[Stei94g]

Ralf Steinmetz; Multimedia Operating Systems: Resource Reservation, Scheduling, File Systems, and Architectures; Technical Report 43.9402, IBM European Networking Center Heidelberg, Germany, 1994; Vorabversion zu [Stei95a und Stei95b].

Multimedia Operating Systems: Resource Reservation, Scheduling, File Systems, and Architectures

Ralf Steinmetz

Technical Report 43.9402

IBM European Networking Center Vangerowstraße 18 • 69115 Heidelberg • Germany

Phone: +49-6221-59-3000 • Fax: +49-6221-59-3400

steinmet @ dhdibmip_bitnet steinmet @ heidelbg.vnet

Abstract:

Multimedia applications demand for the processing of audio and video data such that humans perceive these media in a natural -error free and non-artificial- way. This continuous-media data has its origin at sources like microphones, cameras and files. From these sources the data is transferred to destinations like loudspeakers, video windows and files located at the same computer or at a remote station. On the way from source to sink the digital data is processed by at least some type of move, copy or transmit operation. Therefore in this data manipulation process there are always many resources which are under control of the operating system. This resource management and the respective scheduling must be performed according to the realtume demands of multimedia applications.

This paper outlines the main features of the operating systems capable of handling multimedia data, namely resource management, scheduling and file system issues. For all these topics the most relevant approaches and alternatives are surveyed. This survey concludes with a presentation of the most well known architectures of operating system extensions for multimedia.

1. Introduction

The operating system is the shield of the computer hardware to all other software components. It provides a comfortable environment for the execution of programs, and it ensures an effective utilization of the computer hardware. The operation system offers various services related to the essential resources of a computer: CPU, main memory, storage and all input and output devices.

.....

۲

Multimedia applications demand for the processing of audio and video data such that humans perceive these media in a natural -error free and non-artificial- way. This continuous-media data has its origin at sources like microphones, cameras and files. From these sources the data is transferred to destinations like loudspeakers, video windows and files located at the same computer or at a remote station. On the way from source to sink the digital data is processed by at least some type of move, copy or transmit operation. Therefore in this data manipulation process there are always many resources which are under control of the operating system.

The integration of discrete and continuous multimedia data demands for additional services that have to be provided by many operating system components. The major aspect in this context is real-time processing of continuous-media data.

The process management must take into account the timing requirements imposed by the handling of multimedia data. Appropriate scheduling methods should be applied. In contrast to the traditional real-time operating systems, multimedia operating systems also have to consider tasks without hard timing restrictions under the aspect of fairness.

To obey timing requirements, single components are conceived as resources that are reserved prior to execution. This concept of resource reservation has to cover all resources on a data path, i.e., all resources that deal with continuous media. It also may affect parts of the application that process continuous-media data. In distributed systems, for example, the resource management also comprises network capacity [HVWW94].

The communication and synchronization between single processes must reflect the restrictions through real-time requirements and the timing relations among different media.

The main memory is available as a shared resource to single processes. In multimedia systems, the memory management has to provide access to data with a guaranteed timing delay and efficient data manipulation functions. For instance, physical data copy operations have to be avoided due to their negative impact on performance; buffer management operations (such as known from communication system) should be used.

The database management is an important component in multimedia systems. But, database management abstracts from the details of storing data on secondary media storage. Therefore the database management should rely on file management services provided by the multimedia operating system in order to access single files and file systems. For example, the incorporation of a CD-ROM XA file system as an integral part of a multimedia file system allows transparent and guaranteed continuous retrieval of audio and video data to any application using the file system; the database system is one of those applications. However often database systems implement their own access to stored data.

Since the operating system shields devices from applications programms it has to provide services for the device management. In multimedia systems the important issue is the integration of audio and video devices in a similar way as any other input/output device. The addressing of a camera can be done in a similar way as the addressing of a keyboard in the same system, although most of the present systems do not apply this technique.

2

÷.,

Product information dealing with operating system extensions towards the integration of multimedia [IBM92, Micr91, DrMu92, IBM92a, IBM92b] typically provide a detailed description of application interfaces. In this paper we concentrate on basic concepts and internal tasks of a multimedia operating system.

As the essential aspect of any multimedia operating system is the notion of "real-time", the following paragraph starts with a detailed study of this idea in its relationship to multimedia. Subsequently the concept of resource management is discussed. The process management contains a brief presentation of traditional real-time scheduling algorithms. Further, their suitability and adaptability towards continuous-media processing is examined. The section of file systems outlines disk access algorithms, data placement and structuring. The subsequent paragraphs illustrate interprocess communication and synchronization, memory management, and device management. We conclude this presentation of operating system issues with a discussion on typical system architectures which comprise real-time and non real-time environments.

2. Real-Time

The notion of "real-time" has been developed independently from the discussion of continuous-media processing. Therefore, the next paragraph starts with a general definition of realtime. Later-on it shows the relevance of real-time for multimedia data and processes.

2.1 The Notion of 'Real-Time'

The German national institute for standardization DIN, similar to the American ANSI, defines a real-time process in a computer system. "A real-time process is a process which delivers the results of the processing in a given time-span". Programs for the processing of the data must be available during the entire run-time of the system. The data may require processing at an a priori known point in time, or it may be demanded without any previous knowledge [DIN85].

The system must enforce externally defined time constrains. Internal dependencies and their related time limits are implicitly considered. External events occur – depending on the application – deterministically (at a predetermined instant) or stochastically (randomly, with an unknown timing). The real-time system has the permanent task to receive information, occurring spontaneously or in periodic time intervals, from the environment and/or to deliver it to the environment in time.

The main characteristic of real-time systems is the correctness of the computation. This correctness does not only apply to an errorless computation, but it also depends on the time the result is presented [ShRa89]. Hence, a real-time system can not only fail because of massive hardware or software failures, but also because the system is unable to execute its critical workload in time [KrLe91]. Deterministic behavior of the system refers to the adherence of, in advanced defined, time spans for the manipulation of data, i.e. a guaranteed response time exists. Speed and efficiency are not -as often assumed by mistake- the main characteristics of a real-time system. In a petro-chemical plant for example, the result is not only unacceptable when the engine of a vent responds too fast, but also when it responds with a large delay. Another example is the playback of a video sequence in a multimedia system. The result can only be accepted when the video is presented neither too fast nor too slow. Timing and logical dependencies among different related tasks, processed at the same time, have also to be considered. These dependencies refer to the internal and external restrictions. In context of multimedia data streams, this refers to the processing of synchronized audio and video data where the relation between the two media has to be considered.

Deadlines

3

A *deadline* represents the latest time for the presentation of a processing result. It marks the border between normal (correct) and anomalous (failing) behavior. In real-time system we have both, hard and soft deadlines.

The term *soft deadline* is often used for deadline which cannot be exactly determined and where failing does not produce an unacceptable result. We understand a soft deadline as a deadline which in some cases is missed and the miss may be tolerated as long as (1) not too many deadlines are missed and/or (2) they are not missed by much. Such soft deadlines are only reference points with a certain acceptable tolerance. For example start and arrival times of planes or trains where the deadlines can vary about 10 minutes can be considered as soft deadlines.

Whereas soft deadlines may be violated, hard deadlines should never. A hard deadline violation is a system failure. Hard deadlines are determined by the physical characteristics of realtime processes. Failing such a deadline causes costs which can be measured in terms of money (e.g., inefficient use of raw materials in a process control system), or human and environmental terms (e.g. accidents due to untimely control in a nuclear power plant or fly-by-wire avionics systems) [Jeff90].

Characteristics of Real-Time Systems

The necessity for deterministic and predictable behavior of a real-time system leads to required processing guarantees for time-critical tasks. Such guarantees cannot be assured for events that occur at random intervals with unknown arrival times, processing requirements or deadlines. Further, all given guarantees are only valid under the premise that no processing machine collapses during the run-time of real-time processes. A real-time system is distinguished by the following features (c.f. [ShRa89]):

- Predictably fast response to time-critical events and accurate timing information. For example, in the control system of a nuclear power plant the response to a malfunction must occur within a well-defined period in order to avoid a potential disaster.
- A high degree of schedulability. Schedulability refers to the degree of resource utilization at, or below, the deadline of each time-critical task can be taken into account.
- Stability under transient overload. Under system overload the processing of critical tasks must be ensured. These critical tasks are vital to the basic functionality provided by the system.

Management of manufacturing processes and the control of military systems are the main application areas for real-time systems. Such process control systems are responsible for realtime monitoring and control. Real-time systems are also used as command and control systems in fly-by-wire aircraft, automobile anti-lock braking systems and the control of nuclear power plants [KrLe91]. New areas for real-time systems are computer conferencing and multimedia in general, the topic of our work.

2.2 Real-Time and Multimedia

Audio and video data streams consist of single, periodically changing values of continuousmedia data, e.g., audio samples or video frames. Each logical data unit (LDU) has to be presented at a well determined deadline. Jitter is only allowed before the final presentation to the user. A piece of music, for example, has to be played back with constant speed.

To fulfill these timing requirements of continuous media, the operating system must use realtime scheduling techniques. These techniques have to be applied to all system resources involved in the continuous-media data processing. The entire *end-to-end* data path is involved. The CPU is just one of these resources, all components have to be considered including main memory, storage, I/O devices and networks.

Real-Time Requirements in Multimedia Systems

Traditional real-time scheduling techniques, used for command and control systems in application areas such as factory automation or aircraft piloting, have a high demand for security and fault-tolerance. The requirements derived from these demands somehow counteract to realtime scheduling efforts applied to continuous media. Multimedia systems which are not used in traditional real-time scenarios have different (in fact, more favorable) real-time requirements:

- The fault-tolerance requirements of multimedia systems are usually less strict than of those real-time systems that have a direct physical impact. A shorttime failure of a continuous-media system will not directly lead to the destruction of technical equipment or constitute a threat to human life. Please note, that this is a statement which applies in general but not always. The support of remote surgery, e.g., by video and audio has stringent delay and correctness requirements.
- For many applications missing a deadline in a multimedia system is although it should be avoided - not a severe failure. It may even be unnoticed: If an uncompressed video frame (or parts of it) is not available on time it can be simply omitted. The human viewer will hardly notice it, assuming this does not happen for a contiguous sequence of frames. For audio, requirements are more stringent because the human ear is more sensitive to audio gaps than the human eye is to video jitter.
- A sequence of digital continuous-media data results from periodically sampling a sound or image signal. Hence, in processing the data units of such a data sequence, all time-critical operations are periodic. Schedulability considerations for periodic tasks are much easier than for sporadic ones [Mok84].
- The bandwidth demand of continuous media is not always *that* stringent, it must not be a priori fixed but it may eventually be lowered. As some compression algorithms are capable of using different compression ratios leading to different qualities the required bandwidth can be negotiated. If not enough capacity for full quality is available the application may also accept a reduced quality (instead of no service at all). The quality may also be adjusted dynamically to the available bandwidth, *e.g.*, by changing encoding parameters. This is known as scalable video.

In a traditional real-time system, timing requirements result from the physical characteristics of the technical process to be controlled, *i.e.*, they are provided externally. Some applications must meet external requirements, too. A distributed music rehearsal is a futuristic example: Music played by one musician on an instrument connected to his workstation has to be made available to all other members of the orchestra within a few milliseconds, otherwise the underlying knowledge of a global unique time is disturbed. If human users are involved just in the input or only the output of continuous media, delay bounds are more flexible. Consider the playback of a video from a remote disk. The actual delay of a single video frame to be transferred from the disk to the monitor is unimportant. Frames must only arrive in a regular fashion. The user will notice any difference in delay only as start delay (i.e., for the first video frame to be displayed).

3. Resource Management

Multimedia systems with integrated audio and video processing will be at the limit of their capacity, even with data compression and utilization of new technologies. Current computers do not allow "processing" of these data according to their deadlines without any reservation and real-time process management. "Processing" in this context refers to any kind of manipulation and communication of data.



Figure 1: Window of insufficient resources

This stage of development is known in [ATWG90] as "the window of insufficient resources" (see Figure 1). With CD-DA (Compact Disc Digital Audio) quality highest audio requirements are satisfied. In video technology the required data transfer rate will ascend with the development of digital HDTV and larger TV-screens. Therefore, no redundancy of resource capacity can be expected in the near future. In a multimedia system the given timing guarantees for the processing of continuous media must be adhered at every system component, i.e., at every hardware and software component on the data path. The actual requirements depend upon the type of media and the nature of the supported applications [StMe92]. A video image should not be presented late because the communication system has been busy with a transaction from a database management system. In any realistic scenario we encounter several multimedia applications which concurrently make use of shared resources. Hence, even high band width networks and huge processing capabilities require the use of real time mechanisms in order to provide guaranteed data delivery. The concept of integration does, on the other hand, not allow to solve this problem by rendering the system for traditional applications.

Thus, in an integrated distributed multimedia system several applications compete for system resources. This shortage of resources requires their careful allocation. The system management must employ adequate scheduling algorithms to serve the requirements of the applications. Thereby the resource is first allocated and then managed.

The "resource management" in distributed multimedia systems covers several computers including the involved communication networks. It allocates all resources involved in the data transfer process between sources and sinks. For instance, a CD-ROM XA device has to be allocated exclusively, each CPU on the data path has to provide 20% of its capacity, the network has to allocate a certain amount of its bandwidth and the graphic processor has to be reserved up to 50% for such a process. The required throughput and a certain delay is guaranteed. At the connection establishment phase the resource management ensures that the new "connection" does not violate performance guarantees already provided to existing connections.

Applied to operating systems, this model covers the CPU (including process management), memory management, the file system and device management. Therefore we chose to detail

this issues for all resources in a generic notion of resources in the following paragraphs. The resource reservation is identical for all resources whereas the management is different for each.

3.1 Resources

A resource is a system entity required by tasks for manipulating data. Each resource has a set of distinguishing characteristics classified using the following scheme:

- A resource can be active or passive. An *active resource* is the CPU or a network adapter for protocol processing, it provides a service. A *passive resource* is the main memory, bandwidth of communication ????? or a file system (whenever we do not take care of the processing of the adapter); it denotes some system capability required by active resources.
- A resource can be either used **exclusively** by one process at the time or shared between various processes. Active resources are often exclusive, passive resources can usually be shared among processes.
- A resource that exists only once in the system is known as a single, otherwise it is a multiple resource. In a transputer based multiprocessor system the individual CPU is a multiple resource.

Each resource has a capacity which results from the ability of a certain task to perform using the resource in a given time-span. In this context **capacity** refers to CPU capacity, frequency range or, for example, the amount of storage. For real-time scheduling only the temporal division of resource capacity among real-time processes is of interest.

Process management belongs to the category of active, shared, and most often to single resources. A file system on an optical disc with CD-ROM XA format is a passive, shared, single resource.

3.2 Requirements

The requirements of multimedia applications and data streams have to be served by the single components of a multimedia system. The resource management maps these requirements onto the respective capacity. The transmission and processing requirements of local and distributed multimedia applications can be specified according to the following characteristics:

- 1. The *throughput* is determined by the needed data rate a connection needs in order to satisfy the application requirements. It also depends on the size of the data units.
- 1. We distinguish between local and global (end-to-end) delay:
 - a. The *delay* "at the resource" is the maximum time span for the completion of a certain task at this resource.
 - b. The *end-to-end delay* is the total delay for a data unit to be transmitted from the source to its destination. For example, the source of a video telephone is the camera the destination is the video window on the screen of the partner.
- 3. The *jitter* (or delay jitter) determines the maximum allowed variance in the arrival of data at the destination.
- 4. The *reliability* defines error detection and correction mechanisms used for the transmission and processing of multimedia tasks. Errors can be ignored, indicated and/or corrected. It is important to notice that error correction through re-transmission is rarely appropriate for time-critical data because the re-transmitted data will usually arrive late.

Forward error correction mechanism are more useful. In terms of reliability, we also mean the CPU errors due to unwanted delays in processing a task which exceed the demanded deadlines.

In accordance with communication systems these requirements are also known as "quality of service parameters" (QoS).

3.3 Components and Phases

One possible realization of resource allocation and management is based on the interaction between clients and the respective resource managers. The *client* selects the resource and requests a resource allocation by specifying its requirements through a QoS specification. This is equivalent to a workload request. First the *resource manager* checks its own resource utilization, and decides if the reservation request can be served or not. All existing reservations are stored, this way their share in terms of the respective resource capacity is guaranteed. Moreover, this component negotiates the reservation request with other resource managers if necessary.

The following example of a distributed multimedia system illustrates this generic scheme. During the connection establishment phase the QoS parameters are usually negotiated between the requester (client application) and the addressed resource manager. The negotiation starts in the simplest case with specification of the QoS parameters by the application. The resource manager checks whether these requests can be guaranteed or not. A more elaborate method is to optimize single parameters. In this case two parameters are determined by the application (e.g., throughput and reliability), the resource manager then calculates the best achievable value for the third parameter (e.g., delay). To negotiate the parameters for end-to-end connections over one or more computer networks, resource reservation protocols like ST-II are employed [Topo90]. Here, the resource managers of the single components of the distributed system allocate the necessary resources.



Figure 2: Components grouped for the purpose of video data transmission

In the following case shown in Figure 2 two computers are connected over a LAN. The transmission of video data between a camera connected to a computer *Server* and the screen of the computer *User* involves for all depicted components a resource manager.

This example illustrates that in addition to the individual resource managers there must exist a protocol for coordination between these services is necessary, such as ST-II.

Phases of the Resource Reservation and Management Process

A resource manager provides components for the different phases of the allocation and management process:

- 1. Schedulability Test: The resource manager checks with the given QoS parameters (e.g. throughput and reliability) if there is enough remaining resource capacity available to handle this additional request.
- 2. Quality of Service Calculation: After the schedulability test the resource manager calculates the best possible performance (e.g. delay) the resource can guarantee for the new request.
- 3. *Resource Reservation:* The resource manager allocates the required capacity in order to meet the QoS guarantees for each request.
- 4. **Resource Scheduling:** Incoming messages from connections are scheduled according to the given QoS guarantees. For process management, for instance, the allocation of the resource is done by the *scheduler* at the moment the data arrives for processing.

With respect to the last phase for each resource a scheduling algorithm is defined. The schedulability test, QoS calculation and resource reservation depend upon this algorithm used by the scheduler.

3.4 Allocation Scheme

Reservation of resources can be made either in a *pessimistic* or in an *optimistic* way:

- The *pessimistic approach* avoids resource conflicts by making reservations for the worst case, i.e. resource bandwidth for the longest processing time and the highest rate which might ever be needed by a task is reserved. Resource conflicts are therefore avoided. This leads potentially to an underutilization of resources. In a multimedia system the remaining processor time (i.e. the time reserved for traffic but not used) can be used by discrete media tasks. This method results in a guaranteed quality of service.
- With the optimistic approach resources are reserved according to an average workload only. This means that the CPU is only reserved for the average processing time. This approach may overbook resources with the possibility of unpredictable packet delays. QoS parameters are met as far as possible. Resources are highly utilized, though an overload situation may result in failure. To detect an overload situation tohandle it accordingly an monitor can be implemented. The monitor may, for instance, preempt processes according to their importance.

The optimistic approach is considered to be an extension of the pessimistic approach. It requires that additional mechanism to detect and solve resource conflicts are implemented.

3.5 Continuous Media Resource Model

This section specifies a model frequently adopted to define QoS parameters and hence the characteristics of the data stream. It is based on the model of *Linear Bounded Arrival Processes (LBAP)* as described in [Ande93]. In this model a distributed system is decomposed into a chain of resources traversed by the messages on their end-to-end path. Examples of such resources are single schedulable devices (such as CPU) or combined entities (such as networks).

The data stream consists of LDUs. In this context, we call them *messages*. In a first step the data stream itself is characterized as strictly periodic, irregular with a definite maximum message size. Various data streams are independent of each other.

A closer inspection shows a possible variance of the message rate, the maximum rate is welldefined. This variance of the data rate results in an accumulation of messages (*burst*) where the maximal range is defined by the maximum allowed number of messages.

In the LBAP model a burst of messages consists of messages arrived *ahead of schedule*. LBAP is a message arrival process at a resource defined by three parameters.

- M = Maximum message size (byte/message)
- R = Maximum message rate (message/second)
- B = Maximum Burstiness (message)

Example

The LBAP model is discussed in terms of a specific example:

Two workstations are interconnected by a LAN. A CD-player is attached to one workstation. Single channel audio data is transferred from the CD-player of this workstation over the network to the other computer. At this remote station, the audio data is delivered to a speaker. The audio signal is sampled with the frequency of 44.1 kHz. Each sample is coded with 16 bit. This results in a data rate of

 $R_{byte} = 44100 \text{ Hz} \times \frac{16 \text{ bit}}{8 \text{ bits/byte}} = 88200 \text{ bytes/s}$

The samples on a CD are assembled to frames. These frames are the audio messages to be transmitted. 75 of these audio messages are transmitted per second (R) according to the CD-format standard. Therefore we encounter a maximum message size of.

$$M = \frac{88200 \text{ bytes/s}}{75 \text{ messages/s}} = 1176 \text{ bytes/message}$$

Up to 12000 bytes are assembled into one packet and transmitted over the LAN. In a packet of 12000 byte transmitted over the LAN we will never encounter more messages than

It obviously follows:

- M = 1176 bytes/message
- R = 75 messages/s
- B = 10 messages

Burst

In the calculation below it is assumed that, because of lower adjacent data rate, a burst never exceeds a maximum data rate. Hence, bursts do not succeed one another. During a time interval of length t, the *maximal number of messages* arriving at a resource must not exceed

$$\overline{M} = B + R \times t(message)$$

e.g. assume t = 1 s

 \overline{M} = (10 messages + 75 messages/s x 1 s = 85 messages)

The introduction of *Burstiness B* allows for short time violations of the rate constraint. This way programs and devices that generate burst of messages can be modelled. Bursts are, e.g., generated when data is transferred from disks in a bulk transfer mode or in our example when messages are assembled to larger packets.

Maximum Average Data Rate

The maximum average data rate of the LBAP is:

 $\vec{R} = M \times R(bytes/s)$

e.g.

 \overline{R} = (1176 bytes/message × 75 messages/s = 88200 bytes/s)

Maximum Buffer Size

It is guaranteed that messages are processed according to their rate. Messages which arrive "ahead of schedule" have to be queued. For delay \leq period the *buffer size* is:

 $S = M \times (B+1)$ (byte)

e.g.

S = 1176 byte/message x 11 message = 12936 byte

Logical Backlog

The function b(m) represents the logical backlog of messages. This is the number of messages which already have arrived "ahead of schedule" at the arrival of message m. Let a_i be the actual arrival time of message $m_i; 0 \le i \le n$: then b(i) is defined by:

$$b(m_0) = 0 \text{ message}$$

$$b(m_i) = max(0 \text{ messages}, b(m_{i-1}) - (a_{i-1})R + 1 \text{ message})$$

e.g.:

$$a_{i-1} = 1.00 \text{ s}; a = 1.013 \text{ s}; b (m_{i-1}) = 4 \text{ messages}$$

$$b(m_i) = max(0 message, 4 message - (1.013 s - 1.00 s) \times 75 messages/s + 1 message)$$

= 4 messages

Logical Arrival Time

The logical arrival time defines the earliest time a message m_i can arrive at a resource when all messages arrive according to their rate. The logical arrival time of a message can then be defined as:

$$l(m_{i}) = a(n_{i}) + \frac{b(m_{i})}{R}$$

e.g.

$$I(m_i) = 1.013 \text{ s} + \frac{4 \text{ messages}}{75 \text{ messages/s}} = 1.06 \text{ s}$$

Equivalently it can be computed as:

$$I(m_0) = a_0$$

 $I(m_i) = \max(a_{i-1}) + \frac{1}{B})$

e.g.

$$I(m_{i-1}) = 1.053 \text{ s}$$

 $I(m_i) = max(1.01\overline{3} \text{ s}, 1.05\overline{3} \text{ s} + \frac{1 \text{ message}}{75 \text{ messages/s}}) = 1.0\overline{6} \text{ s}$

Guaranteed Logical Delay

The guaranteed logical delay of a message m denotes the maximum time between the logical arrival time of m and its latest valid completion time. It results from the processing time of the messages and the competition among different sessions for resources, i.e. the waiting time of the message. If a message arrives "ahead of schedule" the actual delay is the sum of the logical delay and the time by which it arrives too early, it is then larger than the guaranteed logical delay. It can also be smaller than the logical delay when it is completed "ahead of schedule". The deadline d(m) is derived from the delay for the processing of a message m at a resource. The deadline is the sum of the logical arrival time and its logical delay.

Workahead Messages

If a message arrives "ahead of schedule" and the resource is in an idle state, the message can be processed immediately. The message is then called a *workahead message*, the process is a workahead process. A maximum *workahead time A* can be specified (e.g., from the application) for each process. This results in a maximum *workahead limit W*

$$W = A \times R$$

e.g.

$$A = 0.04 \text{ s}$$

0.04 s x 75 messages/s = 3 messages

If a message is processed "ahead of schedule" the logical backlog is greater than the actual backlog. A message is *critical* if its logical arrival time has passed. Throughout the rest of the chapter the LBAP-model is assumed to apply to the arrival processes at each resource. The resource must ensure that the arrival processes at the output interface obeys the LBAP-parameters.

4. Process Management

The process management deals with the resource *main processor*. The capacity of this resource is specified as processor capacity. The process manager maps single processes onto resources according to a specified scheduling policy such that all processes meet their requirements.

In most systems a process under control of the process manager can adopt one of the following states:

- In the initial state no process is assigned to the program. The process is in the *idle* state.
- If a process is waiting for an event, i.e. the process lacks one of the necessary resources for processing, it is in the *blocked* state.
- If all necessary resources are assigned to the process, it is *ready to run*. The process only needs the processor for the execution of the program.
- A process is *running* as long as the system processor is assigned to it.

The process manager is the scheduler. This component transfers process into the ready-to-run state by assignment of a position in the respective queue of the dispatcher, it is the essential part of the operating system kernel. The *dispatcher* manages the transition from *ready to run* to *run*. In most operating systems the next process to run is chosen according to a priority policy. Between processes with the same priority the one with the longest *ready time* is chosen.

Today and in the near future existing operating systems must be considered to be the base of continuous-media processing on workstations and personal computers. I the next years there will be certainly no *new* developed multimedia operating systems which will be accepted in the market, therefore already existing *multitasking* systems must cope with multimedia data handling. The next paragraph provides a brief description of real-time support typically available in such systems.

4.1 Real-Time Process Management in Conventional Operating Systems: An Example

UNIX with the various variants, Windows-NT, Apple's System 7 and OS/2 are, respectively will be, the most widely installed operating systems with multitasking capabilities on Personal Computers (including the Power PC) and workstations. Although some of them are enhanced with special priority classes for real-time processes, this is not sufficient for multimedia applications. In [NHNW93], for example, the SVR4 UNIX scheduler which provides a real-time static priority scheduler in addition to a standard UNIX timesharing scheduler is analyzed. For this investigation three applications have been chosen to run concurrently; "typing" as an interactive application, "video" as continuous-media application and a batch program. The result was that only through trial and error a particular combination of priorities and scheduling class assignments might be found that works for a specific application set, i.e. additional features have to be provided for the scheduling of multimedia data processing. In order to be more specific let us have a deeper look into real-time capabilities of one of these systems, namely OS/2. On basis of this system the available real-time support is demonstrated.

Th**r**eads

OS/2 was designed as a time-sharing operating system without taking serious real-time applications into account. An OS/2 thread can be considered as a light weight process: It is the dispatchable unit of execution in the operating system. A thread belongs to exactly one address space (called process in the OS/2 terminology). All threads share the resources allocated by the respective address space. Each thread has its own execution stack, register values and dispatch state (either executing or ready to run). Each thread belongs to one of the following priority classes.

- The time-critical class is reserved for threads that require immediate attention.
- The *fixed-high* class is intended for applications that require good responsiveness without being time critical.
- The regular class is used for the executing of normal tasks.
- The *idle-time* class contains threads with the lowest priorities. Any thread in this class is only dispatched if no thread of any other class is ready to execute.

Priorities

Within each class 32 different priorities (0, ..., 31) exist. Through time-slicing threads of equal priority have equal chances to execute. A context switch occurs whenever a thread tries to access an already otherwise allocated resource. The thread with the highest priority is dispatched and the time-slice is started again. At the expiration of the time slice, OS/2 can preempt the dispatched thread if other threads of equal or higher priority are ready to execute. The time slice can be varied between 32 msec and 65536 msec. The goal at the determination of the time slice duration is to keep the number of context switches low and to get a fair and efficient schedule for the whole run-time of the system. Threads of the regular class may be subject of a dynamic rise of priority as a function of the waiting time.

Threads are preemptive, i.e. if a higher-priority thread becomes ready to execute, the scheduler preempts the lower-priority thread and assigns the CPU to the higher-priority thread. The state of the preempted thread is recorded so that execution can be resumed later.

Physical Device Driver as Process Manager

In OS/2 applications with real-time requirements can run as physical device drivers (PDD) at ring 0 (kernel mode). These PDDs can be made non-preemptable. An interrupt that occurs on a device (e.g., packets arriving at the network adapter) can be serviced from the PDD immediately. As soon as an interrupt happens on a device, the PDD gets control and does all the work to handle the interrupt. This may also include tasks which could be done by application processes running in ring 3 (user mode). The task running at ring 0 should (but must not) leave the kernel mode after 4 msec.

PDD programming is complicated mainly due to difficult testing and debugging. PDD is bound to its device; it only handles requests from its device regardless to any other events happening in the system. Different streams that request real-time scheduling can only be served by their PDDs. They run in competition with each other without the possibility to coordinate or manage them by any higher instance. This is insufficient for a multimedia system where messages can arrive at different adapter cards. Internal time-critical system activities can not be controlled and managed through PDD's. Therefore they can not be considered and accounted for during scheduling decisions. The execution of real-time processes with PDD's is only a reasonable solution for a system where streams arrive at only one device and no other activity in the system has to be considered. Operating system extensions for continuous-media processing can be implemented as PDD's: In this approach a real-time scheduler and the process management run as PDD being activated by a high resolution timer. In principle, this is the implementation scheme of the OS/2 Multimedia Presentation Manager which represents the multimedia extension to OS/2.

Enhanced System Scheduler as Process Manager

Time-critical tasks can also be processed together with normal applications running in ring 3, the user level. The critical tasks can be implemented by threads running in the priority class time-critical with one of the 32 priorities within this class. Each real-time task is assigned to one thread. A thread is interrupted if another thread with higher priority requires processing. Non time-critical applications run as threads in the regular class. They are dispatched by the operating system scheduler according to their priority.

The main advantage of this approach is the control and coordination of all time-critical threads through a higher instance, the system scheduler. This instance, running with a higher priority than all other threads, controls and coordinates threads according to the adapted scheduling algorithm and the respective processing requirements. It can observe the run-time behavior of single threads. Another entity, the resource manager, determines feasible schedules, takes care of quality of service calculating and resource reservation. The competition for the CPU is regulated. The employment of an internal scheduling strategy and resource management allows the provision of processing guarantees. Yet it requires that the native scheduler is enhanced.

Meta-Scheduler as Process Manager

The normal priority driven system scheduler is used to schedule all tasks. A meta-scheduler is employed to assign priorities to real-time tasks. I.e., this meta-scheduler considers only tasks with real-time requirements. Non time-critical tasks are processed when no time-critical task is ready for execution. In an integrated system the process management of continuous data processes will not be realized as a meta-scheduler it rather will be part of the system process manager itself. This meta-scheduler approach is also applied in many UNIX systems.

4.2 Real-time Processing Requirements

Continuous-media data processing has to occur in exactly predetermined - usually periodic intervals. Operations on this data recurs over and over and have to be completed at certain deadlines. The real-time process manager determines a schedule for the resource CPU that allows to make reservations and to give processing guarantees. The problem is to find a feasible schedule which schedules all time-critical continuous-media tasks in a way that each of them can meet its deadlines. This must be guaranteed for all tasks in every period for the whole run time of the system. In a multimedia system, continuous and discrete media data are processed concurrently.

For scheduling of multimedia tasks two conflicting goals have to be considered:

- An uncritical process should not suffer from starvation because time-critical processes are executed. Multimedia applications rely as much on text and graphics as on audio and video. Therefore not all resources should be occupied by the time-critical processes and their management processes.
- On the other hand a time-critical process must never be subject to priority inversion. The scheduler has to ensure that any priority inversion (also between time-critical processes with different priorities) is avoided or reduced as far as possible.

Apart of the overhead caused by the schedulability test and the connection establishment, the costs for the scheduling of every message has to be minimized. They are more critical because they occur periodically with every message during the processing of real-time tasks. The overhead generated by the scheduling and the operating system is part of the processing time and therefore has to be added to the processing time of the real-time tasks. Thus, it is favorable to keep them low. It is particularly difficult to observe the timing behavior of the operating system and its influence on the scheduling and the processing of time-critical data. It can lead to time garbling of application programs. Therefore, operating systems in real-time systems can not be viewed detached from the application programs and vice-a-versa.

4.3 Traditional Real-Time Scheduling

The problem of real-time processing is widely known in computer science [HaSh89, Oper89, ShGo90, TiKo91]. Some real-time scheduling methods are employed in operations research. They differ from computer science real-time scheduling because they operate in a static environment where no adaptation to changes of the workload is necessary [WoCr87].

The goal of traditional scheduling on time sharing computers is optimal throughput, optimal resource utilization, and fair queueing. In contrast, the main goal of real-time tasks is to provide a schedule that allows all, respectively as many time-critical processes as possible to be processed in time, according to their deadline. The scheduling algorithm has to map tasks onto resources such that all tasks meet their time requirements. Therefore, it must be possible to show, or to proof, that a scheduling algorithm applied to real-time systems fulfills the timing requirements of the task.

There are several attempts to solve real-time scheduling problems. Many of them are just variations of basic algorithms. In order to find the best solutions for multimedia systems, first two basic algorithms are analyzed "Earliest Deadline First Algorithm" and "Rate Monotonic Scheduling", and their advantages and disadvantages are elaborated. In the next section a system model is introduced, the relevant expressions are explained in their context.

4.4 Real-Time Scheduling: System Model

All scheduling algorithms to be introduced are based on the following system model for the scheduling of real-time tasks. Their essential components are the resources (as discussed previously), the tasks and the scheduling goals.

A task is a schedulable entity of the system, and it corresponds to the notion of a thread in the previous description. In a hard real-time system, a task is characterized by its timing constraints as well as by its resource requirements. In the considered case, only periodic tasks without precedence constraints are discussed, i.e. the processing of two tasks is mutually inde-

pendent. For multimedia systems this can be assumed without any major restriction. Synchronized data, for example, can be processed by a single process.



Figure 3: Characterization of periodic tasks

The time constraints of the periodic task T are characterized by the following parameters (s, e, d, p) as described in [LeMe80]:

- s: Starting point
- e: Processing time of T
- d: Deadline of T
- p: Period of T
- *r*: Rate of T (r = 1/p)

Whereby $0 \le e \le d \le p$ (see Figure 3). The starting point s is the first time where the periodic task requires processing. Afterwards, it requires processing in every period with a processing time of e. At s+(k-1)p the task T is ready for the k-processing. The processing of T in period k has to be finished at s+(k-1)p+d. For continuous-media tasks it is assumed that the deadline of the period (k-1) is the ready time of period k, this is known as congestion avoiding deadlines: The deadline for each message (d) coincides with the period of the respective periodic task (p).

Tasks can be preemptive or non-preemptive. A preemptive task can be interrupted by the request of any task with a higher priority. Processing is continued in the same state later on. A **non-preemptive** task can not be interrupted until it volunterily yields the processor. Any high-priority task has to wait until the low priority task is finished. The high-priority task is then subject to priority inversion. In the following all tasks processed on the CPU are considered as preemptive unless otherwise stated.

In a real-time system the scheduling algorithm must determine a schedule for an exclusive, limited resource that is used by different processes concurrently such that all of them can be processed without violating any deadlines. This notion can be extended to a model with multiple resources (e.g. CPU) of the same type. It can also be extended to cover different resources such as memory and bandwidth for communication, i.e., the function of a scheduling algorithm is to determine, for a given task set, whether or not a schedule for executing the tasks on an exclusive bounded resource exists, such that the timing and the resource constraints of all tasks are satisfied (*planing goal*). Further, it has to calculate a schedule if one exists. A scheduling algorithm is said to guarantee a newly arrived task if the algorithm can find a schedule where the new task and all previously guaranteed tasks can finish processing in every period over the whole run-time to their deadlines. If a scheduling algorithm guarantees a task, it ensures that

the task finishes processing prior to its deadline [CSRa88]. To guarantee tasks it must be possible to check the schedulability of the newly arrived tasks.

A major performance metric for a real-time scheduling algorithm is the guarantee ratio. The guarantee ratio is the total number of guaranteed tasks versus the number of tasks which could be processed.

Another performance metric is the *processor utilization*. This is the amount of processing time used by guaranteed tasks versus the total amount of processing time [LiLa73]:

$$U = \sum_{i=1}^{n} \frac{e_i}{p_i}$$

4.5 Earliest Deadline First Algorithm

The Earliest Deadline First (EDF) algorithm is one of the best known algorithms for real-time processing. At every new ready state, the scheduler selects among the tasks that are ready and not fully processed that one with the earliest deadline. The requested resource is assigned to the selected task. At any arrival of a new task (according to the LBAP model), EDF must be computed immediately leading to a new order - i.e. the running task is preempted and the new task is scheduled according to its deadline. The new task is processed immediately if its deadline is earlier than the one of the interrupted task. The processing of the interrupted task is continued according to the EDF algorithm later on. EDF is not only an algorithm for periodic tasks but also for tasks with arbitrary requests, deadlines and service execution times [Dert74]. In this case, no guarantee about the processing of any task can be given.

EDF is an optimal, dynamic algorithm: Le., it produces a valid schedule whenever one exist. A dynamic algorithm schedules every instance of each incoming task according to its specific demands. Tasks of periodic processes have to be scheduled in each period again. With n tasks, which have arbitrary ready-times and deadlines, the complexity is $\Theta(n^2)$.

For a dynamic algorithm like EDF, the upper bound of the processor utilization is 100%. Compared with any static priority assignment EDF is optimal in a sense that if a set of tasks that can be scheduled by any static priority assignment it also can be scheduled by EDF. With a priority driven system scheduler, each task is assigned a priority according to its deadline. The highest priority is assigned to the task with the earliest deadline, the lowest to the one with the furthest. With every arriving task, priorities might have to be adjusted.

Applying EDF to the scheduling of continuous-media data on a single processor machine with priority scheduling, process priorities are likely to be rearranged quit often. A priority is assigned to each task ready for processing according to its deadline. Common systems usually provide only a restricted number of priorities. If the computed priority of a new process is not available, the priorities of other processes have to be rearranged until the required priority is free. In the worst case, the priorities of all processes have to be rearranged. This may cause a considerable overhead. The EDF scheduling algorithm itself makes no use of the previously known occurrence of periodic tasks.

EDF is used by different models as basic algorithm. An extension of EDF is the *time-driven* scheduler (TDS). Tasks are scheduled according to their deadline. Further, the TDS is able to handle overload situations. If an overload situation occurs the scheduler aborts tasks which can not meet their deadlines any more. If there still is an overload situation the scheduler removes tasks which have a low "value density". The value density corresponds to the importance of a task for the system.

In [LLSY91 another priority driven EDF scheduling algorithm is introduced. Here every task is divided in to a *mandatory* and an *optional* part. A task is terminated according to the deadline of the mandatory part even if it is not completed at this time. Tasks are scheduled with respect to the deadline of the mandatory parts. A set of task is said to be schedulable if all tasks can meet the deadlines of their mandatory parts. The optional parts are processed if the resource capacity is not fully utilized. Applying this to continuous-media data the method can be used in combination with the encoding of data according to their importance. Take, e.g., a single uncompressed picture in a bitmap format. Each pixel of this monochrome picture is coded with 16 bit. The processing of the 8 most significant bits is mandatory whereas the processes can be scheduled. In an overload situation the optional parts are aborted. This implementation leads to a decreased quality by media scaling. During QoS requirement specification the tasks accepted or were informed that scaling may occur. In such a case scaling QoS parameter(s) can be introduced which reflects the respective implementation. Therefore this approach avoids errors and improves system performance at the expense of media quality.

4.6 Rate Monotonic Algorithm

The Rate Monotonic scheduling principle was introduced by Liu and Layland in 1973 [LiLa73]. It is an optimal, static, priority-driven algorithm for preemptive, periodic jobs. Optimal in this context means that there is no other *static* algorithm that is able to schedule a task set which can not be scheduled by the rate monotonic algorithm. A process is scheduled by a static algorithm at the beginning of the processing. Subsequently, each task is processed with the priority calculated at the beginning. No further scheduling is required. The following five assumptions are necessary prerequisites to apply the rate-monotonic algorithm:

- 1. The requests for all tasks with deadlines are periodic. Le. with constant intervals between consecutive requests.
- 2. The processing of a single task has to be finished before the next task of the same data stream becomes ready for execution. Deadlines consist of run-ability constrains only. Le. each task must be completed before the next request occurs.
- 3. All tasks are independent, Le., the requests for a certain task do not depend on the initiation or completion of requests for any other task.
- 4. Run-time for each request of a task is constant. Run-time denotes the maximum time which is required by a processor to execute the task without interruption.
- 5. Any non-periodic task in the system has no required deadline. Typically they initiate periodic task or are tasks for failure recovery. They usually displace periodic tasks.

Further work has shown that not all of these assumptions are mandatory to employ the ratemonotonic algorithm [LSST91, SKG091].

Static priorities are assigned to tasks once at connection set up phase according to their request rates. The priority corresponds to the importance of a task relatively to other tasks. Tasks with higher request rates will have higher priorities. The task with the shortest period gets the highest priority and the task with the longest period the lowest priority.

The rate-monotonic algorithm is a simple method to schedule time critical, periodic tasks on the respective resource. A task will always meet its deadline if this can be proven to be true for the longest *response time*. The response time is the time span between the request and the end



Figure 4: An Example of critical instants

of processing of a task. This time span is maximal when all processes with a higher priority request to be processed at the same time. This case knows as *critical instant* (see Figure 4). In this figure the priority of a is, according to the rate-monotonic algorithm, higher than the one of b, and the one of b is higher than the one of c. The *critical time zone* is the time interval between the critical instant and the completion of a task.

4.7 EDF and Rate Monotonic: Context switches

Consider an audio and a video stream scheduled according to the rate monotonic algorithm. Let the audio stream have a rate of 1/75 s/sample and the video stream a rate of 1/25 s/frame. The priority assigned to the audio stream is then higher than the priority assigned to the video stream. The arrival of messages from the audio stream will interrupt the processing of the video stream. If it is possible to complete the processing of a video message that requests processing at the critical instant before its deadline, the processing of all video messages to their deadlines is ensured, thus a feasible schedule exists.



Figure 5: Rate monotonic versus EDF: Context switches in preemptive systems

If more than one stream is processed concurrently in a system it is very likely that there might be more context switches with a scheduler using the rate monotonic algorithm than one using EDF. Figure 5 shows an example.

4.8 EDF and Rate Monotonic: Processor Utilizations

The processor utilization of the rate monotonic algorithm is upper bounded. It depends on the number of tasks which are scheduled, their processing times, and their periods. There are two issues to be considered:

- 1. The upper bound of the processor utilization which is determined by the critical instant.
- For each number n of independent tasks t(j) a constellation can be found where the maximum possible processor utilization is minimal. The least upper bound of the processor utilization is the minimum of all processor utilizations over all sets of tasks t(j); j ∈ {1, ..., n}. that fully utilize the CPU. A task set fully utilizes the CPU when it is not possible to raise the processing time of one task without violating the schedule.

Following this assumptions, [LiLa73] gives an estimation of the maximal processor utilization where the processing of each task to its deadline is *guaranteed* for any constellation. A set of *m* independent, periodic tasks with fixed priority order will always meet its deadline if:

$$U(m) = m \times \left(2^{\frac{1}{m}}-1\right) \ge \frac{e_1}{p_1} + \dots + \frac{e_m}{p_m}$$

According to [LeSh86] and [LiLa73] for large m the least upper bound of the processor utilization is U=ln2. Hence it is sufficient to check if the processor utilization is less or equal to the given upper bound to find out if a task set is schedulable or not. Most existing systems check this by simply comparing the processor utilization with the value of ln 2.



Figure 6: Rate monotonic versus EDF: Processor utilization

With EDF, a processor utilization of 100% can be achieved because all task are scheduled dynamically according to their deadlines. Figure 6 shows an example where the CPU can be utilized to 100% with EDF but where rate monotonic scheduling fails.

.

-

.

.

The problem of underutilizing the processor is aggregated by the fact that in most cases the average task execution time is considerably lower than the worst case execution time. Therefore scheduling algorithms should be able to handle transient processor overload. The rate monotonic algorithm on average ensures that all deadlines will be met even if the bottleneck utilization is well above 80%. With one deadline postponement, the deadlines on average are met when the utilization is over 90%. [SSLe89] mentions an achieved utilization bound for the Nowy's Insertial Navigation System of 88%.

As described above, a static algorithm schedules a process once at the beginning of the processing, single tasks are not explicitly scheduled afterwards. Whereas a dynamic algorithm schedules every incoming task according to its specific demands. Since the rate monotonic algorithm is an optimal static algorithm no other static algorithm can achieve a higher processor utilization.

4.9 Extensions to the Rate Monotonic Scheduling

There are several extensions to this algorithm. One of them divides a task in to a mandatory and an optional part. The processing of the mandatory part delivers a result which can be accepted by the user. The optional part only refines the result. The mandatory part is scheduled according to the rate monotonic algorithm. For the scheduling of the optional part, other, different policies are suggested [ChLi88, LLNa87, ChLi89].

In some systems there are aperiodic tasks next to periodic ones. To meet the requirements of periodic tasks and the response time requirements of aperiodic requests, it must be possible to schedule both, aperiodic and periodic tasks. If the aperiodic request is an aperiodic continuous stream (e.g. video images as part of a slide show), we have the possibility to transform it into a periodic stream. Every timed data item can be substituted by n items. The new items have the duration of the minimal life span. The number of streams is increased but since the life span is decreased the semantic remains unchanged. The stream is now periodical because every item has the same life span [Herr90]. If the stream is not continuous, we can apply a sporadic server to respond to aperiodic requests. The server is provided with a computation budget. This budget is refreshed t units of time after it has been exhausted. Earlier refreshing is also possible. The budget represents the computation time reserved for aperiodic tasks. The server is only allowed to preempt the execution of periodic tasks as long as the computation budget is not exhausted. Afterwards it can only continue the execution with a background priority. After refreshing the budget, the execution can resume at the servers assigned priority. The sporadic server is especially suitable for events that occur rarely but must be handled quickly (e.g. a telepointer in a CSCW-application) [ShGo90, SSLe89, Spru90].

The rate monotonic algorithm is, for example, applied in real-time systems and real-time operating systems by the NASA and the European Space Agency [ShRa89]. It is particularly suitable for continuous-media data processing because it makes optimal use of their periodicity. Since it is a static algorithm there is nearly no rearrangement of priorities and hence -in contrast to EDF- no scheduling overhead to determine the next task with the highest priority. Problems emerge with data streams which have no constant processing time per message as specified in MPEG-2 (e.g. a compressed video stream where one of five pictures is a full picture and all others are up-dates of a reference picture). The simplest solution is to schedule this tasks according to their maximum data rate. In this case the processor utilization is decreasing. The idle time of the CPU can be used to process non-time-critical tasks. In multimedia systems, for example, this is the processing of discrete media.

4.10 Other Approaches for In-Time Scheduling

Apart of the two in detail discussed methods, further scheduling algorithms have been evaluated towards their suitability for the processing of continuous-media data. In the following paragraphs the most significant approaches are briefly described and the reasons for their "non suitability, compared to EDF and rate-monotonic, towards their employment in multimedia systems are enumerated.

Least Laxity First (LLF): The task with the shortest remaining laxity is scheduled first [CrWo90, LeSh86]. The laxity is the time between the actual time t and the deadline minus the remaining processing time. The laxity in period k is:

$$I_{k} = (s + (k-1)p + d) - (t+e)$$

LLF is an optimal, dynamic algorithm for exclusive resources. Furthermore it is also an optimal algorithm for multiple resources if the ready-times of the real-time tasks are the same. The laxity is a function of deadline, processing-time and the current time. Thereby the processingtime can not be exactly specified in advance. At the calculation of the laxity the worst-case is assumed. Therefore the determination of the laxity is inexact. The laxity of waiting processes is dynamically changing over time. During run-time of a task, another task may get a lower laxity. This task has then to preempt the running task. Consequently, tasks can preempt each other several times without dispatching a new task. This may cause numerous context switches. At each scheduling point (when a process becomes ready to run or at the end of a time slice) the laxity of each task has to be newly determined. This leads to an additional overhead compared with EDF. Since we have only a single resource to schedule there is no advantage in the employment of LLF compared with EDF. Future multimedia systems might be multi-processor systems; here LLF might be of advantage.

Deadline Monotone Algorithm: If the deadline of tasks are less than their period $(d_i < p_i)$, the prerequisites of the rate monotonic algorithm are violated. In this case, a fixed priority assignment according to the deadlines of the tasks is optimal.

A task T_i gets a higher priority than a task T_j if $d_i < d_j$. No effective schedulability test for the deadline monotone algorithm exists. In order to determine the schedulability of a task set, it has to be checked for each task if it meets its deadline in the worst case. In this case, all tasks require execution to their critical instant [LeWh82, LSST91]. Tasks with a deadline shorter than their period, for example, arise during the measurements of temperature or pressure in control systems. In multimedia systems deadlines equal to period lengths can be assumed.

Shortest Job First (SJF): The task with the shortest remaining computation time is chosen for execution [CrWo90, Fren82]. This algorithm guarantees that as many tasks as possible meet their deadlines under an overload situation if all of them have the same deadline. In multimedia systems, where the resource management allows overload situations, this might be a suitable algorithm.

Apart of the most important real-time scheduling methods discussed above other might be employed for the processing of continuous-media data (an on-line scheduler for tasks with unknown ready times is introduced by [HoLe88]; in [HaSh89] a real-time monitoring system is presented where all necessary data to calculate an optimal schedule is available). In most multimedia systems with preemptive tasks, the rate-monotonic algorithm in different variations is employed. So far, no other scheduling technique has been proven to be at least as suitable for multimedia data handling as the EDF and rate monotonic approaches.

4.11 Preemptive versus Non-Preemptive Task Scheduling

Real-time tasks can be distinguished into preemptive and non-preemptive tasks. If a task is non-preemptive, it is processed and not interrupted until it is finished or requires further resources. The contrary of non preemptive tasks are preemptive tasks. The processing of a preemptive task is interrupted immediately by a request of any higher priority task. In most cases, where tasks are treated as non-preemptive, the arrival times, processing times, and deadlines are arbitrary and unknown to the scheduler until the task actually arrives. The best algorithm is the one which maximizes the number of completed tasks. In this case it is not possible to provide any processing guarantees or to do resource management.



To guarantee the processing of periodic processes and to get a feasible schedule for a periodic

Figure 7: Preemptive versus non-preemptive scheduling

task set, tasks are usually treated as preemptive. One reason is that high preemptability minimizes priority inversion. Another reason is that for some non-preemptive task sets no feasible schedule can be found, whereas for preemptive scheduling it is possible. Figure 7 shows an example where the scheduling of preemptive tasks is possible but non-preemptive tasks can not be scheduled.

In [LiLa73] Liu and Layland show that a task set of *m* periodic, preemptive tasks with processing times e_i and request periods $p_i \forall i \in (1, ..., m)$ is schedulable

• with fixed priority assignment if

$$\sum \frac{e_i}{P_i} \le \ln 2$$

• and for deadline driven scheduling if

∑
$$\frac{e_i}{p_i}$$
≤1

Here the preemptiveness of tasks is a necessary prerequisite to check their schedulability.

The first schedulability test for the scheduling of *non-preemptive* tasks was introduced by Nagarajan and Vogt in [NaV092]. Assume, without loss of generality, that task *m* has highest priority and task 1 lowest. They proof that a set of *m* periodic streams with periods p_i , dead-lines d_i , processing times e_i and $d_i \leq p_i \forall i \in (1, ..., m)$ is schedulable with the non-preemptive fixed priority scheme if

$$d_{m} \geq e_{m} + max_{(1 \leq i \leq m)}e_{i}$$

$$d_{j} \geq e_{j} + max_{(1 \leq j \leq m)}e_{j} + \sum_{j=i+1}^{m}e_{j}F(d_{i}-e_{j}, P_{j})$$
where $F(x, y) = ceil(\frac{x}{y}) + 1$.

This means that the time between the logical arrival time and the deadline of a task t_i has to be larger, or equal to the sum of the own processing time e_i and the processing time of any higher priority task that requires execution during that time interval plus the longest processing time of all lower and higher priority tasks $(\max_{1 \le j \le m} e_j)$ that might be serviced at the arrival of the task t_i .



Figure 8: Deadline requirements for non-preemptive scheduling

The schedulability test is an extension of Liu's and Layland's. Given m periodic tasks with periods p_i and the same processing time E per message. Let $d_i = p_i + E$ be the deadline for task t_i . Then the streams are schedulable

with the non-preemptive rate monotonic scheme with:

$$\sum \frac{1}{\rho_j} \times E \le \ln 2$$

• with deadline-based scheduling, the same holds with:

$$\sum \frac{1}{\rho_j} \times E \le 1$$

Consequently, non-preemptive continuous-media tasks can also be scheduled. However, the scheduling of non-preemptive tasks is less favorable because the number of schedulable task sets is smaller compared to preemptive tasks.

4.12 Scheduling of Continuous Media Tasks: Prototype Operating Systems

Most multimedia operating systems apply one of the above discussed methods. In some systems the scheduler is replaced by a real-time scheduler. Therefore, these systems can be viewed as new operating systems. They are usually not compatible to existing systems and applications. Other systems apply a *meta-scheduler* based on an existing process manager. Only these systems will have a commercial impact in short and medium term because they allow to run already existing applications.

ARTS

The "Advanced Real-Time Technology Operating System" is a real-time operating system for a distributed environment with one real-time process manager. It was developed on SUN3 work-stations, connected with a real-time network based on the IEEE.802.5 Token Ring and Ethernet, by the computer science department of the Carnegie Mellon University. To solve the scheduling problems the *time-driven scheduler* (TDS) with a priority inheritance protocol was adopted. This priority inheritance protocol was used to prevent unbounded priority inversion among communication tasks. Tasks with hard deadlines are scheduled according to the rate-monotonic algorithm. The system is also provided with other scheduling methods for experimental reasons [MeTo90].

YARTOS

"Yet Another Real-Time Operating System" was developed at the University of North Carolina at Chapel Hill as an operating system kernel to support conferencing applications [JSPo91]. An optimal, preemptive algorithm to schedule tasks on a single processor was developed. The scheduling algorithm results from the integration of a synchronization scheme to access shared resources with the EDF algorithm. Here, a task has two notions of deadline, one for the initial acquisition of the processor, and one for execution of operations on resources. To avoid priority inversion, tasks are provided with separate deadlines for performing operations on shared resources. It is guaranteed that no shared resource is accessed simultaneously by more than one task. Further, a shared resource is not longer occupied by a single task than absolutely necessary.

Split-Level Scheduling

The split level scheduler was developed within the DASH-project at the University of California at Berkeley. Its main goal is to provide a better support for multimedia applications [Ande93]. The applied scheduling policy is deadline/workahead scheduling. The LBAP-model is used to describe the arrival processes. Critical processes have priority over all other processes and they are scheduled according to the EDF algorithm preemptively. Interactive processes have priority over workahead processes as long as they do not become critical. The scheduling policy for workahead processes is unspecified but may be chosen to minimize context switching. For non real-time processes, a scheduling strategy like UNIX time-slicing is chosen.

Three Class Scheduler

This scheduler was developed as part of a video-on-demand file servicer at DEC, Littleton. The design of the scheduler is based on a combination of weighted round-robin and rate monotonic

scheduling [RVG*93]. Three classes of schedulable tasks are supported. The isochronous class with the highest priority applies the rate-monotonic algorithm, the real-time and the generalpurpose class are using the weighted round-robin scheme. A general-purpose task is preemptive and runs with a low priority. The real-time class is suitable for tasks that require guaranteed throughput and bounded delay. The isochronous class supports real-time periodic tasks that require performance guarantees for throughput, bounded latency and low jitter. Real-time and isochronous tasks can only be preempted in "preemption windows". The scheduler executes tasks from a ready queue in which all isochronous tasks are arranged according to their priority. At the arrival of the task the scheduler determines whether the currently running task has to be preempted. General-purpose tasks are immediately preempted, real-time tasks are preempted in the next preemption window, isochronous tasks are preempted in the next preemption window if their priority is lower than the one of the new task. Whenever the queue is empty the scheduler alternates between the real-time and general purpose classes using a weighted round-robin scheme.

Meta-Scheduler

To support real-time processing of continuous-media meta-scheduler for the operating systems AIX [WBV094] and OS/2 [MSSt92] was developed at the European Networking Center of IBM in Heidelberg. Both are based on the LBAP-model. According to the rate-monotonic algorithm rates are mapped onto system priorities.

Experience with the Meta-Scheduler Approach

In this paragraph the employment of the OS/2 meta-scheduler is discussed [MSSt92]. The experiences show the limits of this approach. For example, each single process in the system is able to run with a priority initially intended for real-time tasks. These processes are not scheduled by the resource manager and therefore violate the calculated schedule. A malicious process can block the whole system simply by running with the highest priority without giving up control.

The management of scheduling algorithms requires exact time measurement. In OS/2, for example, it is not possible to measure the exact time a thread is using the CPU. Any measurement of the processing time includes interrupts. If a process is interrupted by another process it also includes the time needed for the context switch. The granularity of the OS/2 system timers are insufficient for the processing of real-time tasks. Hence the rate control is inaccurate because it is determined by the granularity of the system timer.

To achieve full real-time capabilities at least the native scheduler of the operating system would have to be extended. The operating system should be enhanced by a class of fast nonpreemptive threads and the ability to mask interrupts for a short period of time. Priorities in this thread class should only be assigned to threads that are already registered by the resource manager. This class should be reserved exclusively for selected threads and monitored by a system component with extensive control facilities. Performance enhancement of the scheduler itself incorporating some mechanisms of real-time scheduling like earliest deadline first would be another solution. The operating system should in any case provide a time measurement tool that allows the measurement of pure CPU-time and a timer with a finer granularity. This may be achieved through a timer chip.

5. File Systems

The file system is said to be the most visible part of an operating system. Most programs write or read on files, their program code as well as user data stored in files. The organization of the

file system is an important factor for the usability and convenience of the operating system. A file is a sequence of information held as unit for storage and use in computer systems [Krak88]. Files are stored in secondary storage, they can be used by different applications. The lifespan of files is usually longer than the execution of a program. In traditional file systems, the information types stored in files are sources, objects, libraries and executables of programs, numeric data, text, payroll records etc. [PeSi83]. In multimedia systems, the stored information also covers digitized video and audio with their related real-time "read" and "write" demands. Therefore, additional requirements in the design and implementation of file systems have to be considered.

The file system provides access and control functions for the storage and retrieval of files. From the user's viewpoint, it is important how the filesystem allows to organize and structure files. The internals, which are more important in our context, i.e., the organization of the file system deal with the representation of information in files, their structure and organization in secondary storage. Because of its importance for multimedia, disk scheduling is also presented in this context.

The next section starts with a brief characterization of traditional file systems and disk scheduling algorithms. Subsequently, different approaches to organize multimedia files and disk scheduling algorithms for the use in multimedia systems are discussed.

5.1 Traditional File Systems

The two main goals of traditional files systems are: (1) to provide a comfortable interface for file access to the user, and (2) to make efficient use of storage media. Whereas the first goal is still an area of interest for research (e.g. indexing for file systems [Salz91] and intelligent file systems for the content-based associative access to file system data [GiT091]), the structure, organization and access of data stored on disk, have been extensively discussed and investigated over the last decades. In order to understand the specific multimedia developments in this area this paragraph gives a brief overview on files, file systems organizations and file access mechanisms. Later, disk scheduling algorithms for file retrieval are discussed.

. ...

File Structure:

We commonly distinguish between two ways of file organization:

In sequential storage, each file is organized as a simple sequence of bytes or records. Files are stored consecutively on the secondary storage media as shown in Figure 9. They are separated from each other by a well defined "end of file" bit pattern, character, or character sequence. A file descriptor is usually placed at the beginning of the file and is, in some systems, repeated at the end of the file. Sequential storage is the only possible way to organize the storage on tape, but it can also be used on disks. The main advantage is its efficiency for sequential access as well as for direct access [Krak88]. Disk access time for reading and writing is minimized.

Additionally, for further improvement of performance with caching the file can be read ahead of the user program [Jans85]. In systems where file creation, deletion, and size modification occur frequently, sequential storage has major disadvantages. Secondary storage is splitted, fragmented, through creation and deletion operations, and files can not be extended without copying the whole files into a larger space. By a process the files may be copied such that all files are located adjecently, i.e., without any "holes" between them.

In non-sequential storage the data items are stored in a non-contiguous order. There exist mainly two approaches:



Figure 9: Contiguous and non-contiguous storage

- One way is to use linked blocks, where physical blocks, containing consecutive logical locations, are linked using pointers. The file descriptor must contain the number of blocks occupied by the file, the pointer to the first block, and it may also have the pointer to the last block. A serious disadvantage of this method is the cost of the implementation for random access because all prior data has to be read. In MS-DOS a similar method is applied. A file allocation table (FAT) is associated with each disk. One entry in the table represents one disk block. The directory entry of each file holds the block number of the first block. The number in the slot of an entry refers to the next block of a file. The slot of the last block of a file contains an end-of-file mark [Tane87].
- Another approach is to store block information in mapping tables. Each file is associated with a table where, apart from the block numbers, information like owner, file size, creation time, last access time etc. are stored. Those tables usually have a fixed size which means that the number of block references is bounded. Files with more blocks are referenced indirectly by additional tables assigned to the files. In UNIX, with each file a small table (on disk) called an i-node is associated (see Figure 10). The indexed sequential approach is an example for multi-level mapping; here logical and physical organization is not clearly separated [Krak88].



Figure 10: The UNIX i-node [Tane87]

Directory Structure:

Files are usually organized in directories. Most of today's operating systems provide tree structured directories where the user can organize the files according to his personal needs. In multimedia systems, it is important to organize the files in a way that allows an easy, fast, and contiguous data access

Disk Management:

Disk access is a slow and costly transaction. In traditional systems, a common technique to reduce disk access are block caches. Using a block cache, blocks are kept in memory because it is expected that future read or write operations access this data again, thus, performance increased due to shorter access time. Another way to increase performance is to reduce disk arm motion. Blocks, that are likely to be accessed in sequence, are placed together on one cylinder. To refine this method, the rotational positioning can be taken into account. Consecutive blocks are placed on the same cylinder but in an interleaved way as shown in Figure 11. Another important issue is the placement of the mapping tables (e.g. I-nodes in UNIX) on the disk. If they are placed near the beginning of the disk then the distance between them and the blocks will be, on average, half the number of cylinders. To improve this, they can be placed in the middle of the disk. Hence the average seek time is roughly reduced by a factor of two. In the same way, consecutive blocks should be placed on the same cylinder. The use of the same cylinder for the storage of the mapping tables and the referred blocks also improves the performance.



Figure 11: Interleaved and non-interleaved storage

Disk Scheduling:

Whereas strictly sequential storage devices (e.g. tapes) do not have a scheduling problem, for random access storage devices every file operation may require movements of the read/write head. This operation, known as 'to seek' is very time consuming, i.e., a seek time in the order of 250 ms for CD's is still state of the art. The actual time to read or write a disk block is determined by

- the seek time (the time required for the movement of the read/write head).
- the latency time or rotational delay (the time during which the transfer can not be proceed until the right block or sector rotates under the read/write head),
- the actual data transfer time needed for the data copy from disk into main memory,

Usually the seek time is the largest factor of the actual transfer time. Most systems try to keep the cost of seeking low by applying special algorithms for the scheduling of disk read/write operations. The access of the storage device is a problem greatly influenced by the file allocation method. For instance, a program, reading a contiguously allocated file, generates requests which are located close together on a disk, thus head movement is limited. Linked or indexed files with blocks, which are widely scattered, cause many head movements. In multiprogramming systems, where the disk queue may often be non-empty, fairness is also a criterion for scheduling. Most systems apply one of the following scheduling algorithms:



- First-Come-First-Served (FCFS): With this algorithm the disk driver accepts requests one at the time, and it serves them in the incoming order. This is an easy to program and an intrinsically fair algorithm. However, it is not optimal with respect to head movements because it does not consider the location of the other queued requests. This results in a high average seek time. Figure 12 shows an example of the application of FCFS to a request of 3 queued blocks.
- Shortest-Seek-Time First (SSTF): At every point in time when a data transfer is requested, SSTF selects among all requests that one with the minimum seek time from the current head position, therefore the head is moved to the closest track in the request queue. This algorithm was developed to minimize seek time and it is in this sense optimal. SSTF is a modification of Shortest Job First (SJF), and like SJF, it may cause starvation of some requests. Request targets in the middle of the disk will get immediate service on the expense of requests in the innermost and outermost disk areas. Figure 13 demonstrates the operation of the SSTF algorithm.
- SCAN: Like SSTF, scan orders the requests to minimize seek time. In contrast to SSTF, it takes the direction of the current disk movement into account. It first serves all requests in one direction until it does not have any request in this direction anymore. The head movement is then reversed and service is continued. SCAN provides a very good seek time, because the edge tracks get better service times. Note, that middle tracks still get a better service then edge tracks. When the head movement is reversed it first serves tracks that have recently been serviced, where the heaviest density of requests, assuming a uni-



Figure 13: SSTF disk scheduling



Figure 14: SCAN disk scheduling

form distribution, is at the other end of the disk. Figure 14 shows an example of the SCAN algorithm.

C-SCAN: C-SCAN also moves the head in one direction, but it offers fairer service with
more uniform waiting times. It does not alter the direction as in SCAN, instead it scans in
cycles, always increasing or decreasing, with one idle head movement from one edge to
the other between two consecutive scans. The performance of C-SCAN is somewhat less
than the one of SCAN. Figure 15 shows the operation of the C-SCAN algorithm.

Traditional file systems are not designed for their employment in multimedia systems. They do, for example, not consider requirements like real-time which are important to the retrieval of stored audio and video. To serve this requirements, new polices in the structure and organization of files, and in the retrieval of data from the disk have to be



Figure 15: C-SCAN disk scheduling

applied. The next section outlines the most important developments in this area.

5.2 Multimedia File Systems

Compared to the increased performance of processors and networks, storage devices have become only marginally "faster" [Mull91]. The effect of this increasing speed mismatch is the search for new storage structures, and storage and retrieval mechanisms with respect to the file system. Continuous-media data is different to discrete data in:

- **Real-Time Characteristics:** As mentioned above, the retrieval, computation, and presentation of continuous media is time dependent. The data has to be presented (read) before a well defined deadline with small jitter only. Thus, algorithms for the storage and retrieval of such data have to consider time constraints, and additional buffers to smooth the data stream, have to be provided.
- File Size: Compared to text and graphics, video and audio have very large storage space requirements. Since the file system has to store information ranging from small unstructured units like text files to large, highly structured data units like video and associated audio. It has to organize the data on disk in a way that efficiently uses the limited storage. For example, the storage requirements of uncompressed CD quality stereo audio are 1.4 Mbits/s, low but acceptable quality compressed video still requires about 1Mbit/s using, e.g., MPEG-1.
- *Multiple Data Streams:* A multimedia system has to support different media at a time. It does not only have to ensure that all of them get a sufficient share on the resources, it also has to consider tight relations between different streams arriving from different sources. The retrieval of a movie, for example, requires the processing and synchronization of audio and video.

There are different ways to support continuous media in file systems. Basically it can be distinguished between two approaches. With the first approach, the organization of files on disk remains as it is. The necessary real-time support is provided through special disk scheduling algorithms and sufficient buffer to avoid jitter. In the second approach, the organization of audio and video files on disk is optimized for their use in multimedia systems. Especially scheduling of multiple data streams still remains an issue of research.

In this section, the different approaches are discussed and examples of existing prototypes are introduced. First, a brief introduction of the different storage devices employed in multimedia systems is given. Then the organization of files on disks is discussed. Subsequently different disk scheduling algorithms for the retrieval of continuous media are introduced.

5.2.1 Storage Devices

The storage subsystem is a major component of any information system. Due to the immense storage space requirements of continuous media, conventional magnetic storage devices are often not sufficient anymore. Tapes, still in use in some traditional systems, are inadequate for multimedia systems because they can not provide independent accessible streams, and random access is slow and expensive.

Apart from common disks with large capacity, some multimedia applications, such as kiosk systems, use CD-ROMs to store data. In general, disks can be characterized in two different ways.

- First, how information is stored on them. There are rewriteable (magnetic and optical) disks, write-once (WORM) disks and read-only disks like CD-ROMs.
- The second distinctive feature is the way of recording. It is distinguished between magnetic and optical disks. The main difference between them is the access time and their track capacity. The seek time on magnetic disks is typically above 10 ms whereas on optical disks 200 ms is a common lower bound. Magnetic disks have a constant rotation speed (constant angular velocity, CAV). Thus, while the densitz varies the storage capacity is the same on inner and outer tracks. Optical disks have varying rotation speed (constant linear velocity, CLV) and hence, the storage density is the same on the whole disk.

Therefore different algorithms for magnetic and optical disks are necessary. File systems on CD-ROMs are defined at ISO 9660, it is considered to be closely related to CD-ROMs and CD-ROM-XA. Very few variations are possible. Hence, we will focus the description on algorithms applicable to magnetic storage devices.

5.2.2 File Structure and Placement on Disk

Whereas in conventional file systems the main goal of the file organization is to make efficient use of the storage capacity (i.e. to reduce internal and external fragmentation) and to allow arbitrary deletion and extension of files, in multimedia systems the main goal is to provide a constant and timely retrieval of data. Internal fragmentation occurs when blocks of data are not entirely filled. On average the last block of a file is only half utilized. The use of large blocks leads to a larger waste of storage due to this internal fragmentation. External fragmentation mainly occurs when files are stored in a contiguous way. After the deletion of a file the gap can only be filled by a file with the same or a smaller size. Therefore there are usually small fractions between two files that are not used, storage space for continuous media is wasted.

As mentioned above, the goals for multimedia file systems can be achieved through providing enough buffer for each data stream and the employment of disk scheduling algorithms, especially optimized for real-time storage and retrieval of data. The advantage of this approach (where data blocks of single files are scattered) is flexibility. External fragmentation is avoided and the same data can be used by several streams (via references). Even using only one stream this might be of advantage, for instance, it is possible to access one block twice, e.g. when a phrase in a sonata is repeated. However, due to the large seek operations during playback, even with optimized disk scheduling, large buffers have to be provided to smooth jitter at the data retrieval phase. Therefore, there are also long initial delays at the retrieval of continuous media.

Another problem in this context is the restricted transfer rate. With upcoming disk arrays, which might have 100 and more parallel heads, the projected seek and latency times of less than 10 ms, and a block size of 4 Kbytes at a transfer rate of 0.32 Gigabit/s will be achieved. But this is, for example, not sufficient for the simultaneous retrieval of four or more production level MPEG-2 videos compressed in HDTV-quality that may require transfer rates of up to 100 Mbit/s. [Stei94a, Stei94b]

Approaches, which use specific disk layout, take the specialized nature of continuous-media data into account to minimize the cost of retrieving and storing streams. The much greater size of continuous-media files and the fact that they will usually be retrieved sequentially, because of the nature of operation performed on them (such as play, pause, fast forward etc.), are reasons for an optimization of the disk layout. The own application related experience has shown that continuous-media streams predominantly belong to the write-once-read-many nature and streams that are recorded at the same time are likely to be played back at the same time (e.g. audio and video of a movie) [LoSh93]. Hence, it seems to be reasonable to store continuous-media data in large data blocks contiguously on disk. Files that are likely to be retrieved together are grouped together on the disk. Thus, interference due to concurrent access of these files are minimized. With such a disk layout the buffer requirements and the seek times decrease.

The disadvantage of the contiguous approach is external fragmentation, and copying overhead during insertion and deletion. To avoid this without scattering blocks in a random manor over the disk, a multimedia file system can provide constrained block allocation of the continuous media. In [GeCh92] different placement strategies were compared. The size of the blocks (M) and the size of the gaps (G) between them can be derived from the requirement of continuity. The size is measured in terms of sectors. We assume that the data transfer rate r_{dt} is the same as the disk rotation rate (sectors/s). The continuity requirement in this case is met if the time to skip over a gap and to read the next media block does not exceed the duration of its playback $T_{play}(s)$ [RKVi92].

$$T_{play}(s) \ge \frac{M(\text{sectors}) + G(\text{sectors})}{r_{dt}(\text{sectors/s})}$$

Since there are two variables in the equation, the storage pattern (M, G) is not unique. There are several combinations possible to satisfy the above equation. Problems occur if the disk is not sufficiently empty, so that single data streams can not be stored exactly according to their storage pattern. In this case the continuity requirements for each block are not strictly maintained. To serve the continuity requirements, read-ahead and buffering of a determined number of blocks has to be introduced. See, e.g., [RaVi91, RKVi92, ViRa93] for a detailed description of this storage method.

Some systems using scattered storage make use of a special disk space allocation mechanism to allow fast and efficient access. Abbott performed the pioneer work in this field [Abbo84]. He was especially concerned about the size of single blocks and their positions on disk. Another topic to be considered is the placement of different streams. With interleaved place-

35

. Sugar
ment all n'th block of each stream are in close physical proximity on disk. A contiguous interleaved placement is possible as well as a scattered interleaved placement. With interleaved data streams synchronization is much easier to handle. On the other hand, the insertion and deletion of single parts of data streams become more complicated.

In [KWYe94] a layout algorithm was developed and analyzed which provides a uniform distribution of media blocks on the disk after copying or writing audio and video files. It takes into account that further files will be merged. Therefore a set of non-filled gaps is left. This uniform distribution is achieved by storing new blocks at the center of existing - so far - non-filled gaps. With this 'central merging method' gaps are successively split into two new equal gaps. It was shown that the mean efficiency of the secondary storage usage with this algorithm is about 75 % without violation of any real-time constraint [KWYe94].

5.2.3 Disk Scheduling Algorithms

The main goal of traditional disk scheduling algorithms is to reduce the cost of seek operations, to achieve a high throughput, and to provide fair disk access for every process. The additional real-time requirements, introduced by multimedia systems, make traditional disk scheduling algorithms, such as described before, inconvenient for multimedia systems. Systems without any optimized disk layout for the storage of continuous media depend far more on reliable and efficient disk scheduling algorithms than others. In case of contiguous storage, scheduling is only needed to serve requests from multiple streams concurrently. In [LoSh93], a round robin scheduler is employed that is able to serve hard real-time tasks. Here, additional optimization is provided through the close physical placement of streams that are likely to be accessed together.

The overall goal of disk scheduling in multimedia systems is to meet the deadlines of all timecritical tasks. Closely related is the goal to keep the necessary buffer space requirements low. As many streams as possible should be served concurrently, but aperiodic requests should also be schedulable without delaying their service for an infinite amount of time. The scheduling algorithm has to find a balance between time constrains and efficiency.

Earliest Deadline First

Let us first look at the EDF scheduling strategy as described above for CPU scheduling but used for the file system issue as well. Here the block of the stream with the nearest deadline would be read first. The employment of EDF as shown in Figure 16 in the strict sense results in poor throughput and an excessive seek time. Further, as EDF is most often applied as a preemptive scheduling scheme the costs for preemption of a task and scheduling another task are considerably high. The overhead caused by this is in the same order of magnitude as at least one disk seek. Hence EDF has to be adapted or combined with file system strategies.

SCAN-Eaeliest Deadline First

The SCAN-EDF strategy is a combination of the SCAN and the EDF mechanisms [ReWy93]. The seek optimization of SCAN and the real-time guarantees of EDF are combined in the following way: Like in EDF, the request with the earliest deadline is always served first. Among requests with the same deadline the specific one, that is first according to the scan direction, is served first. Among the remaining requests this principle is repeated until no request with this deadline is left.



Figure 16: EDF disk scheduling

Since the optimization only applies for requests with the same deadline, its efficiency depends on how often it can be applied (i.e., how many requests have the same or a similar deadline). To increase this probability the following tricky technique can be used: All requests have release times that are multiples of the period p. Hence, all requests have deadlines that are multiples of the period p. Therefore the requests can be grouped together and be served accordingly. For requests with different data rate requirements, in addition to SCAN-EDF, the employment of a periodic fill policy is proposed [YeVa92] to let all requests have the same deadline. With this policy, all requests are served in cycles. In every cycle each request gets an amount of service time that is proportional to its required data rate. The cycle length is equal to the sum of the service times of all requests. Thus, in every cycle all requests can be given a deadline at the end of the cycle.

SCAN-EDF can easily be implemented. Therefore EDF has to be modified slightly. If D_i is the deadline of task *i* and N_i is the track position, then the deadline can be modified to be $D_i + f(N_i)$. Thus the deadline is deferred. The function f() converts the track number of *i* into a small perturbation of the deadline as shown in the example of Figure 17. It has to be small enough so that $D_i + f(N_i) \le D_j + f(N_j)$ holds for all $D_i \le D_j$. For $f(N_i)$ it was proposed to choose the following function [ReWy93]:

$$f(N_i) = \frac{N_i}{N_{max}}$$

where N_{max} is the maximum track number on disk. Other functions might also be appropriate.

We enhanced this mechanism by proposing a more accurate perturbation of the deadline which takes into account the actual position of the head (N). This position is measured in terms of block numbers and the current direction of the head movement (see also Figure 18 and Figure 19):

1: If the head moves toward N_{max}, i.e., upwards, then



Figure 17: SCAN-EDF disk scheduling with $N_{max}=100$ and $f(N_i)=N_i/N_{max}$



1A: for all blocks $N_{\rm i}$ located between the actual position N and N_{max} the perturbation of the deadline is

$$f(N_i) = \frac{N_i - N}{N_{max}}$$
 for all $N_i \ge N$

1B: for all blocks N_i located between the actual position and the first block (no. 0)

$$f(N_i) = \frac{N_{max} - N_i}{N_{max}}$$
 for all $N_i < N$

2: If the head moves downwards towards the first blocks then



Figure 19: Accurate EDF-SCAN algorithm, head moves downwards

2A: for all blocks located between the actual position and Nmax

$$f(N_i) = \frac{N_i}{N_{max}}$$
 for all $N_i > N$

2B: for all blocks located between this first block with the block number 0 and the actual position.

$$f(N_i) = \frac{N - N_i}{N_{max}}$$
 for all $N_i \le N$

Our algorithm is more computing intensive than those with the simple calculation of [ReWy93]. In cases with only a few equal deadlines our algorithm provides improvements and the expenses of the calculations can be tolerated. In situations with many, i.e., typically more than 5) equal deadlines the simple calculation provides sufficient optimization and additional calculations should be avoided.

SCAN-EDF was compared with pure EDF and different variations of SCAN. It was shown that SCAN-EDF with deferred deadlines performs well in multimedia environments [ReWy93].

Group Sweeping Scheduling

With Group Sweeping Scheduling (GSS) requests are served in cycles, in round-robin manor [CKYu93]. To reduce disk arm movements, the set of n streams is divided is into g groups. Groups are served in fixed order. Individual streams within a group are served according to SCAN; therefore it is not fixed at which time or order individual streams within a group are served. In one cycle a specific stream may be the first to be served, in another cycle it may be the last in the same group. A smoothing buffer which is sized according to the cycle time and data rate of the stream assures continuity. If the SCAN scheduling strategy is applied to all streams of a cycle without any grouping then the playout of a stream can not be started until



Figure 20: Group sweeping scheduling as disk access strategy

the end of the cycle of its first retrieval (where all requests are served once) because the next service may be in the last slot of the following cycle. As the data must be buffered in GSS the playout can be started at the end of the group in which the first retrieval takes place. Whereas SCAN requires buffers for all streams, in GSS the buffer can be reused for each group. Further optimizations of this scheme are proposed in [CKYu93]. In this method it is ensured that each stream is served once in each cycle. GSS is a trade-off between the optimization of buffer space and arm movements. In order to provide the requested guarantees for continuous-media data we propose here to introduce a "joint deadline" mechanism: We assign to each group of streams one deadline, the "joint deadline". This deadline is specified as being the earliest one out of the deadline of all streams in the respective group. Streams are grouped in a way that all of them comprise similar deadlines. Figure 20 shows an example of group sweeping scheduling.

Mixed Strategy

In [Abbo84] a mixed strategy was introduced based on the shortest seek (also called greedy strategy) and the balanced strategy. As shown in Figure 21 every time data is retrieved from disk it is transferred into buffer memory allocated for the respective data stream. From there the application process removes it once at a time. The goal of the scheduling algorithm is

- to maximize transfer efficiency by minimizing seek time and latency, and
- to serve process requirements with a limited number of buffer space.

With *shortest seek* the first goal is served, i.e., the process which data block is closest is served first. The *balanced strategy* chooses the process which has the least amount of buffered data for service, because this process is likely to run out of data. The crucial part of this algorithm is



Figure 21: Mixed disk scheduling strategy

the decision which of the two strategies has to be applied (shortest seek or balanced strategy). For the employment of shortest seek two criteria must be fulfilled: The number of buffers for all processes should be balanced (i.e. all processes should nearly have the same number of buffered data) and the overall required bandwidth should be sufficient for the number of active processes, so that none of them will try to read immediately data out of an empty buffer. In [Abbo84] the *urgency* is introduced as an attempt to measure both. The urgency is the sum of the reciprocals of the current "fullness" (amount of buffered data). This number measures both, the relative balance of all read processes and the number of them. If the urgency is large then the balance strategy will be used, if it is small it is safe to apply the shortest seek algorithm.

Continuous Media File System

The CMFS Disk Scheduling is a non preemptive disk scheduling scheme designed for the Continuous Media File System (CMFS) at UC-Berkeley [AOG091]. Different policies can be applied in this scheme. Here the notion of the slack time *H* is introduced. The slack time is the time during which CMFS is free to do non-real-time operations or workahead for real-time processes, because the current workahead of each process is sufficient so that no process would starve even if it would not be served for H seconds. The considered real-time scheduling policies are:

• The Static/Minimal Policy is based on the minimal workahead-augmenting set (WAS). A process p_i reads a file at a determined rate R_i . To each process a positive integer M_i is assigned which denotes to the time overhead required to read a block covering, e.g., the seek time. The CMFS performs a set of operations (i.e. disk operations required by all processes) by seeking the next block of a file and reading M_i blocks of this file. Note, we consider only read operations; the same also holds with minor modifications for write operations. This seek is done for every process in the system. The data read by a process during this operation "lasts"

$$\frac{M_i \times A}{R_i}$$

where A is the block size in byte. The WAS is a set of operations where the data read for each process "lasts longer" than the worst-case time to perform the operations (i.e. the sum of the read operations of all processes is less than the time read data lasts for a process). A schedule is derived from the set that is workahead-augmenting and feasible (i.e. the requests are served in the order given by the WAS). The "Minimal Policy", the minimal WAS, is the schedule where the worst-case elapsed time needed to serve an operation set is the least (i.e. the set is ordered in a way that reduces time needed to perform the operations, e.g. by reducing seek times). The Minimal Policy does not consider buffer requirements. If there is not enough buffer this algorithm causes an buffer overflow. The "Static Policy" modifies this schedule such that no block is read if this would cause a buffer overflow for that process. With this approach starvation is avoided but its use of short operations causes high seek overhead.

- With the *Greedy Policy* a process is served as long as possible. Therefore it computes at each iteration the slack time *H*. The process with the smallest workahead is served. The maximum number *n* of blocks for this process is read; *n* is determined by *H* (the time needed to read *n* blocks has to be less or equal than *H*) and the currently available buffer space.
- The Cyclical Plan Policy distributes the slack time among processes in order to maximize the slack time. It calculates H and increases the minimal WAS with H milliseconds of additional reads; an additional read for each process is done immediately after the regular read determined by the minimal WAS. This policy distributes workahead by identifying the process with the smallest slack time and schedules an extra block for it; it is done until H is exhausted. The number of block reads for the least-workahead is determined. This procedure is repeated every time the read has completed.

The Aggressive version of the Greedy and the Cyclical Plan Policy calculates H of all processes except the least-workahead process that is by both policies immediately served. If the buffer size limit of a process is reached all policies skip to the next process. Non-real-time operations are served if there is enough slack time. First performance measurements of the above introduced strategy showed that Cyclical Plan increases system slack faster at low values of the slack time (which is likely to be the case at system set up). With a higher system slack time, apart of the Static/Minimal Policy, all policies perform about the same.

All of the disk scheduling strategies described above have been implemented and tested in prototype file systems for continuous media. Their efficiency depends on the design of the entire file system, the disk layout, tightness of deadlines, and last not least on the application behaving. It is not yet common sense, which algorithm is the "best" method for the storage and retrieval of continuous-media files. Further research must show which algorithm serves the timing requirements of continuous media best and ensures that aperiodic and non-real-time requests are efficiently served.

5.2.4 Data Structuring

Continuous-media data is characterized by consecutive, time dependent logical data units. The basic data unit of a motion video is a *frame*. The basic unit of audio is a *sample*. Frames contain the data associated with a single video image, a sample represents the amplitude of the analog audio signal at a given instance. Further structuring of multimedia data was suggested in the following way [RaVi91, Rang93, StFr92]: A *strand* is defined as an immutable sequence of continuously recorded video frames, audio samples, or both. I.e., it consists of a sequence of blocks which contain either video frames, audio samples or both. Most often it includes headers and further information related to the compression used. The file system holds primary indices in a sequence of "Primary Blocks". They contain mapping from media block numbers to their disk addresses. In "Secondary Blocks" pointers to all Primary Blocks are stored. The

"Header Block" contains pointers to all secondary blocks of a strand. General information about the strand like recording rate, length etc. is also included in the header block.

Media strands that together constitute a logical entity of information (e.g. video and associated audio of a movie) are tied together by synchronization to form a multimedia rope. A rope contains the name of its creator, its length and access rights. For each media strand in this rope the strand ID, rate of recording, granularity of storage, and corresponding block-level is stored (information for the synchronization of the playback start for all media at the strand interval boundaries). Editing operations on ropes manipulate pointers to strands only. Strands are regarded as immutable objects because editing operations like insert or delete may require substantial copying which can consume significant amount of time and space. Intervals of strands can be shared by different ropes. Strands that are not referenced by any rope can be deleted, its storage is reclaimed [RaVi91]. The following interfaces are the operations that file systems provide for the manipulation of ropes:

- RECORD [media] \rightarrow [requestID, mmRopeID] A multimedia rope, represented by mmRopeID and consisting of media strands, is recorded until a STOP operation is issued.
- PLAY [mmRopeID, interval, media] \rightarrow requestID This operation plays a multimedia rope consisting of one or more media strands.
- STOP [requestID] This operation stops the retrieval or storage of the corresponding multimedia rope.
- Additionally the following operations are supported: INSERT [baseRope, position, media, withRope, withInterval] REPLACE [baseRope, media, baseInterval, withRope, withInterval] SUBSTRING [baseRope, media, interval] CONCATE [mmRopeID1, mmRopeID2] DELETE [baseRope, media, interval]

Example: INSERT baseRope:Rope1, position: 3, media:AudioVisual, withRope: Rope2,

withInterval: from:0, length:2



Figure 22: INSERT Operation

Figure 22 provides an example of the INSERT operation whereas Figure 23 shows the REPLACE operation.

The storage system is divided into two layers:

- The *rope server* is responsible for the manipulation of multimedia ropes. It communicates with applications, allows the manipulation of ropes, and communicates with the underlying storage manager to record and play back multimedia strands. It provides the rope abstraction to the application. The rope access methods were designed similar to UNIX file access routines. Status messages about the state of the play or record operation are passed to the application. The *storage manager* is responsible for the manipulation of strands. It places the strands on disk to ensure continuous recording and playback. The interface to the rope server includes four primitives for manipulating strands:
 - 1. "PlayStrandSequence" takes a sequence of strand intervals and displays the given time interval of each strand in sequence.
 - 2. "RecordStrand" creates a new strand and records the continuous-media data either for a given duration or until StopStrand is called.
 - 3. "StopStrand" terminates a previous PlayStrandSequence or RecordStrand instance.
 - 4. "DeleteStrand" removes a strand from storage.

Example: REPLACE[baseRope:Rope1, media:video, baseInterval:[start:0, length: 3], withRope: Rope2, withInterval: [start:0, length:3]]



Figure 23: REPLACE Operation

The experimental Video File Server introduced in [Rang93] supports integrated storage and retrieval of video. The "Video Rope Server" presents a device independent directory interface to users (Video Rope). A Video Rope is characterized as a hierarchical directory structure constructed upon stored video frames. The "Video Disk Manager" manages a frame oriented motion video storage on disk, including audio and video components.

6. Further Operating System Issues

6.1 Interprocess Communication and Synchronization

In multimedia systems interprocess communication refers to the exchange of different data between processes. This data transfer has to be very efficient because continuous media require the transfer of a large amount of data in a given time span. For the exchange of discrete media data the same mechanisms are used than in traditional operating systems. Data interchange of continuous media is closely related to the memory management and is discussed in the respective section. Synchronization guarantees timing requirements between different processes. In the context of multimedia this is an especially interesting aspect. Different data streams, database entries, document portions, positions, processes, etc. have to be synchronized. Thus, synchronization is important for various components of a multimedia system and therefore it is not included in this discussion on operating systems.

6.2 Memory Management

The memory manager assigns physical resource *memory* to a single process. Virtual memory is mapped onto actually available memory. With *paging* less frequently used data is swapped between main memory and external storage. Pages are transferred back into the main memory when data on them is required by a process. Note, continuous-media data must not be temporary paged out of the main memory. If a page of virtual memory containing code or data required by a real-time process is not in real memory when it is accessed by the process, a page fault occurs, meaning that the page has to be read from disk. Page faults affect the real-time performance very seriously, so they must be avoided. A possible approach is to lock code and/ or data into real memory. However, care should be taken when locking code and/or data into real memory. Real memory is a very scarce resource to the system. Committing real memory by pinning (locking) will decrease overall system performance. The typical AIX kernel will not allow more than about 70% of real memory to be committed to pinned pages [IBM91].

The transmission and processing of continuous data streams by several components require a very efficient data transfer restricted by time constrains. Memory allocation and release functions provide well defined access to shared memory areas. In most cases, no real processing of data but only a data transfer is necessary. For example, the camera with a digitalization process is the source and the presentation process is the sink. The essential task of the other components is the *exchange* of continuous-media data with relatively high data rates in real-time. The processing involves to compute, to add, to interpret and to strip headers. This is well known in communications [McRo93]. The actual implementation can either be with external devices and dedicated hardware in the computer or can be realized with software components.

Early prototypes of multimedia systems incorporate audio and video based on external data paths only. Memory management, in this case, has a switching function only, i.e., to control an external switch.

A first step towards integration was the incorporation of the external switch function into the computer. Therefore, some dedicated adapter cards, that are able to switch data streams with varying data rates, were employed.

A complete integration can be achieved with a full digital approach within the computer, i.e., to offer a pure software solution. Data is transmitted between the single components in realtime. Copy operations are - as far as possible - reduced to the exchange of pointers and the check of access rights. This requires the access of a shared address space. Data can also directly be transferred between different adapter cards. The transfer of continuous-media data takes place in a real-time environment. This exchange is controlled but not necessarily executed by the application. The data transfer has to be performed by processes running in a real-time environment. The application running in a non-real-time environment generates, manipulates and consumes this data streams at an operating system interface.

6.3 Device Management

The device management and the actual access to a device allows the operating system to integrate all hardware components. The physical device is represented by an abstract device driver. The physical characteristics of devices are hidden. In a conventional system such devices are, e.g. a graphics adapter card, disk, keyboard, and mouse. In multimedia systems additional devices like cameras, microphones, speakers, and dedicated storage devices for audio and video have to be considered. In most existing multimedia systems such devices are not often integrated by device management and the respective device drivers.

Existing operating system extensions for multimedia usually provide one common systemwide interface for the control and management of data streams and devices. In Microsoft Windows and OS/2 this interface is known as the Media Control Interface (MCI). The multimedia extensions of Microsoft Windows, for example, provide the following classes of function calls:

- System commands are not forwarded to the single device driver (MCI driver), they are served by a central instance. An example for such a command is the query concerning all devices connected to the system ("Sysinfo").
- Each device driver has to be able to process *compulsory commands*. For instance the query for specific characteristics ("capability info") and the opening of a device ("open") are such commands.
- Basic commands refer to characteristics that all devices have in common. They can be supported by drivers. If a device driver processes such a command it has to consider all variants and parameters of the command. A data transmission is typically started by the basic command "play".
- Extended commands may refer to both, device types and special single devices. The "seek" command for the positioning on an audio CD is an example.

On the basis of a controllable camera the required concepts are explained in more detail. A camera has functions to adjust the focal length, focus and position. An abstraction of the functionality provided by the physical camera as an video input device covers the following layers, which relate to different components in a multimedia system:

- The application has access to a *logical camera* without knowledge about the specific control functions of the actually employed camera. The focal length is adjusted in millimeters. The driver translates a specific "set focal length command" into a sequence of camera hardware control commands and passes them to the control logic. The provision of such an abstract interface and the transformation into hardware dependent commands is a task of the device management of a multimedia operating system.
- Different input device classes have similar characteristics. The zoom operation of a camera can be applied in a similar way to the presentation of a still image. The still image would be *zoomed*. If, for instance, the image is stored on a photo CD with a given resolution the zoom operation would, according to the requested focal length, result in the presentation of images with different resolution respectively special parts of the image only. This kind of abstraction is part of the programming environment of a multimedia system and not task of an operating system, although in some cases it is performed by the operating system. The basic commands define several operations supported by different devices. The basic command used for the start of a data transmission between the camera and the video window of an application called play command in this context can be used in a second realization for file transfer as a kind of copy command.

To complete the description of the camera control, the positioning of the camera is discussed in this paragraph. To change the position of the camera the application specifies the target coordinates in a polar coordinate system. Yet, a concrete camera control can only execute commands like "move swivel slope head in a specific direction with a defined speed". The direction can be "left" or "right", respectively "up" or "down". Eight different speed levels are given but it is only possible to change the speed in steps of maximum two levels. During acceleration consecutively commands with speed level 2, 4, 6, 8 must be executed. It is the task of the camera driver to perform the mapping of coordinates into this positioning controlled by time and speed.

In order to define the required application interface the selectable control class can be subdivided into four function categories [RSSS90]:

1. Defined, compulsory, and generic:

All operations that must be provided for each device driver regardless to its specific functionality belong to this category. This corresponds to the above mentioned commands of the MCI.

2. Defined, compulsory, and device specific:

All functions and parameters specified in this category have to be provided by the device driver. Therefore, it exists a defined interface in the respective operating system. For example a camera driver has to be able to answer an inquiry for an eventual existing auto focus mechanism.

3. Defined, but not compulsory:

For each device type, a set of functions is defined which covers all so fare known possibilities. The functions can not be provided by all different devices and drivers. In case of the camera, such functions are, e.g., to position and to adjust the focal length because not every camera has these facilities. The interface is defined having in mind what is possible and meaningful. If such a function is employed although it is not supported by the implementation a well defined error handling mechanism applies. The application can handle this errors, and therefore it is independent of the connected physical devices.

4. Not defined, and not compulsory:

We have to be aware that there always will be unpredictable "new" devices and special developments. Hence, the operating system provides a fourth category of functions to cover all these calls.

An unambiguous definition of these categories allows an easier integration of devices into the programming environment. The multimedia extensions of todays operating systems have a device management with a first step of functional distinction towards the above outlined categories.

7 System Architecture

The employment of continuous media in multimedia systems also imposes additional, new requirements to the system architecture. A typical multimedia application does not require processing of audio and video to be performed by the application itself. Usually data is obtained from a source (e.g. microphone, camera, disk, network) and is forwarded to a sink (e.g. speaker, display, network). In such a case the requirements of continuous-media data are satisfied best if it takes "the shortest possible path" through the system, i.e. to copy data directly from adapter to adapter. The program than merely sets the correct switches for the data flow by connecting sources to sinks. Hence the application itself never really touches the data as it is the case in traditional processing. A problem with direct copying from adapter to adapter is the control and the change of quality of service parameters. In multimedia systems such an adapter to adapters systems this connection is static. This architecture of low-level data

streaming corresponds with proposals for using additional new busses for audio and video transfer within a computer. It also enables a switch-based rather then a bus-based data transfer architecture [Finn91][HaMc91]. Note, in practice we encounter header and trailers surrounding continuous-media data coming from devices and being delivered to the devices. In the case of compressed video data, e.g. the MPEG-2 program stream contains several layers of headers compared with the actual group of pictures to be displayed.

Most of todays multimedia systems have to coexist with conventional data processing. They share hardware and software components. For instance, the traditional way of protocol processing is slow and complicated. In high speed networks protocol processing is the bottleneck because it can not provide the necessary throughput. Protocols like VMTP, NETBLT and XTP try to overcome this drawback but research in this area has shown that throughput in most communication systems is not bounded by protocol mechanisms but by the way they are implemented [CJRS89]. Time intensive operations are, for example, physical buffer copying. Since the memory on the adapter is not very large and it may not store all related compressed images, data has to be copied at least once from adapter into main memory. Further copying should be avoided. An appropriate buffer management allows operations on data without performing any physical copy. In operating systems like UNIX the buffer management must be available in both, the user and the kernel space. The data need to be stored in shared memory to avoid copying between user and kernel space. For further performance improvement protocol processing should be done in threads with upcalls, i.e. the protocol processing for an incoming message is done by a single thread. A Development to support such a protocol process management is, for example, the x-Kernel.

The architecture of the protocol processing system is just one issue to be considered in the system architecture of multimedia supporting operating systems. Multimedia data should be delivered from the input device (e.g. CD-ROM) to an output device (e.g. a video decompression board) across the fastest possible path. The paradigm of streaming from source to sink is an appropriate way of doing this. Hence the multimedia application opens devices, establishes a connection between them, starts the data flow, and returns to other duties.



Figure 24: Real-time (RTE) and non real-time environments (NRTE)

As stated above, the most dominant characteristic of multimedia applications is to preserve the temporal requirement at the presentation time. Therefore multimedia data is handled in a realtime environment (RTE), i.e., its processing is scheduled according to inherent timing requirements of multimedia data. On a multimedia computer the RTE will usually coexist with a nonreal-time environment (NRTE). The NRTE deals with all data that has no timing requirements. Figure 24 shows the approached architecture. Multimedia I/O devices are in general accessed from both environments. Data such as a video frame, for example, is passed from the RTE to the display. The RTE is controlled by related functions in the NRTE. The establishment of communication connections at the start of a stream must not obey timing requirements, but the data processing for established connections has to. All control functions are performed in the NRTE. The application usually calls only these control functions and is not involved in the active continuous-media data handling. Therefore the multimedia application itself typically runs in the NRTE and is shielded from the RTE. In some scenarios users may want applications to "process" continuous-media data in an application specific way. In our model such an application comprises a module running as stream handler in the RTE. The rest of the applications run in the NRTE, both using the available stream control interfaces. System programs such as communication protocol processing and database data transfer programs make use of this programming in the RTE. Whereas applications like authoring tools and media presentation programs are relieved from the burden of programming in the RTE, they just interface and control the RTE services. Applications determine processing paths which are needed for their data processing and the control devices and paths.

To reduce data copying buffer management functions are employed in the RTE as implementation means for data transfer. This buffer management is located "between" the stream handlers. Stream handlers are all entities in the RTE which are in charge of multimedia data. Typical stream handlers are filter and mixing functions, but also parts of the communication subsystem described above can be treated in the same way. Each stream handler has endpoints for input and output through which data units flow. The stream handler consumes data units from one or more input endpoints and generates data units through one or more output endpoints.

Multimedia data usually "enters" the computer through an input device, a source, and "leaves" it through an output device, a sink (where storage can serve as an I/O device in both cases). Sources and sinks are implemented by a device driver. Applications access stream handlers by establishing sessions with them. A session constitutes a virtual stream handler for exclusive use by the application which has created it. Depending on the required QoS of a session, an underlying resource management subsystem multiplexes the capacity of the underlying physical resources among the sessions. To manage the RTE data flow through the stream handlers control operations are used which belong to the NRTE. These functions make up the *stream management system* in the multimedia architecture. There are operations provided by all stream handlers (e.g. operations to establish sessions and to connect their endpoints) and operations specific to individual stream handler (they usually determine the content of a multimedia stream and apply to particular I/O devices).

Some applications which are all in the NRTE have the need to correlate discrete data such as text and graphics with continuous streams or to post-process multimedia data (e.g. to display the time stamps of a video stream like a VCR). These applications need to obtain segments of multimedia at the stream handler interface. With a grab function the segments are copied to the application as if stream duplication took place. Due to this operation the data units loose their temporal properties because they enter the NRTE. Applications that have to generate or transform multimedia data keeping the real-time characteristics must use a stream handler included in the RTE, which performs the required processing.

The synchronization of streams is a function that is provided by the stream management subsystem. Synchronization is specified on a connection basis and can be expressed using the notions of *clock* or *logical time systems*. It determine points in time at which the processing of data units shall start. For regular streams, the stream rates can be used to relate data units to synchronization points. Sequence numbers can accomplish the same. Time stamps are a more versatile means for synchronization as they can also be used for non periodic traffic. Synchronization is often implemented by delaying the execution of a thread or by delaying the receive operation on a buffer exchanged between stream handlers.

Many operating systems already provide extensions to support multimedia applications. In the next paragraphs three of these multimedia extensions are presented.

7.1 UNIX Based Systems

In the Unix operating system the applications in the user space generally make use of system calls in the NRTE. Either the whole operating system or a part of it is also located in the NRTE and in the kernel space. Extensions to the operating system providing real-time capabilities make up the RTE part of the kernel space (see Figure 25).



Figure 25: NRTE and RTE in UNIX systems

The actual implementation of the RTE varies substantially:

- SUN OS does not yet provide RTE.
- AIX includes real-time priorities. This feature provides the basis for the RTE in the AIX based Ultimedia server.
- The IRIX operating system on Silicon Graphics Workstations has real-time capabilities, i.e., it includes a RTE.

7.2 QuickTime

QuickTime is a software extension to the Macintosh System. It provides the capability to capture, store, manage, synchronize, and display continuous-media data. A more detailed description can be found in [DrMu92]. It introduces digitized video as standard data type into the system, and it allows an easier handling of other continuous media like audio and animation. Standard applications are enhanced by multimedia capabilities. Apple has announced Quick-Time to be available for other operating systems like Windows and UNIX as well. An integration of future hardware and software developments is possible.

The standard data type of QuickTime is a movie. All kinds of continuous-media data are stored in movie documents. Additionally time information like the creation and modification date. duration etc. is also kept in the movie document. With each movie, a poster frame is associated that appears in the dialog box. Other information like current editing selection, spatial characteristics (transformation matrix, clipping region) and a list of one or more tracks are associated with the movie. A track represents a stream of information (audio or video data) that flows in parallel to every other track. With each track information like creation and modification data, duration, track number, spatial characteristics (transformation matrix, display window, clipping region), a list of related tracks, volume and start time, duration, playback rate, and data reference for each media segment is stored. A media segment is a set of references to audio and video data including time information (creation, modification, duration), language, display or sound quality, an media data type, and data pointers. Future releases will have, apart from audio and video tracks, "custom tracks" such as a subtitle track. All tracks can be viewed or heard concurrently. The tracks of a movie are always synchronized. Since movies are documents they can not only be played (including pausing, stepping through etc.) but also be edited, operations like cut, copy and paste are possible. Movie documents can be part of other documents. QuickTime is scalable, hardware components like accelerator or compressor/decompressor cards can be employed.



Figure 26: QuickTime Architecture

The QuickTime architecture comprises three major components (see Figure 26):

The *Movie Toolbox* offers a set of services to the user that allows to incorporate movies into applications. These applications may directly manipulate characteristics of audio and video data of movies. The movie is integrated in the desktop environment. Movie data can be imported and exported with the system clipboard and a movie can be edited within the Movie Toolbox.

The second component known as the *Image Compression Manager* provides a common interface for compression and decompression of data which is, independent of the implementation, to and from hard disk, CD-ROM, and floppy. It offers a directory service in order to select the correct compression component. Different interface levels for different application requirements are available. The compression techniques are a proprietary image compression scheme, a JPEG implementation, and a proprietary video compressor for digitized video data (leading to a compression ratio of 8:1 and if temporal redundancies are also removed to a ratio of 25:1). An animation compressor can compress digital data in lossy and lossless (error-free) mode. A graphics compressor is also available. The pixel depth conversion in bits per pixel can be used as a filter to be applied in addition to other compressors.

The Component Manager provides a directory service related to the components. It is the interface between the application and various system component. It shields developers from having to deal with the details of interfacing with specific hardware. In the Component Manager object oriented concepts (e.g. hierarchical structure, extensible class libraries, inheritance of component functionality, instance-based client/server model) are applied. Thus, applications are independent of implementations, they can easily integrate new hardware and software, and they can adapt to the available resources. The components managed by the Component Manager are the "Clock", the "Image Compressor" and "Image Decompressor", the "Movie Controller", the "Sequence Grabber", "Sequence Grabber Channel" and the "Video Digitizer". Further, application defined components can be added.

There is a simple resource management scheme applied to the local environment only: In case of scarce resources audio is prioritized over video, i.e., audio playback is maintained (if possible) whereas single video frames might be skipped. If an application calls the Movie Toolbox during playback, there are the following possibilities to handle these calls.

- The commonly used mode is a preemptive calling sequence, where the application returns to the system after each update. This might cause a jerky movie output.
- With a non-preemptive calling sequence, the application does not return to the system while a movie is played. This counteracts the multitasking capability.
- The high-performance controlled preemptive calling sequence is a compromise, where the application gives up the control to the Movie Toolbox for a specified time period (e.g., 50 ms).

As additional resource management scheme for better performance, it is recommended to turn off the virtual memory while playing QuickTime movies. If it is on, it will cause the sound to skip and it will lower the frame rate during the playback of a movie. However, no RTE exists.

The concept of components in QuickTime allows for easy extension without effecting applications. It attempts to form a hierarchical structure of functionality by components. The movie controller component eases user interface programming. A disadvantage of QuickTime is that there is no clear layering of abstractions for programmers and that the functionality of managers and components sometimes overlap.

7.3 Windows Multimedia Extensions

The Microsoft Windows Multimedia Extensions (WME) are an enhancement to the Windows programming environment. They provide high-level and low-level services for the development of multimedia applications for application developers, using the extended capabilities of a multimedia personal computer [Micr91].

The following services for multimedia applications are provided by the WME:

• A Media Control Interface (MCI) for the control of media services. It comprises an



Figure 27: MS Windows Multimedia Extensions architecture

extensible string-based and message-based interface for the communication with MCI device drivers. The MCI device drivers are designed to support the playing and recording of waveform audio, the playing of MIDI (Musical Instrument Digital Interface) files, the playing of compact disc audio from a CD-ROM disc drive, and the controlling of some video disc players.

- A Low-Level API (Application Programming Interface) provides access to multimedia related services like playing and recording of audio with waveform and MIDI audio devices. It also supports the handling of input data from joysticks and precise timer services.
- A multimedia file I/O service provides buffered and unbuffered file I/O. It also supports the standard IBM/Microsoft Resource Interchange File Format (RIFF) files. These services are extensible with custom I/O procedures that can be shared among applications.
- The most important device drivers are available for multimedia applications are:
 - An enhanced high-resolution video display driver for Video 7 and Paradise VGA cards providing 256 colors, improved performance, and other new features.
 - A high-resolution VGA video display driver allowing the use of a custom 16-color palette as well as the standard palette.
 - A low resolution VGA video display driver providing 320-by-320 resolution with 256 colors.
- The Control Panel Applets that allows the user to change display drivers, to set up a screen saver, to install multimedia device drivers, to assign waveform sounds to system

alerts, to configure the MIDI Mapper, and to calibrate joysticks. A MIDI Mapper supports the MIDI patch service, that allows MIDI files to be authored independently of enduser MIDI synthesizer setups.

Figure 27 shows the rough architecture of MS Windows Multimedia Extensions:

The MMSYSTEM library provides the Media Control Interface services and low-level multimedia support functions. The communication between the low-level MMSYSTEM functions and multimedia devices such as waveform, MIDI, joystick, and timer is provided by the multimedia device drivers. The high-level control of media devices is provided by the drivers for the Media Control Interface.

The main concepts of the architecture of the Multimedia Extensions are extensibility and device independence. They are provided by a translation layer (MMSYSTEM) that isolates applications from device drivers and centralizes device independent code, run-time linking that allows the MMSYSTEM translation layer to link to the drivers it needs, and a well defined and consistent driver interface that minimizes the development of specialized code and makes the installation and upgrade process easier.

7.4 OS/2 Multimedia Presentation Manager/2

The Multimedia Presentation Manager/2 (MMPM/2) is part of IBMs Operating System/2 (OS/2). OS/2 is a well suited platform for multimedia, because it supports, e.g., preemptive multi-tasking, priority scheduling, overlapped I/O, and demand-paged virtual memory storage.



Figure 28: The architecture of the OS/2 Multimedia Presentation Mangager/2 Figure 28 provides an overview of the architecture.

The *Media Control Interface (MCI)* is a device-independent programming interface that offers commands similar to an entertainment system. The following list comprises a selection of typical MCI-commands:

- "Open", "close", "status of a device" are provided for all devices.
- For playback and recording device-dependent "play", "record", "resume", "stop", "cue", and "seek" commands exist.
- "Set cue point" allows for synchronization.

"Get table of contents of a CD-ROM" is an example of a device specific command.

A *logical device* in MMPM/2 is a logical representation of the functions available from either a hardware device, a hardware device with software emulation or a software emulation only. The actual implementation is not relevant to an application because the MCI provides this device independence.

Examples for logical devices are an "Amplifier-Mixer Device" similar to a home stereo amplifier-mixer, a "Waveform Audio Device" to record and play digital audio, a sequencer device for MIDI-sounds, a "CD Audio Device" that provides access to audio compact discs (CD-DA), a "CD-XA Device" to support CD-ROM/XA discs and a "Videodisc Device" to control video disc players which deliver analog video and audio signals.

The *Multimedia I/O functions* enable media drivers and applications to access and manipulate data objects that are stored in memory or on a file system. Storage system I/O processes handle the access to specific storage devices. File format I/O processes manage the access to data stored in file formats like "RIFF Waveform" and "BitMap". They use the services of the storage system I/O processes.

The implementation of data streaming and synchronization is supported by the Stream Programming Interface (SPI). It provides access to the SyncStream Manager that coordinates and manages the data buffers and synchronization data. Pairs of stream handlers implement the transport of data from a source to a sink.

Ease of use is supported in MMPM/2 on several levels. The installation of programs and setup of devices is supported by unified graphical user interfaces that centralize these functions for easy access. Also a style guide for applications ensures that there is a common look and feel of applications that correspond to this guide. There is a high flexibility because the application developers and device providers can integrate their own logical devices, I/O processes and stream handlers. So, new media devices, data formats, etc. can be integrated in MMPM/2 and can be used by every application using the Media Control Interface.

OS/2 with MMPM/2 is a platform that has some basic operating mechanisms to support the processing and presentation of multimedia information as it is needed in multimedia application scenarios. It incorporates a RTE implemented as a set of device drivers. MMPM/2 is an advanced platform for the development of these multimedia applications by providing the media and stream abstractions.

Finally it should be pointed out that MMPM/2 and WME look very similar and have many concepts in common.

8. Concluding Remarks

In this paper we addressed the major issues of operating systems related to multimedia data processing namely, resource management, scheduling, and file systems. This discussion includes the most relevant existing architectures of such systems.

The concepts employed by current multimedia operating systems have been initially used in real-time systems and were adapted to the requirements of multimedia data. Todays operating systems incorporate these functions either as device driver or as extensions based on the existing operating system scheduler and file systems. In a next step an integration of real-time processing and non real-time processing in the native system kernel can be expected.

Acknowledgment

The author gratefully acknowledges the many valuable advices and the work performed by Andreas Mauthe in this context. Klara Nahrstedt and Lars Wolf suggested many improvements to all sections which we incorporated into this final version. Thank You!

9 References

[Abbo84] C. Abbott; Efficient Editing of Digital Sound on Disk; Journal Audio Eng. Soc. Vol. 32, No. 6, June 1984, pp. 394-402. D. P. Anderson, L. Delgrossi, R. G. Herrtwich; Process Structure and Schedul-[ADHe91] ing in Real-Time Protocol Implementations; Kommunikation in verteilten Systemen, GI/ITG-Fachtagung, Proceedings, Springer Verlag, Mannheim, February 1991, pp. 83-95. [AHSc90] D. P. Anderson, R. G. Herrtwich, C. Schaefer: SRP; A Resource Reservation Protocol for Guaranteed-Performance Communication in the Internet; Technical Report No.90-006, International Computer Science Institute, February 1990. [Ande93] D. P. Anderson; Meta-Scheduling for Distributed Continuous Media; ACM-Transactions on Computer Systems, Vol. 11, No. 3, August 1993. M. Andrews; Guaranteed Performance for Continuous Media in a General [Andr89] Purpose Distributed System; Master Thesis, University of California, Berkeley, October 1989. D. P. Anderson, G. Homsy; A Continuous Media I/O Server and Its Synchroni-[AnHo91] zation Mechanism; IEEE-Computer, Vol. 24, No. 10, October 1991, pp. 51-57. [AOGo91] D. P. Anderson, Y. Osawa, R. Govindan; Real-Time Disk Storage and Retrieval of Digital Audio/Video Data; Technical Report No. UCB/CSD 91/646, Computer Science Division, University of California, Berkeley, September 1991. D.P. Anderson, S. Tzou, R. Wahbe, R. Govindan, M. Andrews; Support for [ATWG90] Continuous Media in the DASH System; Proc. of the 10th ICDCS, Paris, May 1990. [BIFi87] J. Blazewicz, G. Finke; Minimizing Mean Weighted Execution Time Loss on Identical and Uniform Processors; Information Processing Letters, Vol. 24, March 1987, pp. 258-263. [ChLi88] J.-Y. Chung, J. W. S. Liu; Algorithms for Scheduling Periodic Jobs to Minimize Average Error, IEEE-Real-Time Systems Symposium, Huntsville, 1988, pp. 142-151. [ChLi89] J.-Y. Chung, J. W. S. Liu; Performance of Algorithms for Scheduling Periodic Jobs to avoid Timing Faults; Proceedings of 22'nd Hawaii International Conference on System Sciences, Hawaii, 1989, pp. 683-692. D. D. Clark, Van Jacobson, J. Romkey, H. Salwen; An Analysis of TCP Process-[CJRS89] ing Overhead; IEEE Communications Magazine, June 1989, pp. 23-29. [CKYu93] M.-S. Chen, D. D. Kandlur, P. S. Yu; Optimization of the Group Sweeping Scheduling (GSS) with Heterogeneous Multimedia Streams; Proceedings of the first ACM International Conference on Multimedia, Anaheim, California, 1993, pp. 235-241. D. W. Craig, C. M. Woodside; The Rejection Rate for Tasks with Random Arriv-[CrWo90] als, Deadlines and Preemptive Scheduling; IEEE-Transaction on Software Engineering, Vol. 16, No. 10, October 1990, pp. 1198-1208. S.-C. Cheng, J. A. Stankovic, K. Ramamritham; Scheduling Algorithms for [CSRa88] Hard Real-Time Systems - A Brief Survey; Hard Real-Time Systems (J. A.

	Stanko Press, 1	vic, K.Ramamritham, Hrsg), Washington, DC, IEEE Computer Society 1988, pp.150-178.
[Dert	74] M. L. cessing pp. 807	Dertouzos: Control Robotics; <i>The Procedural Control of Physical Pro-</i> ;; Information Processing 74 - North Holland Publishing Company, 1974, 7-813.
[DIN	85] Germa nology Kõln, I	n national standardization body (Deutsches Institut fur Normung); Termi- in Computing (Informationsverarbeitung-Begriffe); DIN 43000, Berlin- Beuth, 1985.
[DrM	[u92] D. L. 1 1992.	Drucker, M. D. Murie; QuickTime Handbook; Hayden, Carmel, USA,
[Ferr	90] D. Fen nationa 1990.	rari; Client Requirements for Real-Time Communication Services; Inter- al Computer Science Institute, Technical Report 90-007, Berkely, March
[Finn	91] Gregor Archite 1991, j	cy G. Finn; An Integration of Network Communication with Workstation ecture; ACM Computer Communication Review, Vol. 21, No. 5, October op. 18-29.
[Fox9	91] E. A. puter,	Fox; Advances in Interactive Digital Multimedia Systems; IEEE Com- Vol. 24, No.10, October 1991, pp. 9-22.
[Fren	82] S. Frei the Jol	nch; Sequencing and Scheduling: an Introduction to the Mathematics of Shop, Ellis Horwood Limited, Chichester, 1982.
[GeC	h92] J. Gen Storag No.1, J	amell, S. Christodoulakis; <i>Principles of Delay Sensitive Multimedia Data</i> <i>e and Retrieval</i> ; ACM Transactions on Information Systems, Vol.10, January 1992.
- [GiTo	D. W. Operat Castle pp. 20	Gifford, J. W. O'Toole; Intelligent File Systems for Object Respositories; ing Systems of the 90s and Beyond, International Workshop Dagstuhl , Germany, July 1991, Proceedings, Springer-Verlag, Heidelberg, 1991, -24.
[GoA	An91] R. Gov media No. U	vindan, D. P. Anderson; Scheduling and IPC Mechanisms for continuous ; University of California, Computer Science Division, Technical Report CB/CSD 91/622, Berkeley, March 1991.
[Gre	i88] I. Gre Morga	if, et al.; Computer Supported Cooperative Work: A Book of Readings; in Kaufman Publisher, San Mateo, 1988.
[HaN	Mc91] M. Ha Revie	yter, Derek McAuley; <i>The Desk Area Network</i> ; ACM Operating Systems w, Vol. 25, No. 4, October 1991, pp. 14-21.
[HaS	Sh89] D. Ha with I Monio	ban, K. G. Shin; Application of Real-Time Monitoring to Scheduling Tasks Random Execution Times; IEEE-Real-Time Systems Symposium, Santa ca, 1989, pp. 172-180.
[Her	r90] R. G. <i>Media</i> pp. 11	Herrtwich; Time Capsules: An Abstraction for Access to Continuous- a Data; IEEE-Real-Time Systems Symposium, Lake Buena Visat, 1990, -20.
[HeS	St91] R. G. munil gung,	Herrtwich, R. Steinmetz; Towards Multimedia: Why and How; Telekom- cation und multimediale Anwendungen der Informatik, GI-21. Jahresta- Proceedings, Springer Verlag, Darmstadt, October 1991, pp. 327-342.

[HHSt91]	D. Hehmann, R. G. Herrtwich, R. Steinmetz; Creating HeiTS: Objectives of the Heidelberg High-Speed Transport System; Technical Report No. 439102, IBM-ENC Heidelberg, 1991.
[HLPa91]	J. Hyman, A. A. Lazar, G. Pacifics; <i>MARS: The Magnet II Real-Time Schedul-</i> ing Algorithm; acm-Press, Computer Communication Review, Vol. 21, No. 4, SIGCOMM '91, Zürich, September 1991, pp. 285-294.
[HeWo92]	A System Software Structure for Distributed Multimedia Systems; 5. European SIGOPS Workshop, Le Mont Saint-Michel, France, September 1992.
[HoLe88]	K. S. Hong, J. YT. Leung; On-Line Scheduling of Real-Time of Tasks; IEEE- Real-Time-System Symposium, Huntsville, 1988, pp. 244-258.
[HVWW94]	R. G. Herrtwich, C. Vogt, H. Wittig, L. Wolf; <i>Resource Mangament for Distrib- uted Multimedia Systems</i> ; Technical Report No. 43.9403 IBM European Net- working Center, IBM Heidelberg, 1994.
[Iaco88]	E. Iacobucci: OS/2 Programer's Guide; McGraw-Hill, Berkeley, 1988.
[IBM89]	IBM Corp.; IBM Operating System/2 Programming Tools and Information Ver- sion 1.2: Control Program Programming Reference; IBM, September 1989.
[IBM91]	IBM Corp.; AIX Version 3.1: RISC System/6000 as a Real-Time System; IBM- International Technical Support Center, Austin, March 1991.
[IBM92]	IBM Corp.; IBM OS/2 Programming Guide, Volume I, Control Program Inter- face; IBM, March 1992.
[IBM92a]	IBM Corp.; The OS/2 Multimedia Advantage; IBM, 1992.
[IBM92b]	IBM Corp.; Multimedia Presentaion Manager/2, Programming Guide; IBM, 1992.
[Jans85]	P. A. Janson; Operating Systems, Structures and Mechanisms; Academic Press Inc. Orlando, 1985.
[Jeff89]	K. Jeffay; Analysis of a Synchronization and Scheduling Discipline for Real- Time Tasks with Preemption Constraints; IEEE-Real-Time Systems Sympo- sium, Santa Monica, 1989, pp. 295-305.
[Jeff90]	K. Jeffay: Scheduling Sporadic Tasks with Shared Resources in Hard-Real-Time Systems; University of North Carolina at Chapel Hill, Department of Computer Science, Technical Report TR90-039, November 1990.
[JSPo91]	K. Jeffay, D. L. Stone, D. E. Poirier; YARTOS: Kernel Support for Efficient, Pre- dictable Real-Time Systems; Proceedings of IFAC, Workshop on Real-Time Programming, Atlanta, Pergamon Press, May 1991.
[JSSm91]	K. Jeffay, D. L. Stone, F. D. Smith; <i>Kernel Support for Live Digital Audio and Video</i> ; Second International Workshop on Network and Operating System Support for Digital Audio and Video, Heidelberg, December 1991.
[Koga90]	M. S. Kogan; OS/2 2.0 Overview; OS/2 Notebook; Microsoft Press, Washinton, 1991.
[Krak88]	S. Krakowiak; Principles of Operating Systems; MIT Press, Cambridge, Massa- chusetts, 1988.
[KrLc91]	C. M. Krishna, Y. H. Lee; Real-Time Systems; IEEE-Computer, May 1991, pp.

10	11
10	-11.
~ v	

- [KWYe94] A. Karmouch, R. Wang, T. Yeap; Design and Analysis of a Storage Retrieval Model for Audio and Video Data; Technical Report, Multimedia Information Systems, Department of Electrical Engineering, University of Ottawa, Canada, 1994.
 - [LaEf91] B. Lamparter, W. Effelsberg; X-Movie: Transmission and Presentation of Digital Movies under X; Second International Workshop on Network and Operating System Support for Digital Audio and Video, Heidelberg, November 1991.
 - [Law173] E. L. Lawler; Optimal Sequencing of a Single Machine Subject to Precedence Constraints; Management Science, Vol.19, No. 5, January 1973, pp. 544-546.
 - [LeMe80] J. Y.-T Leung, M. L. Merrill; A Note on Preemptive Scheduling of Periodic Real-Time Tasks; Information Processing Letters, Vol. 11, No. 3, Nov. 1980, pp. 115-118.
 - [LeSh86] J. P. Lehoczky, L. Sha; *Performance of Real-Time Bus Scheduling Algorithms*; ACM Performance Evaluation Review, Vol. 14, No. 1, May 1986, pp. 44-53.
 - [Letw88] G. Letwin; Inside OS/2; Microsoft Press, Washington, 1988.
 - [LeWh82] J. Y.-T. Leung, J. Whitehead; On the Complexity of Fixed-Priority Scheduling of Periodic Real-Time Tasks; Performance Evaluation 2, 1982, pp. 237-350.
 - [LiLa73] C. L. Liu, J. W. Layland; Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment; Journal of the ACM, Vol.20, No.1, Jan. 1973, pp. 46-61.
 - [LLNa87] J. W. S. Liu, K.-J. Lin, S. Naturajan; Scheduling Real-Time, Periodic Jobs Using Imprecise Results; IEEE-Real-Time Systems Symposium, San Jose, 1987, pp. 252-260.
 - [LLSY91] J. W. S. Liu, K.-J. Lin, W.-K. Shin, A. C. Yu; Algorithms for Scheduling Imprecise Computations; IEEE-Computer, May 1991, pp. 58-68.
 - [LoSh93] P. Lougher, D. Shepherd; *The Design of a Storage Service for continuous media*; The Computer Journal, Vol. 36, No. 1, 1993, pp. 32-42.
 - [LSST91] J. P. Lehoczky, L. Sha, J.K. Strosnider, H. Tokuda; Fixed Priority Scheduling Theory for Hard Real-Time Systems; in Foundations of Real-Time Computing, Scheduling and Resource Management, Kluwer Academic Publisher. Norwell, 1991, pp. 1-30.
 - [McRo93] B. C. McKellar, J. Roos; Buffer Management in Communication Systems; Technical Report 43.9312 IBM European Networking Center, IBM Heidelberg, 1993.
 - [MeT090] C. W. Mercer, H. Tokuda; *The ARTS Real-Time Object Model*; IEEE-Real-Time System Symposium, Lake Buena Vista, 1990, pp. 2-10.
 - [MeT091a] C. W. Mercer, H. Tokuda; Priority Consistency in Protocol Architectures; Second International Workshop on Network and Operating System Support for Digital Audio and Video, Heidelberg, November 1991, pp. 22-27.
 - [MeTo91b] C. W. Mercer, H. Tokuda; An Evaluation of Priority Consistency in Protocol Architectures; Carnegie-Mellon University, Pittsburgh, 1991.

[Micr91]	Microsoft Windows; Multimedia: Programmer's Workbook; Microsoft Press, 1991.
[Mok84]	A. K. Mok; The Design of Real-Time Programing Systems Based on Process Models; IEEE-Real-Time Systems Symposium, Austin, 1984.
[MSSt92]	A. Mauthe, W. Schulz, R. Steinmetz; Inside the Heidelberg Multimedia Operat- ing System Support: Real-Time Processing of continuous media in OS/2; Tech- nical Report No. 43.9214, IBM European Networking Center, IBM Heidelberg, 1992.
[Mull91]	S. J. Mullender; Systems of the Nineties - Distributed Multimedia Systems: Systems of the 90s and Beyond; International Workshop Dagstuhl Castle, Germany, July 1991, Proceedings, Springer Verlag, Heidelberg, 1991, pp. 273-278.
[NaCh91]	A. D. Narasimhalu, S. Christodoulakis; <i>Multimedia Information Systems: The Unfolding of a Reality</i> ; IEEE-Transaction on Computer, Vol. 24, No.10, October 91, pp. 6-8.
[NaVo92]	R. Nagarajan, C. Vogt; Guaranteed-Performance of Multimedia Traffic over the Token Ring; Technical Report No.439201, IBM-ENC, Heidelberg, 1992.
[NHNW93]	J. Nieh, J. G. Hanko, J. D. Northcutt, G. A. Wall; SVR4UNIX Scheduler Unac- ceptable for Multimedia Applications; Proceedings of the 4th International Workshop on Network and Operating System Support for Digital Audio and Video, Lancaster, U.K., 1993, pp 35-47.
[Oper89]	R. Levin, et al.; Operating Systems Review; ACM Press, Operating Systems Review, Vol.23, No.3, July 1989.
[OrHa91]	R. Orfali, D. Harkey; Client-Server Programming with OS/2, Extended Edition; Van Nostrand Reinhold, New York, 1991.
[PeSi83]	J. Peterson, A. Silberschatz; Operating System Concepts; Addison-Wesley Pub- lishing Company, 1983.
[Rang93]	P. V. Rangan; Video Conferencing, File Storage, and Management in Multmedia Computer Systems; Computer Networks and ISDN Systems, March 1993.
[RaVi91]	P. V. Rangan, H. M. Vin; <i>Designing File Systems for Digital Video and Audio</i> ; Proceedings of the 13th ACM Symposium on Operating Systems Principles, Monterey California, Operating Systems Review, Vol. 25, No. 5, October 1991, pp 81-94.
[ReWy93]	A. L. N. Reddy, J. Wyllie; Disk Scheduling in a Multimedia I/O System; Proceedings of the first ACM International Conference on Multimedia, Anaheim, California, 1993, pp. 225-233.
[RKVi92]	P. V. Rangan, T. Kaeppner, H. W. Vin; <i>Techniques for Efficient Storage of Digi- tal Video and Audio</i> ; Proceedings of 1992 Workshop on Multmedia Information Systems, Tempe, Arizona, February 1992.
[RSLe88]	R. Rajkumar, L. Sha, J. P. Lehoczky; Real-Time Synchronization Protocols for Multiprocessors; IEEE-Real-Time Systems Symposium, Huntsville, 1988, pp. 259-269.
[RSSS90]	J. Rückert, B. Schöner, R. Steinmetz, H. Schmutz; A Distributed Multimedia Environment for Advanced CSCW Applications; IEEE Multimedia '90, Bor-

deaux, France, 15.-17. November 1990.

- [RVG*93] K. K. Ramakrishnan, L. Vaitzblit, C. Gray, U.Vahalia, D. Ting, P. Tzelnic, S. Glaser, W. Duso; Operating System Support for a Video-On-Demand File Service; Proceedings of the 4th International Workshop on Network and Operating-System Support for Digital Audio and Video, Lancaster, U.K., 1993, pp 225-236.
- [Salz91] J. H. Salzer; File Sytems Indexing and Backup; Operating Systems of the 90s and Beyond, International Workshop Dagstuhl Castle, Germany, July 1991, Proceedings, Springer-Verlag, Heidelberg, 1991, pp. 13-19.
- [ShGo90] L. Sha, J. B. Goodenough; *Real-Time Scheduling Theory and ADA*; IEEE-Transaction on Computer, Vol. 23, No. 4, April 1990, pp. 53-64.
- [ShRa89] L. Sha, R. Rajkumar; Real-Time Systems, A Tutorial of the Rate-Monotonic Scheduling Framework with Bus-Related Issues; Futurebus+ P896.3, Draft 4.0, 1989.
- [SKG091] L. Sha, M. H. Klein, J. B. Goodenough; Rate Monotonic Analysis for Real-Time Systems; in Foundations of Real-Time Computing, Scheduling and Resource Management, Kluwer Academic Publisher, Norwell, 1991, pp. 129-156.
- [SSLe89] B. Sprunt, L. Sha, J. Lehoczky; Aperiodic Task Scheduling for Hard Real-Time Systems; The Journal of Real-Time Systems, Vol.1, 1989, pp. 27-60.
- [Stei93] R. Steinmetz; Multimedia Technology: Fundamentals and Introduction (in German); Springer-Verlag, 1993.
- [Stei94a] R. Steinmetz; Data Compression in Multimedia Computing: Principles and Techniques; acm journal 'Multimedia Systems', Vol.1, No.4, February 1994, pp.166-172.
- [Stei94b] R. Steinmetz; Data Compression in Multimedia Computing: Standards and Systems; acm journal 'Multimedia Systems', Vol.1, No.5, March 1994.
- [StFr91] R. Steinmetz, C. Fritzsche; *Abstractions for Continuous-Media Programming*; Computer Communications, Vol.15, No.6, July/August 1992.
- [StMe92] Steinmetz, T. Meyer; *Modelling Distributed Multimedia Applications*; IEEE International Workshop on Advanced Communications and Applications for High Speed Networks (IWACA), München, March 1992.
- [StNa94] R. Steinmetz, K. Nahrstedt; *The Fundamentals in Multimedia Systems*; Prentice-Hall, Inc., to appear in November 1994.
- [StRa89] J. A. Stankovic, K. Ramamritham; The Spring Kernel: A new Paradigma of Real-Time Operating Systems; ACM Press, Operating System Review, Vol. 23, No.3, July 1989, pp. 54-71.
- [Tane87] A. S. Tanenbaum; *Operating System, Design and Implementation*; Prentice-Hall, Inc., Englewood Cliffs, 1987.
- [Tesh91] T. Teshima; Constructing Image Databases for the Collection of Art Museums; Telekommunikation und multimediale Anwendungen der Informatik, GI-21. Jahrestagung, Darmstadt, Proceedings, Springer Verlag, October 1991, pp. 50-56.
- [THDu91] W. Tawbi, E. Horlait, S. Dupuy; High Speed Protocols: State of the Art in Mul-

	timedia Applications; Second International Workshop on Network and Operat- ing System Support for Digital Audio and Video, Heidelberg, November 1991, pp 349-358.
[TiKo91]	A. M. Van Tilborg, G. M. Koob, (Eds.); Foundations of Real-Time Computing, Scheduling and Resource Management; Kluwer Academic Publisher, Norwell, 1991.
[ToMe89]	H. Tokuda, C. W. Mercer; ARTS: A Distributed Real-Time Kernel; ACM Press, Operating Systems Review, Vol. 23, No. 3, July 1989, pp. 29-53.
[Topo90]	C. Topolocic (Ed.): Experimental Internet Stream Protocol, Version 2 (ST II); Internet Network Working Group, RFC 1190, October 1990.
[TWWa87]	H. Tokuda, J. W. Wendorf, HY. Wang; Implementation of a Time-Driven Scheduler for Real-Time Operating Systems; IEEE-Real-Time Systems Symposium, San Jose, 1987, pp. 271-280.
[Vest90]	S. Vesttal; The Accuracy of Predicting Rate Monotonic Scheduling Perfor- mance; Proceedings of TRI-ADA '90, December 1990.
[VHNa92]	C. Vogt, R. G. Herrtwich, R. Nagarajan; <i>HeiRAT: The Heidelberg Resource Administration Technique, Design Philosophy and Goals</i> ; Technical Report No. 43.9213, IBM-ENC, Heidelberg, 1992.
[Vide90]	Video Logic; <i>Multimedia Development Tools & Resources</i> ; VideoLogic Limited, Hertfordshire, 1990.
[ViRa93]	H.M. Vin, P. V. Rangan; Techniques for Efficient Storage of Digital Video and Audio; Computer Communications, Vol.16, March 1993, pp.168-176.
[VZSR91]	H. M. Vin, P. T. Zellweger, D. C. Swinehart, P. V. Rangan; <i>Multimedia Conferencing in the Etherphone Environment</i> ; IEEE-Computer, Vol. 24, No. 10, October 1991, pp. 69-79.
[WBV094]	L. Wolf, W. Burke, C. Vogt; CPU Scheduling in Multimedia Systems; Technical Report No. 439404, IBM European Networking Center, IBM Heidelberg, 1992.
[WoCr87]	C. M. Woodside, D. W. Craig; Local Non-Preemptive Scheduling Policies for Hard Real-Time Distributed Systems; IEEE-Real-Time Systems Symposium, San Jose, 1987, pp. 12-17.
[WoHe94]	L. Wolf, R. G. Herrtwich; The System Architecture of the Heidelberg Transport System; ACM Operating Systems Review, Vol. 28, No. 2, April 1994, pp. 51-64.
[WWV92]	H. Wittig, L. Wolf, C. Vogt, CPU Utilization of Multimedia Processes; The Heidelberg Predictor of Execution Times (HeiPOET), Technical Report No. 439405, 1994.
[YeVa92]	J. Yee, P. V. Vatarya; Disk Scheduling policies for Real-Time Multimedia Appli- cations; Technical Report, University of California, Berkeley, August 1992.
[ZhKe91]	H. Zhang, S. Keshav; <i>Comparison of Rate-Based Service Desciplines</i> ; ACM- Press, Computer Communication Review, Vol. 21, No. 4, SIGCOMM'91, Zürich,s September 1991, pp. 113-122.

List of European Networking Center Technical Reports

-	TR 43.9214	Andreas Mauthe Werner Schulz Ralf Steinmetz	Inside the Heidelberg Multimedia Operating System Support: Real-Time Processing of Continuous Media in OS/2
	TR 43.9215	Emanuel Farber Thomas Schütt	The Heidelberg High Speed Transport System: First Performance Results
	TR 43.9301	R. Oechsle M. Graf	The Internet Protocol Family over ATM
	TR 43.9302	Luca Delgrossi Frank O. Hoffmann	A Detailed Tour of ST-II for the Heidelberg Transport System
	TR 43.9303	Luca Delgrossi Ralf G. Herrtwich Frank O. Hoffmann	An Implementation of ST-II for the Heidelberg Transport System
	TR 43.9304	Derick Jordaan Martin Paterok Carsten Vogt	Layered Quality of Service Management in Heterogeneous Networks
	TR 43.9305	L. Delgrossi; Ch. Halstrick D. Hehmann; R. Herrtwich O. Krone; J. Sandvoss; C. Vogt	Media Scaling for Audiovisual Communication with the Heidelberg Transport System
~	TR 43.9306	Barbara Twachtmann Ralf G. Herrtwich	Multicast in the Heidelberg Transport System
	T'R 43.9307	R. Steinmetz	Compression Techniques in Multimedia Systems: A Survey
	TR 43.9308	Martin Bever Ulrich Schäffer Claus Schottmüller	ISO OSI FTAM and High Speed File Transfer: No Contradiction
	TR 43.9309	Ralf Steinmetz	Videodatenkompression für verteilte Multimedia-Anwendungen
	TR 43.9310	Ralf Steinmetz Clemens Engler	Human Perception of Media Synchronization
-	TR 43.9311	Luca Delgrossi; Chr. Halstrick Ralf Herrtwich; F. Hoffmann Jochen Sandvoss; B.Twachtmann	Reliability Issues in Multimedia Transport
	TR 43.9312	Brian Craig McKellar Jan Roos	Buffer Management in Communication Systems

TR 43.9313	Robert Erfle Johannes Rückert Ingo Barth; Gabriel Dernler Franz Fabian; Kurt Rothermel Frank Sembach	Multimedia Document Handling - A Survey of Concepts and Methods
TR 43.9314	Martin Paterok Stefan Kaetker Carsten Vogt Luca Delgrossi, Hartmut Wittig	An SNMP MIB for the ST-II Protocol and the Heidelberg Resource Administration Technique
TR 43.9315	Luca Delgrossi Ralf Herrtwich Carsten Vogt, Lars Wolf	Reservation Protocols for Internetworks: A Comparison of ST-II and RSVP
TR 43.9316	Thomas Käppner Lars Wolf	Architecture of HeiPhone: A Testbed for Audio/Video Teleconferencing
TR 43.9317	Ralf Herrtwich Luca Delgrossi Frank Hoffmann, Sibylle Schaller	Receiver-Initiated Communication with ST-II
TR 43.9318	Michael Altenhofen Juergen Dittrich Rainer Hammerschmidt, Ralf Her Carsten Kruschel, Ansgar Kuecke Thomas Steinig, Kathrin Werner,	Implementing the BERKOM Multimedia Collaboration Service rrtwich, Thomas Kaeppner s, Florin Spanachi Joerg Winckler
TR 43.9319	Lars Wolf Ralf Herrtwich	The System Architecture of the Heidelberg Transport System
TR 43.9401	Jochen Sandvoss Jörg Winckler Hartmut Wittig	Network Layer Scaling: Congestion Control in Multimedia Communication with Heterogenous Netoworks
TR 43.9402	Ralf Steinmetz	Multimedia Operating Systems: Resource Reservation, Scheduling, File Systems, and Architectures
TR 43.9403	Carsten Vogt	Quality-of-Service Calculation for Multimedia Streams variable Bit Rate
TR 43.9404	Heinrich J. Stüttgen	Network Evolution and Multimedia Communication



Technische Universi

Technische Universität Darmstadt Fachgebiet Multimedia Kommunikation Rundeturmstr. 10 64283 Darmstadt 87/20

Phone: +49 6151 16-6150

and the second se

Communications Lab

·

oFG Kolloquium igitale Signalraibeilung" anglover ...3.84

Petri Netze als Beschreibungs- und Analysehilfen zur Strukturierung von Algorithmen auf Mehrrechnersysteme

K.Hoffmann, R.Steinmetz Institut für Übertragungstechnik Technische Hochschule Darmstadt, Merckstr.25, 6100 Darmstadt

Einleitung

Petri Netze sind ein Hilfsmittel zur Spezifikation, Modellierung und Entwicklung von sehr unterschiedlichen Systemen auf verschiedenen Abstraktionsebenen: Von der Programm- und Hardwarebeschreibung /Cap Dspl 82/ bis zur Darstellung chemischer Prozesse, juristischer Beweisführung und wirtschaftlicher Produktionsabläufe. Die besonderen Vorteile liegen in der Verarbeitung paralleler Prozesse, in dem fundierten mathematischen Hintergrund, nicht nur in der Simulation, sondern auch in der Analysierbarkeit und in den Möglichkeiten zur graphischen Darstellung /Peterson 81, Reisig 82/.

Mit Unterstützung eines geeigneten Petri Netzes werden hier digitale Signalverarbeitungsanlagen u.A. auf ihre Verklemmungsfreiheit hin untersucht. Gleichzeitig werden Aussagen zur vorteilhaften Strukturierung (Aufteilung) der

Algorithmen für das spätere Prozeßscheduling bei Mehrrechnerbetrieb getroffen. Das System soll mehrere Algorithmen verarbeiten, die von mehreren Datenströmen gespeist werden, Ergebnisse liefern und mit der Umgebung Kommunikationsprotokolle abwickeln. Die Algorithmen sollen als Prozesse untereinander Informationen austauschen (Bild 1).



KH 84

Hier sollen weder eine Parallelisierung von Algorithmen /Brent 74, Coleman Hughes Powell 81/, eine Analyse paralleler Ausdrücke /Bernstein 66/, noch Schedulingprobleme /Ecker 77/ betrachtet werden.

Die Systementwicklung

Ausgehend von der Aufgabenstellung werden die zu implementierenden Algorithmen definiert. Anschließend werden sie in der Petri Netz Beschreibungssprache PNDL kodiert. Auch die Systeminternen Abhängigkeiten müssen in PNDL klar spezifiziert werden. Für die weitere Strukturierung sind folgende Erläuterungen notwendig:

Eine logisch zusammenhängende Sequenz von zeitlich sequentiellen Abläufen in einem System heißt Prozeß P. Ein Algorithmus wird einmal beschrieben und kann in etlichen Inkarnationen (als Prozesse Pl,P2..Pk) gleichzeitig auf einem System laufen. Die Algorithmen Al,A2,..,An , deren Interkommunikation und der Kommunikation mit anderen Systemen bilden das System. Dieses läßt sich dann stets in eine Menge von Prozessen Pl,P2,..,Pm aufteilen. Dabei kann ein Prozeß aus mehreren Algorithmen oder ein Algorithmus aus mehreren Prozessen bestehen.

Bei der Spezifikation in PNDL wird jeder Algorithmus durch verschiedene Sprachelemente beschrieben und kann noch weiter untergliedert werden (entsprechend 'Program Refinement'). Diese Darstellungsweise bildet das statische Petri Netz. Das Netz kann nun analysiert werden, wobei gegebenenfalls enthaltene Fehler (Verklemmungen) eliminiert und zeitkritische Pfade erkannt werden. In einem weiteren Schritt werden bei verschiedenen Datenkonstellationen an den Eingängen alle Prozesse (alle Inkarnationen) analysiert (dynamisches Petri Netz). Dabei erhält man zusätzlich Hinweise auf ein optimales Scheduling, die dem Prozeßmanager (Teil des Steuerwerks) des späteren Systems als Prozeßverteilungshinweise dienen können. Das Ziel ist dabei eine gleichmäßige Auslastung von Prozessoren, Sicherheitsaspekte oder die Minimierung der systeminternen Kommunikation /Hoffmann Steinmetz 83/. Dann können abhängig oder unabhängig von der Petri Netz Ebene Simulationsverfahren angewendet werden um die Optimierung der Realisierung zu erreichen. .

× •

Ein Petri Netz Modell und die Petri Netz Beschreibungssprache PNDL

Es existieren in der Literatur eine Vielzahl von grundsätzlich unterschiedlichen Modellen, die alle unter dem Oberbegriff 'Petri Netz' zusammengefaßt sind. Die beiden einfachen und allgemein bekannten Arten sind die Bedingungs-Ereignis (B/E) und die Stellen-Transitionen (S/T) Netze /Reisig 82/. Ein Bedingungs-Ereignis Netz B/E ist ein 3-tupel: B/E = (Netz,Init,Schalt) mit:

-Netz = (Bedingungen, Ereignisse, Flußrelation) ist ein 3-tupel und ein gerichteter Graph mit:

-Bedingungen: Nichtleere endliche Menge B von Bedingungen (Knotentyp 1 des Graphs).

-Ereignisse: Nichtleere endliche Menge E von Ereignissen (Knotentyp 2 des Graphs), E und B seien disjunkte Mengen.

-Flußrelation: Teilmenge von $\{B \times E\} \bigcup \{E \times B\}$ (Kanten des Graphs ohne Schlingen).

-Init : Initialisierung des B/E Netzes = B \rightarrow {0,1} . Marken werden auf die Stellen abgebildet.

-Schalt : Menge der Schaltregeln. Siehe hierfür z. Bsp. /Reisig 82/ oder /Peterson 81/.

Die Vorteile der B/E Netze liegen in ihrer einfachen Analysierbarkeit. Soll dagegen ein komplexes Problem als B/E System darstellt werden, so bedeutet dies einen großen, wenn nicht gar unmöglichen manuellen Kodierungs- und Abstraktionsaufwand. Die Fehlerquelle bei der Umsetzung des Problems in das B/E Netz ist nicht vernachlässigbar! Diese Problematik ist vergleichbar mit der Programmierung in höheren Programmiersprachen und in Assemblersprache. Deshalb wurde für komplexere Problemstellungen ein 'höheres' Petri Netz Modell mit einer eigenen Beschreibungssprache PNDL entwickelt. Eine der Erweiterungen gegenüber B/E und S/T Netzen liegt wie bei Prädikat-Ereignis (Pr/E) Netzen /Genrich Lautenbach 79/ in der Individualität der Marken.

Durch eine Petri Netz Transformation wird ein B/E Netz generiert oder es werden die nicht abbildbaren Konstrukte des beschriebenen Systems aufgezeigt (Bild 2). Der Anwender wird dann die Analysemöglichkeiten des B/E Netzes ausnutzen, indem er versucht, seine Systembeschreibung an der nicht transformierbaren Stelle geeignet zu ändern.



Bild²2

In PNDL werden die Marken wie Variable einer Programmiersprache mit Namen und Typ erfaßt:

Zum Beispiel: (Terminalsymbole in Großschrift)

TOKEN TYPE messagetype IS RECORD tln IS INTEGERSUBSET (1..1000)

signal IS ELEMENTOFSET (wait,do)

END RECORD

VAR inmessage, outmessage IS messagetype

Den Stellen (Bedingungen) werden Attribute zugeordnet: Namen, minimale und maximale Verweilzeiten, Warteschlangenorganisationen, Prioritäten für bestimmte Marken, Kapazitäten in Abhängigkeit der Markenwerte und Markentypen.

Den Aktionen (Übergänge) werden Stellen und Transitionseigenschaften zugeordnet:

Namen, minimale und maximale Ausführungszeiten, Boolsche Verknüpfungen als Schaltregeln (B/E, S/T und Pr/E Netze haben nur UND Verknüpfungen). .
Zum Beispiel:

ACTIONS ... getsignal: ACTIVATED FROM signalfifo NAMED inmessage IF (inmessage.signal = wait) OR (inmessage.tln = 110) THEN (SEND police VAR outmessage = inmessage) AND (GET FROM other TYPE ANY)

Gegenüber anderen Petri Netz Modellen (zum Beispiel B/E Netzen) wurde so eine komfortablere Beschreibungsmethode entwickelt. PNDL wurde für Untersuchungen der Prozeßinterkommunikation mit Nachrichten entwickelt /Hoffmann Steinmetz 83/. Andere Entwicklungen mit anderen Zielsetzungen sind: /Cap Dspl 82/, /Epsilon 82/, /Fun 83/ und /Necon 83/.

Beispiel: Sicherungsanlage

Die Sicherungsanlage (Bild 3) bestehe aus n Fühlern (F), die z.Bsp. unbefugtes Betreten von Räumen feststellen. Ein Prozeß Fühler (PF) lege den Schwellenwert zwischen Normal- und Ausnahmezustand fest, er aktiviere bei Schwellenwertüberschreitung einen Prozeß Koordinator (PK) (Bild 4), indem er ihm die Meldung <alarm> sendet. Außerdem reagiere der Prozeß Fühler auf Anfragen des eigenen Prozeß Koordinators. Die Prozeß Koordinatoren kommunizieren untereinander, um das Ausmaß des aufgetretenden Alarmfalls zu beurteilen. Bei entsprechender Alarmsituation werde ein Prozeß Polizei (PP) durch eine Meldung aktiviert, dieser gehöre aber nicht mehr zum Sicherungssystem. Der Prozeß Überwacher (PÜ) diene der Wartung und Initialisierung des Systems.



Bild 4

Bild 3

Ausschnitte aus dem in PNDL beschriebenen System: SYSTEM sicherungsanlage; ANALYSE INVARIANTEN, SHOW_DEADLOCKS, REACHABILITYGRAPH; TOKEN TYPE pf_pk_type IS ELEMENT_OF_SET (Alarm, Antwort_Alarm, Antwort_Ruhe) ... VAR inmessage IS pf_pk_type ... ENVIRONMENT ... (* Fühler, Mensch am PÜ,... *) ... PROCESS pk1 PLACES ... entrypoint = ... ACTIONS ... alarm_von_PF_auswertung: ACTIVATED_FROM entrypoint NAMED inmessage IF firstalarm THEM SEND PROCESS pk2, .. pkn PORT entrypoint VAR outmessage = auch_alarm? ... END SYSTEM sicherungsanlage.

Durch die Netz Transformation (Bild 2) ergibt sich ein B/E Netz. Der für die Analyse interessante Teil sei hier beschrieben (die Nachbedingungen haben für das Beispiel keine Bedeutung und werden deshalb nicht explizit formuliert): .. IF (frei UND Eigenalarm) THEN (sende auch Alarm? UND warte); IF (warte UND empfange Antwort) THEN (werte aus UND frei); IF (frei UND empfange auch Alarm?) THEN Klärung; IF Klärung THEN (sende Antwort UND frei); .. .



Die Zustände dieses Teilsystems lassen sich durch einen Zustandsvektor darstellen: Eigenalarmi frei Auswertung Antwort Warte Klärung auch Alarm?

Mit der in Bild 5 gezeigten Initialisierung ergibt sich dieser Ausschnitt (Bild 6) aus dem Zustandsgraphen. Die Analyse zeigt die enthaltene Verklemmung auf: Wenn beide Prozesse durch einen Eigenalarm aktiviert sind und sich gegenseitig abfragen (auch Alarm?), dann sind die Bedingungen Wartel und Warte2 gesetzt. Nur durch ein anderes Netz oder zusätzliche Marken läßt sich diese Verklemmung auflösen. Im Fall des Sicherungssystems wurden Zeitbedingungen (Timer) eingeführt.



dingungen (Timer) eingeführt. Durch die bisher unerwähnten Sicherheitsanforderungen werden jeweils ein Prozeß_Fühler und der dazugehörige Prozeß_Koordinator in einer Rechnereinheit zusammengefaßt. Der Prozeß_Polizei ist extern, der Prozeß_Überwacher kann auf einer PF-PK Rechnereinheit laufen oder eine eigenständige Rechnereinheit beanspruchen.

Literatur

/Bernstein 66/ : Analysis of Programs for Parallel Processing,

- IEEE Transactions on Computers, Oktober 1966, S.757-763.
- /Brent 74/: The Parallel Evaluation of General Arithmetic Expressions, Journal of the ACM, April 1974, S.201-206.

/Cap Dspl 82/: Rammig, CAP/DSPL Preliminary Language Reference Manual, Forschungsbericht Nr.129, Abt.Informatik, Universität Dortmund, 1982.

/Coleman Hughes Powell 81/: A Method for Syntax Directed Design of Multiprograms, IEEE Transactions on Software Engineering, März 1981, S.189-196.

/Ecker 77/: Organisation von parallelen Prozessen, Reihe Informatik/23, BI Mannheim, 1977.

/Epsilon 82/: Jensen, Kyng, EPSILON a System Description Language, Aarhus University, 1982.

/Fun 83/: Godbersen, Simulation with FUN, Angewandte Informatik, Mai 1983, S.213-219.

/Genrich Lautenbach 79/: The Analysis of Distributed Systems by Means of Predicate/Transition Nets, in Lecture Notes in Computer Science 84, Springer Verlag, Berlin Heidelberg New York, 1979.

/Hoffmann Steinmetz 83/: Einige Überlegungen zur Minimierung der systeminternen Kommunikation in informationsverarbeitenden Systemen unter Anwendung von Petri-Netzen, Forschungsbericht Nr.113, Institut für Übertragungstechnik und Elektroakustik, T.H.Darmstadt, 1983.

/Necon 83/: Scheschonk, der Netz-Editor des NECON-Systems, Workshop: "Petrinetzanwendungen bei der Systemkonstruktion, Bad Honnef, Januar 1983.

/Peterson 81/: Petri Net Theory and the Modeling of Systems, Prentice Hall, Englewood Cliffs New York, 1981.

/Reisig 82/: Petrinetze, Springer Verlag, Berlin New York Heidelberg, 1982.

. ¥

IM VORTRAGSRAUM DER MENSA DER UNIVERSITÄT HANNOVER CALLINSTRASSE **23 3**000 HANNOVER

13. MARZ 1984

DIGITALE SIGNALVERARBEITUNG

DES DFG-SCHWERPUNKTPROGRAMMS

KOLLOQUIUM