Real-Time Database Scheduling: Design, Implementation, and Performance Evaluation

Sang H. Son, Prasad Wagle, and Seog Park†

Department of Computer Science, University of Virginia, Charlottesville, VA 22903, USA

† Department of Computer Science, Sogang University, Seoul, Korea

ABSTRACT

A real-time database system has timing constraints associated with transactions and the database. To ensure that a real-time database system completes as many transactions as possible without violating their timing constraints, its scheduling strategy should be dynamic and use information about the timing constraints associated with transactions and the database. This paper presents an intelligent dynamic scheduling algorithm for transactions in real-time database systems. The scheduling algorithm uses timing information about transactions and the database to enhance the system's ability to meet transaction deadlines. The scheduling algorithm is implemented in a simulated pulse detection system, and its performance is demonstrated by a series of experiments.

1. Introduction

A Real-time database system is a database system that supports real-time computing. *Real-time computing* is a type of computing where the correctness of the system's response depends not only on the logical result of the computation, but also on the time at which the results are produced [Stan88]. The timing constraint on the system's response is called *deadline*. Traditional real-time systems have concentrated on systems which have hard deadlines. If a system misses a hard deadline, the consequences can be disastrous. On the other hand, if the system misses a soft deadline, there may still be some value for computing the response of the system. Real-time systems are assuming an increasingly important role in our society. Examples of current real-time computing systems are command and control systems, aircraft avionics, robotics, network management, and program trading.

This work was supported in part by the Office of Naval Research under contract # N00014-88-K-0245, By IBM Federal Systems Division, and by the Center for Innovative Technology under contract # CIT-INF-90-011.

DATABASE SYSTEMS FOR ADVANCED APPLICATIONS '91
Ed. A. Makinouchi
©World Scientific Publishing Co.

Most of the complex real-time computing applications need to access large amount of data. Thus, we need database systems which are cognizant of the requirements of real-time computing, i.e. *real-time database systems* [Son88]. Transactions in real-time database systems must be scheduled in such a way that they can be completed before their corresponding deadlines expire. For example, both the update and query on the tracking data of a missile must be processed within the given deadlines: otherwise, the information provided could be of little value. In such a system, transaction processing must satisfy not only the database consistency constraints but also the timing constraints.

Concurrency control algorithms in database systems control the interaction among the concurrent transactions in order to prevent them from destroying the consistency of the database. Serializability is a widely accepted notion of the definition of correctness for concurrency control in database systems. A *scheduler* in database systems is entrusted with the task of enforcing the serializability constraints. It accepts database operations from transactions and schedules them appropriately for the data manager. To satisfy both consistency and real-time constraints in real-time database systems, there is a need to integrate concurrency control algorithms with real-time scheduling algorithms.

Real-time scheduling algorithms address the problem of meeting the specified timing constraints. Satisfying the timing constraints of real-time systems demands the scheduling of system resources according to some well-understood algorithms so that the timing behavior of the system is understandable, predictable, and maintainable. The goal of most scheduling problems is to find optimal static schedules which minimize the response time for a given task set. In many real-time systems, however, there is generally no incentive to minimize the response time other than meeting the deadline. Real-time systems are often highly dynamic requiring on-line, adaptive scheduling algorithms. In these cases, the goal is to schedule as many jobs as possible, subject to meeting the task timing constraints. Alternative schedules and/or error handlers are

required and must be integrated with the on-line scheduler.

A real-time database scheduling algorithm must maximize both concurrency and resource utilization subject to three constraints: data consistency, transaction correctness, and transaction deadlines [Son90c]. This requirement can be satisfied by the intelligent integration of two forms of scheduling protocols: concurrency control protocols and real-time scheduling protocols. It is not very straightforward to integrate the two protocols, since while the term "scheduling" is used in both database systems and real-time system schedulers, the processing of scheduling has different assumptions and objectives in the two environments [Son90]. The notion of a schedule as formalized in database systems does not include time, so there is no way to measure the timeliness of a database schedule. On the other hand, the notion of a schedule as formalized in real-time systems does not enforce any consistency metric on data resources.

Satisfying timing constraints while preserving data consistency requires scheduling and concurrency control protocols to accommodate timeliness of transactions as well as data consistency requirements. In real-time database systems, timeliness of a transaction is usually combined with its criticality to take the form of the priority of the transaction. Therefore, proper management of priorities and conflict resolution in real-time transaction scheduling are essential for predictability and responsiveness of real-time database systems. In this paper, we present a dynamic scheduling algorithm developed for real-time database systems. The algorithm employs the notion of *dynamic priority* based on a number of important parameters that affect scheduling decisions. We first discuss a few basic concepts and issues associated with real-time database scheduling in the next section. The proposed algorithm is then discussed in detail, together with experimental study which illustrates the performance of the algorithm.

2. Basic Concepts

2.1. Validity Constraints

Deadlines are timing constraints associated with transactions. There exist another kind of timing constraints which are associated with transactions and data objects in the database. In a database, there may be some data objects which get old or out-of-date if they are not updated within a certain period of time. To quantify this notion of age we associate with each data object a degree of validity which decreases with time. The validity curve associated with each data object is a plot of the degree of validity of the data object with respect to the time elapsed after the object was last modified.

If w is the time of last modification of a data object, we can calculate the validity of the data object at time t from its validity curve. A transaction may require that any data object it reads does not have a degree of validity less than the minimum degree of validity. This constraint could be either hard or soft, like deadlines. Scheduling decision could be made more intelligent by incorporating this validity information about transactions and data objects they read.

2.2. Static versus Dynamic Scheduling

It is possible to statically guarantee real-time constraints by pre-calculating all possible schedules of transactions off-line. There are two reasons why this approach is infeasible. First, the task of finding all possible schedules of transactions is NP-hard [Stan90]. Therefore, the task becomes computationally intractable when there are a large number of simultaneously active transactions. Second, the demands on a real-time database system can change frequently. For example, aperiodic transactions, by their very nature, can be activated at unpredictable times. Therefore, a dynamic scheduling strategy is needed to make the system more flexible. Also, to make "intelligent" scheduling decisions, the scheduling strategy should use as much timing information as possible about transactions and the data objects they access.

A scheduler in database systems accepts database operations from transactions and schedules them appropriately for the data manager [Bern87]. In a distributed system, each site has its own scheduler which can receive database operations from transaction managers at different sites. In conventional database systems, the scheduler is entrusted with the task of enforcing the serializability constraints. In real- time database systems, it is also necessary to take into account the timing constraints associated with the transactions and the database while making scheduling decisions.

3. The Scheduling Algorithm

In this section, we present a dynamic scheduling algorithm for transactions in real-time database systems. The scheduling strategy uses timing and validity information about transactions and data objects to calculate dynamic priorities of transactions. These priorities are then used to make scheduling decisions at all places where transactions contend for scarce resources.

3.1. Information for Intelligent Scheduling

A transaction can be represented as a tuple (SP, RS, WS, A, D, E, V_{min}). The elements of the tuple are described below.

(1) System priority (SP):
: This is the static component of the dynamic priority associated with a transaction. It is a measure of the value to the system of completing the transaction within its timing constraints. For example, transactions dealing with emergency situations should have a higher priority than routine transactions.

(2) Read set (RS):
: This is the set of data objects which the transaction reads.

(3) Write set (WS)
: This is the set of data objects which the transaction writes.

(4) Arrival time (A):
: This is the time at which the transaction arrives in the system.

(5) Deadline (D):
: This is the time before which the transaction has to finish its execution. The transaction specifies whether the deadline is hard or soft.

(6) Runtime estimate (E):
: This is the estimate of the processing time required by a transaction. This includes the time required for CPU as well as I/O operations.

(7) Minimum Validity(V_{min}):
: This is the minimum degree of validity required of all objects read by the transaction. The transaction specifies whether this validity constraint is hard or soft.

The above information about the transaction is available to the system before the transaction is started and remains constant throughout the transaction execution. Since the scheduling strategy is dynamic, it needs information about the transaction which varies with time. The information which varies with time is described below.

(8) Read set validity(RSV):
: This is the degree of validity of data objects in the transaction's read set. The degree of validity of a data object can be calculated from its validity curve. The validity curve of a data object defines a function of the degree of validity of the data object with respect to the time elapsed after the data object was last modified. Therefore, if we know the time the object was last modified, we can calculate the degree of validity of the data object at the current time from the validity curve.

(9) Processing time(P):
: This is the processing time already received by a transaction. This includes the time required for CPU as well as I/O operations.

(10) Current time(C):
: This is the time at which the scheduling decision is made.

3.2. Design Issues

Before implementing any scheduling strategy, it is important to consider the overhead it requires. Obviously, a complicated scheduling strategy requires more time. This factor can be crucial in deciding whether it is of any practical benefit to use the extra information about transactions and the database in the scheduling strategy.

For instance, if the database is disk-resident and the transactions are I/O intensive, the time required for I/O operations would be large compared to the time required for doing CPU operations. In that case, it would not make a big difference whether or not we use a complicated scheduling policy at the CPU level. The bottleneck in this case would be the data objects and it would be imperative to schedule the database operations in an intelligent way. But if the database is memory resident and the transactions are CPU intensive then it would become necessary to use the extra information about transactions in the scheduling decision at the CPU level. Example 1 shows a scenario which illustrates a situation where an intelligent scheduling strategy at the CPU level would be helpful.

Example 1: Assume that transactions execute CPU and I/O instructions alternately. Let the time required for one session of CPU computation be 10 time units and the time required for one I/O operation be 2 time units (if there is no blocking). Let the transactions to be scheduled (T_1 and T_2) have the characteristics given below. This situation can arise if both T_1 and T_2 wait for some other transaction (say, T_3) to release a data object d_1. Assume that T_3 releases the data object at time 5. Thus, the scheduling decision has to be made at time 5.

Transaction	A	E	D	Operations
T_1	0	12	30	read(d_1)
T_2	5	12	20	read(d_1)

According to the FCFS scheduling strategy, T_1 is scheduled first and it completes at time 17. T_2 starts at time 15, but since it requires 12 time units to complete, it misses its deadline at time 20. The execution sequence is shown in Fig. 3.1.

If the system is intelligent enough to follow the elaborate scheduling strategy to be discussed in Section 3.3, T_2 would be scheduled first. (According to the least slack method of assigning priorities, T_2 has a higher priority than T_1, because the slack of T_2 is less than the slack of T_1.) In that case both transactions would meet their deadlines as shown in Fig. 3.2.

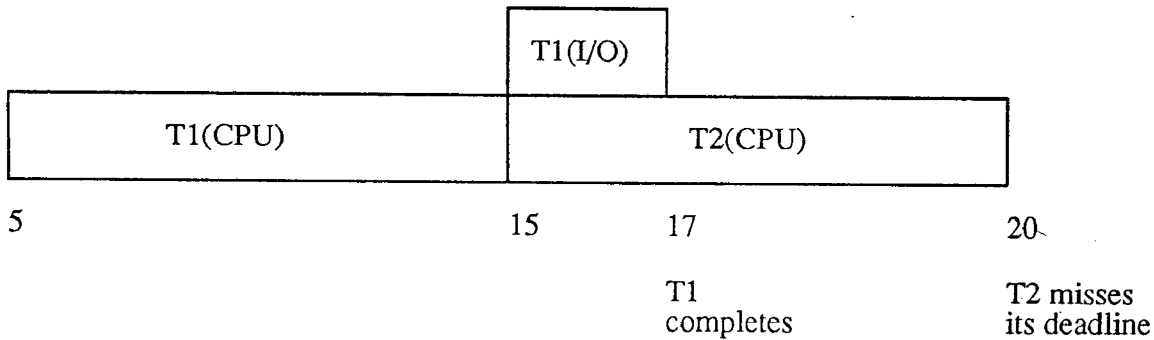

Fig. 3.1 FCFS Scheduling

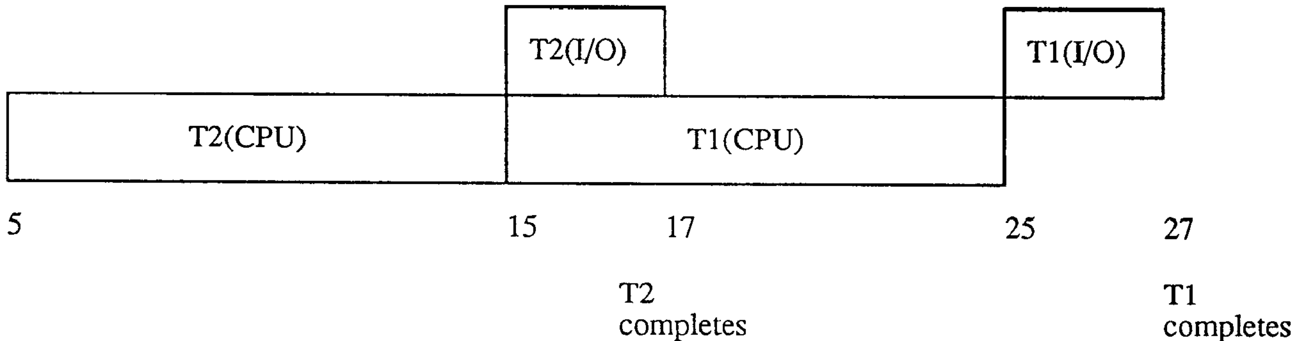

Fig. 3.2 Intelligent Scheduling

An issue involved in designing a scheduling strategy is whether or not to allow preemption. The scheduling decision at the CPU level normally allows preemption. However, if we allow preemption at the data object level, we may have to abort the preempted transaction for maintaining consistency of the database. When preemption is not allowed, the scheduling decision has to be made whenever a transaction relinquishes a resource or when a transaction requests a resource which is not being used. When preemption is allowed the scheduling decision has to be made whenever a transaction either requests or relinquishes a resource.

3.3. A Real-Time Database Scheduler

A scheduling algorithm for transactions in real-time database systems can be decomposed into three components: (1) determining eligibility, (2) assigning dynamic priorities, and (3) making scheduling decisions on granting the resource. Each component of the proposed algorithm is described in detail in the following sections.

3.3.1. Determining Eligibility

Before making a scheduling decision we have to decide whether the transactions involved are eligible for scheduling i.e. whether it is of any use to the system to start processing those transactions. If a transaction is ineligible for scheduling we abort it immediately.

We assume that, if a transaction misses a hard deadline, it is ineligible for scheduling and should be aborted. If a transaction misses a soft deadline, it is still eligible for scheduling. We also check whether it is possible for the transaction to finish before its deadline:

$$(D - C) \geq (E - P)$$

If it is not possible, and the deadline in question is hard, we consider the transaction ineligible for scheduling. However, if the deadline is soft, the transaction remains eligible for scheduling.

The steps taken in incorporating validity constraints are similar to those taken for deadlines. If a transaction misses a hard validity constraint then it is ineligible for scheduling and should be aborted. If the validity constraint missed is soft, then we continue executing the transaction at a different priority. We also

check, for each data object read by the transaction, whether its degree of validity is greater than the minimum validity level expected by the transaction. If that is not the case, and the validity constraint of the transaction is hard, we consider the transaction ineligible for scheduling. However, if the validity constraint is soft, the transaction remains eligible for scheduling.

3.3.2. Assigning Dynamic Priorities

The dynamic priority of a transaction is a number calculated by the scheduler while making the scheduling decision. It is a measure of the importance, to the over-all goals of the system, of scheduling that transaction before others at that point in time [Son89]. Since this measure may change with time, it has to be calculated dynamically when a scheduling decision has to be made.

Dynamic priority (DP) is a weighted sum of the following factors:

(1) System priority (SP): It is the static component of dynamic priority.

(2) Slack with respect to deadline (SDL): It is the amount of time the transaction can be delayed and still meet its deadline. It is calculated as follows:

$$SDL = D - C - (E - P)$$

(3) Slack with respect to minimum validity constraints (SV): It is the amount of time the transaction can be delayed and still be completed without violating its validity constraints.

$$SV = Min\ \{\ t \mid \text{For each data object } d \text{ read with } V_d(T + t) \geq V_{min}\ \}$$

where, $V_d(T + t)$ is the degree of validity of an object d at time $(T + t)$, assuming no updates between time T and $(T + t)$.

Dynamic Priority (DP) is calculated as follows:

$DP := DP_1 + DP_2 + DP_3$
where,
$DP_1 := w_1 * SP$
$DP_2 := w_2 * SDL$
$DP_3 := w_3 * SV$

The factors involved in determining the dynamic priority of a transaction have constraints closely related to the characteristics of real- time transactions. First, $w_1 > 0$, since if SP increases, DP should increase. Also, if SDL > 0 then $w_2 < 0$, since if SDL decreases then DP should increase. If SDL < 0, then the transaction has already missed its deadline. Note that since the transaction is still eligible for scheduling, the deadline missed must have been soft. At this point, there are two options available to us. We could reason as follows: Since the transaction has missed its deadline (soft), it should be finished as soon as possible, and hence its priority must be increased. In that case, $w_2 < 0$. However, we might reason that since the transaction has already missed its deadline, its priority should be reduced so that it does not interfere with other transactions in the system which are nearing their deadlines. In that case, $w_2 > 0$. Similar discussion applies to w_3 and SV.

The relative values of w_1, w_2, w_3 depend on the high level goals of the system. For example, some systems may aim at minimizing the number of transactions that miss their deadline, in which case w_1 would not be very high. Some systems might require that absolutely none of the higher priority transactions be aborted, in which case w_1 would very high. Example 2 shows a scenario which illustrates that a scheduling strategy at the CPU level taking validity constraints into account does prevent unnecessary aborts of transactions.

Example 2: Assume that transactions use the CPU and do I/O operations alternately. Let the time required for one session of CPU computation be 10 time units and the time required for one I/O operation be 2 time units (if there is no blocking). Let the transactions to be scheduled (T_1 and T_2) have the characteristics given below.

Transaction	A	E	D	V_{min}	Operations
T_1	0	12	30	100%	read(d_1)
T_2	0	12	25	50%	read(d_1)

Let the validity curve for object d_1 be as shown in Fig. 3.3, and the time it was last modified be 0. Let the weights w_2 and w_3 for calculating dynamic priorities be -1. This implies that, in the formula for calculating dynamic priorities, the slacks with respect to deadline and validity constraints have the same weight.

Case A) Assume that validity constraints are not considered:
In this case, $DP := DP_1 + DP_2$. The slack of T_1 with respect to deadline is 18. The slack of T_2 with respect to deadline is 13. Therefore,

$$DP_2(T_1) = -18 \text{ and } DP_2(T_2) = -13.$$
$$\text{i.e. } DP_2(T_2) > DP_2(T_1).$$

Assuming equal system priorities, $DP(T_2) > DP(T_1)$, implying that T_2 would be scheduled first. The execution would proceed as shown in Fig. 3.4. T_2 would finish its execution at time 12, and then T_1 would start. However, at time 20 the validity of object d_1 would be 50%. This would violate the validity constraint of T_1, which would have to be aborted.

Case B) Assume that validity constraints are considered:

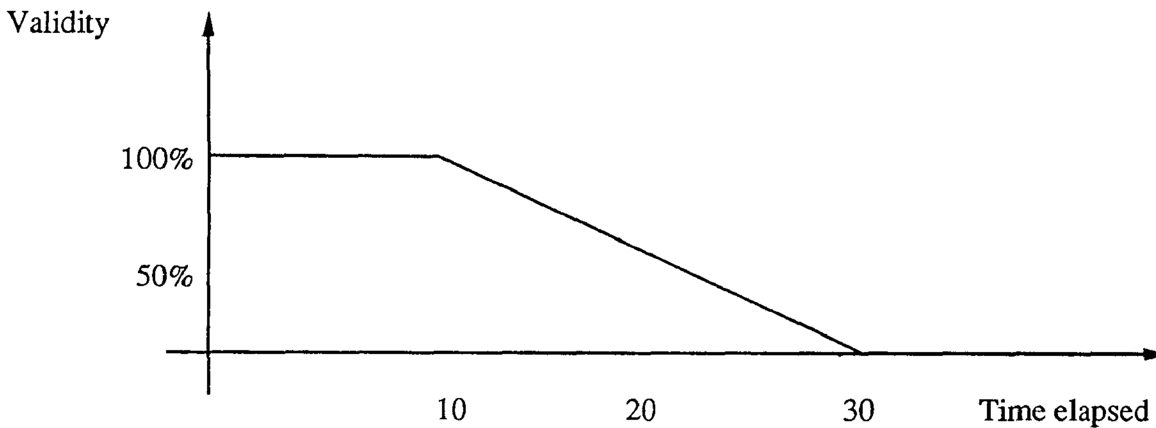

Fig. 3.3 Validity Curve

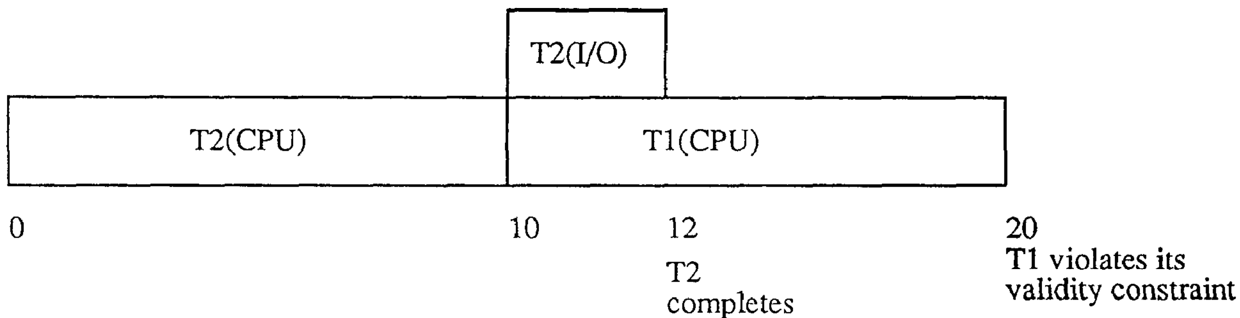

Fig. 3.4 Validity constraints ignored

In this case, DP := $DP_1 + DP_2 + DP_3$. The slack of T_1 with respect to validity constraints is 10. The slack of T_2 with respect to validity constraints is 20. Therefore,

$$DP_3(T_1) = -10 \text{ and } DP_3(T_2) = -20.$$
$$\text{i.e. } DP_2(T_1) + DP_3(T_1) > DP_2(T_2) + DP_3(T_2).$$

Assuming equal system priorities, $DP(T_1) > DP(T_2)$, implies that T_1 would be scheduled first. The execution would proceed as shown in Fig. 3.5. At time 10 the validity of object 1 would be 100%, satisfying T_1's validity constraints. T_1 would finish its execution at time 12, and then T_2 would start. At time 20, the validity of object d_1 would be 50%, satisfying T_2's validity constraints. Thus T_2 would finish its execution at time 22. This example shows that incorporating validity constraints in the scheduling strategy does prevent transactions from being aborted unnecessarily.

3.3.3. Making Scheduling Decisions

The way a scheduling decision is made depends on whether preemption is allowed or not. In the following discussion we assume that the transactions considered have already passed the eligibility test. Let us consider the scheduling algorithms for the two cases:

Case 1. No preemption.

There are more than one transactions requesting a resource and we have to decide the transaction which should be granted the resource. In this case we grant the resource to the transaction with the highest dynamic priority.

Case 2. With preemption.

There is a transaction currently holding a resource and there is another transaction requesting the same resource. We have to decide whether or not to preempt the transaction holding the resource and grant the resource to the transaction requesting it.

Let two transactions T_h and T_r be the resource holder and the requester for a shared resource, respectively. Let $P(T_h)$ and $P(T_r)$ be dynamic priorities of the two transactions. Let $P(T_h^p)$ be the priority of T_h were it

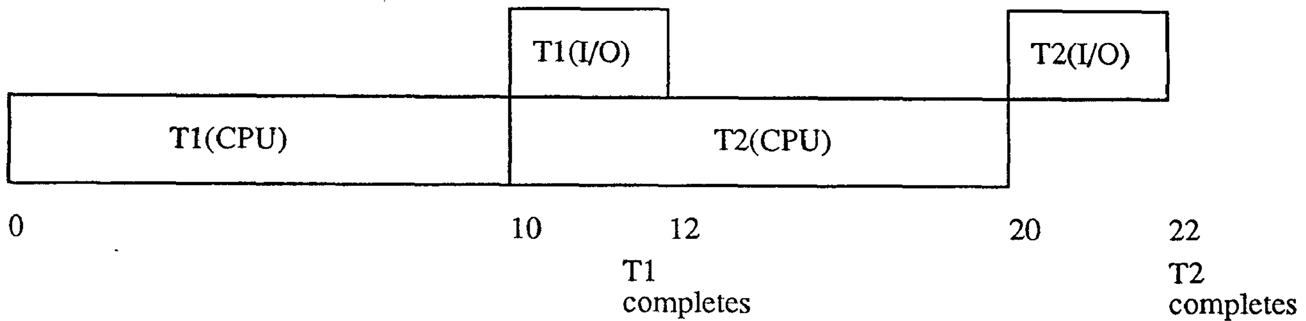

Fig. 3.5 Validity constraints considered

to be preempted by T_r. The algorithm works as follows. It is based on a resolution policy of the conditional restart in [Abbo88].

```
if P(Tr) >  P(Th) and P(Tr) > P(Th^p)  then
     if RemainingTime(Th) > Slack(Tr)
               then preempt Th;
     else   block Tr
                    Th inherits the priority of Tr
     endif;
else block Tr
endif;
```

where $\text{RemainingTime}(T_h) = E(T_h) - P(T_h)$ and $\text{Slack}(T_r) = \text{min. of SDL and SV of } T_r$.

4. Experimentation

4.1. Need for a Real-Life Application

The research on real-time transactions scheduling is still in its infancy. There exists no formal theoretical framework to analyze the performance of the existing scheduling algorithms. For this reason, it is necessary to make an experiment to compare the performance of different scheduling strategies.

Until now, none of the algorithms proposed in previous studies have been evaluated in real systems. [Abbo88] and [Abbo89] present experimental results based on simulation, whereas [Huan89] presents an integrated approach to study real-time transaction processing on a testbed system. In these studies semantically meaningless transactions are randomly generated with random system priorities, resource requirements, and timing constraints. The disadvantage of this approach is that it does not give the researcher a true feel for real-life problems. Also, for any scheduling strategy to be used in industry, it has to be supported by an extensive round of experimentation with a *real-life* application.

4.2. A Pulse Detection System

A pulse detection system is an example of a real-time database system [Hale89]. It is used to detect and track external objects by means of pulses (radar or sonar) received from them. The pulse detection system maintains information about each object in reality in a database of emitter files. It contains a number of simultaneously active transactions with different system priorities, timing constraints, and resource requirements.

The pulse detection system we have implemented runs on a SUN 3/75 Workstation with a color monitor. It is based on the scenario of a battleship surrounded by airborne enemy objects like aircrafts or missiles. It consists of two windows: the *reality* window, and the *operator's console* window

The reality window consists of a stationary battleship at its center and the surrounding enemy objects. Each object has a position and a velocity associated with it. An object is implemented as a process which calculates the new position of the object and displays it in the reality window. The reality window is managed by two modules: `Object` and `Reality`. The module `Object` is responsible for creating objects in reality, continuously updating their positions and detecting collisions. The module `Reality` is responsible for creating the reality window. It has a procedure called `GetPulseData` which simulates the operation of a radar by getting new pulse data of an object in reality.

The operator's console window displays the operator's view of reality as maintained by the pulse detection system. It is supposed to display the most current positions of enemy objects in reality. The operator's console window is managed by the modules: `Detect` and `EmitterFile`. The module `EmitterFile` maintains an emitter file to store information corresponding to each enemy object in reality.

The `Detect` module contains three periodic and two aperiodic transactions. Each transaction is implemented as a process. The following are the periodic transactions with a brief description of what they do.

(1) `Track`: It calls `Reality:GetPulseData` to get a new pulse data of an object in reality. It scans all the emitter files to find an emitter file which correlates with the pulse data received. If it finds such an emitter file, it updates it; else it creates a new emitter file with that pulse data.

(2) `Clean`: It periodically scans the emitter files and deletes emitter files which haven't been updated for a predetermined amount of time assuming that the object which they represents have been destroyed.

(3) `Operator Interaction`: This transaction accepts operator commands. For example, an operator may query the database to find more information about an emitter file, or he may start a transaction to shoot an enemy object.

The operator interaction transaction, in turn, can start two aperiodic transactions, which are:

(1) `Display Information`: This transaction displays information about the object chosen by the operator.

(2) `Shoot Object`: This transaction shoots a missile at the object chosen by the operator.

4.3. The Simulation Feature of the System

Since the pulse detection system we have implemented runs in a simulated environment, it is very important that the experimenter has control over the relative speeds of the transactions being executed and the amount of time a transaction needs to use a resource. To provide this capability, we have created a special software module, called the `Simulation` module. If two or more processes want to use a resource at the same time, a decision has to be made in the `Simulation` as to which process should be granted the resource. This decision is made considering the attributes associated with the different processes according to some scheduling strategy.

Currently each process contending for a shared resource has the following attributes: (1) System priority; (2) Arrival time; (3) Deadline; (4) Run-time estimate; (5) Processing time it has received; and (6) Minimum validity of the data it reads.

The system allows the researcher to choose the scheduling strategy followed, with or without preemption, and examine its effects on the pulse detection system. Currently the following strategies, with or without preemption, are supported: (1) First Come First Served; (2) System Priority; (3) Earliest Deadline First; (4) Least Slack First; and (5) A variant of the scheduling strategy presented in the previous chapter, which will be henceforth referred to as the *Combination* strategy. The *Combination* strategy uses the system priority (SP) and the slack with respect to deadline (SDL) while making its scheduling decisions.

Our intention is to show that the performance of the pulse detection system can be enhanced by the use of intelligent scheduling algorithms. The performance of a scheduling strategy can be judged in two ways: (1) by the visual behavior of the simulated pulse detection system; or (2) by the information about successful completion of transactions displayed each time the scheduling strategy is changed.

4.4. Assumptions

The following are the assumptions made about the simulations.

(1) The consistency of the database is maintained using exclusive locks which are non-preemptible. A more efficient concurrency protocol would be the priority ceiling protocol using shared locks [Sha91].

(2) All transactions have hard timing and validity constraints. When a periodic transaction or an instance of a periodic transaction is started, the run-time estimate and the deadline parameters of the transaction are set.

(3) A transaction cannot use more than one resource at the same time.

4.5. Results of Experiments

To make the differences in the performance of the different scheduling strategies obvious two periodic dummy transactions were added to the system. This is justified, since, real-time systems do have certain background tasks which are not directly connected to the real-time application. The following are the dummy transactions and their characteristics:

(1) Dummy1: Low system priority, Tight deadline.

(2) Dummy2: High system priority, Loose deadline.

The simulation results can be grouped into three cases:

(1) Case 1: Dummy1, but not Dummy2, is activated.

(2) Case 2: Dummy2, but not Dummy1, is activated.

(3) Case 3: Both Dummy1 and Dummy2 are activated.

To quantitatively evaluate the results of a particular scheduling strategy, we calculate its *figure of merit* as follows:

$$\text{figure of merit} = \sum_{\text{Transactiontypes}} (\%\ \text{success})(\text{System Priority})$$

where

$$\%\ \text{success} = \frac{(\text{No. of successful completions})}{(\text{No. of instances started})}$$

The system priorities of the different transaction types is shown in the following table.

Transaction Type	System Priority
Track	2
Clean	1
User Interaction	3
Shoot	3
Display Information	3
Dummy1	0
Dummy2	2

The simulation results based on the above performance metric are summarized in the following tables. The entries in the table are either quantitative (figures of merit) or qualitative (good or bad). The qualitative assessment is done by taking into account the visual behavior of the system.

4.5.1. When Preemption is Allowed

Quantitative Assessment:

Scheduling Strategy	Case 1	Case 2	Case 3
FCFS	300	500	511
System Priority	1159	504	500
Earliest Deadline First	1056	1220	828
Least Slack	305	1036	306
Combination	1114	1350	1194

Qualitative Assessment:

Scheduling Strategy	Case 1	Case 2	Case 3
FCFS	bad	bad	bad
System Priority	good	bad	bad
Earliest Deadline First	bad	good	bad
Least Slack	bad	good	bad
Combination	good	good	good

We observe that the FCFS strategy performs poorly in all the three cases. This is because the FCFS strategy does not possess the requisite intelligence to prevent the dummy transactions from using the resources. This causes the more important transactions to miss their deadline.

In Case 1, the dummy transaction activated has low priority but a tight deadline. The scheduling strategy based on system priority can filter out the dummy transaction. But the earliest deadline first and least slack first strategies do process the dummy transaction, thus causing the system to behave poorly.

In Case 2, the dummy transaction activated has high priority but a loose deadline. The earliest deadline first and least slack first strategies can filter out the dummy transaction. But, the scheduling strategy based on system priority does process the dummy transaction, thus causing the system to behave poorly.

In Case 3, dummy transactions of both kinds are activated. The Combination strategy works well since it uses information about system priority as well as information about the timing constraints while making its scheduling decision.

4.5.2. When Preemption is Not Allowed

Quantitative Assessment:

Scheduling Strategy	Case 1	Case 2	Case 3
FCFS	400	1200	608
System Priority	400	600	600
Earliest Deadline First	400	600	600
Least Slack	400	600	600
Combination	400	600	600

Qualitative Assessment:

Scheduling Strategy	Case 1	Case 2	Case 3
FCFS	bad	bad	bad
System Priority	bad	bad	bad
Earliest Deadline First	bad	bad	bad
Least Slack	bad	bad	bad
Combination	bad	bad	bad

As seen above, in general, scheduling strategies perform poorly when preemption is not allowed. From the output of the simulation runs it is observed that almost all of the track transactions miss their deadlines, implying that the operator's console is empty most of the time. Due to this, the clean transactions trivially complete, since they have no emitter files to clean. But, it is almost impossible to start any transactions to shoot or display information about objects. Thus, the entire purpose of the pulse detection system is defeated.

5. Conclusion

Real-time database systems have timing and validity constraints associated with transactions. To ensure that a real-time database system completes as many transactions as possible without violating their timing and validity constraints, its scheduling strategy should have the following characteristics. First and foremost, the scheduling strategy should be dynamic. Second, it

should use the timing and validity information associated with transactions and the database. Third, the scheduling strategy should be used in every situation where there is a resource contention. Fourth, preemption should be allowed wherever possible.

In this paper, we have presented a dynamic scheduling algorithm for transactions in real-time database systems. The scheduling algorithm uses timing and validity information about transactions and data objects to calculate dynamic priorities of transactions. These priorities are then used to make scheduling decisions when transactions contend for scarce resources. The extra information enables the scheduler to make intelligent decisions so that the system completes as many critical transactions as possible.

To be useful, any scheduling algorithms have to be supported by an extensive round of experimentation with a *real-life* application. The proposed algorithm has been implemented and evaluated using a pulse detection system as a real-life, real-time database application. The experimental results show that scheduling algorithms for real-time database systems can be made more effective by making use of extra information about transactions and the database.

Real-time database systems of tomorrow will be large and complex, since they will be distributed, operate in an adaptive manner in a highly dynamic environment, exhibit intelligent behavior, and be characterized as having catastrophic consequences if the logical or timing constraints of transactions are not met. Meeting the challenges imposed by these characteristics very much depends on a focused and coordinated research efforts in several areas, especially in real-time transaction scheduling. The dynamic scheduling algorithm presented in this paper shows promising characteristics that are important to the problem of real-time transaction scheduling.

References

[Abbo88] Abbott, A. and H. Garcia-Molina, "Scheduling Real-time Transactions: a Performance Evaluation", *Proceedings of the 14th VLDB Conference*, 1988.

[Abbo89] Abbott, A. and H. Garcia-Molina, "Scheduling Real-time Transactions with Disk Resident Data", *Proceedings of the 15th VLDB Conference*, 1989.

[Bern87] Bernstein, P.A., V. Hadzilakos, and N. Goodman, "Concurrency Control and Recovery in Database Systems", Addison-Wesley, 1987.

[Hale89] Haleen, B. A., "SDEX/20 and 43RSS: Navy Standard Operating Systems", *Proceedings of the Workshop on Operating Systems for Mission Critical Computing*, University of Maryland, College Park, Maryland, 1989.

[Huan89] Huang, J. and J. Stankovic, "Experimental Evaluation of Real-Time Transaction Processing", *Proceedings of the Real-Time Systems Symposium*, December 1989.

[Sha91] Sha, L., R. Rajkumar, S. H. Son, and C. Chang, "A Real-Time Locking Protocol," *IEEE Transactions on Computers*, (to appear).

[Son88] Son, S. H., editor, *ACM SIGMOD Record*, *17*, 1, Special Issue on Real-Time Database Systems, March 1988.

[Son89] Son, S. H., "On Priority-Based Synchronization Protocols for Distributed Real-Time Database Systems," *IFAC/IFIP Workshop on Distributed Databases in Real-Time Control*, Budapest, Hungary, Oct. 1989, pp 67-72.

[Son90] Son, S. H. and C. Chang, "Performance Evaluation of Real-Time Locking Protocols using a Distributed Software Prototyping Environment," *10th International Conference on Distributed Computing Systems*, Paris, France, June 1990, pp 124-131.

[Son90b] Son, S. H. and J. Lee, "Scheduling Real-Time Transactions in Distributed Database Systems," *7th IEEE Workshop on Real-Time Operating Systems and Software*, Charlottesville, Virginia, May 1990, pp 39-43.

[Son90c] Son, S. H., "Real-Time Database Systems: A New Challenge," *Data Engineering*, vol. 13, no. 4, Special Issue on Directions for Future Database Research and Development, December 1990, pp 39-43.

[Stan88] Stankovic, J. A., "Misconceptions about Real-Time Computing", *IEEE Computer*, October 1988.

[Stan90] Stankovic, J. A., K. Ramamritham, and D. Towsley, "Scheduling in Real-Time Transaction Systems," *ONR Annual Workshop on Foundations of Real-Time Computing*, Washington, D.C., October 1990, pp 127-144.