Echtzeitbetriebssysteme

Echtzeit-Betriebssystem-Kerne

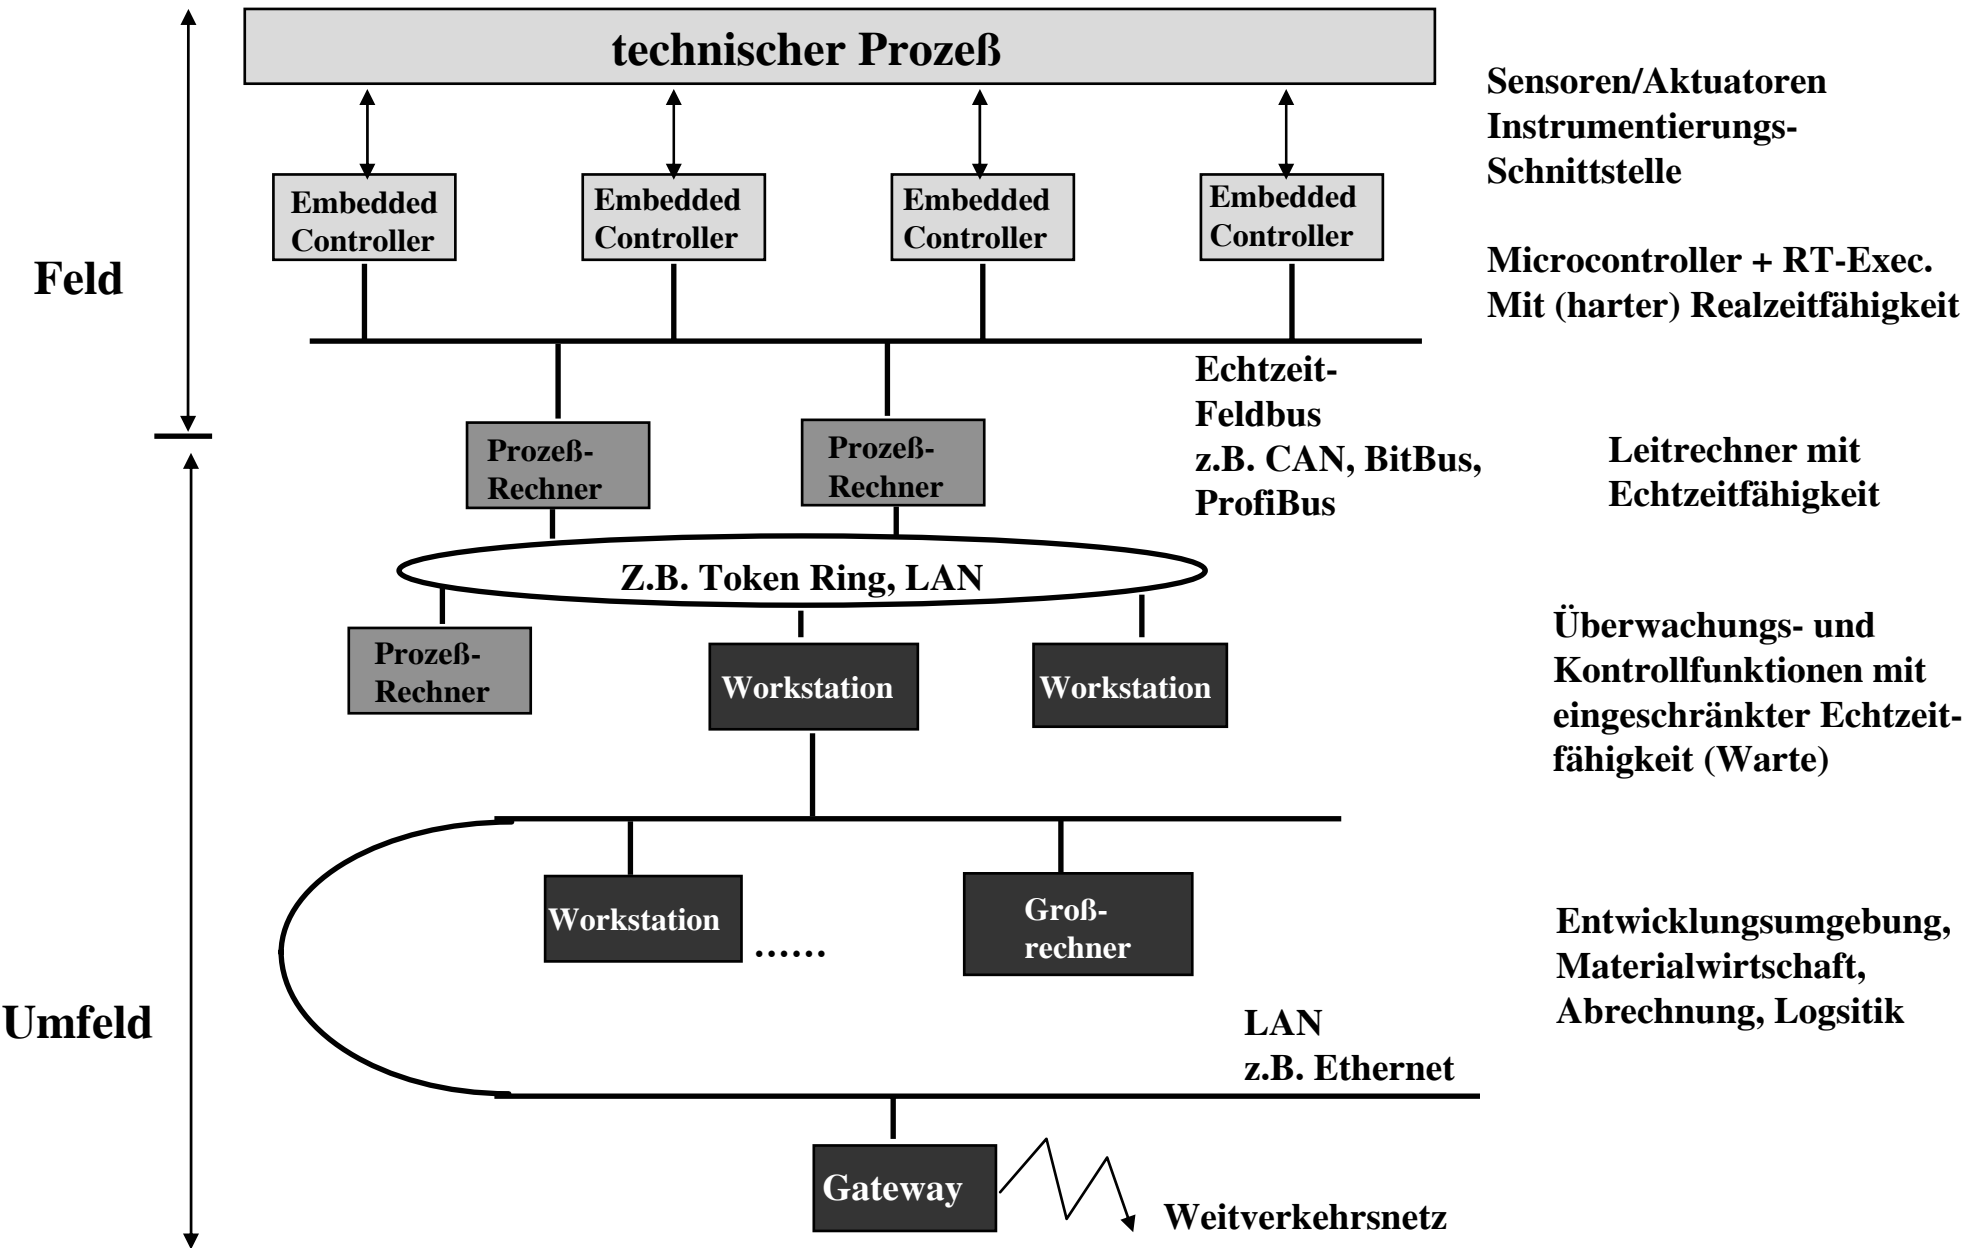**Grundlegende Funktionalität:**

High-Level:
- **Scheduling**
- **Zeitbehandlung**
- **Ein-Ausgabe, Ereignisbehandlung**
- **Kommunikation**
- **Ressourcenverwaltung**

Low-Level:
- **Task Management**
- **Task-Synchronisation**
- **Inter-Task-Kommunikation**
- **Zeitgeber und Uhren**
- **Unterbrechungsbehandlung**
- **Statische/Dynamische Speicherallokation**

Einbettung von Echtzeitsystemen in eine CIM-Umgebung

technischer Prozeß

Sensoren/Aktuatoren
Instrumentierungs-
Schnittstelle

Embedded Controller   **Embedded Controller**   **Embedded Controller**   **Embedded Controller**

Feld

Microcontroller + RT-Exec.
Mit (harter) Realzeitfähigkeit

Echtzeit-
Feldbus
z.B. CAN, BitBus,
ProfiBus

Prozeß-Rechner   **Prozeß-Rechner**

Leitrechner mit
Echtzeitfähigkeit

Z.B. Token Ring, LAN

Prozeß-Rechner   **Workstation**   **Workstation**

Überwachungs- und
Kontrollfunktionen mit
eingeschränkter Echtzeit-
fähigkeit (Warte)

Workstation ...... **Groß-rechner**

Entwicklungsumgebung,
Materialwirtschaft,
Abrechnung, Logsitik

Umfeld

LAN
z.B. Ethernet

Gateway Weitverkehrsnetz

Anforderungen und Technologie

Ebene:	Charakteristische Anforderungen:	Technologien/Betriebssysteme:
Embedded Controller, einfache Steueraufgaben	Harte Echtzeitbedingungen, kurze Reaktionszeiten	Microcontroller, Echtzeitkerne mit eingeschränkter Funktionalität PXROS, RTkernel, VRTX, VXworks EUROS, iRMX, . . .
Leitrechnerebene	Harte Echtzeitbedingungen, kurze Reaktionszeiten, Realisierung konsistenter globaler Zustände, globales, verteiltes Scheduling, Erkennung globaler (kritischer) Zustände	Prozeßrechner oder Workstations, Echtzeitbetriebssysteme für komplexe verteilte Steuerungsaufgaben, Beisp.: OpenVMS, Posix IEEE 1003.13, LynxOS, QNX, RT-Linux, Maruti, SPRING, MARS, RT-Mach,
Warte	Eingeschränkte Echtzeitfähigkeit, graphische Bedienoberfläche	Workstationtechnologie, Touchscreen, GUI, Visualisierungs- werkzeuge (z.B. LabView)
Entwicklungsumgebung	Keine Echtzeitanforderungen, Leistungsfähige, Multi-User- Umgebung zur Programm- Entwicklung, Simulation, Verifikation, Test	High-End-Workstation, hoher Leistungsbedarf, Compiler: ADA, PEARL, C, C++, JAVA Off-line Scheduling Werkzeuge, Simulationsumgebung

Kategorien für Echtzeit- Betriebssysteme

- Prioritätsbasierte Betriebssystemkerne für eingebettete Anwendungen
  Beispiele: VRTX, VxWorks, QNX, OSEK, PXROS, RTkernel, OS9, RTEMS, u8os. . . . .

- Echtzeiterweiterungen für "Timesharing" Betriebssysteme
  Beispiele: RT-Mach, RT-Erweiterungen in POSIX, RED-Linux, RT-Linux, ...

- Echtzeitbetriebssysteme in der Forschung
  Beispiele: MARS, CHAOS, Spring, ARTS, MARUTI, ...

VRTX
u8os

PXROS,
OSEK,

RTKernel,
EUROS, RTOS

Maruti, MARS,
Spring, PS/9

LynxOS, Mach, QNX,
Solaris, RTLinux, NT…

zunehmende
Systemkomplexität

SimpleEmbedded Systems

+ Minimale System-
   Unterstützung
+ Minimaler Platzbedarf
+ Minimaler Overhead

- Aufwendige Off-line Analyse
- Statisches Systemverhalten
- kein Dateisystem
- kein BIOS

| **Anwendung** |
| **Kommunikation** |
| **RT-Executive** |
| **Treiber** |
| **Hardware** |

| **Anwendung** |
| **Kommunikation** |
| **RT-Executive** |
| **BIOS** |
| **Treiber** |
| **Hardware** |

| **Anwendung** |
| **Taskmanagement, (dyn.) Planung, Synchronisation** |
| **RT-Networking** |
| **BIOS** |
| **Treiber** |
| **Hardware** |

| **Anw.1 … Anw.n** |
| **mehrere Benutzer** |
| **Taskmanagement, dyn. Planung, Synchronisation** |
| **Dateisystem** |
| **RT-Networking** |
| **Protection** |
| **virt. Speicher** |
| **BIOS** |
| **Treiber** |
| **Hardware** |

Anpassen der Betriebssystemkerne an die Anwendungsanforderungen

Profile in POSIX (1003.13) :

PSE50 - Minimales Realzeit System Profil:
Geeignet für eingebettete Ein- oder Mehrprozessorsysteme, die ein oder mehrere externe Geräte kontrollieren. Es wird nur ein einzelner Prozeß (nur 1 Adressraum!) mit mehreren Threads unterstützt. Kein Dateisystem, keine Operator-Schnittstelle.

PSE51- Systemprofil für „Real-Time Controller":
Erweiterung zu PSE50. Dateisystem und asynchrone Ein-/Ausgabe.

PSE52 – Spezielles Reazeit-Profil:
Erweiterung zu PSE50 um Systeme mit Speicherverwaltungseinheit. Unterstützt mehrere Prozesse aber kein Dateisystem.

PSE53 – Multi-purpose Realzeit-Profil:
Unterstützt Koexistenz von Echtzeitsystemen und Nicht-Echtzeitsystemen. Für Ein-/ Mehrprozessorsysteme mit MMU, Massenspeicher, Netzwerk etc..

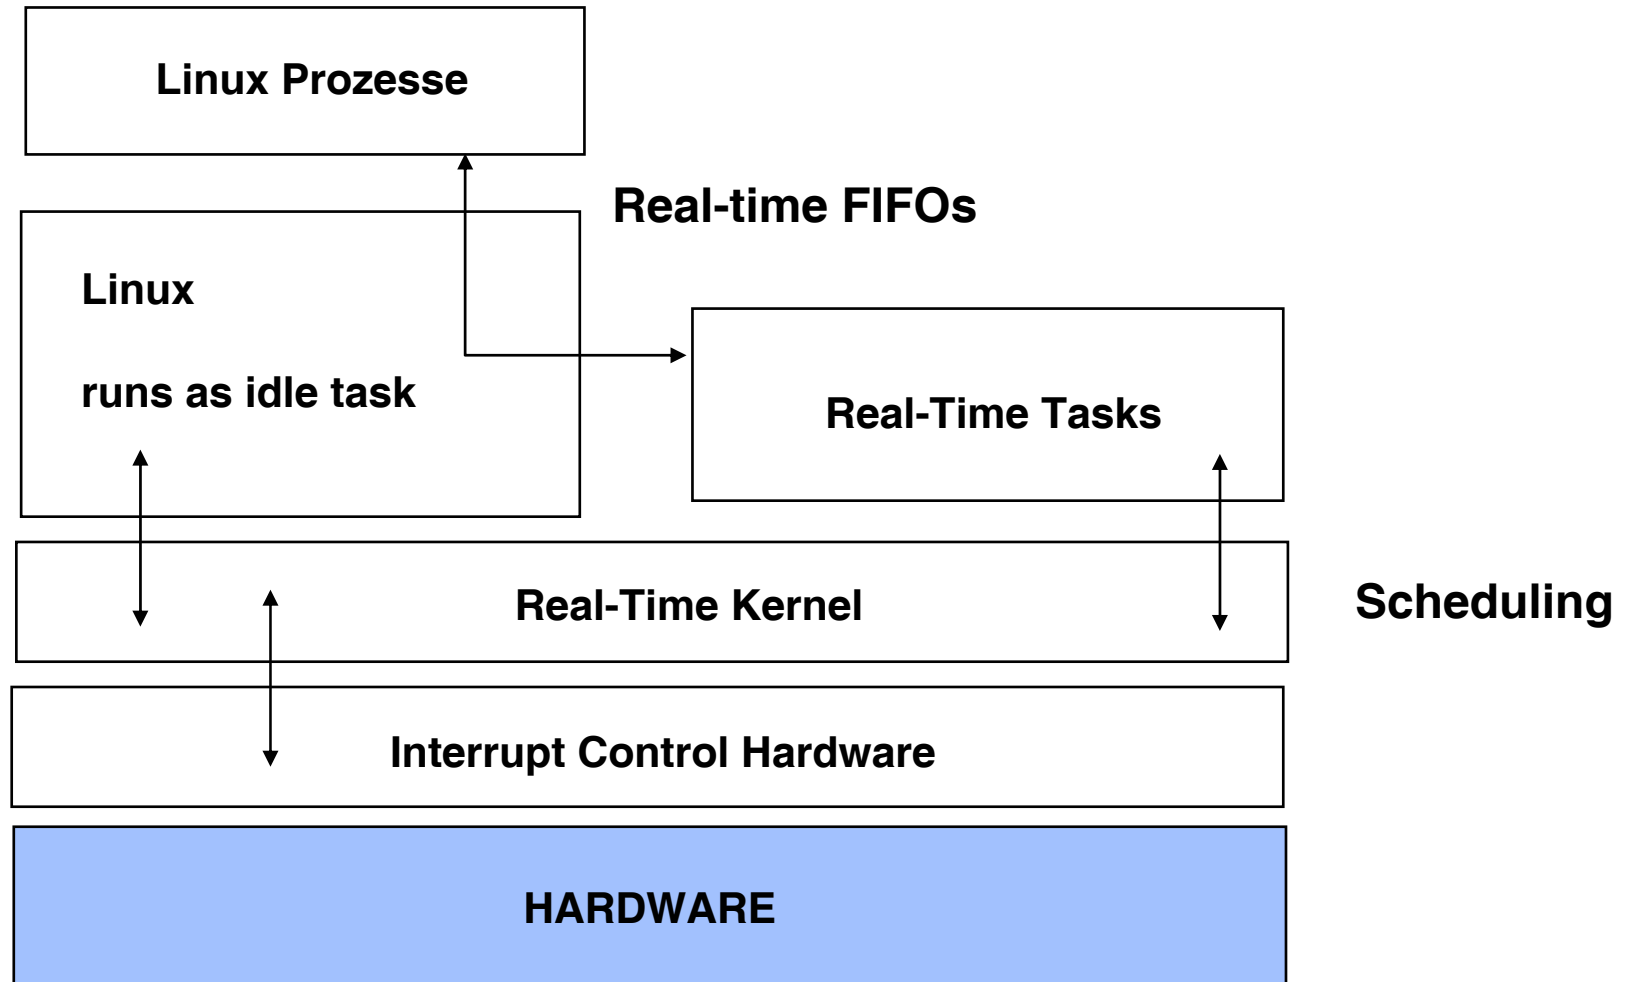# Einbettung von Echzeitkernen in die Systemarchitektur

Beispiel: RTKernel

Anwendungsprogramm	
Standard-Bibliotheken	**RTKernel**	**Laufzeit-System**

MS-DOS

BIOS

Hardware/Betriebsmittel
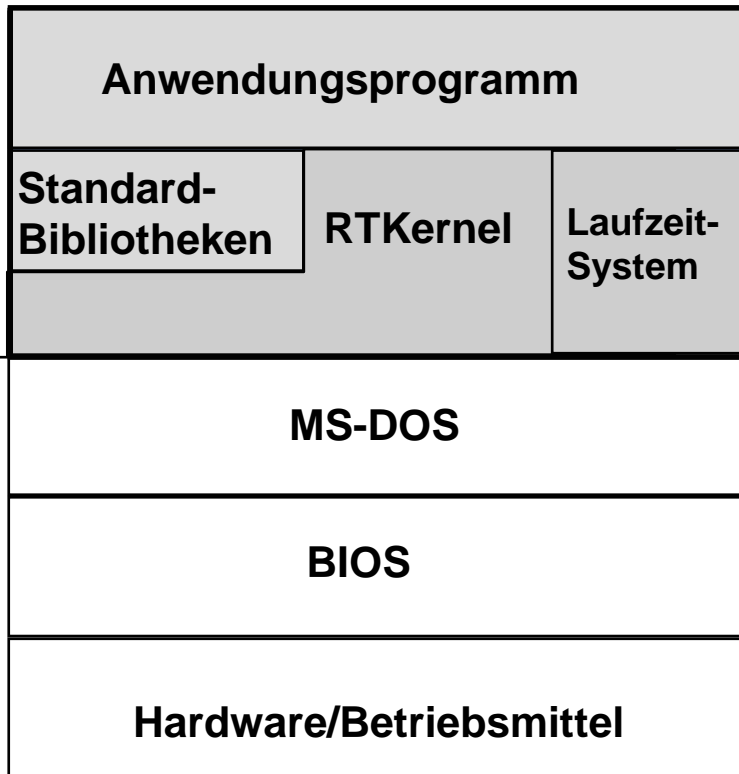
DOS-Anwendung, die durch Hinzubinden der RTKernel Libraries einen echzeitfähige Anwendung darstellt

Eigenschaften::

- **Integration in DOS /BIOS , dadurch sind Basisbetriebssystem-funktionen wie Standard I/O Funktionen (Platte, Floppy, Screen, etc.) und Interruptverarbeitung verfügbar.**
- **Systemannahmen 80x86 Architektur + 8059 Interruptbaustein**
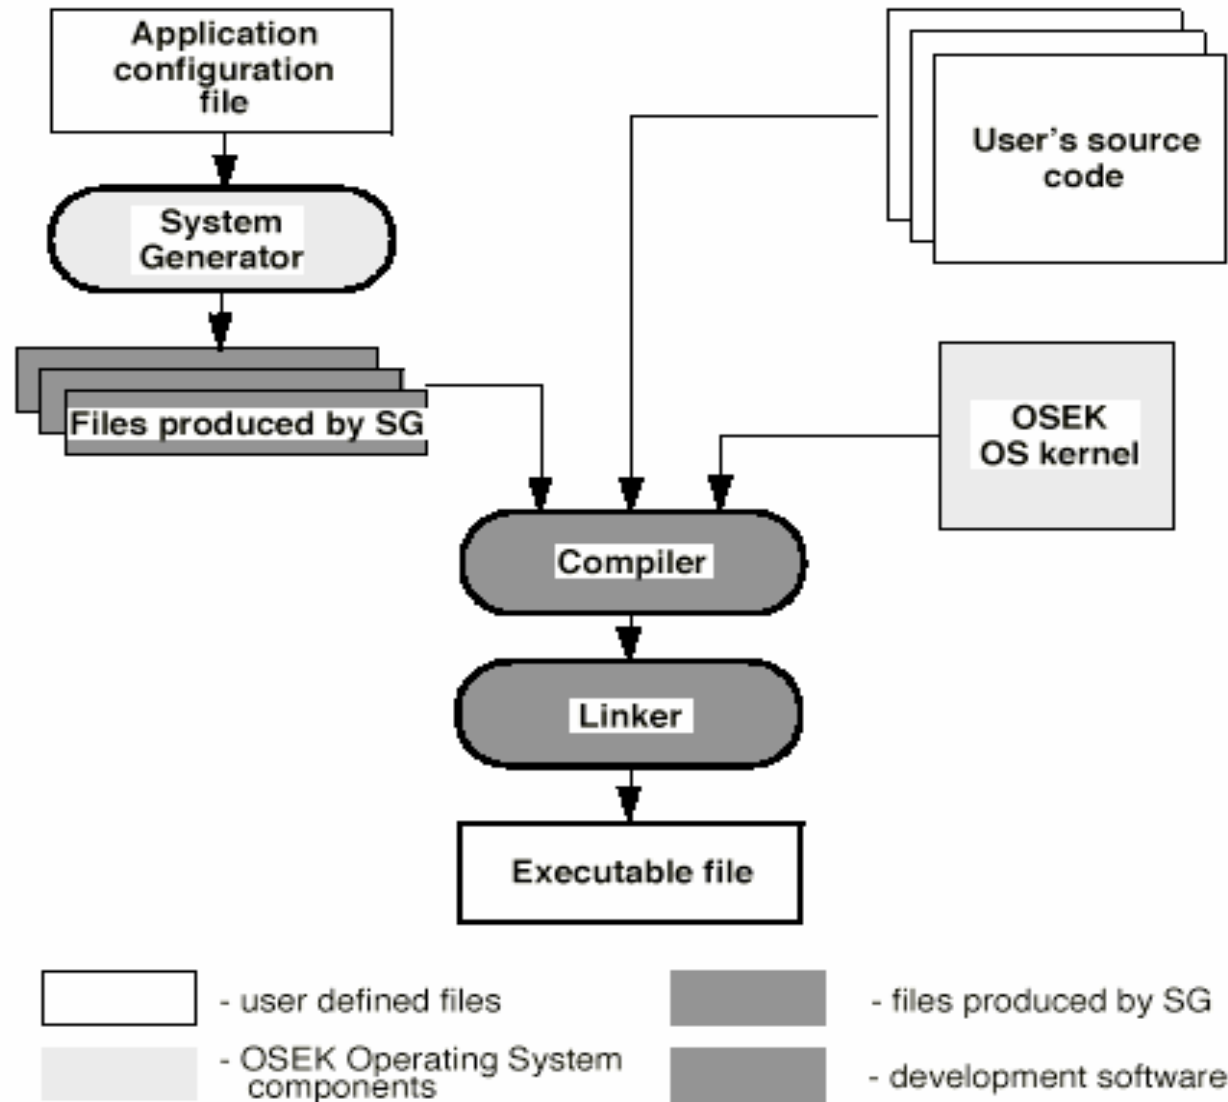- **Programmentwicklung, Debugging, Simulation auf einem System**

Architektur und Einbettung von RTLinux

Linux Prozesse

Real-time FIFOs

Linux

runs as idle task

Real-Time Tasks

Real-Time Kernel

Scheduling

Interrupt Control Hardware

HARDWARE

Schema einer Cross-Entwickungsumgebung:
Konfigurierbarer BS-Kern

P_0 ... P_n BS-Kern   lib

Entwicklungssystem
Editor, Compiler, Binder,Lader,
Simulator, Verifizierer, Scheduling-
Werkzeuge, etc...

Download,
EPROM,
Flash Memory

Embedded Controller

P_0 ... P_n

BS-Kern

Embedded Controller:

• **spezifischer I/O**
• **spezifischer Befehlssatz**
• **spezifische Unterbrechungsbehandlung**
• **variable Zeitgeberfunktionen**

10
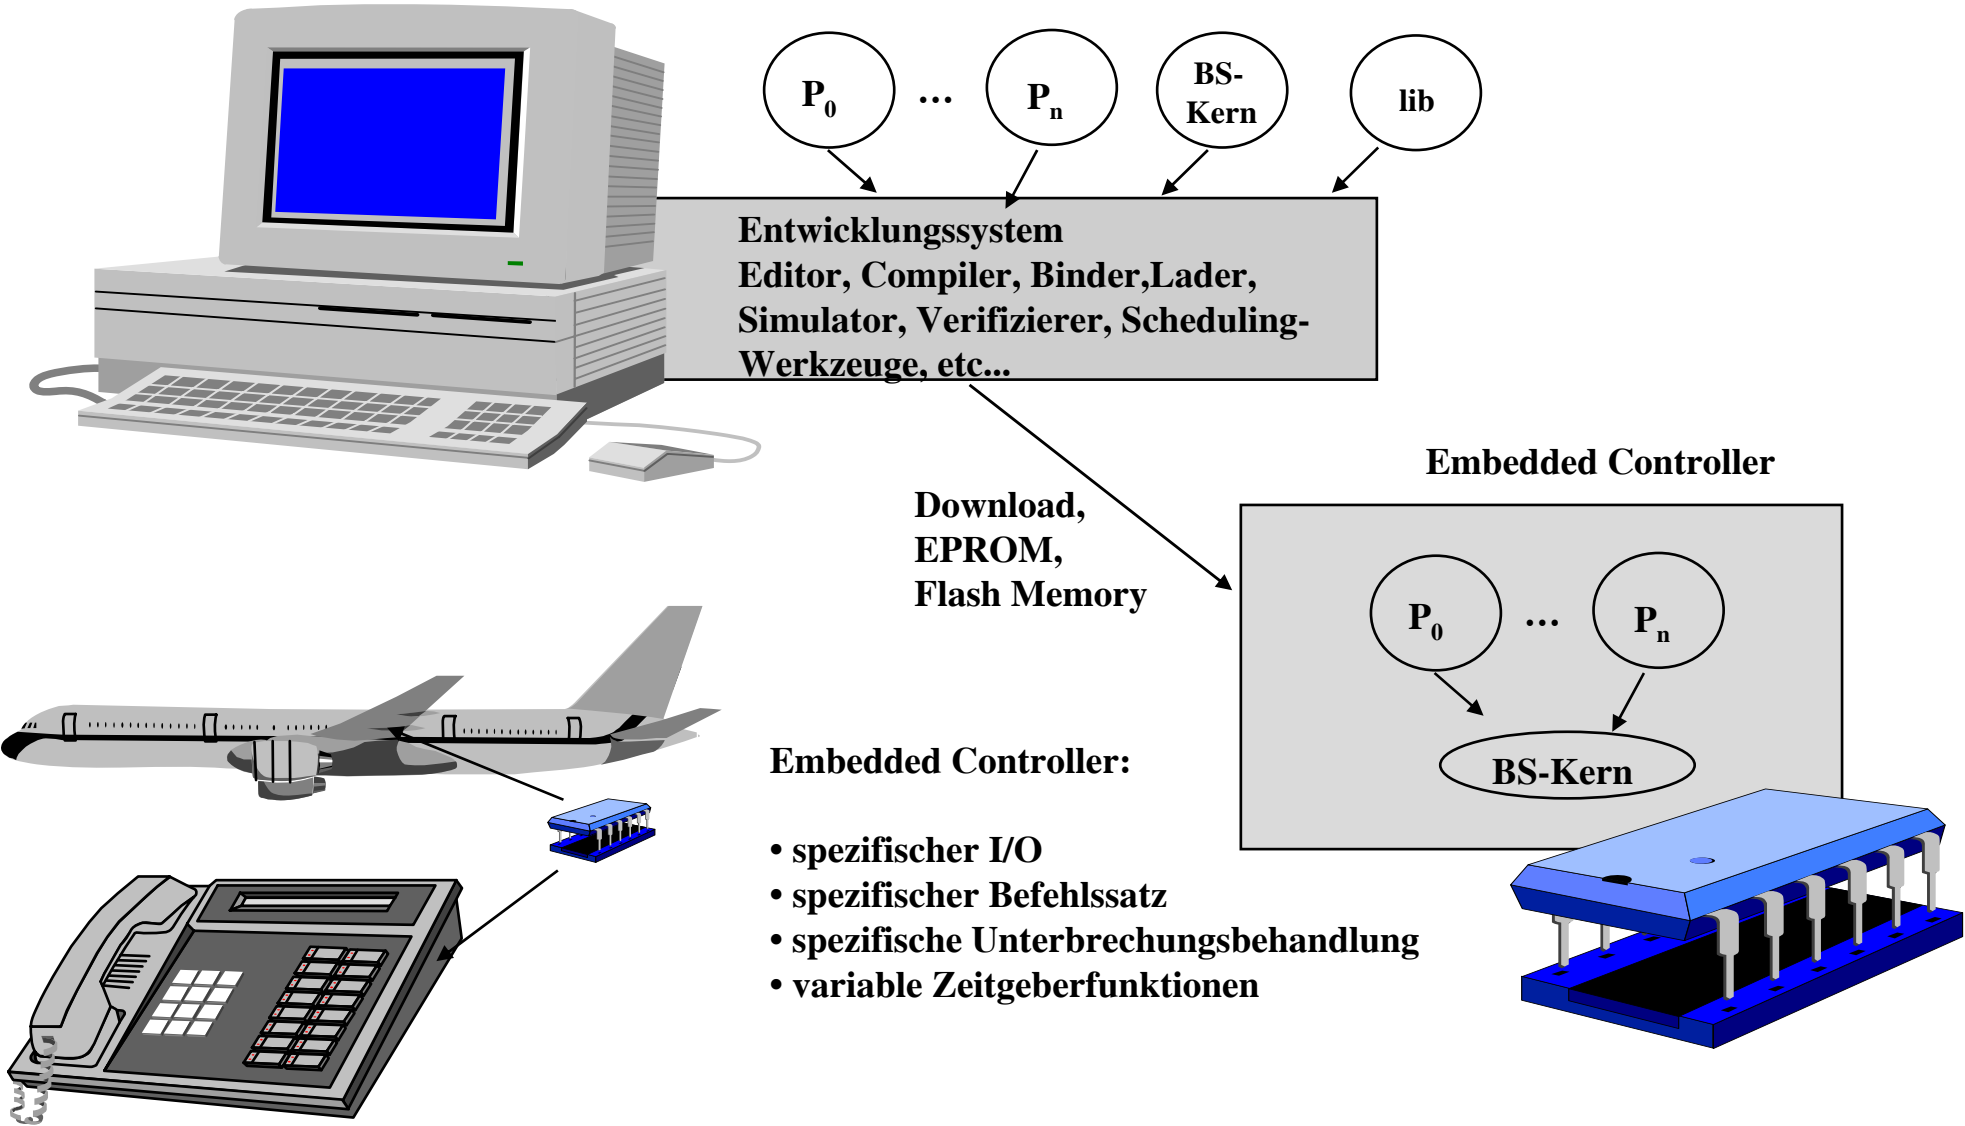
Application building process in OSEK

Beispiel: OSEK Konfigurationsdatei

```
[Property]
TargetMCU = HC08;
ConformanceClass =ECC1;
SimpleScheduler = OFF;
ExtendedStatus = ON;
HookRoutines = ON;
ErrorHandler = ON;
ContextSwitchRoutine = ON;
InterruptMaskControl = OFF;
TaskIndexMethod = OFF;
PersistentNode = OFF;
TaskOwnStack = ON;
TaskBasePage = ON;
StackPool = ON;
SchedulerPolicy = MIXPREEMPT;
CounterSize = 16;
Alarms = ON;
AlarmList = ON;
Resources = ON;
FastResource = OFF;
Events = ON;
StateMessage = ON;
StateMsgDefaultValue = ON;
StateMsgTimeStamp = ON;
EventMessage = ON;
EventMsgTimeStamp = ON;
EventMsgOneToN = OFF;
ActivateOnMsg = ON; /* Message signalling mechanism */
AlarmOnMsg = ON;
SetEventOnMsg = ON;
```

```
[Scheduler]
DefineScheduler( 4, 5, 128, , 64 ) ;
[Interrupt management]
DefineInterrupts( 0x8, 0, 0, 64 );
[User's hook]
DefineHooks ( OSError, OSPreTask, OSPostTask);
[Tasks]
DefineTask (TASKSND, BASIC|PREEMPT|ACTIVATE, 3, TaskSND );
DefineTask (TASKRCV, PREEMPT| BASIC|OWNSTACK, 1,
TaskRCV,,, TaskStack, TASKSTACKSIZE );
DefineTask (TASKPROD, PREEMPT| EXTENDED|ACTIVATE|POOLSTACK, 2,
TaskProd,,, POOL );
DefineTask (TASKCONS, NONPREEMPT|BASIC|POOLSTACK, 4,
TaskCons,,, POOL );
DefineStackPool( POOL, 70, 2 );
[Resources]
DefineResource( MSGACCESS, 1);
[Counters]
DefineSystemTimer( SYSTEMTIMER, –1, 10, 1000000 );
DefineCounter( MSGCOUNTER, 6, 1 );
[Alarms]
DefineAlarm( MSGALARM, SYSTEMTIMER, TASKSND );
DefineAlarm( PRODALARM, SYSTEMTIMER, TASKPROD, TIMEVENT );
DefineAlarm( EVMSGALARM, MSGCOUNTER, TASKCONS );
[State Messages]
DefineStateMessage( MsgA, MSGATYPE, sizeof(MSGATYPE),
WithoutTimeStamp );
ActivateOnMessage( MsgA, TASKRCV );
DefineMessageAlarm( MsgA, MSGALARM, 100 );
[Event Messages]
DefineEventMessage( MsgB, MSGBTYPE, sizeof(int), 5, 1,
WithOverwriteCheck, WithTimeStamp );
```

Fallbeispiel:
OSEK/VDX

OSEK
Offene Systeme und deren Schnittstellen für die
Elektronik im Kraft-fahrzeug

VDX
Vehicle Distributed eXecutive

What is OSEK/VDX?

OSEK/VDX is a joint project of the automotive industry. It aims at an industry standard for an open-ended architecture for distributed control units in vehicles.

A real-time operating system, software interfaces and functions for communication and network management tasks are thus jointly specified.

The term OSEK means:

"Offene Systeme und deren Schnittstellen für die Elektronik im Kraft-fahrzeug"
(Open systems and the corresponding interfaces for automotive electronics).

The term VDX means

„Vehicle Distributed eXecutive".
The functionality of OSEK operating system was harmonised with VDX. For simplicity OSEK will be used instead of OSEK/VDX in the document.

OSEK/VDX partners

The following companies attended and contributed to the OSEK/VDX Technical Committee:

Adam Opel AG,
AFT GmbH,
ATM Computer GmbH,
BMW AG,
Blaupunkt,
Centro Ricerche Fiat,
Denso,
Cummins Engine Company,
Daimler-Benz AG,
Dassault Electronique,
Delco Electronics,
ETAS GmbH & Co KG,
Hella KG Hueck & Co.,
Hewlett Packard France,
IIIT University of Karlsruhe,
ITT Automotive,
Integrated Systems,
LucasVarity,
Magneti Marelli,
Mecel,
Motorola,

National Semiconductor,
NEC Electronics,
NRTT,
Philips Car Systems,
PSA,
Renault,
Robert Bosch GmbH,
Sagem Electronic Division,
SGS Thomson,
Siemens Automotive,
Siemens Semiconductors,
Softing GmbH,
TEMIC,
Texas Instruments,
UTA - United Technologies Automotive,
Valeo Electronique,
Vector Informatik,
Volkswagen AG,
Volvo Car Corporation,
Wind River Systems.

Special support for automotive requirements

Specific requirements for an OSEK operating system arise in the application context of software development for automotive control units. Requirements such as reliability, real-time capability, and cost sensitivity are addressed by the following features:

- **The OSEK operating system is configured and scaled statically. The number of tasks, resources, and services required is statically specified by the user.**

- **The specification of the OSEK operating system supports implementations capable of running on ROM, i.e. the code could be executed from Read-Only-Memory.**

- **The OSEK operating system supports portability of application tasks.**

- **The specification of the OSEK operating system provides a predictable and documented behaviour to enable operating system implementations, which meet automotive real-time requirements.**

- **For each operating system implementation performance parameters must be known.**

OSEK OS - Conformance Classes (CC)

- Generell zwei Gruppen von CC
  - Basic (BCC) erlaubt nur Basic Tasks
    - (Basic Tasks haben keinen *waiting*-Zustand)
  - Extended (ECC) erlaubt auch extended Tasks
    - (Extended Tasks können auf Events warten)

- Aufwärts-kompatible CC's
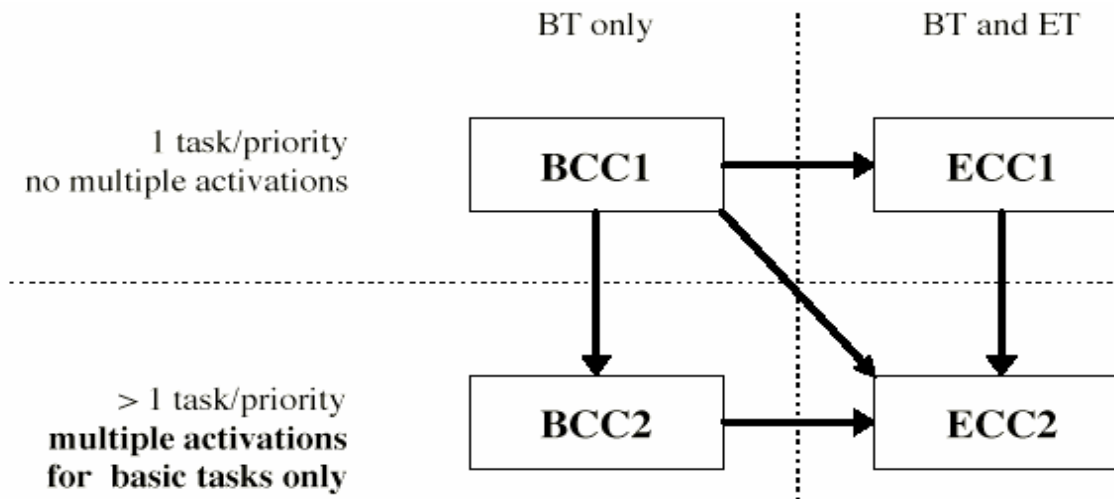
Minimal parameters of implementations

	BCC1	BCC2	ECC1	ECC2
Multiple requesting of task activation	no	yes	BT^9: no ET: no	BT: yes ET: no
Number of tasks which are not in the *suspended* state	≥ 8		≥ 16 (any combination of BT/ET)	
Number of tasks per priority	1	> 1	1 (both BT/ET)	> 1 (both BT/ET)
Number of events per task	—		≥ 8	
Number of priority classes	≥ 8			
Resources	only scheduler	≥ 8 (including scheduler)		
Alarm	≥ 1 (single or cyclic alarm)			
Application Mode	≥ 1			

Basiskomponenten

Unterbrechungsbehandlung

Kommunikation

| OS |
| ISR |
| MESSAGE |

RESOURCE

TASK

COUNTER

EVENT

ALARM

COM

NM

CPU
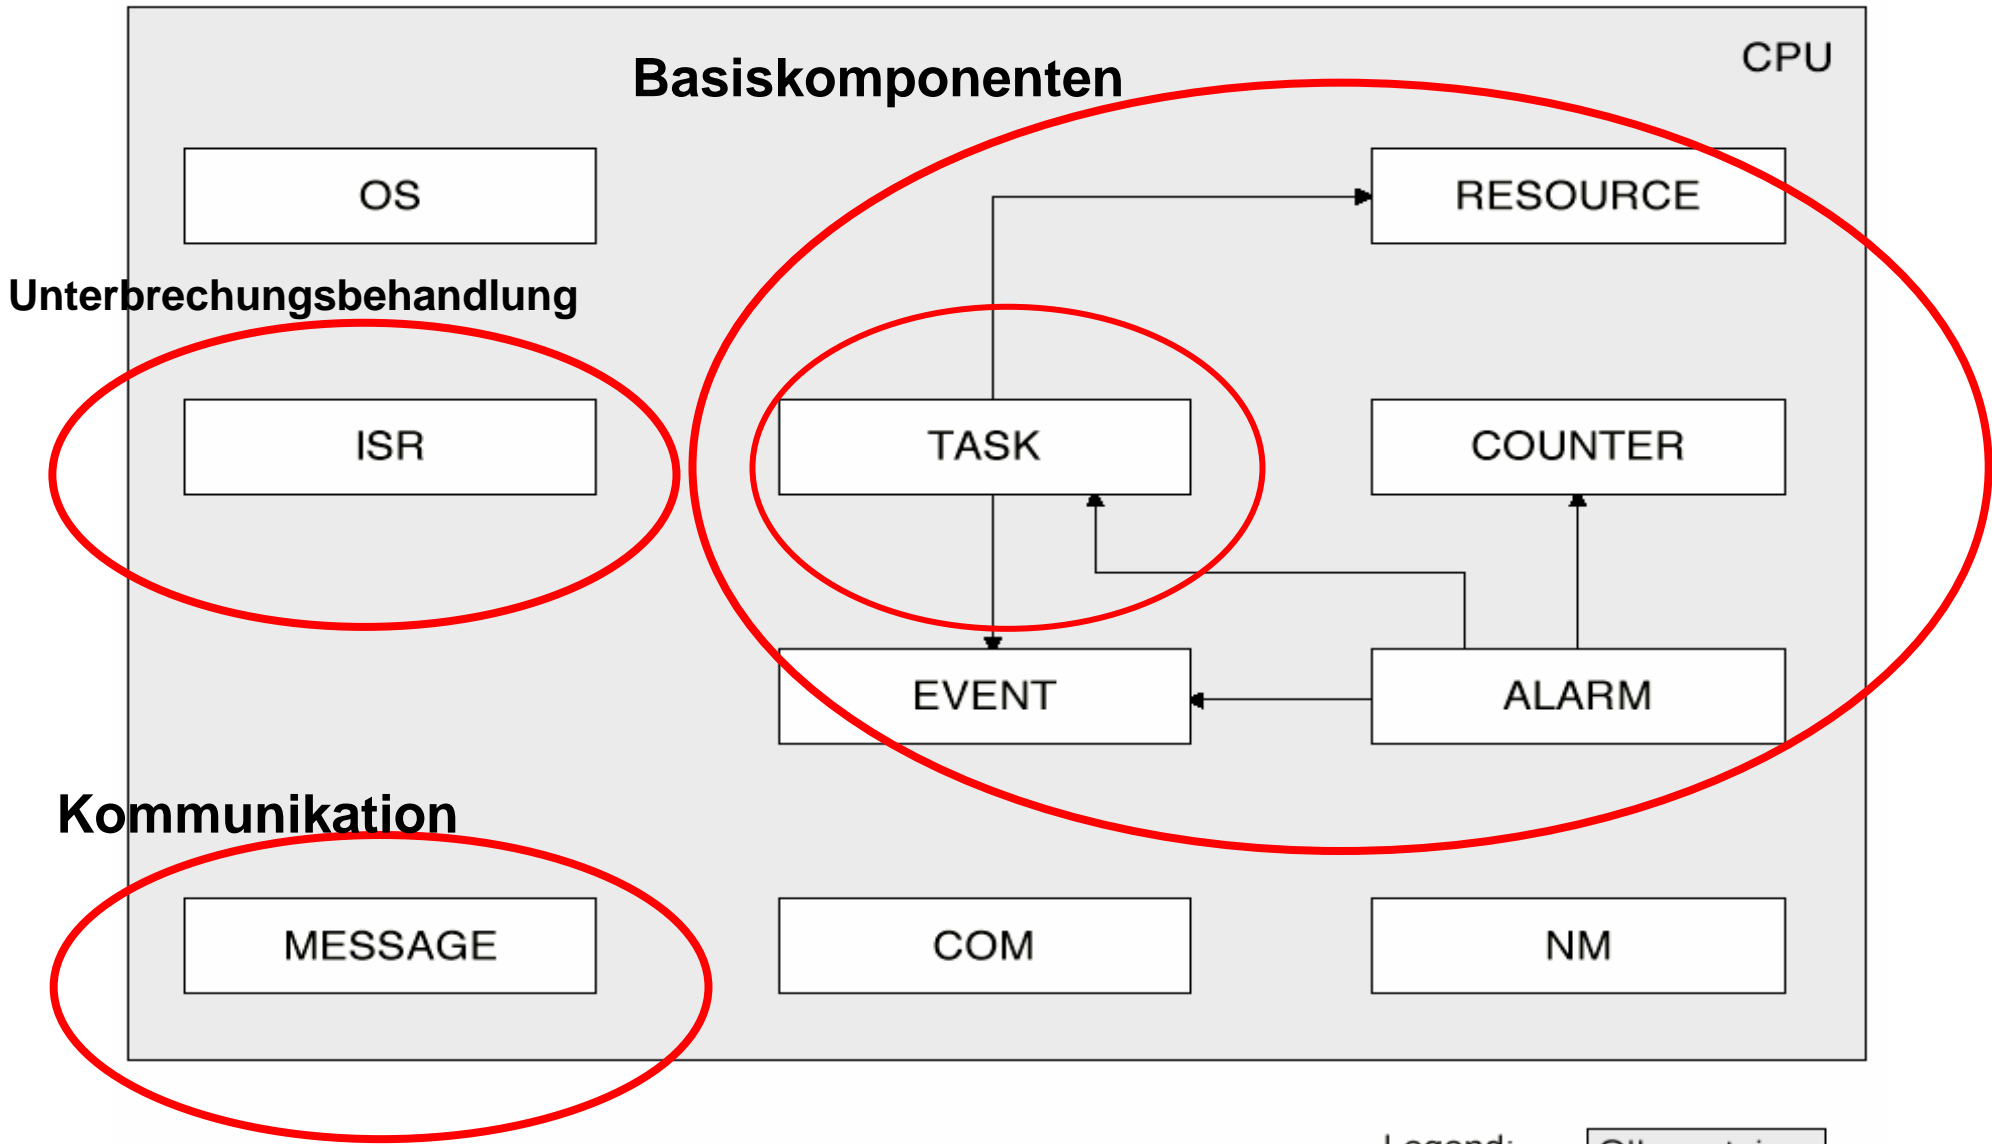
Legend: OIL container

OIL objects

OSEK Komponenten

Beschreibung der Komponenten

CPU: the CPU on which the application runs under the OSEK OS control.

OS: the OSEK OS which runs on the CPU. No standard references are defined in OIL from OS to other system objects but, of course, all system objects are controlled by OS.

ISR: interrupt service routines supported by OS.

RESOURCE: the resource which can be occupied by a task.

TASK: the task handled by the OS.

COUNTER: the counter represents hardware/software tick source for alarms.

EVENT: the event owned by a extended task. Up to 8 events can be assigned to an extended task.

ALARM: the alarm is based on a counter and can activate a task or set an event.

MESSAGE: the message which provides local or network communication.

COM: the communication subsystem. This subsystem is out of the scope of the OSEK OS specification.

NM: the network management subsystem. This subsystem is out of the scope of the OSEK OS specification.

Task management
•Activation and termination of tasks
• Management of task states, task switch

Synchronisation
The operating system supports two means of synchronisation effective on tasks:
• Resource management
  Access control for inseparable operations to jointly used (logic) resources or devices, or for control of a program flow.
• Event control
  Event management for task synchronisation.

Inter-Task Kommunikation

Interrupt management
• Services for interrupt processing

Alarms
• Relative and absolute alarms
• Static (defined at compile-time) and dynamic (defined at run-time) alarms

Error treatment
• Mechanisms supporting the user in case of various errors

Basic Tasks

• **running**

 In the running state, the CPU is assigned to the task, so that its instructions
can be executed. Only one task can be in this state at any point in time, while
all the other states can be adopted simultaneously by several tasks.

• **ready**

All functional prerequisites for a transition into the running state exist, and
the task only waits for allocation of the processor. The scheduler decides
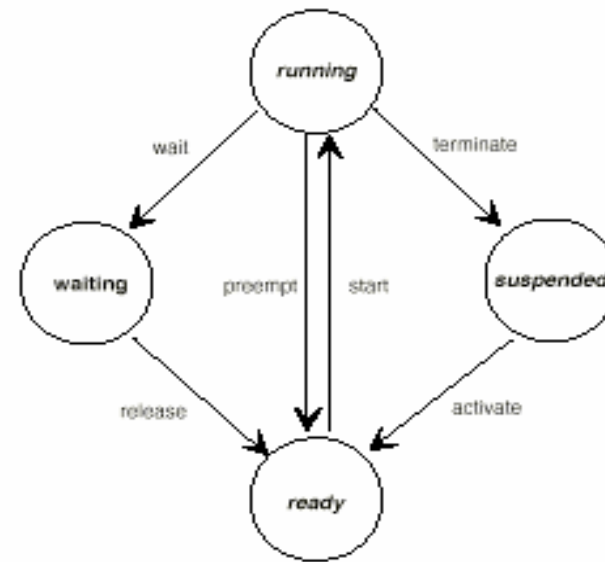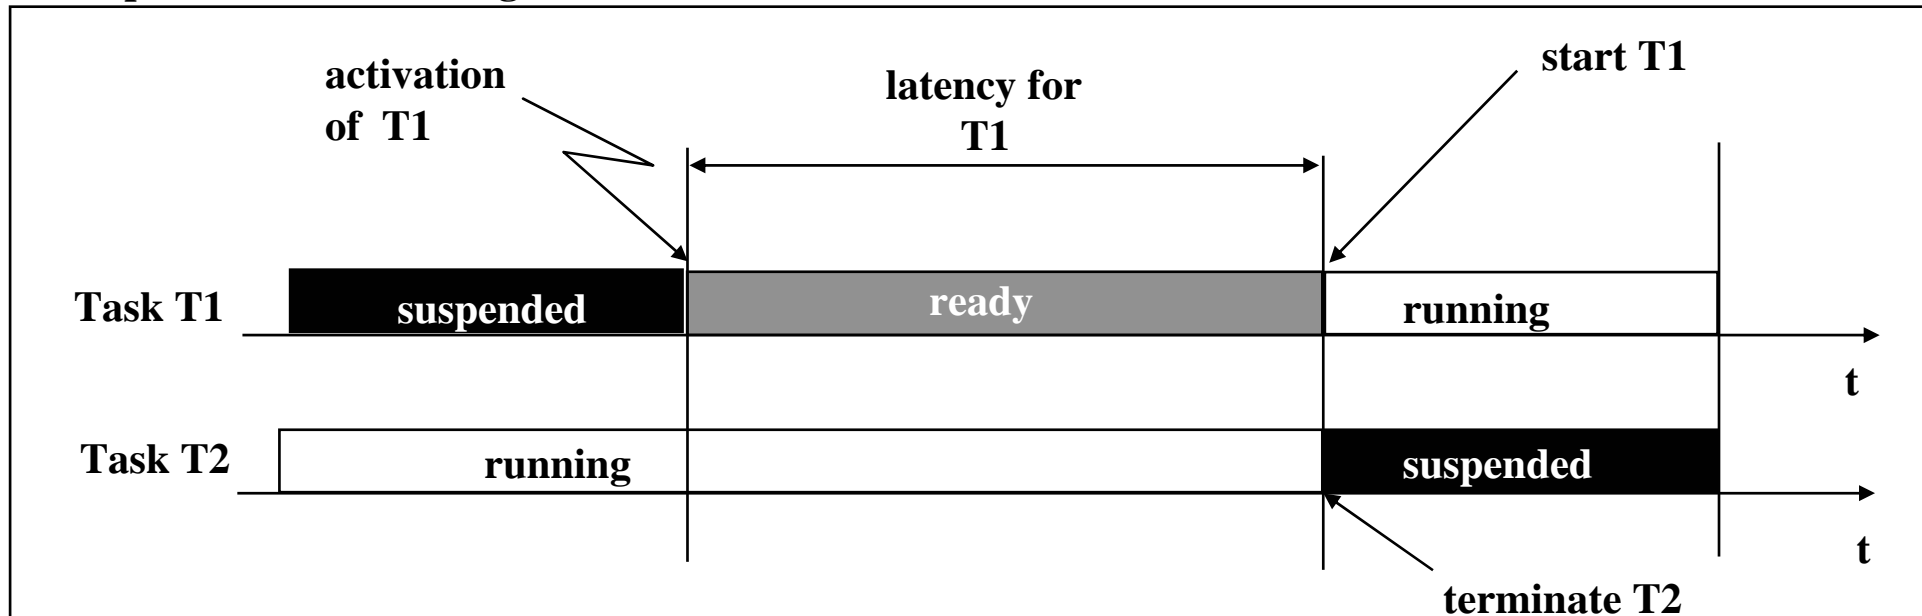which ready task is executed next.

• **suspended**

In the suspended state the task is passive and can be activated.
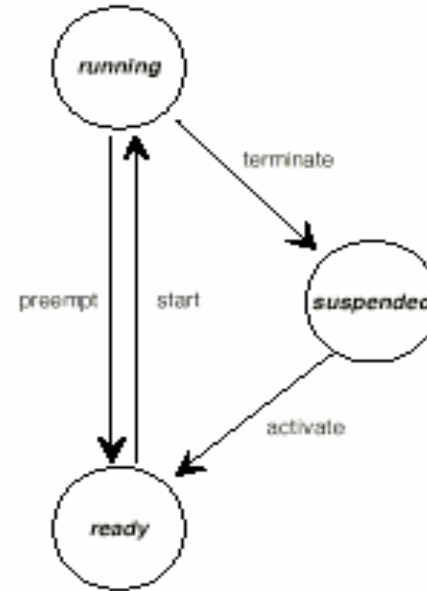
Transition	Former state	New state	Description
activate	suspended	ready[2]	A new task is entered into the *ready* queue by the a system service. The OSEK operating system ensures that the execution of the task will start with the first instruction.
start	ready	running	A *ready* task selected by the scheduler is executed.
preempt	running	ready	The scheduler decides to start another task. The *running* task is put into the *ready* state.
terminate	running	suspended	The *running* task causes its transition into the *suspended* state by a system service.

Extended Tasks

• **running**
In the running state, the CPU is assigned to the task, so that its instructions can be executed. Only one task can be in this state at any point in time, while all the other states can be adopted simultaneously by several tasks.

• **ready**
All functional prerequisites for a transition into the running state exist, and the task only waits for allocation of the processor. The scheduler decides which ready task is executed next.

• **waiting**
A task cannot continue execution because it has to wait for at least one event (see chapter 6, Event mechanism).

• **suspended**
In the suspended state the task is passive and can be activated.

Transition	Former state	New state	Description
activate	suspended	ready	A new task is entered into the *ready* queue by a system service. The OSEK operating system ensures that the execution of the task will start with the first instruction.
start	ready	running	A *ready* task selected by the scheduler is executed.
wait	running	waiting	To be able to continue operation, the *running* task requires an event. It causes its transition into the *waiting* state by using a system service.
release	waiting	ready	At least one event has occurred which a task has *waited* on.
preempt	running	ready	The scheduler decides to start another task. The *running* task is put into the *ready* state.
terminate	running	suspended	The *running* task causes its transition into the *suspended* state by a system service.

non-preemtive Scheduling

activation of T1

latency for T1

start T1

Task T1 | suspended | ready | running |

t

Task T2 | running | | suspended |

t

terminate T2

preemtive Scheduling

latency for T1

activation of T1

terminate T1

Task T1 | suspended | | running | suspended |

t

Task T2 | running | ready | running |

t

start T2
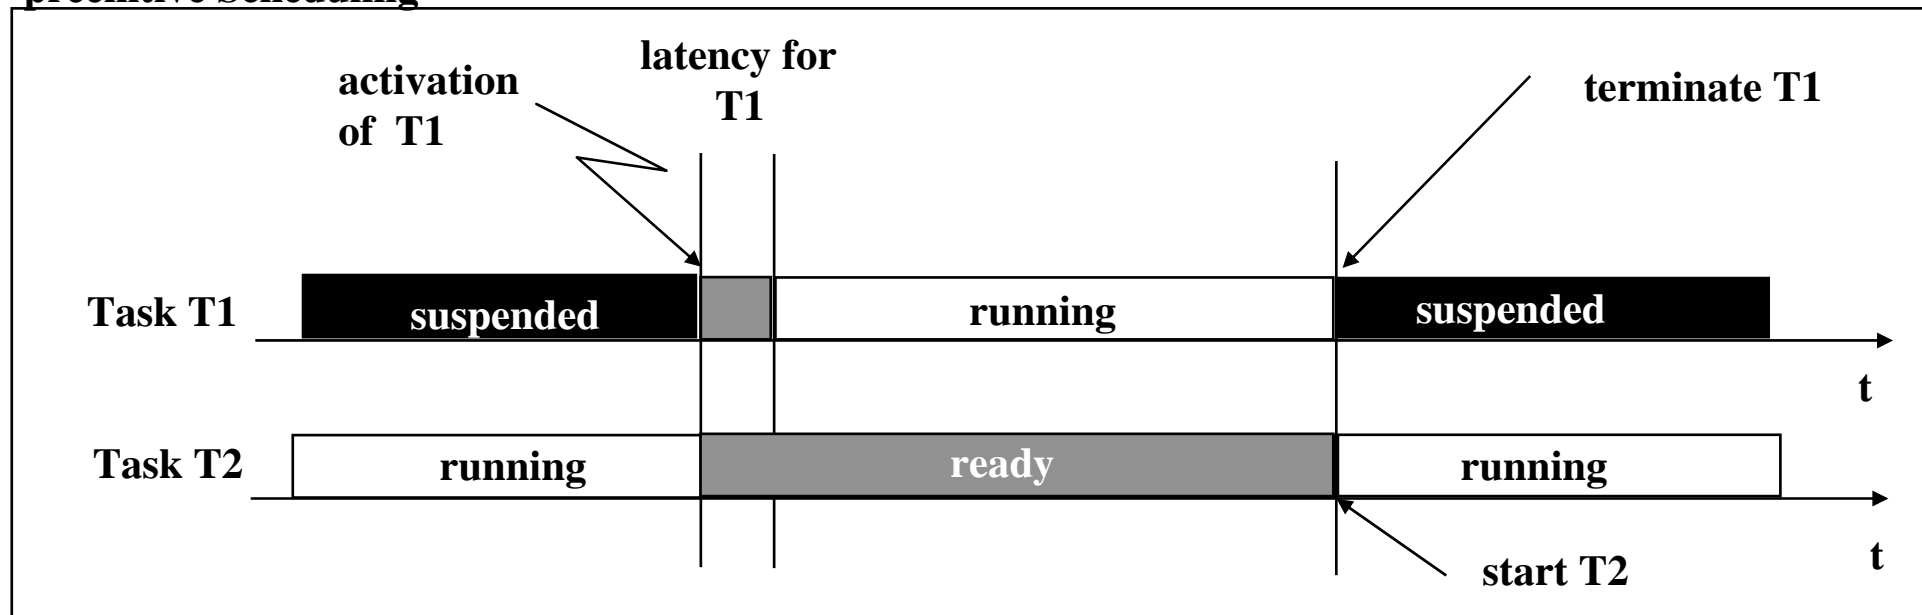
Points of rescheduling (in OSEK)

Non-preemptive task:
- Successful termination of a task.
- Successful termination of a task with explicit activation of a successor task.
- Explicit call of scheduler (system service *Schedule*).
- A transition into the waiting state takes place (system service *WaitEvent* (nur ET)).

Full-preemptive task:
- Successful termination of a task (system service *TerminateTask*).
- Successful termination of a task with explicit activating of a successor task (system service *ChainTask*).
- Activating a task of higher priority (system service *ActivateTask*)
- Explicit wait call, if a transition into the waiting state takes place (extended tasks only, system service *WaitEvent*).
- Setting an event to a waiting task
- message arrival, when task activation or event setting is defined for this message
- alarm expiring, when task activation or event setting is defined for this alarm
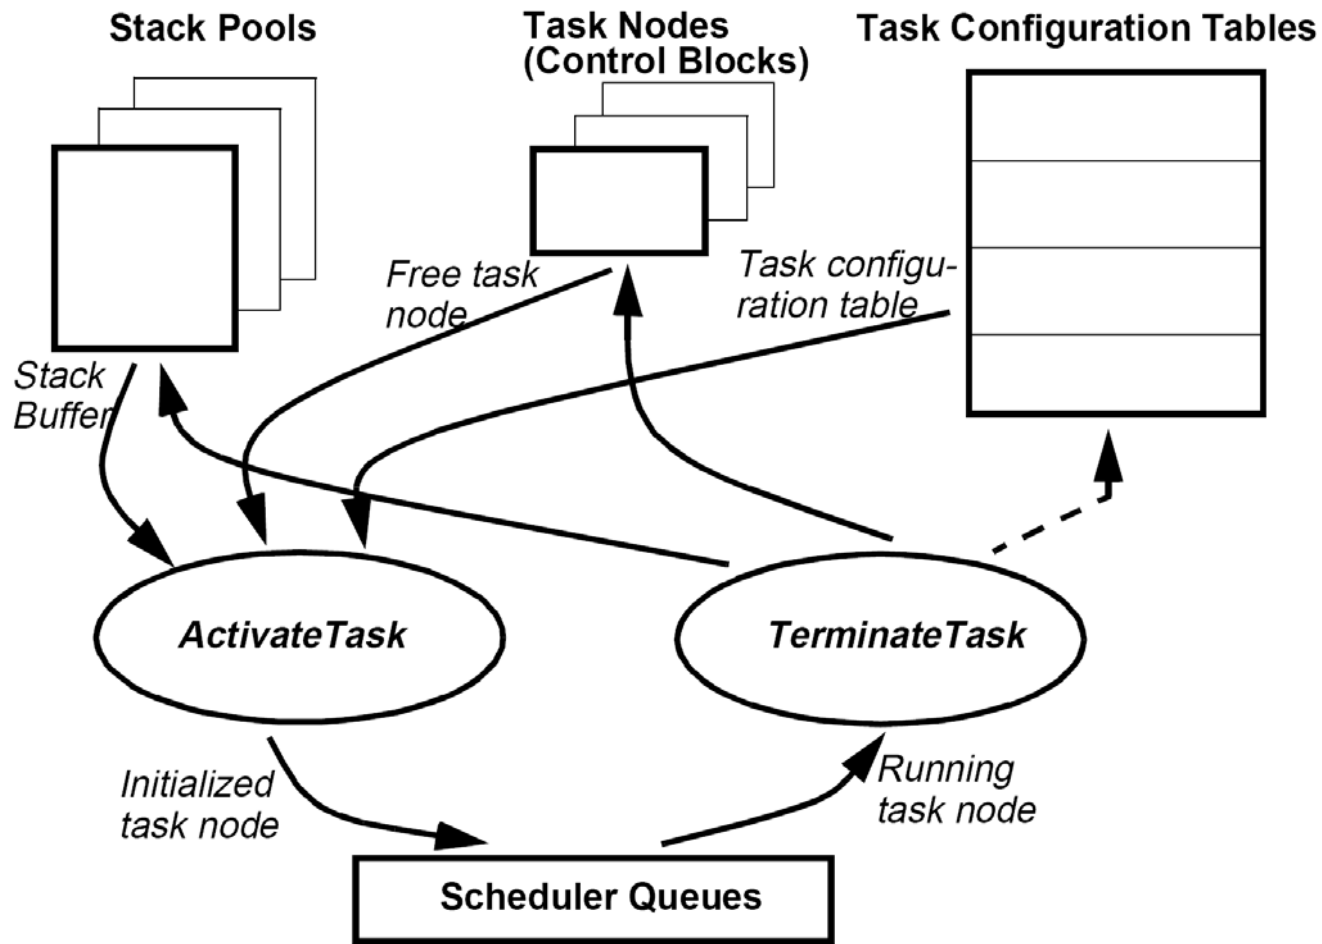
In the case of a full-preemptive system, the user must constantly expect preemption of the running task. If a task fragment must not be preempted, this can be achieved by blocking the scheduler temporarily via the system service *GetResource*.

Scheduler
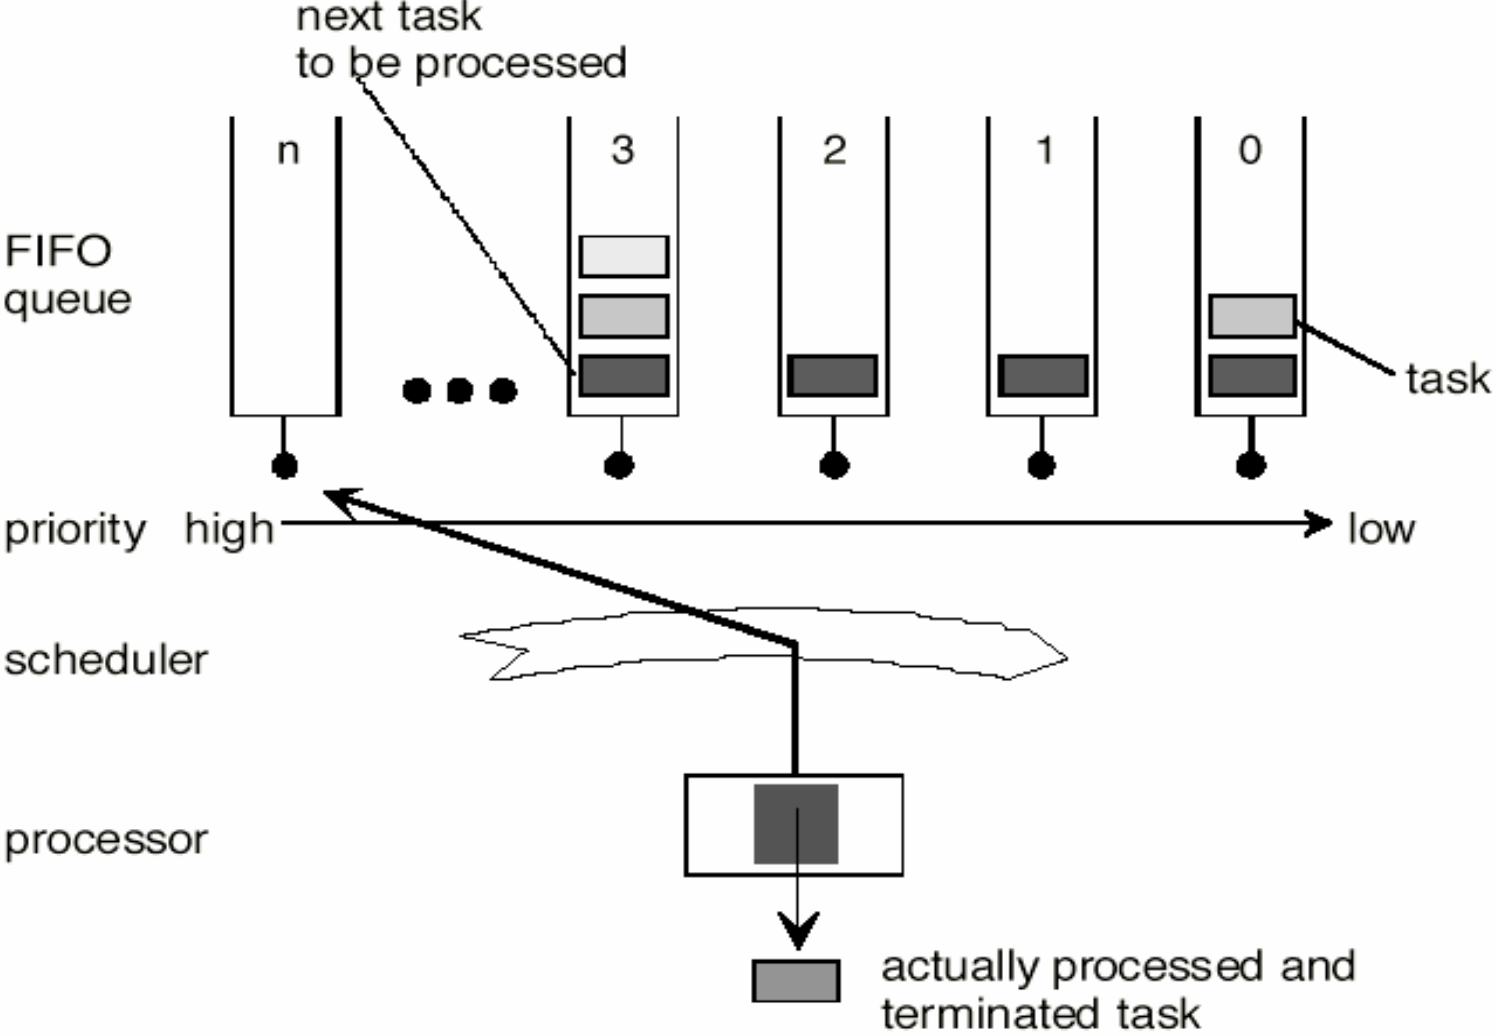
System services related to tasks:

Service Name	Description
ActivateTask	Activates the task, i.e. put it from the *suspended* into the *ready* state
TerminateTask	Terminates the task, i.e. put it from the *ready* into the *suspended* state
ChainTask	Terminates the task and activates a new one immediately
Schedule	Yields control to a higher-priority ready task (if any exists)
GetTaskId	Gets the identifier of the running task
GetTaskState	Gets the status of the specified task

Stack Pools

Task Nodes (Control Blocks)

Task Configuration Tables

Stack Buffer

Free task node

Task configuration table

ActivateTask

TerminateTask

Initialized task node

Running task node

Scheduler Queues

```
DefineTask( <TaskName>, <TaskProperties>, <TaskPriority>,
<EntryPoint> [,<TaskBank>] [,<InterruptMask>]
[,<TaskStack>] [,<TaskStackSize] );
```

Basiskomponenten

CPU

OS

Unterbrechungsbehandlung

ISR

RESOURCE

TASK

COUNTER

EVENT

ALARM

Kommunikation

MESSAGE

COM

NM

OSEK Komponenten

The Event mechanism

- **is a means of synchronisation**

   in conjunction with the wait-state of an extended task

- **is only provided for extended tasks**

   Each extended task has a definite number of events

- **initiates state transitions of tasks to and from the waiting state.**

System services related to Event Processing:

Service Name	Description
SetEvent	Sets events of the given task according to the event mask
ClearEvent	Clears events of the calling task according to the event mask
GetEvent	Gets the current event setting of the given task
WaitEvent	Transfers the calling task into the waiting state until specified events are set

Synchronisation of non-preemtive tasks
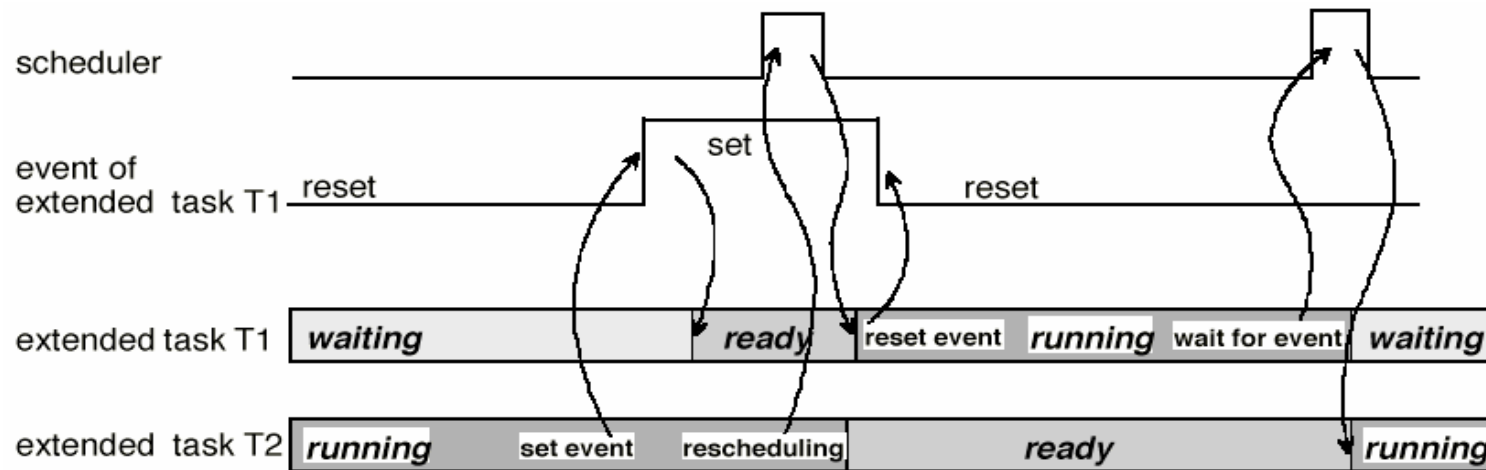
Synchronisation of full-preemtive tasks
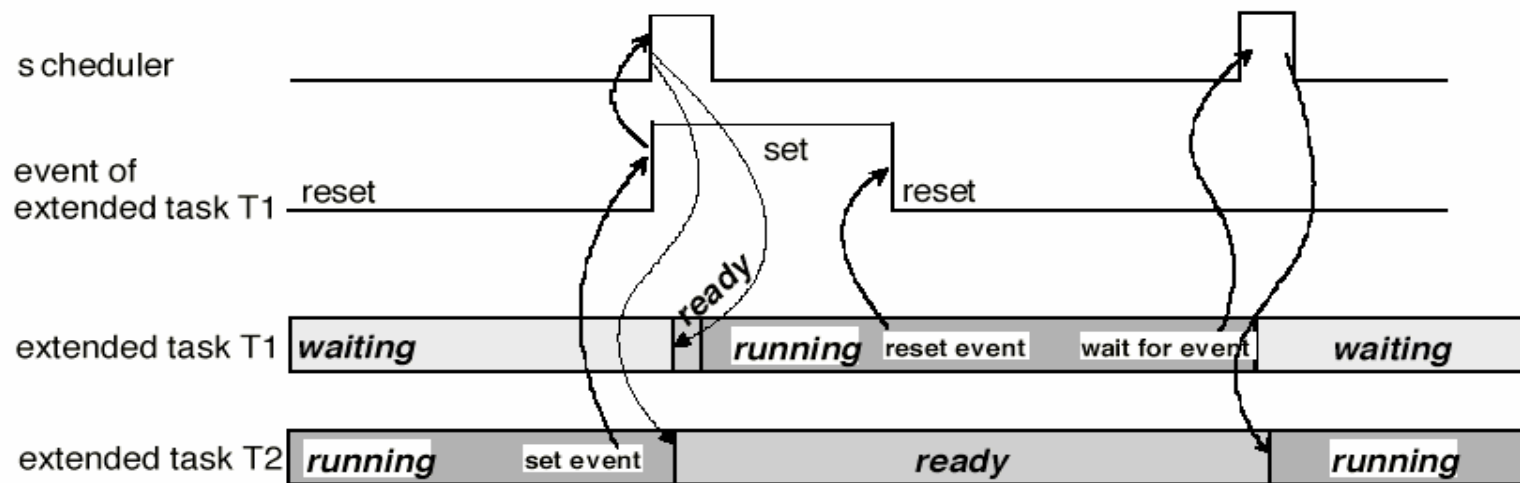
An Alarm may start a task or set an event.

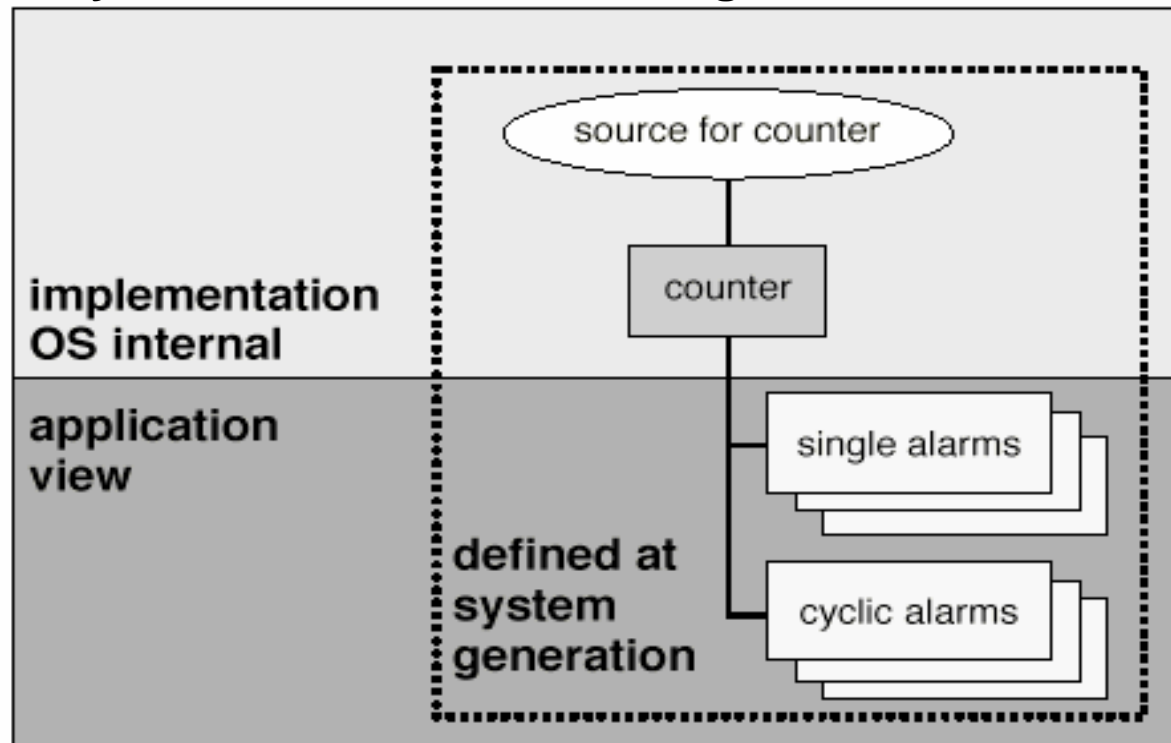**An alarm is statically assigned at system generation time to:**

- **one counter**
- **one task**
- **a notation if the task is to be activated or the event is to be set**
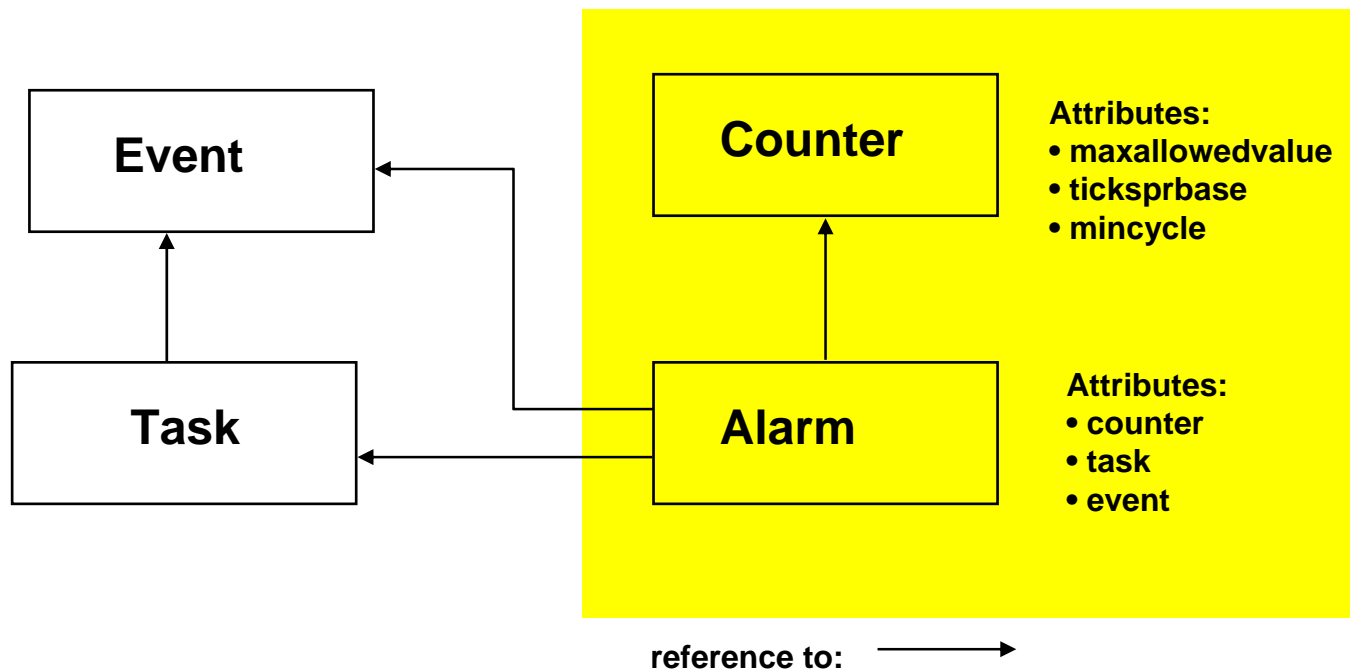
Alarms

- **Processing of reccuring events**
  - **- single alarm and cyclic alarm**
  - **- relative (to a timer value) and absolute points in time**
  - **- time (clock ticks) and event counting**

Layered model of alarm management

Counter and Alarm Management Run-time Services

Service Name	Description
InitCounter	Sets the initial value of the counter
CounterTrigger	Increments the counter value
GetCounterValue	Gets the counter current value
GetCounterInfo	Gets counter parameters
SetRelAlarm	Sets the alarm with a relative start value
SetAbsAlarm	Sets the alarm with an absolute start value
CancelAlarm	Cancels the alarm: the alarm is transferred into the STOP state
GetAlarm	Gets the time left before the alarm expires

Counters and Alarm Generation

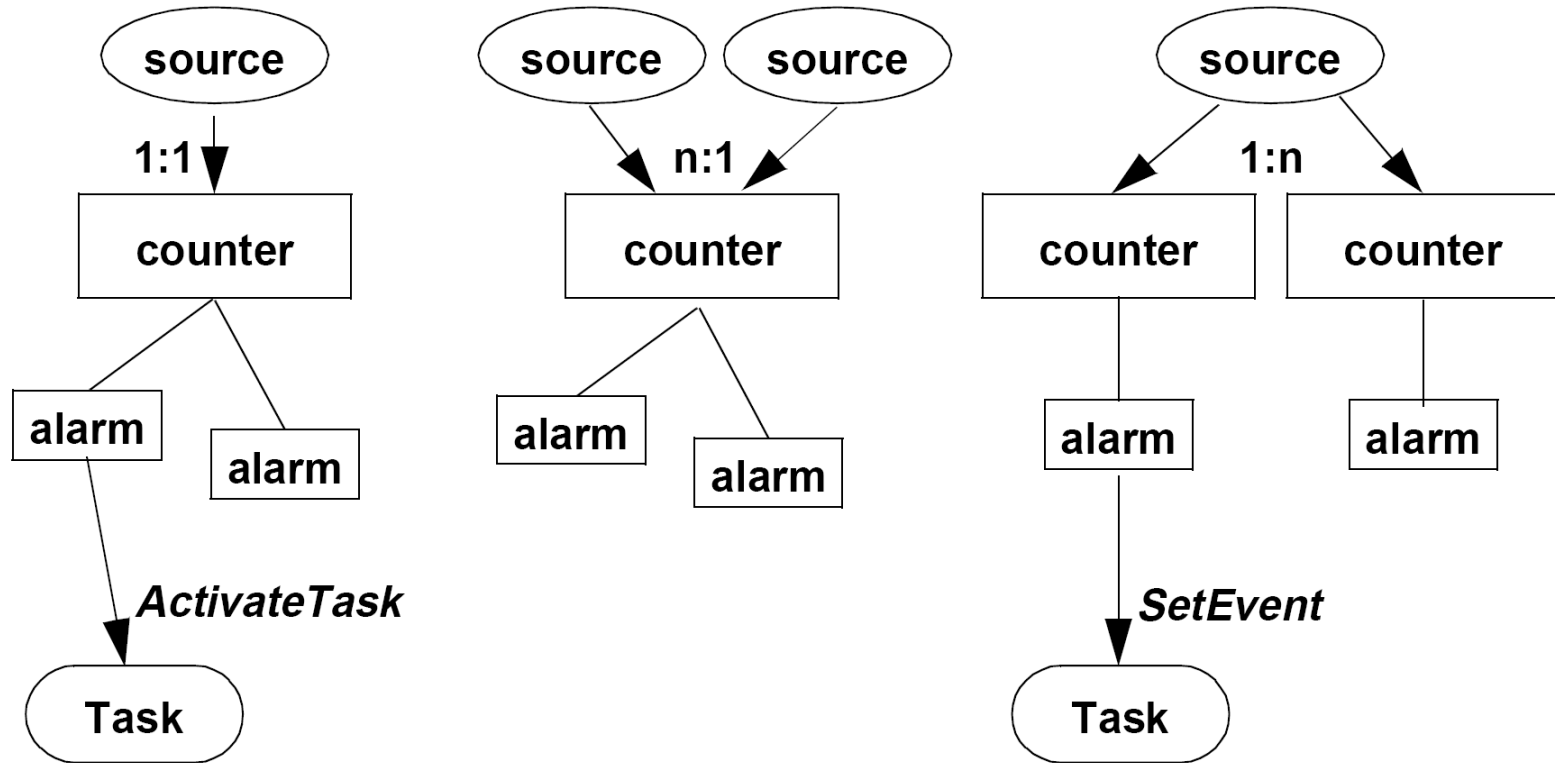DefineCounter(<CounterID>, <maxallowedvalue>, <ticksperbase> [,<mincycle>]);

DefineSystemTimer(<CounterID>, <maxallowedvalue>, <ticksperbase>,
                    <tickduration> [,<mincycle>] [, <HardwareType> [, <HardwareParams>]]);

DefineAlarm(<AlarmID>, <CounterID>, <TaskID> [,<Event>]);

DeclareCounter(CtrRefType <CounterID>);
DeclareAlarm(AlarmRefType <AlarmID>);

Counters and Alarms

source → 1:1 → counter

counter → alarm, alarm

alarm → *ActivateTask* → Task

source, source → n:1 → counter

counter → alarm, alarm

source → 1:n → counter, counter

counter → alarm → *SetEvent* → Task

counter → alarm

Alarms

Relative Alarm (starting from current counter value)
Absolute Alarm

Variations of Absolute Alarm

Basiskomponenten

Unterbrechungsbehandlung

Kommunikation
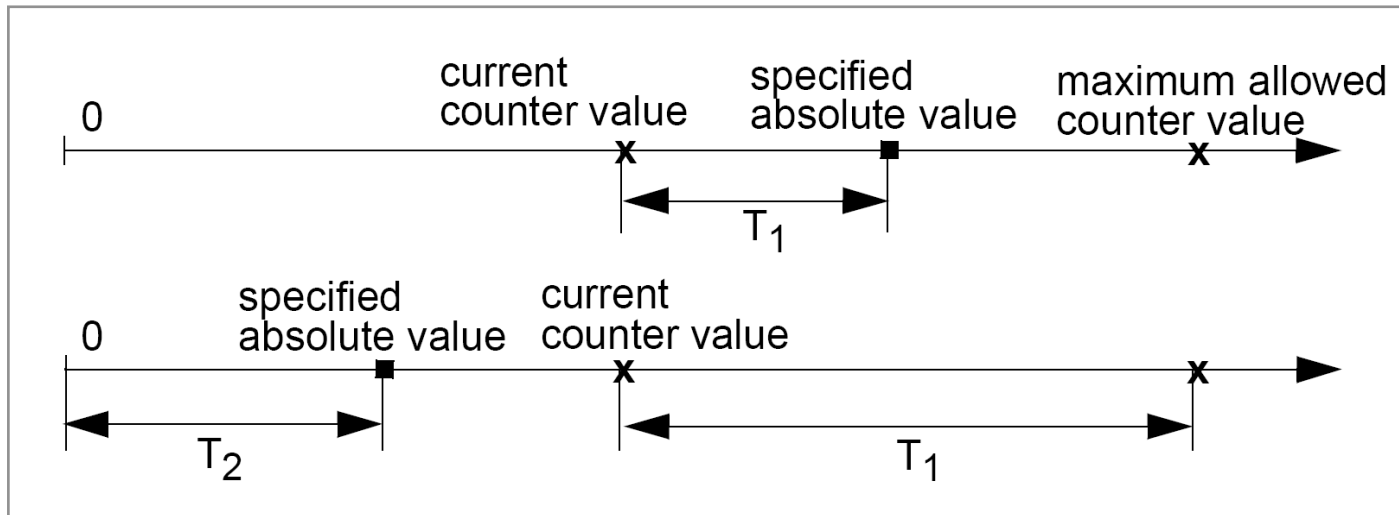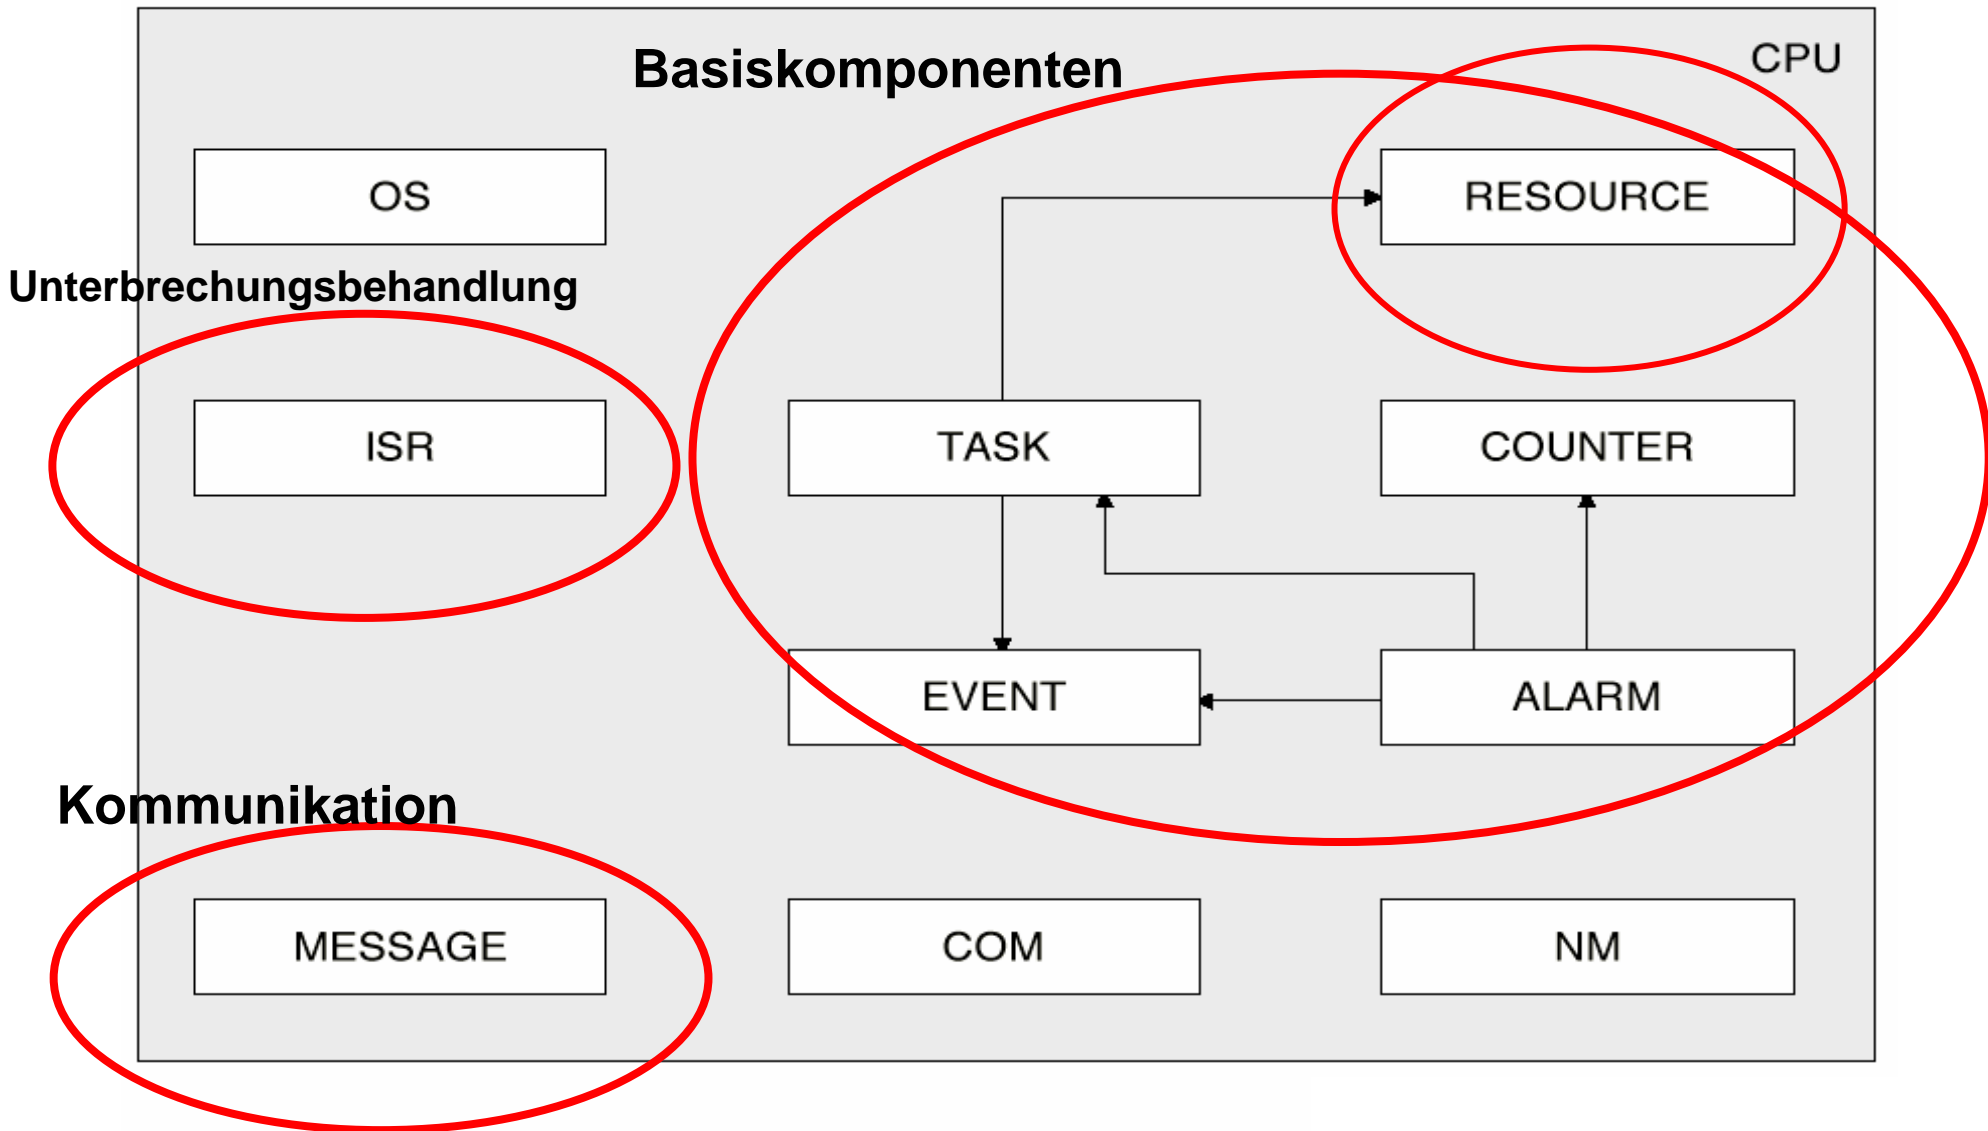
OSEK Komponenten

Resource Management

The resource management is used to coordinate concurrent accesses of several full preemptive tasks with different priorities to shared resources, e.g. management entities (scheduler), program sequences, memory or hardware areas.

Resource management ensures that:
- two tasks cannot occupy the same resource at the same time.
- priority inversion can not occur.
- deadlocks do not occur by use of these resources.
- access to resources never results in a waiting state.

The functionality of resource management is only required in the following cases:
- full-preemptive tasks
- non-preemptive scheduling, if resources are also to remain occupied beyond a scheduling point (in OSEK this only applies to the system service "*Schedule( )*")
- non-preemptive scheduling, if the user intends to have the application code executed under other scheduling policies, too

The resource management is mandatory for all conformance classes.

Basiskomponenten

CPU

OS

Unterbrechungsbehandlung

ISR

RESOURCE

TASK

COUNTER

EVENT

ALARM

Kommunikation

MESSAGE

COM

NM

OSEK Komponenten

Inter-Task Kommunikation

• **Transparent local and network communication by messages.**

• **Messages are encapsulated in message objects which are handled by the OS.**

• **Message objects are defined at system configuration.**

• **A uinique identifier (UID) is assigned to message objects.**

• **Tasks reference messages by this UID.**

• **OSEK distinguishes between queued (event) and unqueued (state) messages.**

• **Task activation and event signalling can be statically defined to be performed at message arrival for notification.**

Inter-Task Kommunikation

OSEK unterscheidet:

      - **State Messages** und

      - **Event Messages**

State Messages **repräsentieren den Wert einer Echtzeitvariablen. Die Empfangsoperation liest die Nachricht, ohne sie zu zerstören. Eine State Message enthält stets den aktuellen Wert der Variablen, d.h. ein alter Wert wird von einem neuen Wert beim „send" überschrieben.**

Event Messages **werden gepuffert, d.h. ein neuer Wert wird in einen neuen Puffer geschrieben. Event Messages werden beim „receive" konsumiert, d.h. zerstörend gelesen.**

Zusammenfassung: Features of the Message Concept

| State Message	Event Message
No buffering	Buffering in a FIFO-queue
No consumption of message	Consumption of messages
Direct access possible in non-preemptive systems (no copies)	No direct access possible (always copies)
Static definition of task activation or event signalling	
Static definition of a message alarm	

Inter-Task Kommunikation

➡ **OSEK unterstützt 1:1 und 1:N Kommunikation für State- und Event-messages.**

➡ **Optional kann statisch festgelegt werden, dass eine Task aktiviert oder notifiziert wird, wenn eine neue Nachricht ankommt.**

➡ **OSEK stellt keine spezielle Systemunterstüzung für das Warten auf Nachrichten zur Verfügung. Hierzu kann der allgemeine „Event-" Mechanismus verwendet werden.**
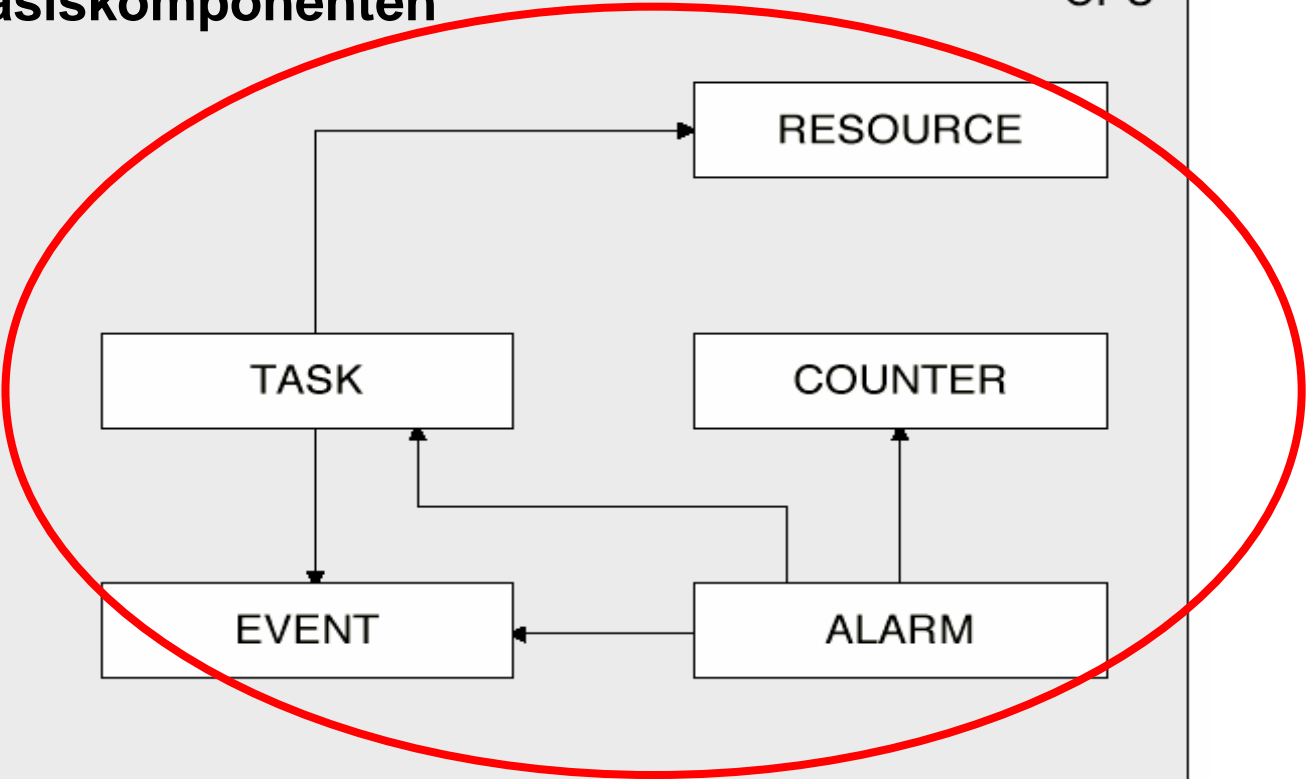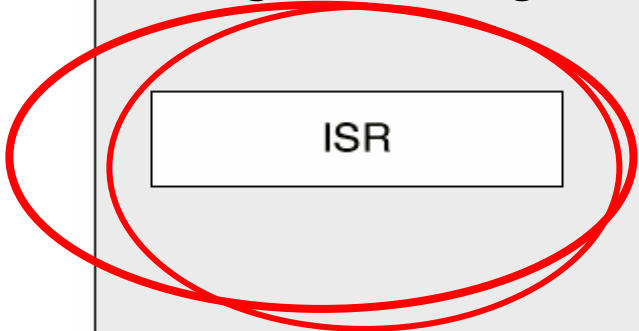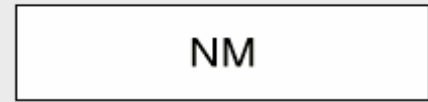
OSEK Komponenten

Interrupt Categories

ISR category 1

ISRs of this type does not use any operating system service. These ISR does not use OS services EnterISR and LeaveISR, consequently, ISR of these type are executed on the current stack. In this case, if ISR uses the stack space for its execution, the user is responsible for the appropriate stack size.

ISR category 2

In ISR category 2 the OSEK Operating System provides the ISR frame to execute more complicated user code. The ISR frame are instructions between OS services EnterISR and LeaveISR. Such ISRs must have the EnterISR call at their beginning to switch to the ISR stack and save the initial interrupt mask. At the end of the ISR the LeaveISR service must be executed to switch back to the task stack and restore the interrupt mask.

ISR category 3

Such ISR's are similar to those of category 2. But the location of the ISR-frame in the code segment is application dependent and user defined.

Interrupt Management

Service Name	Description
Interrupt Management Services	
EnterISR	Registers the switching to the interrupt level and switch context to the ISR stack
LeaveISR	Registers the leaving of the ISR level
EnableInterrupt	Enables interrupts in accordance with the given mask
DisableInterrupt	Disables interrupts in accordance with the given mask
GetInterruptMask	Returns the current state of interrupts
Services allowed for use in ISR	
ActivateTask	Activates the specified task (puts it into the *ready* state)
SendStateMessage	Sends a state message to the specified task
SendEventMessage	Sends an event message to the specified task
CounterTrigger	Increments a counter value

Error Handling, Tracing, Debugging
Hook Routines

The OSEK operating system provides system specific hook routines to allow user-defined actions within the OS internal processing.
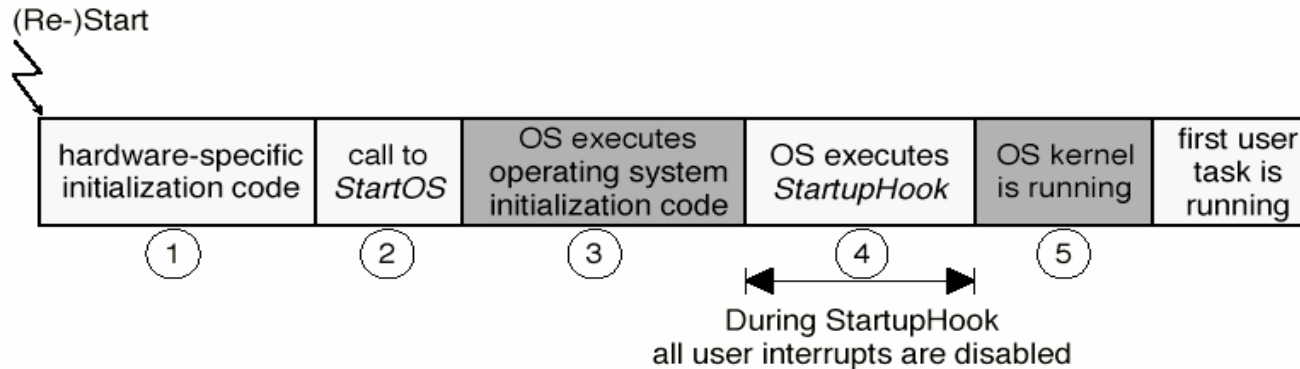
Those hook routines are:
• called by the operating system, in a special context depending on the implementation of the operating system
• higher priority than all tasks
• using an implementation dependent calling interface.
• part of the operating system but user defined
• implemented by the user
• standardised in interface per OSEK OS implementation, but not standardised in functionality (environment and behaviour of the hook routine itself), therefore usually hook routines are not portable
• are only allowed to use a subset of API functions
• optional (the implementation should omit calls to hook routines which do not exist)

In the OSEK operating system hook routines may be used for:
• system start-up. The corresponding hook routine (StartupHook) is called after the operating system start-up and before the scheduler is running.
• system shutdown. The corresponding hook routine (ShutdownHook) is called to a) request a shutdown by the application and b) force a system shutdown in case of a severe error.
• tracing or application dependent debugging purposes as well as user defined extensions of the context switch.
• error handling. The corresponding hook routine (ErrorHook) is called if a system call returns a value not equal to E_OK.

Each implementation of OSEK has to describe the interfaces and conventions for the hook routines, e.g. interfaces of system functions that may be called by the hook routines.

Example: StartupHook

(Re-)Start

hardware-specific initialization code	call to *StartOS*	OS executes operating system initialization code	OS executes *StartupHook*	OS kernel is running	first user task is running
1	2	3	4	5	

During StartupHook
all user interrupts are disabled

1) After a reset, the user is free to execute (non-portable) hardware specific code. The non-portable section ends by detection of the application mode.
2) Call StartOS with the application mode as a parameter. This call starts the operating system.
3) The operating system performs internal start-up functions and
4) calls the hook routine StartupHook, where the user may place initialisation procedures. During this hook routine, all user interrupts are disabled.
5) The operating system enables user interrupt according to the INITIAL_INTERRUPT_DESCRIPTOR, and starts the scheduler.

Example: Debugging with "OSPreTask" and "OSPostTask"
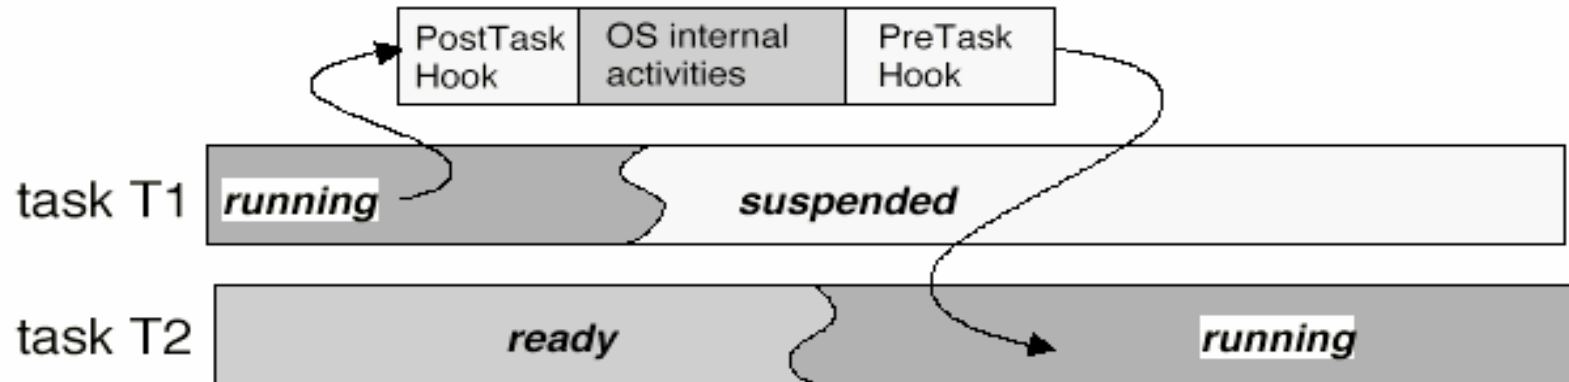
Two hook routines (PreTaskHook and PostTaskHook) are called on task context switches.

These two hook routines may be used for debugging or time measurement (including context switch time). Therefore PostTaskHook is called after leaving the context of the old task, PreTaskHook is called before entering the context of a new task.