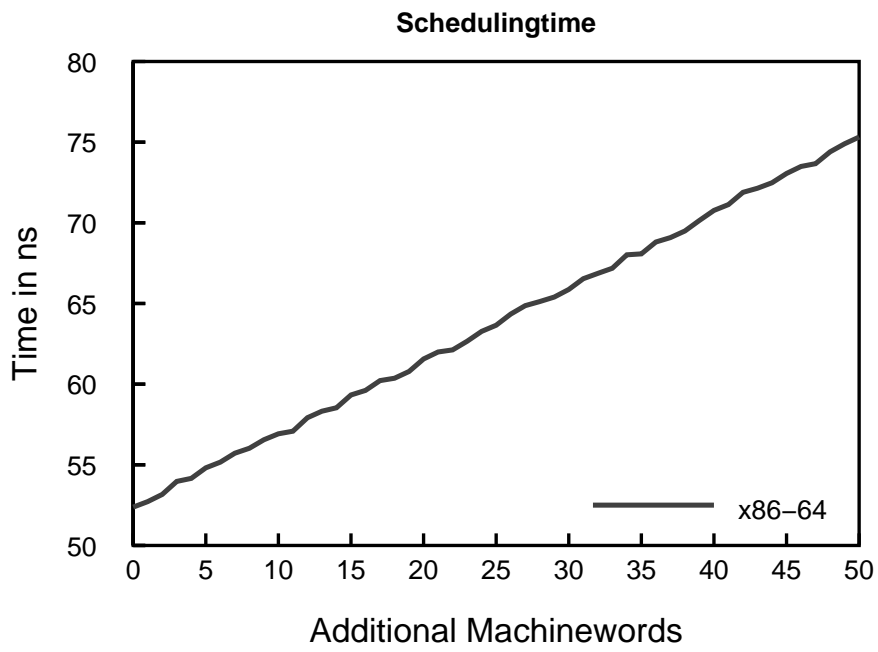


Figure 5.2: Artificially Inflating Context Size

	GetResource	ReleaseResource
Number of executions:	369312353	
Total execution time:	77.8 sec	
Operations/Second:	4.75 M	
Number of CPU cycles:	9928252581	22505181787
Cycles/Execution:	27	61
Theoretical Ops/Second:	15.5 M	6.84 M

Table 5.6: GetResource compared to ReleaseResource

It is not surprising that the `ReleaseResource` system call takes more than twice the time `GetResource` needs. The reason for that is simple: `GetResource` will only increase the task priority or not alter it at all. The currently running task will therefore never be preempted by a `GetResource` system call – no rescheduling point is necessary. When calling `ReleaseResource`, however, it is very well possible the task which is currently in the *running*-state decreases in priority and is therefore preempted (in a full-preemptive system) by a more important task. `ReleaseResource` thus calls `Schedule`, which takes time to come to a scheduling decision.

The overall performance is with 4.75 MOps/s remarkably slower in total than the measurements which were taken without the time stamp counter readout. This can be easily explained: `GetResource` and `ReleaseResource` together take around 88 CPU cycles, while one TSC readout needs 109 cycles to complete. Three readouts are performed per loop (one before `GetResource`, one after `ReleaseResource` and one in between), yielding 327 cycles plus the overhead of cycle counting. So 327 (or around 79%) of the total 415 CPU cycles per loop are used for reading out the TSC – obviously degrading performance.

Chapter 6

Conclusion and Future Work

6.1 Conclusion

The goals of writing a OSEK conform implementation were for one thing to show how it can be done, to point out optimization possibilities and to have a free OSEK implementation available. Through usage of *josek* a clean OSEK conform operating system can be generated within the matter of minutes. It supports multiple architectures (x86-32, x86-64 and AVR) and has been developed so it can easily be ported to even more target platforms. Extending the code generator via use of code hooks is intuitive. Extended tasks, alarms and counters are fully supported resulting in a useful OSEK conform operating system.

During development, attention was paid that the resulting code is well readable – the trade-off between highly optimized and well-readable code was taken in favor of code concinnity. Study of the generated code is easy and reveals interesting points of operating system development. Testing functional components of an OSEK compatible operating system has become simple, as development tools like the GNU Debugger *gdb* or *valgrind* can be used on the resulting code.

josek has become a good basis for running *KESO* on UNIX, which was also one reason for development of an OSEK compliant operating system. As both operating systems aim towards implementation of the OSEK operating system interface, hardly any consultation was necessary as the concept of blackboxing worked surprisingly well.

josek is distributed under the terms of the GNU LGPL-3¹ and can therefore be used and modified without legal shackle as it would be the case with proprietary software.

¹Lesser General Public License

6.2 Future Work

The *josek* code generator can in future be extended to support even more architectures. The generated code could be further improved – this study thesis should give hints on how to get started. Extensive testing of the AVR architecture would be nice, as a perfectly working AVR-port would impressively demonstrate how well OSEK is fit for use in highly embedded systems.

The code package is currently under heavy development because support for a backend different from OSEK is added. Particularly, the *josek* code is changed in order to support the interface of the *EPOS*² operating system. In this context, the project has been renamed to *EPOS-OSEK*. The two projects could be merged together in future so that improvements to one code basis also affect the other and vice versa.

All in all, the *josek* project has to be considered a success. It clearly shows what can be done with considerable effort and has the potential of being used as a starting point for other developers.

²Embedded Parallel Operating System

Bibliography

- [Atm07] Atmel Corporation. *8-bit AVR Microcontroller with 128K Bytes In-System Programmable Flash – ATmega128*. ATMEL, August 2007. <http://www.atmel.com/atmel/acrobat/doc2467.pdf>.
- [BBFT06] Jean-Luc Béchenec, Mikaël Briday, Sébastien Faucou, and Yvon Trinquet. Trampoline – an opensource implementation of the osek/vdx rtos specification. 2006. <http://trampoline.rts-software.org/IMG/pdf/trampoline.pdf>.
- [Fog07] Agner Fog. *Calling conventions for different C++ compilers and operating systems*. Copenhagen University College of Engineering, July 2007. http://www.agner.org/optimize/calling_conventions.pdf.
- [Gud07] Keith Gudger et al. *avr-libc Reference Manual 1.4.6*. May 2007. <http://savannah.nongnu.org/download/avr-libc/avr-libc-user-manual-1.4.6.pdf.bz2>.
- [HS89] Mark D. Hill and Alan Jay Smith. Evaluating associativity in CPU caches. *IEEE Transactions on Computers*, 38(12):1612–1630, 1989. ftp://ftp.cs.wisc.edu/markhill/Papers/toc89_cpu_cache_associativity.pdf.
- [Int07a] Intel Corporation. *Intel 64 and IA-32 Architectures Software Developer’s Manual – Volume 1: Basic Architecture*. August 2007. <http://www.intel.com/design/processor/manuals/253665.pdf>.
- [Int07b] Intel Corporation. *Intel 64 and IA-32 Architectures Software Developer’s Manual – Volume 2A: Instruction Set Reference, A-M*. May 2007. <http://www.intel.com/design/processor/manuals/253666.pdf>.
- [Int07c] Intel Corporation. *Intel 64 and IA-32 Architectures Software Developer’s Manual – Volume 2B: Instruction Set Reference, N-*

- Z. May 2007. <http://www.intel.com/design/processor/manuals/253667.pdf>.
- [ope07] openOSEK.org. *The most common questions answered*. March 2007. <http://www.openosek.org/tikiwiki/tiki-index.php>.
- [OSE04] The OSEK Group. *OSEK/VDK Binding Specification 1.4.2*. July 2004. <http://portal.osek-vdx.org/files/pdf/specs/binding142.pdf>.
- [OSE05] The OSEK Group. *OSEK/VDK Operating System Specification 2.2.3*. February 2005. <http://portal.osek-vdx.org/files/pdf/specs/os223.pdf>.
- [Sch07] Michel Schinz. *Advanced Compiler Construction - Tail Call Elimination*. March 2007. http://lamp.epfl.ch/teaching/archive/advanced_compiler/2007/resources/slides/act-2007-07-tail-calls_6.pdf.
- [SRL90] Lui Sha, Raganathan Rajkumar, and John P. Lehoczky. Priority inheritance protocols: An approach to real-time synchronization. *IEEE Transactions on Computers*, 39(9):1175–1185, 1990.
- [StGDC07] Richard M. Stallman and the GCC Developer Community. *Using the GNU Compiler Collection – Function Attributes*. GNU Press, July 2007. <http://gcc.gnu.org/onlinedocs/gcc-4.2.1/gcc/Function-Attributes.html>.