

Master Thesis

Formal Specification and Verification of Functions of VAMOS Scheduler

Yury Chebiryak



Saarland University, Computer Science Department

Institute for Computer Architecture and Parallel Computing

Prof. Dr. W. J. Paul

Dr. habil. Werner Stephan

October 2005

Contents

| | |
|---|-----------|
| Introduction | 6 |
| 1 A Word on Notation | 7 |
| 1.1 Basic Types | 7 |
| 1.2 Lists | 8 |
| 1.3 Records | 11 |
| 2 The VAMOS* Model | 12 |
| 2.1 The VAMOS* Configuration | 13 |
| 2.1.1 Processes: <i>process_list</i> | 13 |
| 2.1.2 Kernel Data Structures <i>kds</i> | 14 |
| 2.1.3 Global Variables <i>gv</i> | 15 |
| 2.1.4 Connection with the CVM Layer <i>cvm_c</i> | 15 |
| 2.1.5 External Interrupt Handlers <i>external_handlers_list</i> | 15 |
| 2.1.6 Device Drivers <i>devices_list</i> | 15 |
| 2.2 VAMOS System Calls | 16 |
| 3 VAMOS Scheduler | 17 |
| 3.1 Design of the VAMOS Scheduler | 17 |
| 3.2 Clock Interrupt Handler | 18 |
| 3.2.1 Incrementing the Time | 19 |
| 3.2.2 Checking the Timeslice of the Current Process | 20 |
| 3.2.3 Checking for Elapsed IPC Timeouts | 21 |
| 3.3 Switching To Another Process | 24 |
| 3.3.1 Updating the consumed timeslice | 25 |
| 3.3.2 Moving the invoker to the end of its ready list | 26 |
| 3.3.3 Changing the current process | 26 |
| 3.3.4 Remarks | 26 |

| | | |
|----------|--|-----------|
| 3.4 | Changing Scheduling Parameters | 26 |
| 3.4.1 | Updating the scheduling parameters | 29 |
| 3.4.2 | Updating the <i>Ready Lists Array</i> | 29 |
| 3.4.3 | Changing the <i>current_max_prio</i> and the current process | 30 |
| 3.5 | Summary | 30 |
| 4 | The VAMOS Model in Isabelle | 31 |
| 4.1 | Basic Datatypes | 31 |
| 4.2 | The <i>VAMOS Configuration</i> | 32 |
| 4.2.1 | The Record <i>c_conf_t</i> | 32 |
| 4.2.2 | The Record <i>v_data_t</i> | 32 |
| 4.3 | VAMOS Transition Functions | 33 |
| 5 | Implementation | 36 |
| 5.1 | The Programming Language C0 | 36 |
| 5.2 | Data Structures | 37 |
| 5.2.1 | Process Information Block | 40 |
| 5.2.2 | Global Variables | 40 |
| 5.3 | VAMOS System Calls | 40 |
| 6 | Verification Environment | 46 |
| 6.1 | Hoare Logic for Partial Correctness | 46 |
| 6.2 | Abstraction Relations | 47 |
| 6.3 | Imperative Programming Language | 47 |
| 6.3.1 | Assignments | 47 |
| 6.3.2 | Conditional Statement | 48 |
| 6.3.3 | Loop | 48 |
| 6.3.4 | Procedure Call | 49 |
| 6.4 | VCG | 49 |
| 7 | Functional Correctness | 50 |
| 7.1 | Functions on dLists | 50 |
| 7.1.1 | The Function <i>queue_InsertTail()</i> | 50 |
| 7.1.2 | The Function <i>queue_Rotate()</i> | 51 |
| 7.2 | The <i>VamosData?</i> Predicate | 53 |
| 7.3 | Assumptions | 55 |
| 7.3.1 | Assumptions from the <i>VamosData?</i> Predicate | 55 |
| 7.3.2 | Stand-alone Assumptions | 57 |

| | | |
|-------|--|-----------|
| 7.4 | The VAMOS Functions | 57 |
| 7.4.1 | The Function <code>search_next_process()</code> | 58 |
| 7.4.2 | The Function <code>compute_max_prio()</code> | 59 |
| 7.4.3 | The Function <code>wake_up()</code> | 60 |
| 7.4.4 | The Function <code>process_switch_to()</code> | 63 |
| 7.4.5 | The Function <code>process_change_sched_param()</code> | 67 |
| 7.5 | Verification Status | 68 |
| | Conclusion and Future Work | 70 |
| | A CVM and VAMOS Layers | 71 |
| | B C0-implementation | 73 |
| B.1 | The Function <code>queue_InsertTail()</code> | 73 |
| B.2 | The Function <code>queue_Rotate()</code> | 74 |
| B.3 | The Function <code>search_next_process()</code> | 74 |
| B.4 | The Function <code>compute_max_prio()</code> | 75 |
| B.5 | The Function <code>wake_up()</code> | 76 |
| B.6 | The Function <code>process_switch_to()</code> | 77 |
| B.7 | The Function <code>process_change_sched_param()</code> | 79 |
| | C The Function <code>get_my_pid</code> | 81 |
| | Bibliography | 83 |

List of Figures

| | | |
|-----|--|----|
| 3.1 | Scheduler design. | 19 |
| 3.2 | System call <i>processSwitchTo</i> example. | 27 |
| 4.1 | The <i>VAMOS Configuration</i> | 34 |
| 4.2 | System call <i>processSwitchTo</i> modelled in Isabelle. | 35 |
| 5.1 | Implementation of the VAMOS data structures in C0. | 38 |
| 5.2 | Process information block structure implementing the record <i>PrsData.t</i> | 41 |
| 5.3 | Global variables implementing the record <i>gd.t</i> | 42 |
| 5.4 | VAMOS functions. | 44 |
| 5.5 | Static call graph of functions relevant to the VAMOS scheduler. | 45 |
| 7.1 | The function <code>queue_InsertTail()</code> | 52 |
| 7.2 | The function <code>queue_Rotate()</code> | 53 |
| 7.3 | Functional correctness. | 55 |
| 7.4 | The function <code>search_next_process()</code> | 59 |
| 7.5 | Function <code>compute_max_prio()</code> | 61 |
| 7.6 | Annotating the while loop with invariants. | 62 |
| 7.7 | The function <code>wake_up()</code> | 64 |
| 7.8 | The function <code>process_switch_to()</code> | 66 |
| 7.9 | The function <code>process_change_sched_param()</code> | 69 |
| A.1 | The CVM and VAMOS layers. | 72 |
| C.1 | The function <code>process_get_my_pid()</code> | 82 |

Introduction

This master thesis is done in the context of the Verisoft project¹. Verisoft aims at the pervasive formal verification of entire computer systems. In particular, the seamless verification of the academic system is attempted. This system consists of hardware (processor and devices) on top of which runs a microkernel, an operating system, and applications.

In this thesis we focus on the microkernel. The microkernel is divided into two computational models: CVM* (Communicating Virtual Machines) and VAMOS* (Verified Architecture Microkernel Operating System).

The CVM* model represents the low-level functionalities of the microkernel: (i) process switching, (ii) allocation and deallocation of memory, (iii) declaration of interrupt handlers, (iv) I/O with devices, and (v) copying of data between user processes.

Based on CVM* we have the model VAMOS* providing: (i) scheduling mechanisms, (ii) memory management, and (iii) inter-process communication (IPC).

The aim of this thesis is to formally specify the scheduling mechanisms and show the functional correctness of the VAMOS scheduler.

Organization of the document

Chapter 1 presents notation that we use within this thesis. *Chapter 2* outlines the mathematical model VAMOS*. *Chapter 3* describes the design of the scheduling mechanisms, i.e. the VAMOS scheduler and how it operates within the VAMOS* model. In *chapter 4* we come closer to the real world. We present how VAMOS is modelled in theorem proving environment Isabelle, and *chapter 5* concentrates on the implementation of VAMOS using the programming language C0. *Chapter 6* introduces the verification environment that we use in Verisoft project. *Chapter 7* presents specifications of the VAMOS functions relevant to the VAMOS scheduler.

¹<http://www.verisoft.de>

Chapter 1

A Word on Notation

First we introduce notation. Some of these definitions (together with notation) are adopted from [9].

1.1 Basic Types

We use the following basic types:

Boolean numbers: $\mathbb{B} = \{ False, True \}$

Integer numbers: $\mathbb{Z} = \{ \dots, -2, -1, 0, 1, 2, \dots \};$

Natural numbers: $\mathbb{N} = \{ 0, 1, 2, \dots \};$

Process identifier (PID): We identify processes of the VAMOS microkernel using process identifiers (PIDs). The number of processes is currently restricted to 128. We define the set \mathbb{PID} of process identifiers as follows:

$$PID_MAX \in \mathbb{N}, \quad PID_MAX = 128;$$

$$\mathbb{PID} = \{0, \dots, PID_MAX - 1\}.$$

Priority (PRIO): Each process of the VAMOS microkernel has a priority. Currently, there are three different priorities. We define the set \mathbb{PRIO} of priorities as follows:

$$MAX_PRIO \in \mathbb{N}, \quad MAX_PRIO = 3;$$

$$\mathbb{PRIO} = \{0, \dots, MAX_PRIO - 1\}.$$

1.2 Lists

A list is a sequence of elements of a certain type. This type can be a basic or a compound one. We denote a list of elements of type \mathbb{T} by “*list* (\mathbb{T})”. An empty list is denoted by “[]”. For further use we define:

List of one element

A list of one element x of type \mathbb{T} is constructed using square brackets:

$$\begin{aligned} x &: \mathbb{T} \\ [x] &: \textit{list}(\mathbb{T}). \end{aligned}$$

List of more than one element

To denote a list of n elements we enclose them in square brackets and divide by comma:

$$\begin{aligned} a_1, a_2, a_3, \dots, a_n &: \mathbb{T} \\ [a_1, a_2, a_3, \dots, a_n] &: \textit{list}(\mathbb{T}). \end{aligned}$$

For example, list of integers, containing -5 followed by 2, followed by 1 is denoted as follows:

$$[-5, 2, 1] : \textit{list}(\mathbb{I}).$$

Constructor (Cons)

A list is constructed from an element x and a list xs by the infix operator “#”:

$$\begin{aligned} x &: \mathbb{T} \\ xs &: \textit{list}(\mathbb{T}) \\ x \# xs &: \textit{list}(\mathbb{T}). \end{aligned}$$

For example,

$$a_1 \# [a_2, a_3, a_4] = [a_1, a_2, a_3, a_4].$$

In order to describe the VAMOS we would need some functions on lists:

The head of a list

The function *hd* returns the first element of a list. It is undefined for an empty list:

$$\begin{aligned} \textit{hd} : \textit{list}(\mathbb{T}) &\longrightarrow \mathbb{T} \\ \textit{hd}(x \# xs) &\equiv x \end{aligned}$$

For example,

$$hd [a_1, a_2, a_3, a_4] = a_1$$

The tail of a list

The function *tl* returns a list without its first element:

$$tl : list (\mathbb{T}) \longrightarrow list (\mathbb{T})$$

$$tl [] \equiv []$$

$$tl (x \# xs) \equiv xs$$

For example,

$$tl[a_1, a_2, a_3, a_4] = [a_2, a_3, a_4]$$

The length of a list

The function *length* returns the number of elements in a list:

$$length : list (\mathbb{T}) \longrightarrow \mathbb{N}$$

$$length [] \equiv 0$$

$$length (x \# xs) \equiv 1 + length xs$$

For example,

$$length[a_1, a_2, a_3, a_4] = 4$$

Appending two lists

The infix operator “@” appends two lists:

$$@ : list (\mathbb{T}) \times list (\mathbb{T}) \longrightarrow list (\mathbb{T})$$

$$[] @ ys \equiv ys$$

$$(x \# xs) @ ys \equiv x \# (xs @ ys)$$

For example,

$$[a_1, a_2] @ [a_3, a_4] = [a_1, a_2, a_3, a_4]$$

Accessing an element by index

A list can be seen as an array. To access the *i*-th element of a list (*xs*) we use the infix

operator “!”. The operation is undefined for an empty list and for indices greater than length of a list:

$$! : list(\mathbb{T}) \times \mathbb{N} \longrightarrow \mathbb{T}$$

$$(xs) ! i \equiv \begin{cases} hd(xs) & \text{if } (i = 0) \\ (tl(xs)) ! (i - 1) & \text{otherwise} \end{cases}$$

Presence of an element in a list

We use the infix operator “ \in ” to check whether an element is in a list:

$$\in : \mathbb{T} \times list(\mathbb{T}) \longrightarrow \mathbb{B}$$

$$y \in [] \equiv False$$

$$y \in (x \# xs) \equiv \begin{cases} True & \text{if } (x = y) \\ y \in xs & \text{otherwise} \end{cases}$$

Deleting an element from a list

We delete an element with the function *Delete*. We assume only one occurrence of the provided element.

$$Delete : list(\mathbb{T}) \times \mathbb{T} \longrightarrow list(\mathbb{T})$$

$$[] Delete y \equiv []$$

$$(x \# xs) Delete y \equiv \begin{cases} xs & \text{if } (x = y) \\ x \# (xs Delete y) & \text{otherwise} \end{cases}$$

We introduce an extended version of this function. The function takes two lists and deletes all elements of the second list from the first one. It is defined as the infix operator “ \setminus ”:

$$\setminus : list(\mathbb{T}) \times list(\mathbb{T}) \longrightarrow list(\mathbb{T})$$

$$xs \setminus [] \equiv xs$$

$$xs \setminus (y \# ys) \equiv (xs Delete y) \setminus ys$$

Total order in lists

We usually use lists with distinct components. Thus, we can define total order on elements with respect to this list. Obviously, one element precedes another one if its index

is smaller:

$$A : \mathbb{T}, \quad B : \mathbb{T}, \quad a : \mathbb{N}, \quad b : \mathbb{N}, \quad lst : list(\mathbb{T});$$

Let

$$A \in lst, \quad lst ! a = A, \quad B \in lst, \quad lst ! b = B,$$

the total order is defined as follows:

$$A <_{lst} B \iff a < b.$$

1.3 Records

A *record* is a special data type. It is a compound object, containing named components of arbitrary types. These components are called *fields*.

For instance, the record R with the integer field A and the boolean field B is defined by the following notation:

$$R : (A : \mathbb{Z}, B : \mathbb{B})$$

For further use we define:

Accessing a field of a record

In order to access a certain field of a record we use the "." (dot) operator. The field A of the record R presented above is accessed as follows:

$$(A = -5, B = True).A = -5$$

Records equality

Two records are equal if and only if all corresponding fields are equal:

$$R = R' \iff \forall r. (R.r = R'.r)$$

Chapter 2

The VAMOS* Model

This chapter presents the VAMOS* model. The description relies on the terminology introduced in [2].

The VAMOS* model is based on the CVM* model. The CVM* model represents the low-level functionalities of the microkernel: (i) process switching, (ii) allocation and deallocation of memory, (iii) declaration of interrupt handlers, (iv) I/O with devices, and (v) copying of data between user processes.

The model VAMOS* provides: (i) scheduling mechanisms, (ii) memory management, and (iii) inter-process communication (IPC).

We define the VAMOS* model as follows:

$$VAMOS^* = (vamos_conf_t, vamos_input_t, vamos_start_t, \delta_{VAMOS})$$

- **VAMOS* configuration** *vamos_conf_t*: this record defines all configurations of our model;
- **Input signals** *vamos_input_t*: the VAMOS* model has a set of input signals. These are interrupt events and input from devices;
- **initial configuration** *vamos_start_t*: this set defines all valid initial configurations of the VAMOS* model;
- **transition relation** δ_{VAMOS} : defines all possible transitions from a certain VAMOS* configuration.

In the following, we do not consider the whole model, but only the parts affected by the scheduler. Therefore the focus in the forthcoming section is on the VAMOS* configuration.

2.1 The VAMOS* Configuration

The VAMOS* configuration comprises the kernel data structures *kds*, the global variables *gv*, a list of all processes *process_list*, some parts of the CVM configuration *cvm_c*, a list of external handlers *external_handlers_list* and a list of devices *devices_list*:

```

vamos_conf_t = (process_list : list (process_data_t),  kds : kds_t,
                gv : gv_t,  cvm_c : cvm_c_t,
                external_handlers_list : list (PID),  devices_list : list (PID) )

```

2.1.1 Processes: *process_list*

The *process_list* of the VAMOS* configuration contains all processes. Its length is equal to the global constant *PID_MAX*.

```

process_data_t = (timeslice : N,  consumed_timeslice : N,  first_invalid_page : N,
                 pid : PID,  timeout : Z,  priority : PRIO,
                 privileges : B,  state : N,  ipc_length : N,
                 ipc_length2 : N,  ipc_message : N,  ipc_message2 : N,
                 ipc_partner : PID,  ipc_partner2 : PID,  ipc_timeout : Z,
                 ipc_timeout2 : Z,  ipc_send_list : list (PID) )

```

- **timeslice** — The computing time of a process until it is being rescheduled (in clock ticks);
- **consumed_timeslice** — The number of CPU clock ticks already used by a process;
- **first_invalid_page** — The index of the first page in memory not owned by a process;
- **pid** — The process identifier;
- **timeout** — The point in time, until an IPC operation has to happen (in clock ticks);

- **priority** — The priority of a process;
- **privileges** — The boolean variable that indicates whether a process has privileges (i.e. it is a privileged process);
- **state** — The current status of a process. It is either **inactive**, **ready** or **sleeping**;
- **ipc.length** — The length of an IPC message;
- **ipc.length2** — The same for the receive phase of the *ipcSendReceive* operation;
- **ipc.message** — The starting address of an IPC message;
- **ipc.message2** — The same for the receive phase of the *ipcSendReceive* operation;
- **ipc.partner** — The process identifier of the process to communicate with;
- **ipc.partner2** — The process identifier of the process to receive information from (only for the receive phase of the *ipcSendReceive* operation);
- **ipc.send_list** — A list of processes waiting for the IPC communication with this particular process.

2.1.2 Kernel Data Structures *kds*

$$kds.t = \left(\begin{array}{l} ready_lists_array : list(list(PID)), \\ sleeping_list : list(PID), \\ inactive_list : list(PID) \end{array} \right)$$

- **ready_lists_array** contains processes ready to execute, sorted by their priorities. The dimension of this array (i.e. the number of the ready lists) is defined by the global constant *MAX_PRIO*. Ready processes are contained in the ready list indexed by their priorities. If there are no processes of a certain priority, the corresponding ready list is empty;
- **sleeping_list** contains processes waiting for an IPC (sleeping processes);
- **inactive_list** contains unused PIDs.

2.1.3 Global Variables *gv*

$$gv.t = (cup : \mathbb{PID}, \quad current_max_prio : \mathbb{PRIO}, \\ time : \mathbb{N}, \quad next_timeout : \mathbb{Z}, \\ privileged_exists : \mathbb{B}, \quad vamos_pages_used : \mathbb{N})$$

- **cup** — The process identifier of the currently running process;
- **current_max_prio** denotes the maximum priority over all ready processes;
- **time** stores the global time (in clock ticks);
- **next_timeout** stores the closest IPC timeout (the minimum timeout over all processes in the *Sleeping List*);
- **privileged_exists** indicates whether privileged process(-es) already exists;
- **vamos_pages_used** stores the number of memory pages used by the VAMOS microkernel.

2.1.4 Connection with the CVM Layer *cvm_c*

The configuration of the VAMOS* model partially contains the CVM configuration *cvm_c*. It serves to bind the VAMOS layer with the underlying CVM layer. These layers are clearly separated, and the CVM layer is not of interest in the context of this thesis (because the VAMOS scheduler operates only on the VAMOS layer). In order to get the whole picture of the operating system developed in Verisoft refer to [2].

2.1.5 External Interrupt Handlers *external_handlers_list*

A process can be assigned to handle a certain external interrupt. The list *external_handlers_list* indicates processes associated with interrupts. The index in this list corresponds to the interrupt number. These mechanisms are not implemented yet.

2.1.6 Device Drivers *devices_list*

Processes can be assigned with devices and serve as device drivers. These mechanisms are not implemented yet.

2.2 VAMOS System Calls

VAMOS provides several system calls to the user. These system calls change the VAMOS* configuration depending on their input parameters. All VAMOS system calls are presented in Table 2.1. Thorough information about the VAMOS system calls can be found in [3].

| System Call | Description |
|-----------------------------------|--|
| <i>Task Management</i> | |
| processCreate | creates a new process |
| processClone | clones a given process |
| processKill | kills a given process |
| processGetPrivileges | is used to obtain privileges for the currently running process |
| processGetMyPid | returns PID of the currently running process |
| processSwitchTo | current process gives out CPU control voluntarily |
| processChangeSchedulingParam | changes timeslice and priority values for a given process |
| <i>Memory Management</i> | |
| memoryAdd | extends the virtual memory of a process |
| memoryFree | shortens the virtual memory of a process |
| <i>I/O-Devices</i> | |
| ioIn | writes data to device ports |
| ioOut | reads data from device ports |
| <i>Interprocess Communication</i> | |
| ipcSend | sends an IPC-message to a given process |
| ipcSendReceive | sends an IPC-message and waits for reply |
| ipcReceive | receives an IPC-message from a given process |

Table 2.1: VAMOS System Calls, adopted from [2].

Chapter 3

VAMOS Scheduler

VAMOS provides functionalities to serve as base for a multi-tasking operating system. The VAMOS processes are running concurrently and contending for the CPU control. Obviously, only one of them advances at time. Therefore we have to choose a process to execute. This is usually done by the scheduler, which is part of the operating system kernel.

In the following we describe the functionality of the VAMOS scheduler.

3.1 Design of the VAMOS Scheduler

It was conceived that processes of VAMOS can be distinguished in quality of service. For that purpose we have a **priority** associated with every process. A higher priority guarantees better service. Processes ready for execution are sorted in the *Ready Lists Array* by their priorities. Processes with the highest priority are executed in **round robin** fashion. If the ready list of the highest priority is empty, the scheduler executes processes from the ready list of the priority less by 1, and so on. Therefore a process with some priority is not scheduled while there exists a process with the higher priority. This defines the default scheduling policy. Nevertheless, it can be changed using the VAMOS system calls *processSwitchTo* and *processChangeSchedulingParam*.

The VAMOS scheduler is preemptive, i.e. a process can be suspended at an arbitrary instant [15]. Every process has a **timeslice** associated with it. The timeslice is the amount of the CPU time provided to a process for execution. In the VAMOS microkernel time is measured in clock ticks. After the current process has used its whole timeslice, the scheduler selects the next process to execute.

The current process can be preempted even before its timeslice has expired (e.g. if a process with the higher priority awakes).

An overview of the VAMOS scheduler design is presented in Figure 3.1.

As it was mentioned above, the *Ready Lists Array* stores processes that are ready for execution. Processes waiting for an IPC are enqueued in the *Sleeping List*. If a process is killed, it is enqueued to the *Inactive List*. All these lists are part of the kernel data structures *kds*. The global variables *gv* also contain some information relevant for the scheduling decisions: *cup* stores the PID of the currently running process, *current_max_prio* stores the maximum priority over all ready processes, and *next_timeout* stores the minimum IPC timeout over all sleeping processes.

Together with ordinary processes there is a special one, called the *idle* process. It is introduced to avoid the situation that all the ready lists are empty. The *idle* process runs in an endless loop. It does no meaningful work and it has the following properties: the minimum timeslice (1 CPU clock tick), the minimum priority (0), and its state is *ready*. Even privileged processes cannot affect these properties.

Thus, we can state that if the *idle* process is scheduled, there are no processes in the ready lists with higher priorities:

$$\forall c : \text{vamos_conf_t.} \\ \left(c.gv.cup = \text{idle} \implies \left(\forall i : \mathbb{N}. (i > 0 \wedge i < \text{MAX_PRIO} \implies \text{ready_lists_array} ! i = []) \right) \right).$$

Note, that this statement does not hold the other way round — the fact that all the ready lists with higher priorities are empty does not imply that the *idle* process runs, but any with priority 0. That is why the timeslice of the *idle* process is set to the minimum value — in order to lessen the CPU time wasted for the *idle* process while there are other processes requiring execution.

The order of process execution is determined by the clock interrupt handler and two VAMOS system calls. In the following we give a detailed description.

3.2 Clock Interrupt Handler

The core of the VAMOS scheduler is the clock interrupt handler. It is invoked on every clock tick. The clock interrupt handler is in charge of the scheduling decision. In other words, it chooses the process that is executed in the next clock tick.

In general, three main steps should be done to handle a clock interrupt:

- 1) increment the time value;

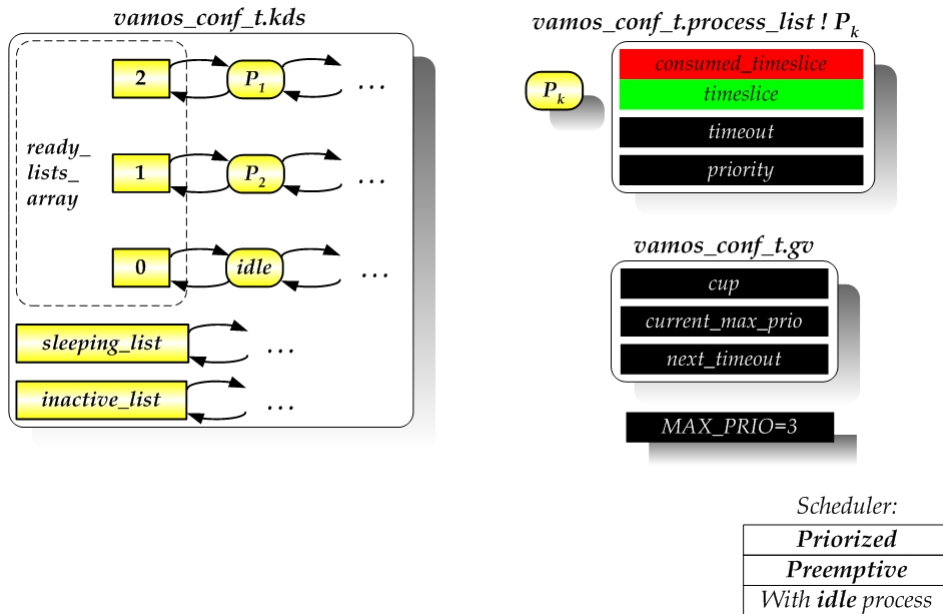


Figure 3.1: Scheduler design.

- 2) check whether the timeslice of the current process is expired;
- 3) check for elapsed IPC timeouts.

On every step we refer to the VAMOS* configuration before the invocation of the clock interrupt handler as c . By updating c we get the VAMOS* configuration c' :

$$c, c' : \text{vamos_conf_t}$$

We emphasize only on updates of the configurations. In other words, we only show the changes, made on c to come to c' .

In the following we present changes done in these steps. Every step is explained in a separated subsection.

3.2.1 Incrementing the Time

The global time counter is incremented by 1:

$$c'.gv.time \equiv c.gv.time + 1$$

3.2.2 Checking the Timeslice of the Current Process

Since the current process used another clock tick we have to check whether the assigned timeslice is spent. Therefore we introduce the predicate *timeslice_complete?*:

$$timeslice_complete? : vamos_conf_t \longrightarrow \mathbb{B};$$

$$timeslice_complete?(c) \equiv$$

$$(((c.process_list ! c.gv.cup).consumed_timeslice + 1) < (c.process_list ! c.gv.cup).timeslice).$$

Depending on this predicate, we consider two cases.

Case 1

If the predicate *timeslice_complete?* does not hold for the configuration *c*, we simply increment the consumed timeslice of the current process:

$$\begin{aligned} (c'.process_list ! c.gv.cup).consumed_timeslice &\equiv \\ &(c.process_list ! c.gv.cup).consumed_timeslice + 1. \end{aligned}$$

Case 2

If the predicate *timeslice_complete?* holds, we have to suspend the current process.

This means: we set its consumed timeslice to zero, put it to the end of its ready list, and change the current process. The first two updates are basic and for the latter we have to consider the position of the current process in its ready list. In general, there are only three possible cases:

- The currently running process is the only element of the ready list indexed by the *current_max_prio*.
In this case the next process to schedule is again the current process. Thus, no update of the configuration is needed;
- The currently running process is the head of the ready list indexed by the *current_max_prio* and this list contains more than one element¹.

¹In fact, this case cannot happen with Simple Operating System (SOS) running on top of the VAMOS. Nevertheless, we consider it here to have all the possible cases covered.

In this case we have to schedule the second process in the ready list:

Let

$$cup_is_head? : vamos_conf_t \longrightarrow \mathbb{B};$$

$$cup_is_head?(c) \equiv \left((hd(c.kds.ready_lists_array ! c.gv.current_max_prio) = c.gv.cup) \right. \\ \left. \wedge (length(c.kds.ready_lists_array ! c.gv.current_max_prio) > 1) \right)$$

in

$$c'.gv.cup \equiv (c.kds.ready_lists_array ! c.gv.current_max_prio) ! 1 \quad \text{if } cup_is_head?(c)$$

- The currently running process is not the head of the ready list indexed by the *current_max_prio* or even not in this ready list. Such situation can happen if the configuration is affected by the *processSwitchTo* system call.

In this case we have to schedule the head of the ready list indexed by the *current_max_prio*:

Let

$$cup_not_head? : vamos_conf_t \longrightarrow \mathbb{B};$$

$$cup_not_head?(c) \equiv \left((hd(c.kds.ready_lists_array ! c.gv.current_max_prio) \neq c.gv.cup) \right. \\ \left. \vee (c.gv.cup \notin (c.kds.ready_lists_array ! c.gv.current_max_prio)) \right)$$

in

$$c'.gv.cup \equiv hd(c.kds.ready_lists_array ! c.gv.current_max_prio) \quad \text{if } cup_not_head?(c)$$

3.2.3 Checking for Elapsed IPC Timeouts

The variable *next_timeout* stores the closest IPC timeout over all processes in the *Sleeping List*. As the global time reaches the *next_timeout* we are sure that at least one timeout of sleeping processes expires. The predicate *timeout_elapsed?* checks whether a timeout expires:

$$timeout_elapsed? : vamos_conf_t \longrightarrow \mathbb{B}$$

$$timeout_elapsed?(c) \equiv (c.gv.next_timeout = c.gv.time).$$

If the predicate holds, we perform the following:

- all the sleeping processes with the elapsed IPC timeouts are woken up;
- the variable *next_timeout* should be recomputed;
- if a process with the priority higher than the *current_max_prio* awakes, we have to update the variable *current_max_prio* and change the current process.

We explain how these actions are done step by step.

Waking up Processes

First, we enclose the sleeping processes with the elapsed IPC timeouts into the *Elapsed List*:

$$filter_not_elapsed : vamos_conf_t \times list(\text{PID}) \longrightarrow list(\text{PID})$$

$$filter_not_elapsed(c, []) \equiv [];$$

$$filter_not_elapsed(c, x \# xs) \equiv \begin{cases} x \# xs & \text{if } ((c.process_list ! x).timeout = c.gv.time) \\ filter_not_elapsed(c, xs) & \text{otherwise} \end{cases};$$

$$elapsed_list(c) \equiv filter_not_elapsed(c, c.sleeping_list).$$

Then it is clear that the *Sleeping List* in the subsequent VAMOS* configuration is defined as follows:

$$rest_sleeping : vamos_conf_t \longrightarrow list(\text{PID})$$

$$rest_sleeping(c) \equiv c.sleeping_list \setminus elapsed_list(c);$$

$$\boxed{c'.kds.sleeping_list \equiv rest_sleeping(c).}$$

Processes from the *Elapsed List* should be enqueued to the corresponding ready lists. First, we define the function *filter_prio* that filters out processes with priorities differing from the given one:

$$filter_prio : list(\text{PID}) \times \text{PRIO} \longrightarrow list(\text{PID})$$

$$filter_prio([], prio) \equiv [];$$

$$filter_prio(x \# xs, prio) \equiv \begin{cases} x \# filter_prio(xs, prio) & \text{if } ((c.process_list ! x).priority = prio) \\ filter_prio(xs, prio) & \text{otherwise} \end{cases}$$

We define the updated *Ready Lists Array* as follows:

$$\forall prio : \text{PRIO}.$$

$$\boxed{\left(c'.kds.ready_lists_array ! prio \equiv (c.kds.ready_lists_array @ filter_prio(elapsed_list(c), prio)) \right)}.$$

We have to update the state of the awaking processes too:

$$\boxed{\forall prcs : \text{PID}. \quad (prcs \in \text{elapsed_list}(c) \implies ((c'.\text{process_list} ! prcs).\text{state} \equiv \text{ready}))}.$$

Updating the Variable *next_timeout*

Since the processes with the elapsed IPC timeouts are deleted from the *Sleeping List*, we have to update the *next_timeout* value²:

$$\text{MinTimeout} : \text{vamos_conf_t} \longrightarrow \mathbb{Z}$$

$$\text{MinTimeout}(c) \equiv \text{Min} \{x : \mathbb{Z} \mid x = (c.\text{process_list} ! prcs).\text{timeout}, \quad prcs \in \text{rest_sleeping}(c)\};$$

$$\boxed{c'.\text{next_timeout} \equiv \text{MinTimeout}(c).}$$

Updating the Variable *current_max_prio*

If at least one of the woken processes has a priority higher than the variable *current_max_prio*, we have to update its value. Also we have to change the current process, regardless the fact that we might have changed it in the previous step³.

First, we compute the highest priority over all awaking processes:

$$\text{MaxPrio} : \text{vamos_conf_t} \longrightarrow \text{PRIO}$$

$$\text{MaxPrio}(c) \equiv \text{Max} \{x : \text{PRIO} \mid x = (c.\text{process_list} ! prcs).\text{priority}, \quad prcs \in \text{elapsed_list}(c)\}.$$

If $\text{MaxPrio}(c)$ is greater than the $c.gv.\text{current_max_prio}$ we perform necessary updates:

$$\boxed{\begin{aligned} &c'.gv.\text{current_max_prio} \equiv \text{MaxPrio}(c) \\ \text{and} \\ &c'.gv.\text{cup} \equiv \text{hd}(\text{filter_prio}(\text{elapsed_list}(c), \text{MaxPrio}(c))). \end{aligned}}$$

²The function *Min* computes the minimum over a set.

³That is because the processes with high priority should be executed as they appear (or wake up).

3.3 Switching To Another Process

Processes with priorities lower than the *current_max_prio* have no opportunity to be executed under the default scheduling policy. They can only be executed if the current process voluntarily relinquish CPU. This mechanism is provided via the VAMOS system call *processSwitchTo*.

This system call is very useful when building a user-level scheduler on top of the VAMOS scheduler. Just imagine a privileged process with maximum priority that gives the CPU control to lower-priority processes. If these lower-priority processes have no privileges, they cannot influence scheduling decisions. Therefore using this scheduling mechanism a user-level scheduler can implement arbitrary scheduling policies.

The process invoking this call voluntarily donates the rest of its timeslice and the CPU control to a target process. The system call gets as a parameter the PID of the target process. Therefore signature of the transition function for this system call is defined as follows:

$$\delta_{switchTo} : vamos_conf_t \times \mathbb{PID} \longrightarrow vamos_conf_t.$$

There are certain restrictions for this system call:

- the invoker is a privileged process;
- the target process has the **ready** status;
- the target process is not the *idle* process.

We define the predicate *success_switch?* comprising these restrictions when trying to switch to the process *prcs*:

$$success_switch? : vamos_conf_t \times \mathbb{PID} \longrightarrow \mathbb{N}$$

$$success_switch?(c, prcs) \equiv$$

$$((c.process_list ! c.gv.cup).privileged) \wedge ((c.process_list ! prcs).status = \mathbf{ready}) \wedge (prcs \neq \mathbf{idle}).$$

In the following we present the changes made by this system call in order to switch to the given process *prcs*. We assume that there are no errors, i.e. the predicate *success_switch?* holds.

In general, we have to perform three steps in order to switch to the given process:

- 1) update the consumed timeslice of the target process (subtract donated timeslice);
- 2) move the invoker to the end of its ready list and set its consumed timeslice to zero;
- 3) change the current process to the target process.

All these steps are performed by the $\delta_{switchTo}$ transition function. We refer to the VAMOS* configuration before the invocation of this system call as c . If switching to the process $prcs$, the system call updates it to the VAMOS* configuration c' :

$$\begin{aligned} c, c' &: vamos_conf_t; \\ c' &= \delta_{switchTo}(c, prcs). \end{aligned}$$

In the following we present changes done in order to obtain c' .

3.3.1 Updating the consumed timeslice

The system call decreases the consumed timeslice of the target process by the rest of the timeslice of the current process. In other words, we extend the potential amount of the running time for the target process, and this extension is bounded by the amount of the rest timeslice the current process has. When donating the rest of the timeslice we have to beware of underflow in natural numbers.

The predicate *underflow?* checks for underflow in the consumed timeslice of the target process:

Let

$$rest_tsl = ((c.process_list ! c.gv.cup).timeslice - (c.process_list ! c.gv.cup).consumed_timeslice)$$

in

$$underflow? : vamos_conf_t \times PID \longrightarrow \mathbb{B}$$

$$underflow?(c, prcs) \equiv (c.process_list ! prcs).consumed_timeslice < rest_tsl.$$

If the predicate holds, we set the consumed timeslice to zero, otherwise we perform the subtraction:

Let

$$prcs_ctsl = (c.process_list ! prcs).consumed_timeslice$$

in

$$(c'.process_list ! prcs).consumed_timeslice \equiv \begin{cases} 0 & \text{if } underflow?(c, prcs) \\ (prcs_ctsl - rest_tsl) & \text{otherwise} \end{cases}$$

3.3.2 Moving the invoker to the end of its ready list

Since the currently running process has completely used its dedicated timeslice, it is suspended: its consumed timeslice is set to zero, and the process is placed at the end of its ready list.

Let

$$cup_prio = (c.process_list ! c.gv.cup).priority;$$

in

$c'.kds.ready_list_array ! cup_prio \equiv$
 $((c.kds.ready_lists_array ! cup_prio) Delete (c.gv.cup)) @ [c.gv.cup],$
 and
 $(c'.process_list ! c.gv.cup).consumed_timeslice \equiv 0.$

3.3.3 Changing the current process

Since the target process is the next computing process, we have to update the configuration accordingly:

$$c'.gv.cup \equiv pres.$$


3.3.4 Remarks

If the target process is already the currently running process, none of these actions are performed.

Figure 3.2 presents a pictorial example of how switching to the process is done.


3.4 Changing Scheduling Parameters

Priority and timeslice of the VAMOS processes are assigned during the process creation (using the *processCreate* system call). These scheduling parameters can be changed even after process creation using the VAMOS system call *processChangeSchedulingParam*. This system call is



1. Update Ready List:

first Delete **cup**, then append to the end;
 donate timeslice, null consumed timeslice of **cup**



2. Switch current process to **prcs**




Figure 3.2: System call `processSwitchTo` example.

very useful when creating a user-level scheduler on top of the VAMOS scheduler. It is available for privileged processes only.

The process invoking this system call attempts to change the priority and the timeslice of some process. New values for priority and timeslice are provided as input parameters together with the PID of the target process. Therefore signature of the transition function for this system call is defined as follows:

$$\delta_{c.s.p} : \text{vamos_conf_t} \times \text{PID} \times \text{PRIO} \times \mathbb{N} \longrightarrow \text{vamos_conf_t}.$$

There are certain restrictions for this system call:

- the invoker is a privileged process;
- the target process has the **active** status (i.e. either **ready** or **sleeping**);
- the target process is not the *idle* process.

We define the predicate *success_change?* comprising these restrictions when trying to change the scheduling parameters of the process *prcs*:

$$\text{success_change?} : \text{vamos_conf_t} \times \text{PID} \longrightarrow \mathbb{N}$$

$$\text{success_change?}(c, \text{prcs}) \equiv$$

$$((c.\text{process_list} ! c.\text{gv.cup}).\text{privileged}) \wedge ((c.\text{process_list} ! \text{prcs}).\text{state} \neq \mathbf{inactive}) \wedge (\text{prcs} \neq \mathbf{idle}).$$

In the following we present the changes made by this system call in order to change priority and timeslice of the given process *prcs*. We assume that there are no errors, i.e. the predicate *success_change?* holds.

In order to switch to the given process we have to perform the following three steps:

- 1) update timeslice and priority of the target process to the provided values;
- 2) move the target process into the appropriate ready list;
- 3) if necessary, change the current process to the target process and update the *current_max_prio*.

All these steps are performed by the $\delta_{c.s.p}$ transition function. We refer to the VAMOS* configuration before the invocation of this system call as *c*. If switching to the process *prcs*, the system call updates it to the VAMOS* configuration *c'*:

$$c, c' : \text{vamos_conf_t};$$

$$c' = \delta_{c.s.p}(c, \text{prcs}, \text{new_prio}, \text{new_tsl}).$$

In the following we present changes done in order to obtain c' .

3.4.1 Updating the scheduling parameters

The timeslice and priority fields of the *process_data.t* record of the target process are updated according to the input variables:

$$\begin{aligned} & (c'.process_list \! prcs).priority \equiv new_prio, \\ \text{and} \\ & (c'.process_list \! prcs).timeslice \equiv new_tsl. \end{aligned}$$

3.4.2 Updating the *Ready Lists Array*

If the target process is ready for execution and its priority is to be changed, we have to enqueue the process into the appropriate ready list. To decide this we introduce the predicate *prio_change?*:

Let

$$prcs_prio = (c.process_list \! prcs).priority,$$

in

$$prio_change? : vamos_conf.t \longrightarrow \mathbb{B};$$

$$prio_change?(c) \equiv ((c.process_list \! prcs).state = \mathbf{ready}) \wedge (new_prio \neq prcs_prio).$$

If the predicate *prio_change?* holds, we remove the target process from its ready list and append it to the ready list, corresponding to the new value of the priority:

$$\begin{aligned} & c'.kds.ready_lists_array \! prcs_prio \equiv ((c.kds.ready_lists_array \! prcs_prio) \text{ Delete } prcs), \\ \text{and} \\ & c'.kds.ready_lists_array \! new_prio \equiv ((c.kds.ready_lists_array \! new_prio) @ [prcs]). \end{aligned}$$

3.4.3 Changing the *current_max_prio* and the current process

If the newly assigned priority of the target process is higher than the *current_max_prio* value, we have to update the *current_max_prio* and switch to that process. In order to check this we introduce the predicate *prio_greater?*:

$prio_greater? : vamos_conf_t \rightarrow \mathbb{B}$

$prio_greater?(c) \equiv ((c.process_list \neq prcs).state = \mathbf{ready}) \wedge (new_prio > c.gv.current_max_prio).$

If the predicate holds, we assign the variable *current_max_prio* to the provided priority and switch the current process:

$$\begin{array}{l}
 c'.gv.current_max_prio = new_prio, \\
 \text{and} \\
 c'.gv.cup = prcs.
 \end{array}$$

3.5 Summary

In this chapter we have outlined the scheduling mechanisms implemented in the VAMOS scheduler and the provided functionality to change the default scheduling policy.

Chapter 4

The VAMOS Model in Isabelle

Having handwritten specification of a computer system is nice, but of no practical use. In the Verisoft project we use the interactive theorem prover Isabelle for verification. Therefore we have to adapt the VAMOS* model to the Isabelle system. This chapter describes how VAMOS* is represented in Isabelle.

4.1 Basic Datatypes

We utilize Isabelle with Higher-Order Logic. It provides the following basic types:

- **natural numbers** — recursive type generated by the constructors `zero` and `successor`. It works well with inductive proofs and recursive function definitions;
- **integer numbers** — lacks induction, but supports true subtraction;
- **lists** — recursive type generated by two constructors (the empty list and the operator that adds an element to the front of a list);
- **records** — compound type, containing named components of arbitrary types;
- **sets** — of any type with usual operations (intersection, union, etc.).

Isabelle's types system does not support subsets like `PID` or `PRIQ` directly. These are modelled as natural numbers \mathbb{N} and we introduce predicates to check the boundaries.

4.2 The VAMOS Configuration

With these basic types the VAMOS* configuration can be mapped to a record in Isabelle. We define the record `v_conf_t` and refer to it as *VAMOS Configuration* (it is depicted on Figure 4.1). In this section we explain its structural differences from the VAMOS* configuration.

The *VAMOS Configuration* consists of two fields: the record `c_conf` of the type `c_conf_t` and the record `v_data` of the type `v_data_t`.

4.2.1 The Record `c_conf_t`

This record type serves for a connection with the CVM layer. It is not presented here in details, because the VAMOS scheduler does not affect this layer.

4.2.2 The Record `v_data_t`

This record type stands for the *VAMOS Data*. It contains kernel information (kernel data structures and global variables) and information about processes. It contains two fields: the record `gd` of the type `gd_t`, and the function `P` that realizes the *process_list* within the VAMOS* model.

The record `gd_t`

This record type contains kernel data structures and global variables. There are only two differences in names: *Sleeping List* is called `wkp` and the *Ready Lists Array* is named `rdy`.

The function `P`

The function `P` realizes *processes_list* from the VAMOS* model. It maps process identifiers to the records of the type `PrCsData_t`.

The record `PrCsData_t`

This record type contains information about a process. Its fields directly correspond to the content of the *t_process_data*, but some names are abbreviated:

- *timeslice* is abbreviated as `tsl`,
- *consumed_timeslice* as `ctsl`,

- *priority* as `pri`,
- *first_invalid_page* as `fip`,
- *privileges* as `prv`,
- *ipc_send_list* as `ipc_sq`.

4.3 VAMOS Transition Functions

In the previous chapter we have presented some transition functions (e.g. $\delta_{switchTo}$) applied to the VAMOS* configuration. We represent them in Isabelle as non-recursive functions, which merely perform record updates to the *VAMOS Configuration*. In general, their signature looks as follows:

$$func_{abstr} : [v_conf_t, input] \longrightarrow v_conf_t.$$

Here the function $func_{abstr}$ is presented in curried notation. So, usually functions take a *VAMOS Configuration*, some input parameters and return a modified configuration. Adaptation to Isabelle notation is merely bookkeeping.

As an example we present function `abstr_switch_to` that corresponds to the function $\delta_{switchTo}$ from the previous chapter. We have omitted some syntactical details (like access to fields of records, etc) to stress only on a whole picture and give an impression how is it done in Isabelle (see Figure 4.2).

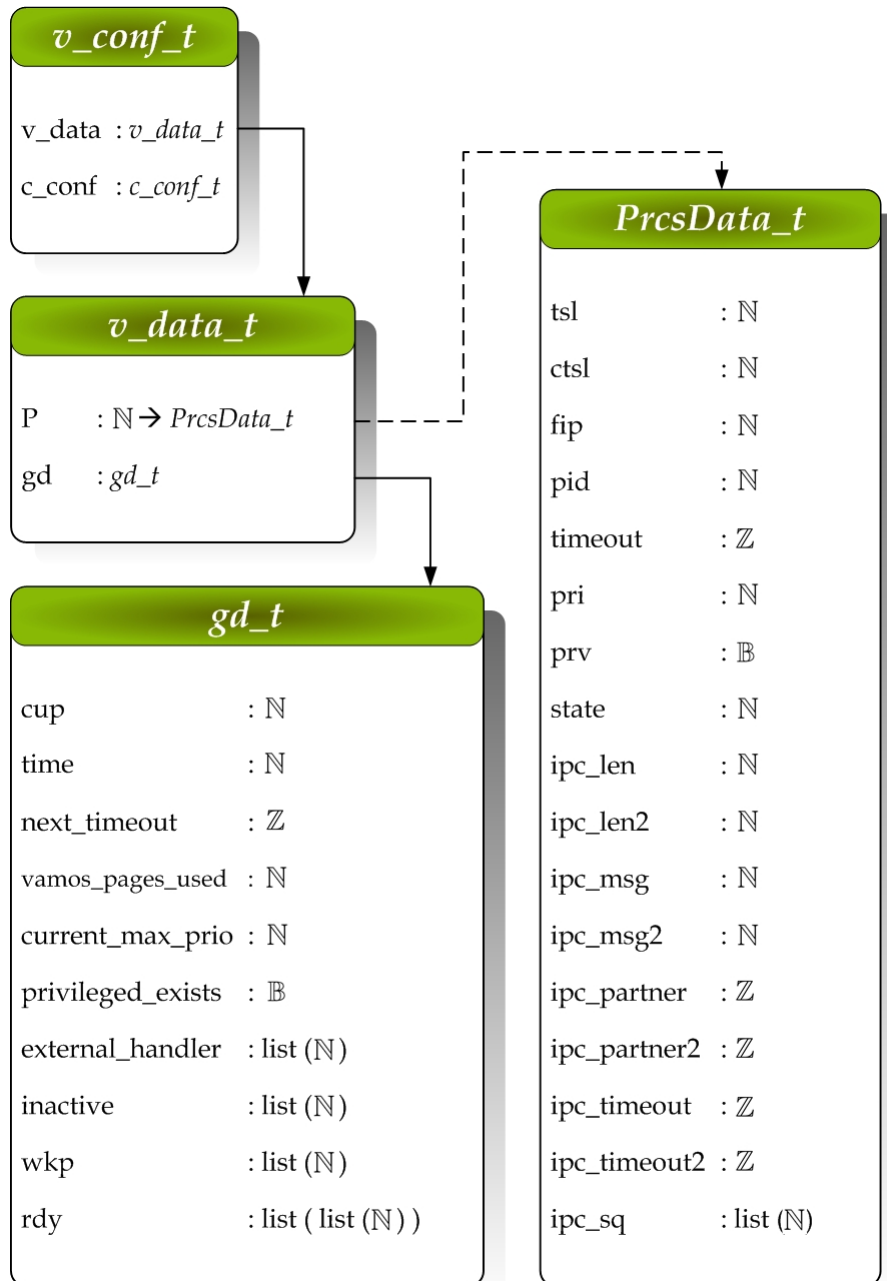


Figure 4.1: The VAMOS Configuration.

```

constdefs  abstr_switch_to :: "[v_conf_t, nat] → v_conf_t"
"
abstr_switch_to v p ==
  let
    current = v.cup;
    current_PrcsData = v.P(current);
    p_PrcsData = v.P(p);
    cup_tsl = current_PrcsData.tsl;
    cup_ctsl = current_PrcsData.ctsl;
    p_ctsl = p_PrcsData.ctsl;
    cup_is_prv = current_PrcsData.prv;
    p_is_ready = (p_PrcsData.state = ready);
    p_is_invalid = (PID_MAX ≤ p ∨ p = 0);
    cup_pri = current_PrcsData.pri
  in
    if (current = p)
      then v
      else if ((¬p_is_invalid) ∧ cup_is_prv ∧ p_is_ready)
        then if (p_ctsl < (cup_tsl - cup_ctsl))
          then v.rdy! cup_pri := (filter (λy. y ≠ current) (v.rdy! cup_pri)) @ [current],
            v.cup := p,
            v.P(current).ctsl := 0,
            v.P(p).ctsl := 0,
          else v.rdy! cup_pri := (filter (λy. y ≠ current) (v.rdy! cup_pri)) @ [current]
            v.cup := p,
            v.P(current).ctsl := 0,
            v.P(p).ctsl := (p_ctsl - (cup_tsl - cup_ctsl))
        else v
"

```

Figure 4.2: System call *processSwitchTo* modelled in Isabelle.

Chapter 5

Implementation

The VAMOS microkernel is located between the Communicating Virtual Machines layer (CVM) and the Simple Operating System layer (SOS). VAMOS uses functionality provided by the CVM layer via so-called CVM-primitives. Functionality provided by the VAMOS microkernel is available for the SOS layer via VAMOS system calls.

The VAMOS microkernel implementation delivers a dispatcher invoking interrupt handlers or functions realizing the services provided to the users. The decision is event-dependent.

The VAMOS operating system is implemented in the C0 programming language [5]. In this chapter we outline the important nuances of the implementation.

5.1 The Programming Language C0

The VAMOS operating system is implemented in C0, which is a subset of the C programming language. C0 differs from standard C, but the grammar is defined unambiguously. The formal C0 semantics is much shorter than the one for C. This makes program verification simpler. Moreover, programs written in C can be translated to C0.

The C0 programming language has the following restrictions [5]:

- no side-effects inside expressions (e.g. function calls, increments like `i++`, etc.);
- the size of arrays is fixed at the compile time;
- in every function is only one `return` statement, which must be the last statement of the function body;
- no pointer arithmetic;

- no pointers to local variables;
- no function pointers;
- only typed pointers are allowed (thus, no `void` pointers);
- no variable declarations in functions after the first statement;
- no initialization during declaration (except declarations of constants);
- only one variable scope inside a function (i.e. no declarations inside nested `{...}` blocks).

The C0 programming language has following built-in types [12]:

- 32-bit signed integers: `int = {-231, ..., 231 - 1}`;
- 32-bit unsigned integers: `unsigned int = {0, ..., 232 - 1}`;
- Boolean: `bool = {true, false}`;
- 8-bit signed integers: `char = {-128, ..., 127}`.

Based on these simple types one can construct compound types[12]:

- Typed pointers: `typ *x`;
- Arrays: `typ a[size]`;
- Structures: `struct styp { typ data }`;

Here `typ` stands for the basic or compound type, `styp` — for the newly declared structural type, and `size` — for an integer constant denoting size of an array.

5.2 Data Structures

The implementation of the VAMOS data structures is depicted in Figure 5.1. In this section we explain how it corresponds to the *VAMOS Configuration* (see Figure 4.1).

Using the type system of C0 we can implement natural numbers \mathbb{N} as **unsigned int**'s, integer numbers \mathbb{Z} as **int**'s, and boolean \mathbb{B} as **bool**'s. Records are implemented as C0-structures.

For what concerns lists, we use the *dList* package developed in the frame of the Verisoft project. With this package we only declare doubly linked list. Then C0-macros declares and defines functions working with these lists. At the time of writing the package includes the following functions [10]:

| Global Variables | | struct pib_t | |
|---------------------|-----------------------|---------------------|-----------------|
| pib_p | current_process | unsigned int | timeslice |
| pib_p | PIB[PID_MAX] | unsigned int | consumed_time |
| unsigned int | time | unsigned int | fip |
| int | next_timeout | unsigned int | pid |
| unsigned int | vamos_pages_used | int | timeout |
| unsigned int | current_max_prio | unsigned int | priority |
| bool | privileged_exists | bool | privileged |
| unsigned int | external_handler[22] | unsigned int | state |
| pib_p | inactive_list | unsigned int | ipc_len_p |
| pib_p | wakeup_list | unsigned int | ipc_len_p2 |
| pib_p | ready_lists[MAX_PRIO] | unsigned int | ipc_msg_p |
| | | unsigned int | ipc_msg_p2 |
| | | int | ipc_partner |
| | | int | ipc_partner2 |
| | | int | ipc_timeout |
| | | int | ipc_timeout2 |
| | | pib_p | queue_next |
| | | pib_p | queue_prev |
| | | pib_p | send_queue |
| | | pib_p | send_queue_next |
| | | pib_p | send_queue_prev |

```

#define PID_MAX 128u
#define MAX_PRIO 3u
typedef struct pib_t *pib_p

```

Figure 5.1: Implementation of the VAMOS data structures in C0.

- `dList_New()` — empty list constructor;
- `dList_InsertHead()` — list constructor (appends provided element to the head of a list);
- `dList_InsertAfter()` — inserts provided element at certain position in a list;
- `dList_GetElement()` — fetches an element from a list (invoker provides the index of the element as an input parameter);
- `dList_Append()` — appends two provided lists;
- `dList_Delete()` — deletes an element from a list;
- `dList_Length()` — returns the length of a list.

Elements of `dLists` are C0-structures. In order to declare a list (and then use functions from the `dList` package) one has to provide the following:

- Declare some structural type `STyp`;
- Within this structural type fields containing valuable information have to be declared. For example, if one declares a list of integers:

```
struct STyp { int i } ;
```

- Supplement this structural type with pointers `prev` and `next`. Their type is "`STyp *`" and they point to neighbors of the element in the list.

For the previous example:

```
struct STyp {
    int i;
    STyp *prev;
    STyp *next };
```

- Declare a pointer that will store the head of a list. The type of this pointer is "`STyp *`".

Thus, on the abstract layer we argue only about elements in a list, but on the implementation side we perform operations with pointers on the heap. The `dList` package not only contains macros to define and declare functions from above, but also their specifications in Isabelle. At the time of writing both partial and total correctness are shown [8].

5.2.1 Process Information Block

Information about the particular processes is contained in the *Process Information Block* (*PIB*) structure `pib_t`. Its members directly correspond to fields of *PrceData.t* record (see Figure 5.2).

Additionally, PIB contains pointers `queue_next` and `queue_prev` to the structure of same type `pib_t`. These pointers are used by functions working on dLists (see above).

Pointer `send_queue` stores to the head of the list of processes waiting for IPC communication with this process, the so-called *Send Queue* (in the *VAMOS Configuration* it is *ipc_sq*). Pointers `send_queue_next` and `send_queue_prev` are used to build this dList.

5.2.2 Global Variables

The global variables of VAMOS implement the *VAMOS Data* record (see Figure 5.3) with some nuances:

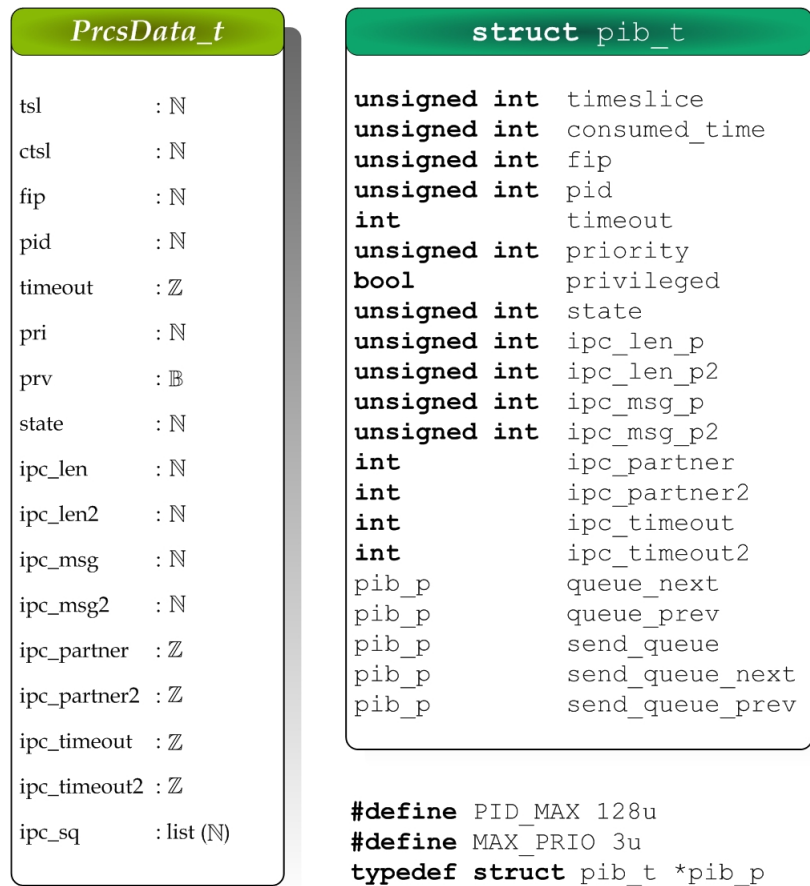
- Instead of the function *P* we have an array of the PIBs;
- The reference `current_process` points to the PIB of the currently running process;
- The list of external handlers is implemented as an array;
- We store the heads of dLists in pointers `inactive` and `wakeup_list` for the *Inactive* and the *Sleeping List* respectively. The array `ready_lists` stores the heads of the ready lists implementing the *Ready Lists Array*. Its dimension is determined by the global constant `MAX_PRIO`.

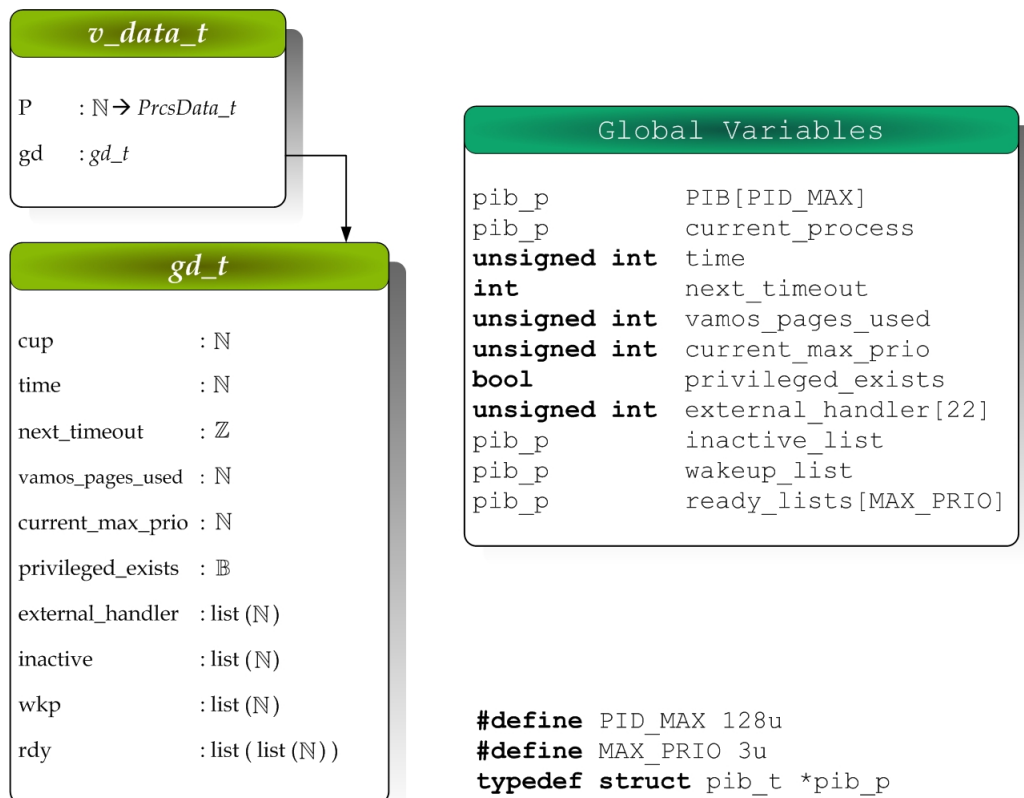
5.3 VAMOS System Calls

VAMOS System Calls are implemented by functions with almost the same names (compare Figure 5.4 with Table 2.1).

The implementation is modularized very well:

- frequent operations are carried out into helper functions (like `compute_max_prio`, `wake_up`, etc.);
- for operations on doubly linked lists functions from the *dList* package are invoked;
- some VAMOS system calls use *CVM-primitives* to change the state of the Virtual Machine.

Figure 5.2: Process information block structure implementing the record *PrdsData_t*.

Figure 5.3: Global variables implementing the record `gd_t`.

All these functions are depicted in Figure 5.4.

The functions `InsertTail` and `Rotate` are not part of the *dList* package. They are specified and verified in the context of this thesis.

Our goal is to specify and verify the VAMOS functions relevant to the VAMOS scheduler.

Since it is unfeasible to verify the function before specifying and verifying functions it calls, we have to consider their static call graph (see Figure 5.5). For example, the function `process_change_sched_param` implementing the system call *processChangeSchedulingParam* invokes: the function `Delete`, which is specified and verified in *dList* package; functions `InsertTail` and `compute_max_prio` that we have to specify and verify.

In the next chapter we specify some VAMOS functions in the order that does not violate calling dependencies.

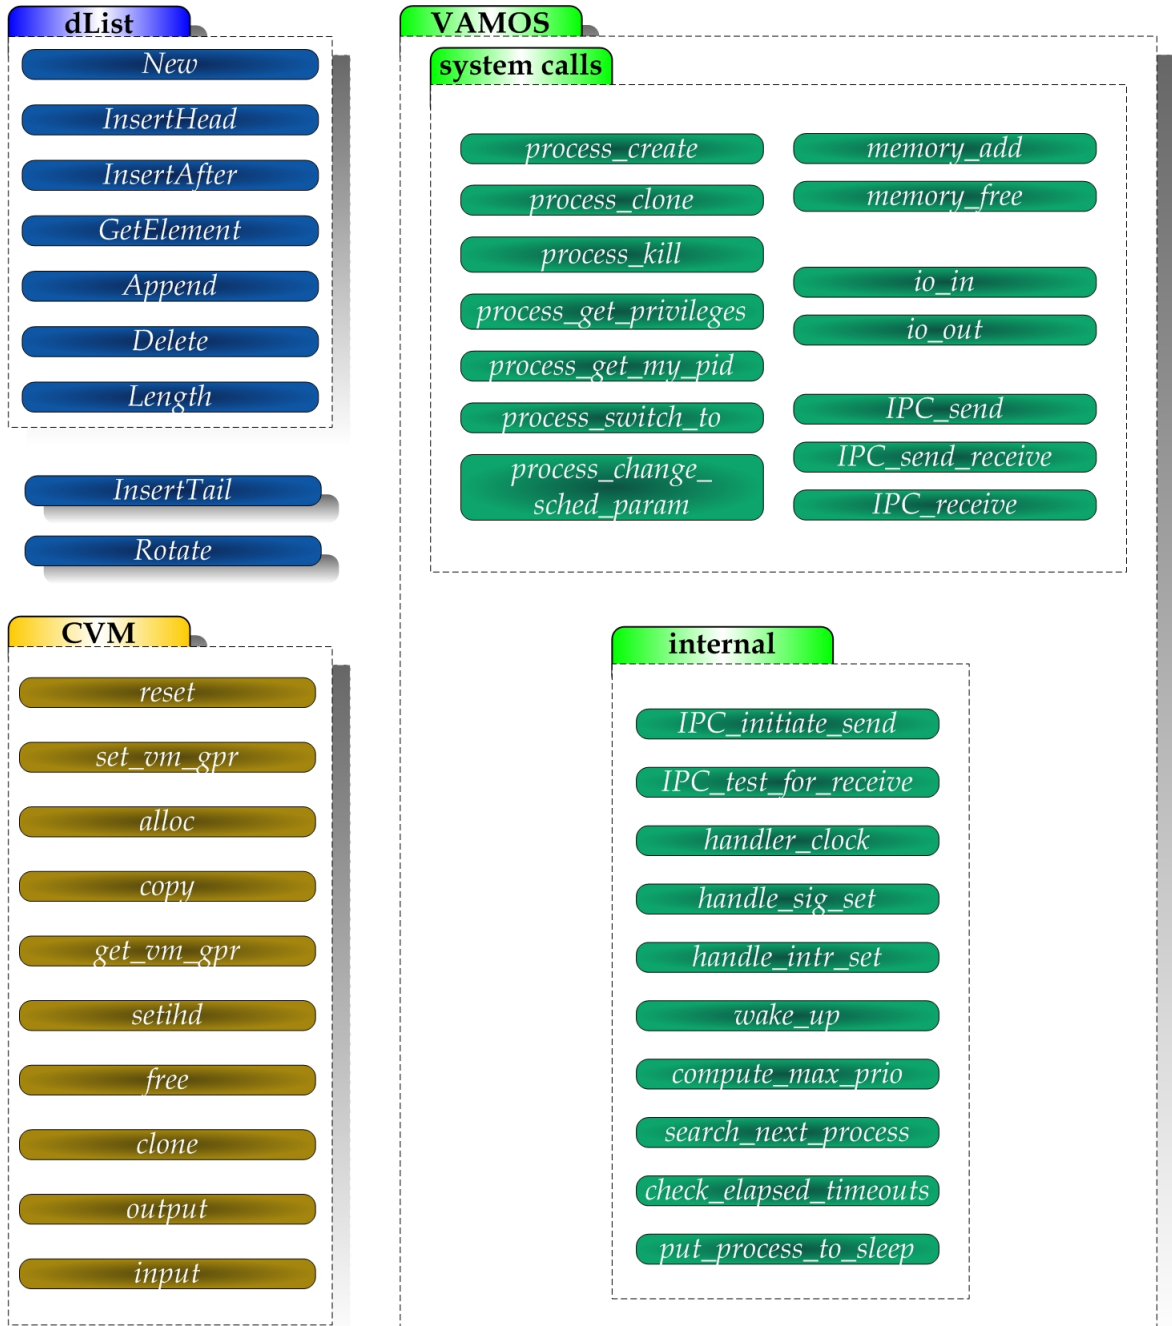


Figure 5.4: VAMOS functions.

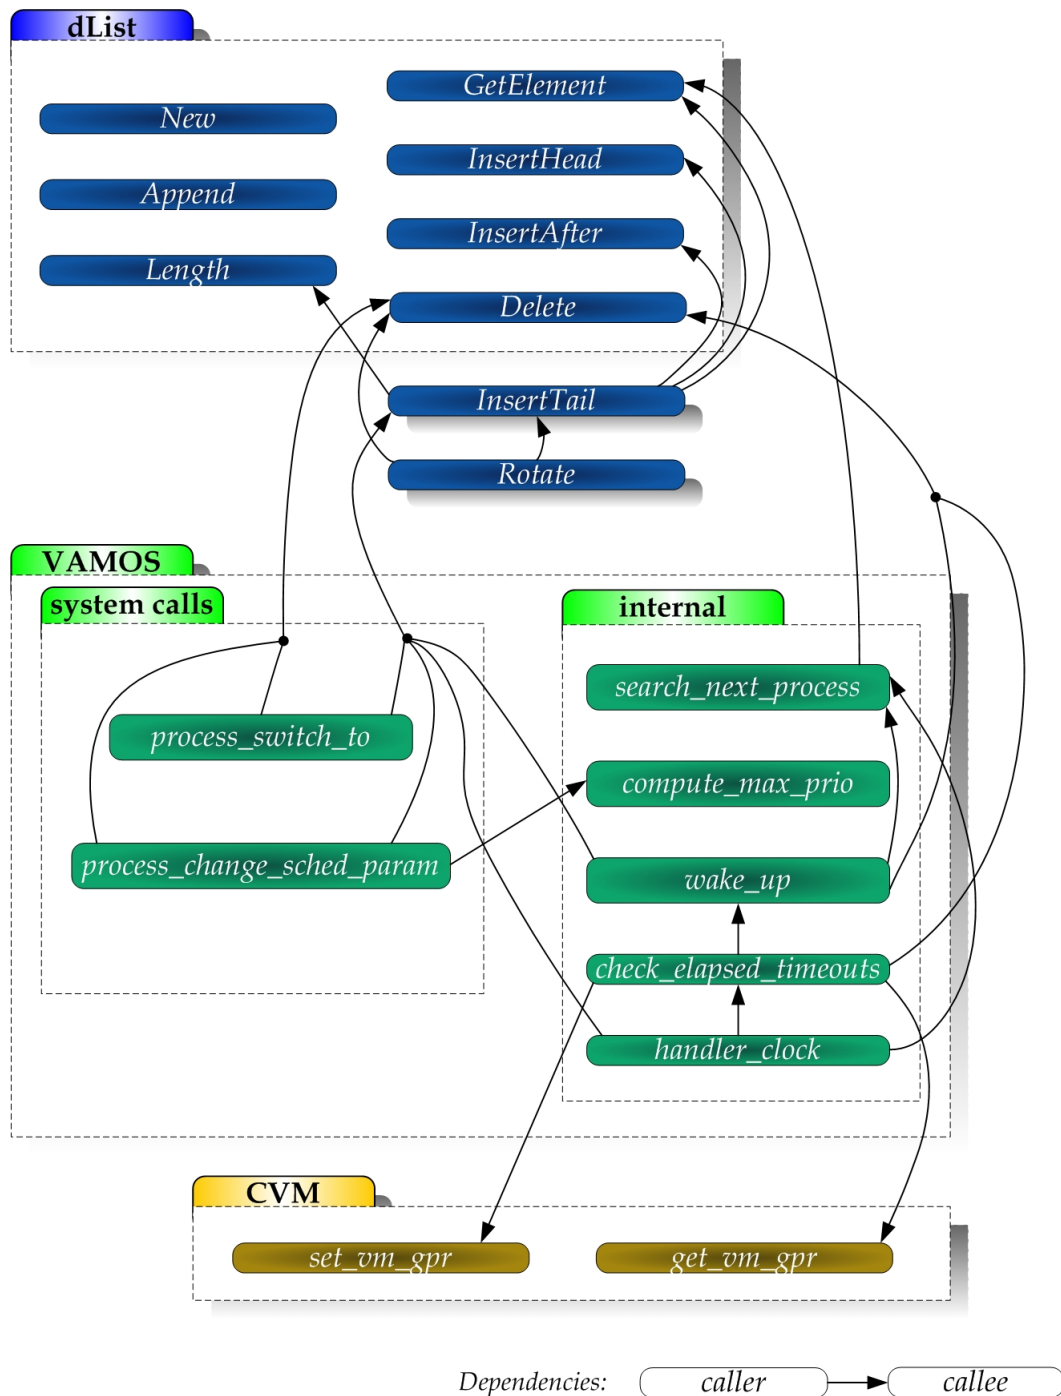


Figure 5.5: Static call graph of functions relevant to the VAMOS scheduler.

Chapter 6

Verification Environment

In the Verisoft project we use verification environment developed by N. Schirmer [11]. This chapter outlines main facilities it provides.

6.1 Hoare Logic for Partial Correctness

In the Verisoft project we use Hoare Logic to verify programs.

Hoare Logic was originally invented by C. A. R. Hoare and in our days widely used for program verification. In Hoare Logic, consequences of the program run are obtained from the program code using so-called inference rules (refer, for instance, to [7]).

Hoare Logic argues on the configuration of the program. Implementation is seen at a coarse-grained level. Implementation data structures and global variables constitute a so-called *state space*. We define the *state space* as a tuple of all the implementation components: global and local variables, and structures with unfolded fields. The actual code is translated into a logical representation, i.e. every statement is mapped to a corresponding inference rule.

In order to formally specify a program in Hoare Logic one has to provide pre- and post-conditions. These are merely assertions on a state space. Then partial correctness is defined as follows¹: a program is correct, if every terminating computation of the program that starts in a state satisfying its precondition terminates in a state satisfying its postcondition.

¹This definition is taken from [12], where you can find more detailed information.

6.2 Abstraction Relations

Pure Hoare Logic argues only on the implementation layer. But we want to relate abstract and implementation layers. This is done using so-called *abstraction relations*. An abstraction relation is merely a predicate taking certain input parameters from both abstract and implementation layers. If it holds, the specified relation exists (this relation is called a *simulation relation*, see [2] for details). Nevertheless, an abstraction relation is not pure simulation relation, because we expand it with validity statements.

6.3 Imperative Programming Language

In fact, abstraction relations bind abstract data structures not directly to the implementation. Instead we using a notion of the state space. In further we relate the *VAMOS Configuration* with the VAMOS state space.

C0-implementation of VAMOS is processed by a translator [8]. This translator processes C0 data structures and returns the VAMOS state space consisting of:

- global variables, whose names are changed to $glob_⟨variable\ name⟩$;
- local variables, whose names are changed to $⟨function\ name⟩_⟨variable\ name⟩$;
- functions, whose names are changed to $fun_⟨function\ name⟩$;
- unfolded C0-structures — tuples of their fields.

For what concerns C0-program itself, it is translated into the imperative programming language. Imperative programs consist of program constructors that change the state space. In the following we present their syntax².

6.3.1 Assignments

The assignment statement of an expression $expr$ to a variable Var has the following syntax:

$$Var ::= expr$$

where

$$expr, Var : \mathbb{T}.$$

²The forthcoming subsections are partially taken from [8], where you can find more detailed information.

An assignment of an expression $expr$ to a field fld of a structure that is pointed to by a reference r is written as³:

$$r \rightarrow fld ::= expr$$

where

$$r : ref,$$

$$expr, r \rightarrow fld : \mathbb{T};$$

6.3.2 Conditional Statement

A conditional statement has the following syntax:

$$IF \textit{bool_expr} \textit{ THEN } \textit{statements}_1 [ELSE \textit{statements}_2] FI$$

where

$$\textit{bool_expr} : \mathbb{B}.$$

Here, $\textit{statements}_1$ and $\textit{statements}_2$ are sequences of program constructs separated by “;”. The code in square brackets can be omitted. Decision is made depending on the boolean expression $\textit{bool_expr}$. If it holds, the sequence $\textit{statements}_1$ is executed. Otherwise, the sequence $\textit{statements}_2$ is executed.

6.3.3 Loop

A while loop has the following syntax:

$$WHILE \textit{bool_expr} DO \textit{statements} OD$$

where

$$\textit{bool_expr} : \mathbb{B}$$

The sequence $\textit{statements}$ is executed as long as the boolean expression $\textit{bool_expr}$ holds.

³As it was mentioned above, C0-structures are unfolded. This results in the heap functions $ref \rightarrow type$ for every type of structures' fields. Here ref is type of pointers. It is similar to the natural numbers type, but additionally contains the constant $Null$ denoting undefined pointer. All definitions and explanations can be found in [8].

6.3.4 Procedure Call

A procedure call has the following syntax:

$$Var ::= CALL \textit{procedure_name} (in_1, in_2 \dots, in_n)$$

where

$$Var : \mathbb{T}$$

Here the value returned by a procedure *procedure_name* is assigned to a variable *Var*. Intuitively, they must have the same type. Input parameters are given in brackets, and have to be separated by a comma.

6.4 VCG

The basic idea about verification condition generator (VCG) for Hoare Logic is simple. Given an imperative program *c*, precondition, postcondition, and some state space *S* we automatically apply the Hoare rules until we end up in a configuration where the program *c* is completely eliminated and a purely logical proof obligation remains. N. Schirmer has implemented a verification condition generator in Isabelle.

The VCG cannot infer final state space of programs with while loops. Therefore we have to annotate such programs with loop invariants. This is done using *INV* construct.

Chapter 7

Functional Correctness

7.1 Functions on dLists

The VAMOS function perform two very frequent operations on dLists. These operations were encapsulated into two helper functions: `queue_InsertTail` and `queue_Rotate`. Their specifications are presented in this section.

7.1.1 The Function `queue_InsertTail()`

This function is intended to enqueue a process into the end of a list.

Signature: $queue_InsertTail_{abstr} : list(\mathbb{PID}) \times \mathbb{PID} \longrightarrow list(\mathbb{PID})$

Let us define the updated list:

$$Rs = queue_InsertTail_{abstr}(Ps, prcs),$$

where

$$prcs : \mathbb{PID},$$

$$Ps, Rs : list(\mathbb{PID}).$$

Precondition: The provided process is not in the target list.

$$prcs \notin Ps.$$

Postcondition: The provided list is updated accordingly:

$$Rs = (Ps @ [prcs]).$$

The implementation of this function in C0 is presented in Appendix B.1. Figure 7.1 presents the function code translated into the imperative programming language and state space for this function. Here the variables marked with the acute prefix (') refer to their values from the current state space (in other words, at this particular point in the program).

This function invokes functions from the *dList* package, that are already specified and verified. Hence proof is simple. With 9 auxiliary lemmata proven the functional verification takes 82 proof steps in Isabelle.

7.1.2 The Function `queue_Rotate()`

This function is intended to enqueue the head of a list to its end (thus, “rotate” the list).

Signature: $queue_Rotate_{abstr} : list(\mathbb{PID}) \longrightarrow list(\mathbb{PID})$

Let us define the updated list:

$$Rs = queue_Rotate_{abstr}(Ps),$$

where

$$Ps, Rs : list(\mathbb{PID}).$$

Precondition: The provided list is not empty.

$$Ps \neq [].$$

Postcondition: The provided list is updated accordingly:

$$Rs = ((tl Ps) @ [hd Ps]).$$

The implementation of this function in C0 is presented in Appendix B.2. Figure 7.2 presents the function code translated into the imperative programming language and state space for this function.

state space

```

res_queue_InsertTail :: "ref"
queue_InsertTail_dummy :: "ref"
queue_InsertTail_result :: "ref"
queue_InsertTail_dummy_int :: "int"
queue_InsertTail_len :: "nat"
queue_InsertTail_p :: "ref"
queue_InsertTail_list :: "ref"

```

procedures

```

fun_queue_InsertTail (queue_InsertTail_list, queue_InsertTail_p | res_queue_InsertTail) =
  "
  'queue_InsertTail_len ::= CALL fun_dlist_pib.t_queue_Length('queue_InsertTail_list);
  IF 'queue_InsertTail_len = 0 THEN
    'queue_InsertTail_result ::= CALL fun_dlist_pib.t_queue_InsertHead
      ('queue_InsertTail_list, 'queue_InsertTail_p)
  ELSE
    'queue_InsertTail_dummy ::= CALL fun_dlist_pib.t_queue_GetElement
      ('queue_InsertTail_list, 'queue_InsertTail_len - 1);;
    'queue_InsertTail_dummy_int ::= CALL fun_dlist_pib.t_queue_InsertAfter
      ('queue_InsertTail_dummy, 'queue_InsertTail_p);;
    'queue_InsertTail_result ::= 'queue_InsertTail_list
  FI;;
  'res_queue_InsertTail ::= 'queue_InsertTail_result
  "

```

Figure 7.1: The function `queue_InsertTail()` translated into the imperative programming language.

state space

```
res_queue_Rotate :: "ref"
```

```
queue_Rotate_result :: "ref"
```

```
queue_Rotate_list :: "ref"
```

procedures

```
fun_queue_Rotate (queue_Rotate_list | res_queue_Rotate) =
```

```
"
```

```
'queue_Rotate_result := CALL fun_dlist_pib_t_queue_Delete
```

```
  ('queue_Rotate_list, 'queue_Rotate_list);;
```

```
'queue_Rotate_result := CALL fun_queue_InsertTail
```

```
  ('queue_Rotate_result, 'queue_Rotate_list);;
```

```
'res_queue_Rotate := 'queue_Rotate_result
```

```
"
```

Figure 7.2: The function `queue.Rotate()` translated into the imperative programming language.

This function invokes the function `dList.Delete()` from the *dList* package and the function `queue.InsertTail()` from the previous section. These are already specified and verified. The functional verification takes 32 proof steps in Isabelle.

7.2 The *VamosData?* Predicate

In order to argue about functional correctness of VAMOS we have to state how the VAMOS functions change the *VAMOS Configuration*. We introduce the main abstraction relation *VamosConf?*. This abstraction relation comprises the following statements:

- 1) it relates abstract and implementation layers;
- 2) it comprises properties of validity of the current *VAMOS Configuration*;
- 3) it relates VAMOS and CVM layers.

The first and second statements are combined into the predicate *VamosData?*. The third one is realized by the predicates *CvmConf?* and *CvmRelVamos?*.

Let v be a *VAMOS Configuration* and S a state space. We assume that the predicate *VamosConf?* holds.

For any VAMOS function f we have its specification f_{abstr} that changes the *VAMOS Configuration* (in the format that is presented in Section 4.3). Applying this function to v we obtain a subsequent *VAMOS Configuration* v' .

Suppose that the function f is implemented in C0 and its body changes the state space S to some state space S' . In the scope of functional correctness we want to assure that v' is a valid *VAMOS Configuration* and that it is related to the state space S' . Therefore we require that the predicate *VamosConf?* holds for the updated *VAMOS Configuration* v' and the state space S' :

Let

$$VamosConf? : v_conf_t \times state_space_t \longrightarrow \mathbb{B};$$

$$v, v' : v_conf_t,$$

$$S, S' : state_space_t,$$

$$VamosConf?(v, S) = True,$$

and function f is defined: $v' = f_{abstr}(v, input_{abstr}),$

and subsequent state space is: $S',$

then $VamosConf? v' S' = True.$

In the frame of this thesis we are only interested in the *VamosData?* predicate, because the VAMOS scheduler does not affect the CVM layer. Henceforth we no longer refer to the whole *VAMOS Configuration*, but only to the *VAMOS Data* part. So, the VAMOS functions take a *VAMOS Data* record d and several input parameters. They usually return updated *VAMOS Data* record d' . We refer to the pair (d, S) as pre-state and to the pair (d', S') as post-state of the function call (see Figure 7.3):

Let $d, d' : v_data_t, S, S' : state_space_t,$

$$VamosData?(d, S) = True,$$

and function f is defined: $d' = f_{abstr}(d, input_{abstr}),$

and subsequent state space is: $S',$

then f is correct if: $VamosData?(d', S') = True.$

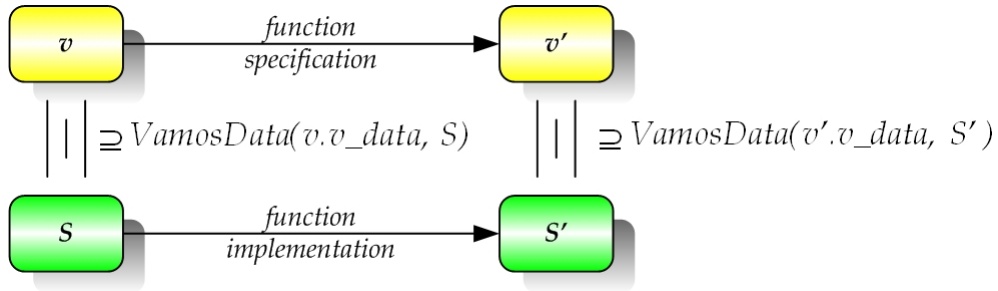


Figure 7.3: Functional correctness.

The *VamosData?* is very complex and is not presented here completely. Instead we emphasize only on those invariants that are relevant to the VAMOS scheduler. Moreover, we stress only on invariants about validity of the *VAMOS Configuration* (thus, nothing about support of the simulation relation).

7.3 Assumptions

In this section we present invariants, which we assume in the pre-state. Most of them are incorporated into the *VamosData?* predicate. Two invariants are stand-alone, because they do not hold for all VAMOS functions. Note that we focus only on invariants relevant to the VAMOS scheduler and stating about validity of the *VAMOS Data*.

7.3.1 Assumptions from the *VamosData?* Predicate

Disjointness of lists of VAMOS

We want to ensure the following:

- 1) a process cannot be in the same list twice (i.e. list contains distinct elements);
- 2) a process cannot be in two VAMOS lists simultaneously.

In fact, the VAMOS lists are implemented using the dList package. From their definition it follows that VAMOS lists are distinct (refer to [10]) and the first goal is obtained.

For the second we remember that every VAMOS process is in one of three states: **ready**, **inactive** or **sleeping**.

So, we relate the state of the process with its list membership:

$$\begin{aligned}
& \forall i : \mathbb{N}. \\
& \left(i < PID_MAX \implies \right. \\
& \quad (d.gd.P(i).state = \mathbf{inactive} \iff i \in d.gd.inactive) \\
& \quad \wedge (d.gd.P(i).state = \mathbf{ready} \iff i \in (d.gd.rdy ! P(i).pri)) \\
& \quad \wedge (d.gd.P(i).state = \mathbf{sleeping} \iff i \in d.gd.wkp) \\
& \quad \left. \wedge (\forall j : \mathbb{N}. j \neq d.gd.P(i).pri \implies i \notin (d.gd.rdy ! j)) \right). \tag{7.1}
\end{aligned}$$

Current process is ready

Another assumption is that the current process is ready. In other words, we want to assure that dead and sleeping processes are not scheduled.

$$d.gd.P(d.gd.cup).state = \mathbf{ready}. \tag{7.2}$$

The Ready Lists Array is not empty

As it was mentioned in Section 3.1 the *idle* process was introduced to overcome the situation, when all the ready lists are empty. The *idle* process state cannot be affected by any of the system calls (e.g. it cannot be killed). The *idle* process has minimum priority. Thus, it is enough to state that the ready list with minimum priority is not empty:

$$d.gd.rdy ! 0 \neq []. \tag{7.3}$$

Length of the Ready Lists Array

The number of the ready lists is equal to the global variable *MAX_PRIO*:

$$length(d.gd.rdy) = MAX_PRIO. \tag{7.4}$$

Priorities of processes are in correct range

Since priorities are modeled as \mathbb{N} we assure they are less than *MAX_PRIO*:

$$\forall prcs : \mathbb{N}. (prcs < PID_MAX \implies d.gd.P(prcs).pri < MAX_PRIO). \tag{7.5}$$

The *current_max_prio* is in correct range

The *current_max_prio* is less than *MAX_PRIO* (just like any other priority):

$$d.gd.current_max_prio < MAX_PRIO. \quad (7.6)$$

7.3.2 Stand-alone Assumptions

We require the predicate *VamosData?* as a precondition for every VAMOS function. Here are presented additional assumptions that cannot be embedded into this predicate, because they do not hold for all VAMOS functions.

The *current_max_prio* is correct

We have to guarantee two things about the *current_max_prio* value:

- the ready list indexed by it is not empty;
- there is no non-empty ready list with higher priority than *current_max_prio*.

We formulate this as follows¹:

$$cmp_correct? : v_data.t \longrightarrow \mathbb{B};$$

$$cmp_correct?(d) \equiv \left(c.gv.current_max_prio = Max\{x : \mathbb{N} \mid x = d.P(prcs).pri, \quad prcs \in Rda\} \right),$$

where

$$Rda \equiv \{x : \mathbb{N} \mid d.P(x).state = \mathbf{ready}\}.$$

The *next_timeout* is correct

We have to guarantee that the *next_timeout* is the minimum timeout over all processes in the *Sleeping List*. Let us define the *nt_correct?* predicate:

$$nt_correct? : v_data.t \longrightarrow \mathbb{B};$$

$$nt_correct?(d) \equiv \left(c.gv.next_timeout = Min\{x : \mathbb{N} \mid x = d.P(prcs).next_timeout, \quad prcs \in d.gd.wkp\} \right)$$

7.4 The VAMOS Functions

We outline here the specifications of the VAMOS functions relevant to the VAMOS scheduler. For every function call we give:

¹Here the function *Max* computes the maximum over a set.

- a short description on what this function calls is intended for;
- its signature — i.e. how it operates on the abstract data structures;
- its precondition — i.e. restrictions on the *VAMOS Configuration* before an invocation of a function call. If a precondition does not hold, the *VAMOS Configuration* is not updated;
- its postcondition — i.e. statements about modifications in *VAMOS Configuration* depending on input parameters.

Functions presented here operate solely on the VAMOS layer, not touching CVM configuration. Thus, the correctness of these function calls can be specified and proven without arguing on the CVM layer. Therefore these function calls are modeled as transition functions operating on the *VAMOS Data*, but not on the whole *VAMOS Configuration*.

The VAMOS functions take a *VAMOS Data* record d and several input parameters $(in_1, in_2, \dots, in_n)$ and return an updated *VAMOS Data* record d' .

We refer to the pair (d, S) as pre-state and to the pair (d', S') as post-state of function call.

Let

$$d : v_data.t, S : state_space.t;$$

and function f is defined:

$$d' = f_{abstr}(d, in_1, in_2, \dots, in_n).$$

7.4.1 The Function `search_next_process()`

The `search_next_process()` function is intended to give the CPU control to the next ready process. That is the head of the ready list indexed by the *current_max_prio*. With assumption *cmp_correct?* we are sure that this concrete ready list is not empty, so we can refer to the first element in it. Therefore this function merely assigns the *cup* to the PID of that process.

There are no input parameters for this function call.

Signature: $search_next_process_{abstr} : v_data.t \longrightarrow v_data.t.$

Let d' be the updated *VAMOS Data*:

$$d' \equiv search_next_process_{abstr}(d).$$

procedures

```

fun_search_next_process( | res_search_next_process ) =
  ”
  ’glob_current_process ::=
    CALL fun_dlist_pib_t_queue_GetElement(’glob_ready_lists ! (’glob_current_max_prio),0));
  ’res_search_next_process ::= 0
  ”

```

Figure 7.4: The function `search_next_process()` translated into the imperative programming language.

Precondition: As precondition we have only assumptions from Section 7.3:

$$VamosData?(d, S) \wedge cmp_correct?(d) \wedge nt_correct?(d).$$

Postcondition: As postcondition we require the current process to be changed accordingly and, of course, the *VAMOS Data* to be valid:

$$\begin{aligned}
 & (d'.gd.cup = hd(d.gd.rdy ! d.gd.current_max_prio)) \\
 & \wedge VamosData?(d', S') \wedge cmp_correct?(d') \wedge nt_correct?(d').
 \end{aligned}$$

The implementation of this function in C0 is presented in Appendix B.3. The functional correctness is shown for the whole *VAMOS* state space, which is very big and is not presented here. Figure 7.4 presents the function code translated into the imperative programming language. It is easy to see, that the current process is changed to the first element in the ready list indexed by the value of the current maximum priority. Since the only invoked function `dList_GetElement` is already verified, functional verification of the function `search_next_process` is easy. It takes 77 proof steps in Isabelle.

7.4.2 The Function `compute_max_prio()`

The `compute_max_prio()` function is used to compute the maximum priority over all ready processes for a certain *VAMOS Data*. This function only computes the value and does not update the *VAMOS Data*. Taking a *VAMOS Data* record as an input parameter it returns a natural number indicating the maximum priority.

Signature: $compute_max_prio_abstr : v_data_t \longrightarrow \mathbb{N}$

Precondition is that the provided *VAMOS Data* record is valid:

$$VamosData?(d, S).$$

Postcondition: As a postcondition we require that this function returns the index of the non-empty list with the highest priority:

$$\left(compute_max_prio_abstr(d) = Max\{i : \mathbb{N} \mid i < MAX_PRIO \wedge d.gd.rdy ! i \neq []\} \right) \\ \wedge VamosData?(d, S')$$

The implementation of this function in C0 is presented in Appendix B.4. Figure 7.5 presents the function code translated into the imperative programming language. The imperative program contains a while-loop. Therefore we have to provide loop invariants. This is done using *INV* construct (see Figure 7.6). Here the variables marked with the prefix σ refer to their values from the pre-state.

After providing loop invariant, we have to prove the following:

- 1) the precondition implies the loop invariant,
- 2) the invariant is maintained inside the loop, and
- 3) the loop invariant together with the loop body imply the postcondition.

The functional verification for this function takes 138 proof steps in Isabelle environment.

7.4.3 The Function `wake_up()`

The `wake_up()` function is intended to remove a process from the *Sleeping List* and place it into the corresponding ready list. Thus, the state of the awaking process is changed to *ready*. If the priority of the awaking process is greater than the *current_max_prio* we have to update the *current_max_prio* and switch to that process.

The only input parameter for this function is the PID of the target process.

Signature: $wake_up_abstr : v_data_t \times \mathbb{N} \longrightarrow v_data_t$

procedures

```

fun_compute_max_prio( | res_compute_max_prio ) =
  ”
  ’compute_max_prio_prio := 0;;
  ’compute_max_prio_i := 3 - 1;;
  WHILE 0 < ’compute_max_prio_i DO
    IF ’glob_ready_lists ! (’compute_max_prio_i) ≠ Null THEN
      ’compute_max_prio_prio := ’compute_max_prio_i;;
      ’compute_max_prio_i := 0
    ELSE ’compute_max_prio_i := ’compute_max_prio_i - 1
  FI
  OD;;
  ’res_compute_max_prio := ’compute_max_prio_prio
  ”

```

Figure 7.5: Function `compute_max_prio()` translated into the imperative programming language.

Let us define:

$$\begin{aligned}
 d' &\equiv \text{wake_up_abstr}(d, \text{prcs}); \\
 CUP &\equiv d.gd.cup; \\
 \text{prcs_is_valid?} &\equiv (\text{prcs} \neq \text{IDLE} \wedge \text{prcs} < \text{PID_MAX}); \\
 \text{prcs_is_active?} &\equiv \overline{(d.P(\text{prcs}).\text{state} = \text{inactive})}; \\
 \text{prcs_is_ready?} &\equiv (d.P(\text{prcs}).\text{state} = \text{ready}); \\
 \text{new_pri_is_valid?} &\equiv (\text{new_pri} < \text{MAX_PRIO}); \\
 \text{cup_is_prv?} &\equiv d.P(CUP).prv;
 \end{aligned}$$

Precondition is that the target process is in the *Sleeping List*:

$$\text{prcs} \in d.gv.wkp \wedge \text{VamosData?}(d, S) \wedge \text{cmp_correct?}(d) \wedge \text{nt_correct?}(d).$$

procedures

```

fun_compute_max_prio( | res_compute_max_prio ) =
" 'compute_max_prio_prio ::= 0;;
  'compute_max_prio_i ::= 3 - 1;;
  WHILE 0 < 'compute_max_prio_i DO
  INV{
    (¬(0 < 'compute_max_prio_i)) ⇒
      (∀j.((j < MAX_PRIO) ∧ ('compute_max_prio_prio < j)) ⇒ ((σglob_ready_lists ! j) = Null))
      ∧ (∀i.((i < MAX_PRIO) ∧ (σglob_ready_lists ! i ≠ Null)) ⇒ (i ≤ 'compute_max_prio_prio))
      ∧ ((σglob_ready_lists ! 'compute_max_prio_prio) ≠ Null)
    ∧ (0 < 'compute_max_prio_i ⇒
      (∀j.((j < MAX_PRIO) ∧ ('compute_max_prio_i < j)) ⇒ (σglob_ready_lists ! j) = Null))
      ∧ ('compute_max_prio_prio = 0))
    ∧ ('res_compute_max_prio = σres_compute_max_prio)
    ∧ ('compute_max_prio_i < MAX_PRIO)
    ∧ (0 ≤ 'compute_max_prio_i)
    ∧ ('compute_max_prio_prio < MAX_PRIO)
    ∧ (0 ≤ 'compute_max_prio_prio)
  }
  IF 'glob_ready_lists!('compute_max_prio_i) ≠ Null
    THEN 'compute_max_prio_prio ::= 'compute_max_prio_i;;
      'compute_max_prio_i ::= 0
    ELSE 'compute_max_prio_i ::= 'compute_max_prio_i - 1
  FI
OD;;
'res_compute_max_prio ::= 'compute_max_prio_prio  "

```

Figure 7.6: Annotating the while loop with invariants.

Postcondition:

$$\begin{aligned}
& (d'.gd.wkp = d.gd.wkp \text{ Delete } prcs) \\
& \wedge \left((d'.gd.rdy \! d.P(prcs).pri) = ((d.gd.rdy \! d.P(prcs).pri) @ [prcs]) \right) \\
& \wedge (d'.P(prcs).state = \mathbf{ready}) \\
& \wedge \left(d.gd.current_max_prio < d.P(prcs).pri \implies \right. \\
& \quad \left. (d'.gd.current_max_prio = d.P(prcs).pri) \wedge (d'.gd.cup = prcs) \right) \\
& \wedge \text{VamosData?}(d', S') \wedge \text{cmp_correct?}(d') \wedge \text{nt_correct?}(d').
\end{aligned}$$

The implementation of this function in C0 is presented in Appendix B.5. Figure 7.7 presents the function code translated into the imperative programming language. This function invokes three other functions:

- 1) The function `dList_Delete()` from the `dList` package. Thus, we can rely on the correctness proofs from this package ;
- 2) The function `queue_InsertTail()` that enqueues a given element to the end of a list. This function was specified and verified in the context of this thesis;
- 3) The function `search_next_process()`.

The function `wake_up()` is not verified completely yet.

7.4.4 The Function `process_switch_to()`

The process invoking this call voluntarily gives the rest of its timeslice and the CPU control to a target process. The PID of the target process is provided as input parameter.

This system call was explained in details and modeled in Section 3.3.

Signature: $process_switch_to_{abstr} : v_data.t \times \mathbb{N} \longrightarrow v_data.t.$

Let us define some abbreviations:

$$\begin{aligned}
d' & \equiv process_switch_to_{abstr}(d, prcs); \\
CUP & \equiv d.gd.cup; \\
prcs_is_valid? & \equiv (prcs \neq \mathbf{idle} \wedge prcs < PID_MAX); \\
prcs_is_ready? & \equiv (d.P(prcs).state = \mathbf{ready}); \\
cup_is_prv? & \equiv d.P(CUP).prv
\end{aligned}$$

procedures

```

fun_wake_up(wake_up_p | res_wake_up) =
”
  'glob_wakeup_list ::= CALL fun_dlist_pib_t_queue_Delete ('glob_wakeup_list, 'wake_up_p);;
  'glob_ready_lists ! ('wake_up_p → 'pib_t_priority) ::=
    CALL fun_queue_InsertTail ('glob_ready_lists ! ('wake_up_p → 'pib_t_priority), 'wake_up_p);;
  'wake_up_p → 'pib_t_state ::= 1;;
  IF 'glob_current_max_prio < ('wake_up_p → 'pib_t_priority) THEN
    'glob_current_max_prio ::= ('wake_up_p → 'pib_t_priority);;
    'wake_up_dummy_int ::= CALL fun_search_next_process()
  FI;;
  'res_wake_up ::= 0
”

```

Figure 7.7: The function `wake_up()` translated into the imperative programming language.

Precondition: the currently running process has privileges, the provided PID is correct, and the target process is ready for execution:

$$\begin{aligned}
 & \text{cup_is_prv?} \wedge \text{pid_is_valid?} \wedge \text{pid_is_ready?} \\
 & \wedge \text{VamosData?}(d, S) \wedge \text{cmp_correct?}(d) \wedge \text{nt_correct?}(d).
 \end{aligned}$$

Postcondition:

$$\begin{aligned}
& (prcs = CUP \implies d' = d) \\
& \wedge \left(prcs \neq CUP \implies \right. \\
& \quad d'.gd.cup = prcs \\
& \quad \wedge d'.P(CUP).ctsl = 0 \\
& \quad \wedge \left((d'.gd.rdy \neq d.P(CUP).pri) = ((d.gd.rdy \neq d.P(CUP).pri) \text{ Delete } CUP) @ [CUP] \right) \\
& \quad \wedge (d.P(prcs).ctsl < (d.P(CUP).tsl - d.P(CUP).ctsl) \implies d'.P(prcs).ctsl = 0) \\
& \quad \wedge (d.P(prcs).ctsl \geq (d.P(CUP).tsl - d.P(CUP).ctsl) \implies \\
& \quad \quad \quad d'.P(prcs).ctsl = (d.P(CUP).tsl - d.P(CUP).ctsl)) \left. \right) \\
& \wedge \text{VamosData?}(d', S') \wedge \text{cmp_correct?}(d') \wedge \text{nt_correct?}(d').
\end{aligned}$$

The implementation of this function in C0 is presented in Appendix B.6. Figure 7.8 presents the function code translated into the imperative programming language. This function invokes the following two functions:

- 1) The function `dList_Delete()`;
- 2) The function `queue_InsertTail()`.

The function `process_switch_to()` function is not verified completely yet.

procedures

```

fun_process_switch_to (process_switch_to_pid | res_process_switch_to) =
  ”
  'process_switch_to_return_result ::= -1;;
  IF 'process_switch_to_pid = ('glob_current_process → 'pib_t_pid) THEN
    'process_switch_to_return_result ::= 0
  ELSE IF (128 ≤ 'process_switch_to_pid) ∨ ('process_switch_to_pid = 0) THEN
    'process_switch_to_return_result ::= -1
  ELSE IF ('glob_pib! ('process_switch_to_pid) → 'pib_t_state) ≠ 1 THEN
    'process_switch_to_return_result ::= -21
  ELSE IF ('glob_pib! ('process_switch_to_pid) → 'pib_t_consumed_time) <
    (('glob_current_process → 'pib_t_timeslice) -
    ('glob_current_process → 'pib_t_consumed_time)) THEN
    'glob_pib! ('process_switch_to_pid) → 'pib_t_consumed_time ::= 0
  ELSE 'glob_pib! ('process_switch_to_pid) → 'pib_t_consumed_time ::=
    ('glob_pib! ('process_switch_to_pid) → 'pib_t_consumed_time) -
    (('glob_current_process → 'pib_t_timeslice) - ('glob_current_process → 'pib_t_consumed_time))
  FI;;
  'glob_current_process → 'pib_t_consumed_time ::= 0;;
  'wake_up_dummy ::= CALL fun_dlist_pib_t_queue_Delete
    ('glob_ready_lists! (('glob_current_process → 'pib_t_priority), 'glob_current_process));
  'wake_up_dummy ::= CALL fun_queue_InsertTail ('wake_up_dummy, 'glob_current_process);
  'glob_ready_lists! (('glob_current_process → 'pib_t_priority)) ::= 'wake_up_dummy;;
  'glob_current_process ::= 'glob_pib! ('process_switch_to_pid);;
  'process_switch_to_return_result ::= 0
  FI
  FI
  FI;;
  'res_process_switch_to ::= 'process_switch_to_return_result
  ”

```

Figure 7.8: The function `process_switch_to()` translated to the imperative programming language.

7.4.5 The Function `process_change_sched_param()`

The invoker of this call changes the priority and timeslice of the target process. The PID of the target process and the new scheduling parameters are provided as input parameters.

This system call was explained in details and modeled in Section 3.4.

Signature: $process_change_sched_param_{abstr} : v_data.t \times \mathbb{N} \times \mathbb{N} \times \mathbb{N} \longrightarrow v_data.t$

Let us define:

$$d' \equiv process_change_sched_param_{abstr}(d, prcs, new_pri, new_tsl);$$

$$CUP \equiv d.gd.cup;$$

$$prcs_is_valid? \equiv (prcs \neq IDLE \wedge prcs < PID_MAX);$$

$$prcs_is_active? \equiv \overline{(d.P(prcs).state = \mathbf{inactive})};$$

$$prcs_is_ready? \equiv (d.P(prcs).state = \mathbf{ready});$$

$$new_pri_is_valid? \equiv (new_pri < MAX_PRIO);$$

$$cup_is_prv? \equiv d.P(CUP).prv;$$

Precondition: the currently running process has privileges, the provided PID is correct, the provided priority is correct, and the target process is active:

$$cup_is_prv? \wedge prcs_is_valid? \wedge new_pri_is_valid? \wedge prcs_is_active?.$$

Postcondition:

$$d'.P(prcs).tsl = new_tsl$$

$$\wedge d'.P(prcs).pri = new_pri$$

$$\wedge (prcs_is_ready? \wedge d.P(prcs).pri \neq new_pri \implies$$

$$((d'.gd.rdy ! d.P(prcs).pri) = ((d.gd.rdy ! d.P(prcs).pri) Delete\ prcs))$$

$$\wedge ((d'.gd.rdy ! new_pri) = ((d.gd.rdy ! new_pri) @ [prcs]))$$

$$\wedge (new_pri > d.gd.current_max_prio \implies$$

$$(d'.gd.current_max_prio = new_pri) \wedge (d'.gd.cup = prcs))$$

$$\wedge VamosData?(d', S') \wedge cmp_correct(d') \wedge nt_correct(d').$$

The implementation of this function in C0 is presented in Appendix B.7. Figure 7.9 presents

the function code translated into the imperative programming language. This function invokes the following three functions:

- 1) The function `dList_Delete()`;
- 2) The function `queue_InsertTail()`;
- 3) The function `compute_max_prio()`.

The function `process_change_sched_param()` is not verified completely yet.

7.5 Verification Status

The current verification status of the VAMOS functions done in the frame of this thesis is summarized in Table 7.1. At the moment of writing, the functions `check_elapsed_timeouts` and `handler_clock` are not specified yet.

| VAMOS Function | Status | | |
|---|---------------|----------------|--------------|
| | Specification | Implementation | Verification |
| <code>queue_InsertTail</code> | + | + | + |
| <code>queue_Rotate</code> | + | + | + |
| <code>search_next_process</code> | + | + | + |
| <code>compute_max_prio</code> | + | + | + |
| <code>wake_up</code> | + | + | - |
| <code>process_get_my_pid</code> | + | + | + |
| <code>process_switch_to</code> | + | + | - |
| <code>process_change_sched_param</code> | + | + | - |
| <code>check_elapsed_timeouts</code> | - | + | - |
| <code>handler_clock</code> | - | + | - |

Table 7.1: Current status of functional verification.

procedures

```

fun_process_change_sched_param ( process_change_sched_param_pid, process_change_sched_param_tsl,
process_change_sched_param_prio | res_process_change_sched_param) =
" IF 'process_change_sched_param_pid = 0  $\vee$  128  $\leq$  'process_change_sched_param_pid THEN
    'process_change_sched_param_return_result ::= -1
ELSE IF ('glob_pib! ('process_change_sched_param_pid)  $\rightarrow$  'pib.t.state) = 3 THEN
    'process_change_sched_param_return_result ::= -22
ELSE IF 3  $\leq$  'process_change_sched_param_prio THEN
    'process_change_sched_param_return_result ::= -6
ELSE 'process_change_sched_param_return_result ::= 0;;
    'process_change_sched_param_element ::= 'glob_pib! ('process_change_sched_param_pid);;
    IF ((''process_change_sched_param_element  $\rightarrow$  'pib.t.state) = 1)  $\wedge$ 
        ((''process_change_sched_param_element  $\rightarrow$  'pib.t.priority)  $\neq$ 
            'process_change_sched_param_prio) THEN
        'glob_ready_lists! ((''process_change_sched_param_element  $\rightarrow$  'pib.t.priority)) ::=
            CALL fun_dlist_pib_t_queue_Delete (
                'glob_ready_lists! ((''process_change_sched_param_element  $\rightarrow$  'pib.t.priority)),
                'process_change_sched_param_element);;
        'glob_ready_lists! ('process_change_sched_param_prio) ::=
            CALL fun_queue_InsertTail(
                'glob_ready_lists! ('process_change_sched_param_prio),
                'process_change_sched_param_element);;
        'glob_current_max_prio ::= CALL fun_compute_max_prio()
    FI;;
    'process_change_sched_param_element  $\rightarrow$  'pib.t.timeslice ::= 'process_change_sched_param_tsl;;
    'process_change_sched_param_element  $\rightarrow$  'pib.t.priority ::= 'process_change_sched_param_prio
    FI
    FI
    FI;;
    'res_process_change_sched_param ::= 'process_change_sched_param_return_result
"

```

Figure 7.9: The function `process_change_sched_param()` translated into the imperative programming language.

Conclusion and Future Work

This thesis is the first step in verification of the VAMOS microkernel. We have presented work done so far in verification of the VAMOS scheduler. We started with introduction of the whole VAMOS model. Then the scheduler design was given with regard to the VAMOS model.

Designed scheduler is implemented in the programming language C0 and partially verified in the theorem proving environment Isabelle. We have outlined both implementation and basis for verification of VAMOS. Specification of the VAMOS functions relevant to scheduler is presented. As the next step all functions composing scheduler should be verified. After functional verification is done certain fairness properties of the VAMOS scheduler can be proven.

On the time of writing the Verisoft project goes on already for two years. A lot of work is done, not only in context of pure verification, but also automation of proofs. The verification environment and Verification Condition Generator developed by N. Schirmer lessen routine work. With the tool developed by S. Tverdyshev even more proof steps can be automated. Proceeding with the automation of proofs we help researchers to focus on essence of proofs.

Appendix A

CVM and VAMOS Layers

Throughout this thesis we refer to the CVM layer and its connections with the VAMOS layer. Here we present this connection as it is modelled in Isabelle (see [Figure A.1](#)).

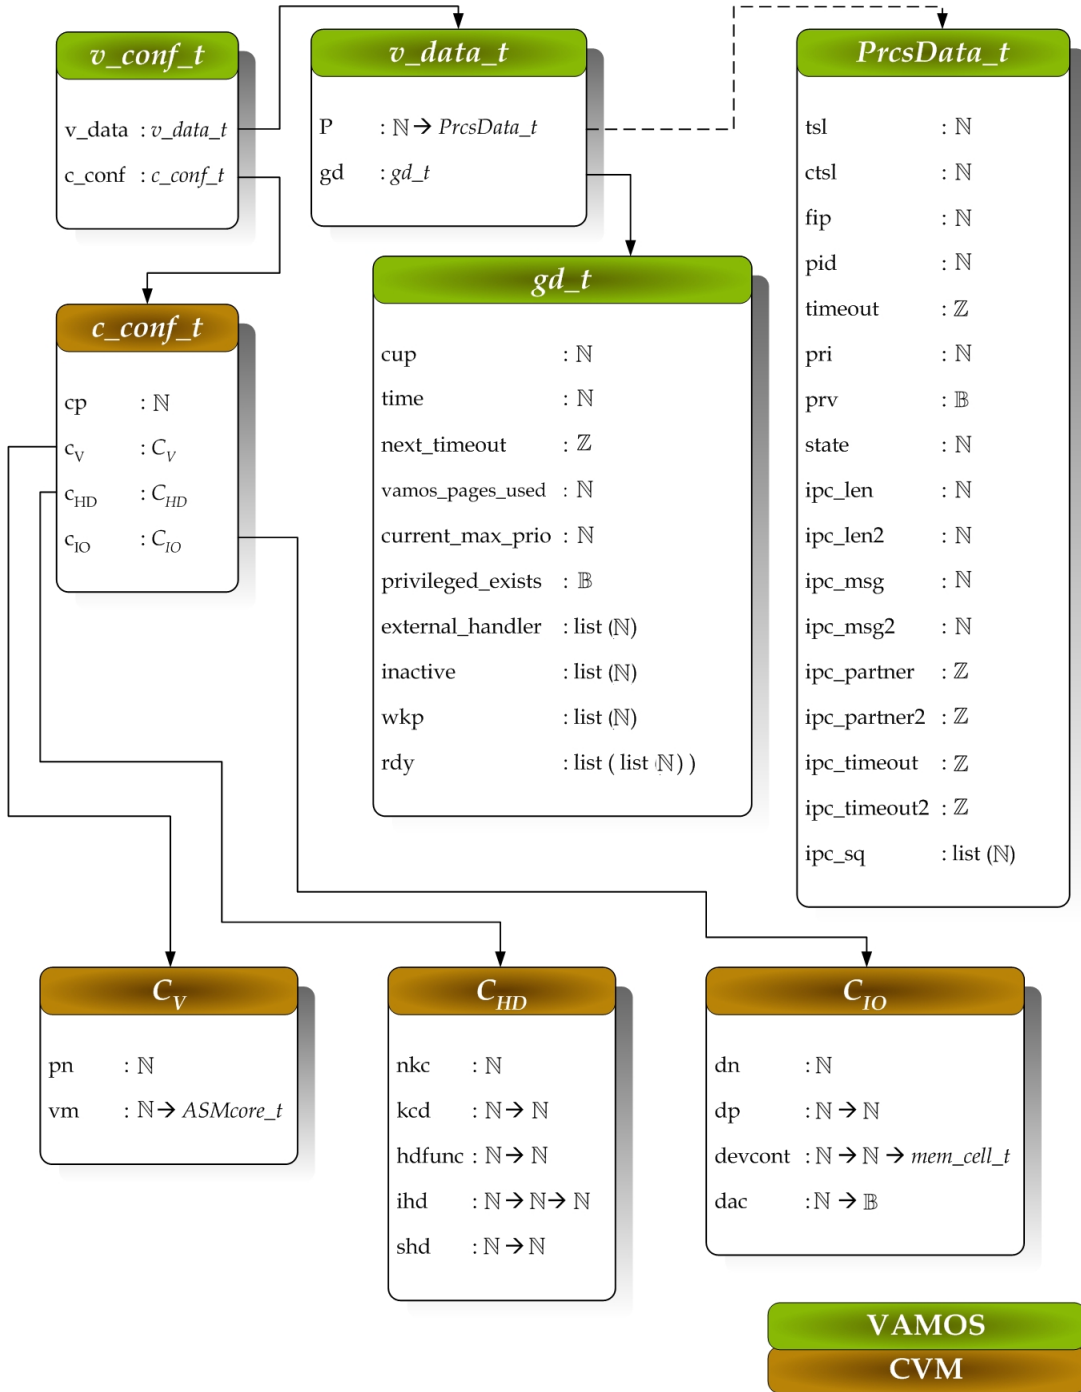


Figure A.1: The CVM and VAMOS layers.

Appendix B

C0-implementation

Here we present how the VAMOS functions are implemented in C0.

B.1 The Function `queue_InsertTail()`

Listing B.1: The function `queue_InsertTail()`.

```
plib_p queue_InsertTail(plib_p list, plib_p p)
{
    unsigned int len;
    int dummy_int;
    plib_p result;
    plib_p dummy;

    len = dlist_plib_t_queue_Length (list);
    if (len == 0u)
    {
        result = dlist_plib_t_queue_InsertHead (list, p);
    }
    else
    {
        dummy = dlist_plib_t_queue_GetElement (list, len-1u);
        dummy_int = dlist_plib_t_queue_InsertAfter (dummy, p);
        result = list;
    }
    return result;
}
```

B.2 The Function `queue_Rotate()`

Listing B.2: The function `queue_Rotate()`.

```
pib_p queue_Rotate(pib_p list)
{
    pib_p result;

    /* First remove the head */
    result = dlist_pib_t_queue_Delete (list, list);

    /* and then insert it into the tail */
    result = queue_InsertTail (result, list);

    return result;
}
```

B.3 The Function `search_next_process()`

Listing B.3: The function `search_next_process()`.

```
int search_next_process()
{
    current_process =
        dlist_pib_t_queue_GetElement (ready_lists[current_max_prio], 0u);

    return 0;
}
```

B.4 The Function `compute_max_prio()`

Listing B.4: The function `compute_max_prio()`.

```
unsigned int compute_max_prio()
{
    unsigned int i;
    unsigned int prio;

    prio = 0u;
    i = MAX_PRIO-1u;

    while( i > 0u )
    {
        if (ready_lists[i]!=NULL)
        {
            prio = i;
            i = 0u;                /* exit from the loop */
        }
        else
        {
            i = i - 1u;
        }
    }
    return prio;
}
```

B.5 The Function wake_up()

Listing B.5: The function wake_up().

```
int wake_up(pib_p p)
{
    pib_p dummy;
    int dummy_int;
    unsigned int len;

    assert(p->state==VAMOS_PROCESS_IPC_SEND | p->state==VAMOS_PROCESS_IPC_RECEIVE);

    // 1. Take out of the wakeup list
    wakeup_list = dlist_pib_t_queue_Delete (wakeup_list, p);

    // 2. Append at the appropriate ready list
    ready_lists[p->priority] = queue_InsertTail (ready_lists[p->priority], p);

    // 3. Set state to ready
    p->state = VAMOS_PROCESS_READY;

    // 4. Update current_max_prio, if priority is bigger than current_max_prio
    //    additionally search next process
    if (p->priority > current_max_prio)
    {
        current_max_prio = p->priority;
        dummy_int = search_next_process();
    }

    return 0;
}
```

B.6 The Function `process_switch_to()`

Listing B.6: The function `process_switch_to()`.

```
int process_switch_to(unsigned int pid)
{
    pib_p dummy;
    int return_result;
    unsigned int len;
    int dummy_int;

    return_result = -1;

    if (pid == current_process->pid)
    {
        return_result = 0;
    }
    else
    {
        if ((pid >= PID_MAX) || (pid == 0u))
        {
            return_result = VAMOS_RC_INVALID_PID;
        }
        else
        {
            if (pib[pid]->state != VAMOS_PROCESS_READY)
            {
                return_result = VAMOS_RC_PROCESS_NOT_READY;
            }
            else
            {
                if (pib[pid]->consumed_time <
                    (current_process->timeslice - current_process->consumed_time))
                {
                    pib[pid]->consumed_time = 0u;
                }
                else
                {
                    pib[pid]->consumed_time = pib[pid]->consumed_time
                        - (current_process->timeslice - current_process->consumed_time);
                }
            }
        }

        // set consumed_time of current_process to 0
        current_process->consumed_time = 0u;
    }
}
```

```
// put current_process at the end of its ready list -- first remove, then put to the
end
ready_lists[current_process->priority] =
    dlist_pib_t_queue_Delete (ready_lists[current_process->priority], current_process);

ready_lists[current_process->priority] =
    queue_InsertTail(ready_lists[current_process->priority], current_process);

// set current_process to the new one
current_process = pib[pid];

return_result = 0;
}
}
}

return return_result;
}
```

B.7 The Function `process_change_sched_param()`

Listing B.7: The function `process_change_sched_param()`.

```

int process_change_scheduling_param (unsigned int pid, unsigned int tssl, unsigned int
    prio)
{
    pib_p      element;
    pib_p      dummy;
    int        return_result;
    unsigned int len;
    int        dummy_int;

    if (pid==0u || pid>=PID_MAX)
    {
        return_result = VAMOS_RC_INVALID_PID;
    }
    else
    {
        if (pib[pid]->state==VAMOS_PROCESS_INACTIVE)
        {
            return_result = VAMOS_RC_PROCESS_INACTIVE;
        }
        else
        {
            if (prio>=MAX_PRIO)
            {
                return_result = VAMOS_RC_INVALID_ARGS;
            }
            else
            {
                return_result = 0;
                element = pib[pid];

                if ( (element->state == VAMOS_PROCESS_READY) && (element->priority != prio) )
                {
                    /* Must switch to different ready list */
                    /* Dequeue */
                    ready_lists[element->priority] =
                        dlist_pib_t_queue_Delete (ready_lists[element->priority], element);

                    /* Insert element in appropriate ready list */
                    ready_lists[prio] = queue_InsertTail (ready_lists[prio], element);
                }
            }
        }
    }
}

```

B.7. THE FUNCTION `PROCESS_CHANGE_SCHED_PARAM()` APPENDIX B. C0-IMPLEMENTATION

```
    /* Recompute the maximum priority */
    current_max_prio = compute_max_prio();
} //end if element->state

/* Update the values of timeslice and priority */
element->timeslice = tsl;
element->priority = prio;

} //end if prio >= MAX_PRIO
} //end if pib[pid]...
} //end if pid == 0u...

return return_result;
}
```

Appendix C

The Function `get_my_pid`

The function `get_my_pid` is not relevant to the VAMOS scheduler, but it was specified and verified in the frame of this thesis. This function implements the system call *processGetMyPid*. This system call is used to get the process identifier of the invoker (i.e., the currently running process). Since processes are restricted to access a microkernel internal information, that is the only way for a process to get his own PID. The returned PID can be used for IPC communication, etc.

This system call has no input parameters.

This system call has no preconditions, and can be invoked by unprivileged processes.

Signature: $process_get_my_pid : v_data.t \rightarrow \mathbb{N}$.

Postcondition: $process_get_my_pid(d) = d.gd.cup$.

The implementation of this function in C0 is presented in Listing C.1. Figure C.1 presents the function code translated into the imperative programming language. The functional verification takes 6 proof steps in Isabelle.

procedures

```
fun_process_get_mypid ( | res_process_get_mypid ) =  
  ”  
  'res_process_get_mypid ::= ('glob_current_process → 'pib.t_pid)  
  ”
```

Figure C.1: The function `process_get_my_pid()` translated into the imperative programming language.

Listing C.1: The function `process_get_my_pid()`.

```
unsigned int process_get_my_pid()  
{  
  return (current_process->pid);  
}
```

Bibliography

- [1] P. BOVET AND M. CESATI. *Understanding the Linux Kernel, 2nd Edition*. December 2002.
- [2] E. ALKASSAR. *Constructing A Formal Framework For Modeling And Verifying A Real Operating System*, Master Thesis, Saarland University, 2005.
- [3] M. GARGANO AND W. J. PAUL. *A CVM specialization: the VAMOS Operating System Kernel*, Verisoft Internal Technical Report #38, May 2004. unpublished.
- [4] M. GARGANO, D. LEINENBACH, AND W. J. PAUL. *CVM: Communicating Virtual Machines*, Interner Technischer Bericht #10, Verisoft Project, May 2004. unpublished.
- [5] D. LEINENBACH. *DIE SPRACHE C0*. *Verisoft Internal Technical Report #2*, June 2004. unpublished.
- [6] D. LEINENBACH, W. J. PAUL, AND E. PETROVA. *Towards the Formal Verication of a C0 Compiler*. In *Proceedings, 3rd International Conference on Software Engineering and Formal Methods (SEFM 2005)*, Koblenz, Germany, Sep 5-9, 2005.
- [7] Z. MANNA AND A. PNUELI. *Temporal Verification of Reactive Systems*. Springer-Verlag, 1995.
- [8] I. MIRONENKO. *Proof of total correctness and absence of runtime faults in Doubly Linked Lists and Strings Libraries*, Diploma Thesis, Saarland University, 2005.
- [9] T. NIPKOW, L.C. PAULSON, AND M. WENZEL. *Isabelle/HOL. A Proof Assistant for Higher-Order Logic*. Springer-Verlag, 2004.
- [10] V. G. NGUIEKOM. *Verikation von doppelt verketteten Listen auf Pointerebene*. Diplomarbeit, Saarland University, Mai 2005.
- [11] N. SCHIRMER. *Verification Environment for Sequential Imperative Programs in Isabelle/HOL*, Verisoft Internal Technical Report #18, June 2004. unpublished.
- [12] A. STAROSTIN. *Formal Verification of a C-Library for Strings*. Diploma Thesis. August, 2005.

- [13] A. S. TANENBAUM. *Modern Operating Systems*. Prentice Hall, 2001.
- [14] A. S. TANENBAUM AND A. S. WOODHULL. *Operating Systems: Design and Implementation*. Prentice Hall, 1997.
- [15] V. TSYBA. *Verification of the L4 Task Scheduler*, Master Thesis, Saarland University, 2003.
- [16] Web Site of the Verisoft project. <http://www.verisoft.de>