



Real-Time Databases and Data Services

KRITHI RAMAMRITHAM

Indian Institute of Technology, Bombay

krithi@cse.iitb.ac.in

SANG H. SON

University of Virginia

son@cs.virginia.edu

LISA CINGISER DIPIPPA

University of Rhode Island

dipippo@cs.uri.edu

1. Introduction

Typically, a real-time system consists of a controlling system and a controlled system. In an automated factory, the controlled system is the factory floor with its robots, assembling stations, and the assembled parts, while the controlling system is the computer and human interfaces that manage and coordinate the activities on the factory floor. Thus, the controlled system can be viewed as the environment with which the computer interacts.

The controlling system interacts with its environment based on the data available about the environment, say from various sensors, for example, temperature and pressure sensors. It is imperative that the state of the environment, as perceived by the controlling system, be consistent with the actual state of the environment. Otherwise, the effects of the controlling systems' activities may be disastrous. Hence, timely monitoring of the environment as well as timely processing of the sensed information is necessary. The sensed data is processed further to derive new data. For example, the temperature and pressure information pertaining to a reaction may be used to derive the rate at which the reaction appears to be progressing. This derivation typically would depend on past temperature and pressure trends and so some of the needed information may have to be fetched from archival storage. Based on the derived data, where the derivation may involve multiple steps, actuator commands are set. For instance, in our example, the derived reaction rate is used to determine the amount of chemicals or coolant to be added to the reaction. In general, the history of (interactions with) the environment are also logged in archival storage.

In addition to the timing constraints that arise from the need to continuously track the environment, timing correctness requirements in a real-time (database) system also arise because of the need to make data available to the controlling system for its decision-making activities. If the computer controlling a robot does not command it to stop or turn on time, the robot might collide with another object on the factory floor. Needless to say, such a mishap can result in a major catastrophe.

Besides robotics, applications such as medical patient monitoring, programmed stock trading, and military command and control systems like submarine contact tracking require timely actions as well as the ability to access and store complex data that reflects the state of the application's environment. That is, data in these applications must be valid, or fresh, when it is accessed in order for the application to perform correctly. In a patient monitoring system, data such as heart rate, temperature, and blood pressure must be collected periodically. Transactions that monitor the danger level of a patient's status must be performed within a specified time, and the data must be accessed within an interval that defines the validity of the data. If not, the computations made by the transactions do not reflect the current state of the patient's health.

A traditional database provides some of the functionality required by these applications, such as coordination of concurrent actions and consistent access to shared data. But they do not provide for enforcement of the timing constraints imposed by the applications. A real-time database (RTDB) has all of the features of a more traditional database, but it can also express and maintain time-constrained data and time-constrained transactions. In the last fifteen years, there has been a great deal of research towards developing real-time databases that provide this functionality. In more recent years, much of the research that has been performed in the area of RTDBs has been applied, and extended in related fields, generally known as real-time data services. This paper surveys the highlights of the research that has been performed to advance the state-of-the-art of RTDBs and real-time data services.

In developing RTDB solutions that provide the required timeliness of data and transactions, there are various issues that must be considered. Below is a discussion of some of the concerns that have been the subject of research in this field.

Data, transaction and system characteristics. A RTDB must maintain not only the logical consistency of the data and transactions, it must also satisfy transaction timing properties as well as data temporal consistency. Transaction timing constraints include deadlines, earliest start times and latest start times. Transactions must be scheduled such that these constraints are met. Data temporal consistency imposes constraints on how old a data item can be and still be considered valid. There are various ways to characterize real-time transactions. If an application requires that the timing constraints on transactions be met in order to be correct, these transactions are considered to be hard. If a firm transaction misses a timing constraint, it no longer has value to the system, and therefore can be aborted. A soft transaction provides some, likely reduced, value to the system after its constraints have been violated. Transactions can also be characterized by their timing requirements. That is, a transaction can be periodic, or aperiodic. Finally, a RTDB system can be static or dynamic. In a static system, all data and requirements are known *a priori*, and the transactions can be designed and analyzed for schedulability up front. A dynamic system does not have this *a priori* knowledge, and therefore must be able to adapt to changing system requirements.

Scheduling and transaction processing. Scheduling and transaction processing techniques that consider data and transaction characteristics have been a major part of the research that has been performed in the field of RTDBs. Many of them consider the

inherent tradeoffs between quality of data vs. timeliness of processing. For example, in a weather monitoring RTDB, it may be necessary to sacrifice image processing time and store lower quality satellite images of areas that have no important weather activity going on. This will allow for more important areas, such as those with hurricane activity, to be more precise. Further, this will also allow for transactions accessing these important images to meet their access timing constraints. For another example, in a programmed stock trading application, a transaction that updates a stock price may be blocked by another transaction that is reading the same piece of data, and has a shorter deadline. If the stock price in question is getting “old” it would be in the interest of its temporal consistency to allow the updating transaction to execute. However, this execution could violate the logical consistency of the data or of the reading transaction. Thus, there is a trade-off between maintaining temporal consistency and maintaining logical consistency. If logical consistency is chosen, then there is the possibility that a stock price may become old, or that a transaction may miss a deadline. If, on the other hand, temporal consistency is chosen, the consistency of the data or of the transactions involved may be compromised.

Distribution. With the advent of high-speed, relatively low-cost networking, many of the applications that require a RTDB are not located on a single computer. Rather, they are distributed, and may require that the real-time data be distributed as well. Issues involved with distributing data include data replication, replication consistency, distributed transaction processing, and distributed concurrency control. When real-time requirements are added, these systems become much more complex. For instance, in a collaborative combat planning application where sensor data is collected by several ships in a region, and various decision-makers are viewing the collected data, it is critical that all collaborators share the same view of the scenario. The degree of distribution is also an issue that researchers have been examining recently. That is, a distributed RTDB that exists on three stationary nodes will have different requirements than a RTDB made up of hundreds of mobile computers each of which can sense local data and share it with the others.

Quality of service and quality of data. It is often not possible to maintain the logical consistency of a database and the temporal consistency as well. Therefore, there must be a trade-off made to decide which is more important. In dynamic systems, where it is not possible to know and control the real-time performance of the database, it may be necessary to accept levels of service that are lower than optimal. It may also be necessary to trade off the quality of the data in order to meet specified timing constraints. The RTDBs must provide mechanisms to specify allowable levels of service and quality of data, and it must also provide a way to adjust these levels when it is necessary.

This paper describes the important research results that have been produced in the above areas. It starts out in Section 2 by defining data freshness and timing properties. Section 3 presents research that has been done in the area of transaction processing for RTDBs. As mentioned above, this is where much of the RTDB focus has been in the past. Quality of service (QoS), and also quality of data (QoD) issues are addressed in Section 4. In Section 5, we discuss some new applications for which RTDB research has been useful. Finally, Section 6 concludes by discussing challenges that still exist for this

research area. It also proposes a research agenda for those interested in doing continuing work in this field.

2. Data Freshness and Timing Properties

The need to maintain consistency between the actual state of the environment and the state as reflected by the contents of the database leads to the notion of temporal consistency: the need to maintain coherency between the state of the environment and its reflection in the database. This arises from the need to keep the controlling system's view of the state of the environment consistent with the actual state of the environment.

A RTDB is composed of real-time objects which are updated by periodic sensor transactions. An object in the database models a real world entity, for example, the position of an aircraft. A real-time object is one whose state may become invalid with the passage of time. Associated with the state is a temporal validity interval. To monitor the states of objects faithfully, a real-time object must be refreshed by a sensor transaction before it becomes invalid, that is, before its temporal validity interval expires. The actual length of the temporal validity interval of a real-time object is application dependent.

Here, we do not restrict the notion of sensor data to the data provided by physical sensors. Instead, we consider a broad meaning of sensor data. Any data item, whose value reflects the time-varying real-world status, is considered a sensor data item.

Sensor transactions are generated by intelligent sensors which periodically sample the value of real-time objects. When sensor transactions arrive at RTDBs with sampled data values, their updates are installed and real-time data are refreshed. So one of the important design goals of RTDBs is to guarantee that temporal data remain fresh, that is, they are always valid. RTDBs that do not keep track of temporal validity of data may not be able to react to abnormal situations in time. Therefore, efficient design approaches are desired to guarantee the freshness of temporal data in RTDBs while minimizing the CPU workload resulting from periodic sensor transactions.

Whereas in the rest of this section we will be dealing with sensor transactions, in particular, with the derivation of their timing properties, it must be pointed out that many RTDBs also possess user transactions, transactions that access temporal as well as nontemporal data. User transactions typically arrive aperiodically. User transactions do not write any temporal data object, but they can read/write non-temporal data and their execution time and data access pattern can be time-varying. For example, in a decision support system a transaction can read different sets of dynamic, that is, temporally varying, data and perform different operations according to the situation.

From here on, $T = \{\tau_i\}_{i=1}^m$ refers to a set of periodic sensor transactions $\{\tau_1, \tau_2, \dots, \tau_m\}$ and $\chi = \{X_i\}_{i=1}^m$ refers to a set of temporal data. All temporal data are assumed to be kept in main memory. Associated with $X_i (1 \leq i \leq m)$ is a validity interval of length α_i : transaction $\tau_i (1 \leq i \leq m)$ updates the corresponding data X_i . Because each sensor transaction updates different data, no concurrency control is considered for sensor transactions. We assume that a sensor always samples the value of a temporal data at the beginning of its period, and all the first instances of sensor transactions are initiated at the same time. C_i, D_i and $P_i (1 \leq i \leq m)$ denote the execution

time, relative deadline, and period of transaction τ_i , respectively. D_i of sensor transaction τ_i is defined as follows:

$$D_i = \text{Deadline of } \tau_i - \text{Sampling Time of } \tau_i.$$

Deadlines of sensor transactions are firm deadlines. The goal of a RTDB designer is to determine P_i and D_i such that all the sensor transactions are schedulable and CPU workload resulting from sensor transactions is minimized. Let us assume a simple execution semantics for periodic transactions: a transaction must be executed once every period. However, there is no guarantee on when an instance of a periodic transaction is actually executed within a period. We also assume that these periodic transactions are preemptable.

Assume the period and relative deadline of a sensor transaction are set to be equal to the data validity interval. Because the separation of the execution of two consecutive instances of a transaction can exceed the validity interval, data can become invalid under this One–One approach. So this approach cannot guarantee the freshness of temporal data in RTDBs.

Example 2.1 Consider Figure 1. A periodic sensor transaction τ_i with deterministic execution time C_i refreshes temporal data X_i with validity interval α_i . The period P_i and relative deadline D_i of τ_i are assigned the value α_i . Suppose $J_{i,j}$ and $J_{i,j+1}$ are two consecutive instances of sensor transaction τ_i . Transaction instance $J_{i,j}$ samples data X_i with validity interval $[T, T + \alpha_i)$ at time T , and $J_{i,j+1}$ samples data X_i with validity interval $[T + \alpha_i, T + 2\alpha_i)$ at time $T + \alpha_i$. From Figure 1, the actual arrival time of $J_{i,j}$ and finishing time of $J_{i,j+1}$ can be as close as $2C_i$, and as far as $2P_i$, that is, $2\alpha_i$ when the period of τ_i is α_i . In the latter case, the validity of data X_i refreshed by $J_{i,j}$ expires after time $T + \alpha_i$. Since $J_{i,j+1}$ cannot refresh data X_i before time $T + \alpha_i$, X_i is invalid from $T + \alpha_i$ until it is refreshed by $J_{i,j+1}$, just before the next deadline $T + 2\alpha_i$.

In order to guarantee the freshness of temporal data in RTDBs, the period and relative deadline of a sensor transaction are each typically set to be less than or equal to one-half of the data validity interval (Ramamritham, 1993; Ho et al., 1997). In Figure 1, the farthest distance (based on the arrival time of a periodic transaction instance and the finishing time of its next instance) of two consecutive sensor transactions is $2P_i$. If $2P_i \leq \alpha_i$, then the freshness of temporal data X_i is guaranteed as long as instances of sensor transaction τ_i do not miss their deadlines.

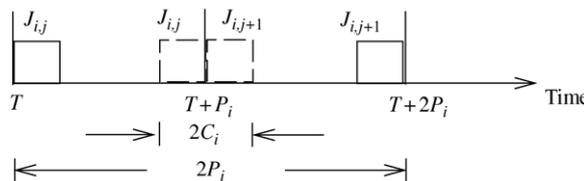


Figure 1. Extreme execution cases of periodic sensor transactions.

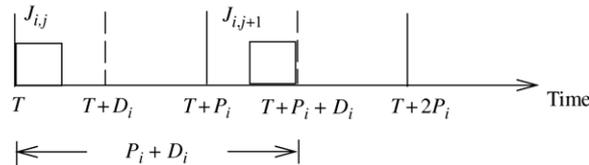


Figure 2. Illustration of More-less approach.

Unfortunately, even though data freshness is guaranteed, this design approach at least doubles CPU workload of the sensor transaction in the RTDBs compared to the One-One approach. Next, we introduce the More-Less approach whose goal is to minimize the CPU workload of sensor transactions while guaranteeing the freshness of temporal data in RTDBs. Recall that, for simplicity of discussion, we have assumed that a sensor transaction is responsible for updating a single temporal data item in the system. In More-Less, the period of a sensor transaction is assigned to be more than half of the validity interval of the temporal data updated by the transaction, while its corresponding relative deadline is assigned to be less than half of the validity interval of the same data. However, the sum of the period and relative deadline always equals the length of the validity interval of the data updated. Consider Figure 2. Let $P_i > \alpha_i/2$, $C_i \leq D_i < P_i$ where $P_i + D_i = \alpha_i$. The farthest distance (based on the arrival time of a periodic transaction instance and the finishing time of its next instance) of two consecutive sensor transactions J_{ij} and J_{ij+1} is $P_i + D_i$. In this case, the freshness of X_i can always be maintained if sensor transactions make their deadlines. Obviously, the load incurred by sensor transaction τ_i can be reduced if P_i is enlarged (which implies that D_i is shrunk). Therefore, we have the constraints $C_i \leq D_i < P_i$ and $P_i + D_i = \alpha_i$ which aim at minimizing the CPU workload of periodic transaction τ_i .

Example 2.2 Suppose there is temporal data X_i with validity interval α_i in a uniprocessor RTDB system. τ_i updates X_i periodically. For simplicity of discussion, we assume that $\delta_i = 0$. Our goal is to assign proper values to P_i and D_i given C_i and α_i so as to reduce the CPU workload resulting from sensor transaction τ_i . Suppose $C_i = \frac{1}{5}\alpha_i$, possible values of P_i , D_i and the corresponding CPU workload according to the three different design approaches are shown in Table 1.

Only half-half and more-less can guarantee the freshness of temporal data X_i if all the sensor transactions complete before their deadlines. We also notice that $U_o < U_m <$

Table 1. Comparison of three approaches.

Approach	D_i	P_i	Workload
One-one	α_i	α_i	$U_o = \frac{C_i}{\alpha_i} = \frac{1}{5}$
Half-half	$\frac{\alpha_i}{2}$	$\frac{\alpha_i}{2}$	$U_h = \frac{2C_i}{\alpha_i} = \frac{2}{5}$
More-less	$\frac{\alpha_i}{3}$	$\frac{2\alpha_i}{3}$	$U_m = \frac{3C_i}{2\alpha_i} = \frac{3}{10}$

U_h (see Table 1). If $P_i = ((N - 1)/N)\alpha_i$, then $D_i = (1/N)\alpha_i$, where $N \geq 2$. The freshness of temporal data in RTDBs is guaranteed if all sensor transactions complete before their deadlines. In such a case, we also notice that $U_m = (NC_i)/((N - 1)\alpha_i)$ and $U_o \leq U_m < U_h$. Theoretically, if $N \rightarrow \infty$, $U_m \rightarrow U_o$.

Unfortunately, how close U_m can get to U_o depends on C_i since $D_i \geq C_i$ implies $(\alpha_i/C_i) \geq N$. As N increases, relative deadlines become shorter and sensor transactions are executed with more stringent time constraints.

Therefore, given a set of sensor transactions in RTDBs, we need to find periods and deadlines of update transactions based on the temporal validity intervals of data such that the CPU workload of sensor transactions is minimized and the schedulability of the resulting sensor transactions is guaranteed. The more-less approach achieves this, as shown in Xiong and Ramamritham (2004). Recent work (Ho et al., 1997) on similarity-based load adjustment also adopts this approach to adjust periods of sensor transactions based on similarity bounds.

Concepts from active databases can be used to maintain the temporal validity of derived objects, where a derived object's value is dependent on the values of associated sensor data and it is updated sporadically. The technique described in Ahmed and Vrbsky (2000) is one such approach and is designed to ensure that maintaining temporal consistency does not result in a significant deterioration of the real-time performance with respect to meeting transaction deadlines.

Before we close this section, we would like to point out that effective scheduling of real-time transactions is greatly hampered by the large variance that typically exists between average-case and worst-case execution times in a database system. This unpredictability arises due to a variety of factors including data-value dependent execution patterns, blocking or restarts due to data contention, and disc latencies that are a function of the access history. Rastogi et al. (2000) attempt to eliminate some of these sources of unpredictability and thereby facilitate better estimation of transaction running times. It employs main-memory database technology and presents new logical and physical versioning techniques that ensure that read-only transactions are completely isolated from update transactions with respect to data contention.

3. Transaction Processing

In this section, issues related to the timely processing of transactions and queries such that they produce logically correct database states as well as outputs. Section 3.1 deals with scheduling and concurrency control issues, Section 3.2 summarizes results in the commit processing area, Section 3.3 presents work on recovery, and Section 3.4 outlines approaches developed to deal with resources such as I/O and buffers.

3.1. Scheduling and Concurrency Control

In this section, we discuss various aspects of transaction and query processing where the transactions and queries have time constraints attached to them and there are different consequences of not satisfying those constraints.

A key issue in transaction processing is predictability. In the context of an individual transaction, this relates to the question: “will the transaction meet its time-constraint”? Section 3.1.1 deals with the processing of transactions that have hard deadlines, that is, deadlines that must be met, and Section 3.1.2 deals with transactions that have soft deadlines, that is, deadlines which can be missed, but the larger the misses, the lower the value accrued to the system.

3.1.1. Dealing with Hard Deadlines

All transactions with hard deadlines must meet their time constraints. Since dynamically managed transactions cannot provide such a guarantee, the data and processing resources as well as time needed by such transactions have to be guaranteed to be made available when necessary. There are several implications.

First, we have to know when the transactions are likely to be invoked. This information is readily available for periodic transactions, but for aperiodic transactions, by definition, it is not. The smallest separation time between two incarnations of an aperiodic transaction can be viewed as its period. Thus, we can cast all hard real-time transactions as periodic transactions.

Second, in order to ensure *a priori* that their deadlines will be met, we have to determine their resource requirements and worst-case transaction execution times. This requires that many restrictions be placed on the structure and characteristics of real-time transactions.

Once we have achieved the above, we can treat the transactions in a manner similar to the way real-time systems treat periodic tasks that require guarantees, that is, by using static table-driven schedulers or preemptive priority-driven approaches. Static table-driven schedulers reserve specific time slots for each transaction. If a transaction does not use all of the time reserved for it, the time may be reclaimed to start other hard real-time transactions earlier than planned. Otherwise, it can be used for soft real-time transactions or left idle. The table-driven approach is obviously very inflexible. One priority-driven approach is the rate-monotonic priority assignment policy. One can apply the schedulability analysis tools associated with it to check if a set of transactions are schedulable given their periods and data requirements. This is the approach discussed in Sha et al. (1988) where periodic transactions that access main memory resident data via read and write locks are scheduled using rate-monotonic priority assignment.

We mentioned earlier that the variance between the worst-case computational needs and actual needs must not be very large. We can see why. Since the schedulability analysis is done with respect to worst-case needs, if the variance is large, many transactions that may be doable in the average case will be considered infeasible in the worst-case. Also, if the table-driven approach is used, a large variance will lead to large idle times.

In summary, while it is possible to deal with hard real-time transactions using approaches similar to those used in real-time systems, many restrictions have to be placed on these transactions so that their characteristics are known *a priori*. Even if one is

willing to deal with these restrictions, poor resource utilization may result given the worst-case assumptions made about the activities.

3.1.2. *Dealing with Soft Transactions*

With soft real-time transactions, we have more leeway to process transactions since we are not required to meet the deadlines all the time. Of course, the larger the number of transactions that meet their deadlines the better. When transactions have different values, the value of transactions that finish by their deadlines should be maximized. The complexity involved in processing real-time transactions comes from these goals. That is to say, we cannot simply let a transaction run, as we would in a traditional database system, and abort it should its deadline expire before it commits. We must meet transaction deadlines by adopting priority-assignment policies and conflict resolution mechanisms that explicitly take time into account. Note that priority assignment governs CPU scheduling and conflict resolution determines which of the many transactions contending for a data item will obtain access. As we will see, conflict resolution protocols make use of transaction priorities and because of this, the priority assignment policy plays a crucial role. We also discuss the performance implications of different deadline semantics.

Scheduling and conflict resolution. Rather than assigning priorities based on whether the transactions are CPU or I/O (or data) bound, real-time database systems must assign priorities based on transaction time constraints and the value they impart to the system. Possible policies are:

- earliest-deadline-first,
- highest-value-first,
- highest-value-per-unit-computation-time-first,
- longest-executed-transaction-first.

It has been shown that the priority assignment policy has significant impact on performance and that when different transactions have different values, both deadline and value must be considered (Huang et al., 1989)

For the purpose of conflict resolution in real-time database systems, various time-cognizant extensions of two phase locking, optimistic, and timestamp based protocols have been proposed in the literature (see, e.g., Abbott and Garcia-Molina, 1992; Buchmann et al., 1989; Haritsa et al., 1990a; Huang et al., 1989, 1991, 1992; Lin and Son, 1990; Son et al., 1991; Stankovic et al., 1991). Some of these are discussed below.

In the context of two-phase locking, when a transaction requests a lock that is currently held by another transaction we must take into account the characteristics of the transactions involved in the conflict. Considerations involved in conflict resolution are the

deadline and value (in general, the priority) of transactions, how long the transactions have executed, and how close they are to completion. Consider the following set of protocols:

- If a transaction with a higher priority is forced to wait for a lower priority transaction to release the lock, a situation known as priority inversion arises. This is because a lower priority transaction makes a higher priority transaction to wait. In one approach to resolving this problem, the lock holder inherits the lock requester's priority whereby it completes execution sooner than with its own priority.
- If the lock holding transaction has lower priority, abort it. Otherwise let the lock requester wait.
- If the lock holding transaction is closer to its deadline, lock requester waits, independent of its priority.

Priority Inheritance is shown to reduce transaction blocking times (Huang et al., 1992). This is because the lock holder executes at a higher priority (than that of the waiting transaction) and hence finishes early, thereby blocking the waiting higher priority transaction for a shorter duration. However, even with this policy, the higher priority transaction is blocked, in the worst case, for the duration of a transaction under strict two phase locking. As a result, the priority inheritance protocol typically performs even worse than a protocol that makes a lock requester wait independent of its priority.

If a higher priority transaction always aborts a low priority transaction, the resulting performance is sensitive to data contention. On the other hand, if a lower priority transaction that is closer to completion inherits priority rather than aborting, then a better performance results even when data contention is high. Such a protocol is a combination of the abort-based protocol proposed for traditional databases and the priority-inheritance protocol proposed for real-time systems (Sha et al., 1990). Said differently, the superior performance of this protocol (Huang et al., 1992) shows that even though techniques that work in real-time systems on the one hand and database systems on the other hand may not be applicable directly, they can often be tailored and adapted to suit the needs of real-time database systems. It should be noted that abort-based protocols (as opposed to wait-based) are especially appropriate for real-time database systems because of the time constraints associated with transactions.

Let us now consider optimistic protocols. In protocols that perform backward validation, the validating transaction either commits or aborts depending on whether it has conflicts with transactions that have already committed. The disadvantage of backward validation is that it does not allow us to take transaction characteristics into account. This disadvantage does not apply to forward validation. In forward validation, a committing transaction usually aborts ongoing transactions in case they conflict with the validating transaction. However, depending on the characteristics of the validating transaction and those with which it conflicts, we may prefer not to commit the validating transaction. Several policies have been studied in the literature (Haritsa et al., 1990a,b; Huang et al., 1991). In one, termed wait-50, a validating transaction is made to wait as

long as more than half the transactions that conflict with it have earlier deadlines. This is shown to have superior performance.

Time-cognizant extensions to timestamp-based protocols have also been proposed. In these, when data accesses are out of timestamp order, the conflicts are resolved based on their priorities. In addition, several combinations of locking-based, optimistic and timestamp-based protocols have been proposed but require quantitative evaluation (Lin and Son, 1990).

Exploiting multiple versions of data for enhanced performance has been addressed in (Kim and Srivastava, 1991). Multiple versions can reduce conflicts over data. However, if data must have temporal validity, old versions which are outdated must be discarded. Also, when choosing versions of related data, their relative consistency requirements must be taken into account: consider a transaction that uses multi-versioned data to display aircraft positions on an air-traffic controller's screen. The data displayed must have both absolute validity as well as relative validity.

Different transaction semantics are possible with respect to discarding a transaction once its deadline is past. For example, with firm deadlines, a late transaction is aborted once its deadline expires (Haritsa et al., 1990b). In general, with soft deadlines, once a transaction's value drops to zero, it is aborted (Huang et al., 1989). On the other hand, in the transaction model assumed in Abbott and Garcia-Molina (1992), all transactions have to complete execution even if their deadlines have expired. In this model, delayed transactions may cause other transactions also to miss their deadlines and this can have a cascading effect.

Experimental studies reported in the literature cited at the end of this article are very comprehensive and cover most aspects of real-time transaction processing, but most have not considered time constraints associated with data.

Database systems in which time validity intervals are associated with the data are discussed in Kuo and Mok (1993, 2000) and Song and Liu (1995). Such systems introduce the need to maintain data temporal consistency in addition to logical consistency. The performance of several concurrency control algorithms for maintaining temporal consistency are studied in Song and Liu (1995). In the model introduced in Song and Liu (1995), a real-time system consists of periodic tasks which are either read-only, write-only or update (read-write) transactions. Data objects are temporally inconsistent when their ages or dispersions are greater than the absolute or relative thresholds allowed by the application. Studies of two-phase locking and optimistic concurrency control algorithms, as well as rate-monotonic and earliest deadline first scheduling algorithms show that the performances of the rate-monotonic and earliest deadline first algorithms are close when the load is low. At higher loads, earliest deadline first outperforms rate-monotonic when maintaining temporal consistency. They also observed that optimistic concurrency control is generally worse at maintaining temporal consistency of data than lock based concurrency control, even though the former allows more transactions to meet their deadlines. It is pointed out in Song and Liu (1995) that it is difficult to maintain the data and transaction time constraints due to the following reasons:

- A transient overload may cause transactions to miss their deadlines.

- Data values may become out of date due to delayed updates.
- Priority based scheduling can cause preemptions which may cause the data read by the transactions to become temporally inconsistent by the time they are used.
- Traditional concurrency control ensures logical data consistency, but may cause temporal data inconsistency.

Motivated by these problems, and taking into account the fact that transactions process data with validity constraints and such data will be refreshed with sensor transactions, the notion of data-deadline is introduced in Xiong et al. (2002b). Informally, data-deadline is a deadline assigned to a transaction due to the temporal constraints of the data accessed by the transaction. Further, a forced wait policy which improves performance significantly by forcing a transaction to delay further execution until a new version of sensor data becomes available is also proposed.

In Kuo and Mok (1993), a class of real-time data access protocols called similarity stack protocols (SSP) is proposed. The correctness of SSP is based on the concept of similarity which allows different but sufficiently timely data to be used in a computation without adversely affecting the outcome. SSP schedules are deadlock free, subject to limited blocking and do not use locks. In Kuo and Mok (2000), weaker consistency requirements based on the similarity notion are proposed to provide more flexibility in concurrency control for data-intensive real-time applications. While the notion of data similarity is exploited in their study to relax serializability (hence increase concurrency), in Xiong et al. (2002b), it is coupled with data-deadline and used to improve the performance of transaction scheduling. The notion of similarity is used to adjust transaction workload by Ho et al. (1997), and incorporated into embedded applications (e.g., process control) in Chen and Mok (1999).

3.2. *Distributed Databases and Commit Protocols*

Many real-time database applications are inherently distributed in nature. These include the intelligent network services database, telecom databases, mobile telecommunication systems and the 1–800 telephone service in the United States. More recent applications include the directory, data feed and electronic commerce services that have become available on the World Wide Web. The performance, reliability, and availability of such applications can be significantly enhanced through the replication of data on multiple sites of the distributed network.

An algorithm for maintaining consistency and improving the performance of replicated DRTDBS is proposed in Son (1987). In this algorithm, a multiversion technique is used to increase the degree of concurrency. Replication control algorithms that integrate real-time scheduling and replication control are proposed in Son and Kouloumbis (1993). In contrast to the relaxed correctness assumed by the latter, conventional one-copy serializability supported by the algorithm called MIRROR reported in Xiong et al.

(2002b). MIRROR (Managing Isolation in Replicated Real-time Object Repositories), is a concurrency control protocol specifically designed for firm-deadline applications operating on replicated real-time databases. MIRROR augments the classical O2PL concurrency control protocol with a state-based realtime conflict resolution mechanism. In this scheme, the choice of conflict resolution method is a dynamic function of the states of the distributed transactions involved in the conflict. A feature of the design is that acquiring the state knowledge does not require inter-site communication or synchronization, nor does it require modifications to the two-phase commit protocol.

The performance of the classical distributed 2PL locking protocol (augmented with the priority abort (PA) and priority inheritance (PI) conflict resolution mechanisms) and of OCC algorithms in replicated DRTDBS was studied in the literature for real-time applications with soft deadlines. The results indicate that 2PL-PA outperforms 2PL-PI only when the update transaction ratio and the level of data replication are both low. Similarly, the performance of OCC is good only under light transaction loads. Making clear-cut recommendations on the performance of protocols in the soft deadline environment is rendered difficult, however, by the following: (1) There are two metrics: missed deadlines and mean tardiness, and protocols which improve one metric usually degrade the other, (2) the choice of the post-deadline value function has considerable impact on relative protocol performance, and (3) there is no inherent load control, so the system could enter an unstable state.

Temporal consistency guarantees are also studied in distributed real-time systems. In Zhou and Jahanian (1998), Distance constrained scheduling is used to provide temporal consistency guarantees for real-time primary-backup replication service.

Let us now consider the transaction commitment process. Once a transaction reaches its commit point, it is better to let it commit quickly so that its locks can be released soon. If commit delays are not high, which will be the case in a centralized database, the committing transaction can be given a high enough priority so that it can complete quickly. The solution is not so easy in a distributed system because of the distribution of the commitment process. Furthermore, since a deadline typically refers to the deadline until the end of the two-phase commit, but since the decision on whether or not to commit is taken in the first phase, we can enter the second phase only if we know that it will complete before the deadline. This requires special handling of the commit process. An alternative is to associate the deadline with the beginning of the second phase, but this may delay subsequent transactions since locks are not released until the second phase. A few papers in the literature (Soparkar et al., 1992; Yoon, 1994) have tried to address the issue of distributed commit processing but have required either relaxing the traditional notion of atomicity or strict resource allocation and resource performance guarantees from the system.

In Haritsa et al. (2000), its authors propose and evaluate a commit protocol that is specifically designed for the real-time domain without these changes. PROMPT allows transactions to optimistically borrow in a controlled manner, the updated data of transactions currently in their commit phase. This controlled borrowing reduces the data inaccessibility and the priority inversion that is inherent in distributed real-time commit processing. A simulation-based evaluation shows PROMPT to be highly successful as compared to the classical commit protocols in minimizing the number of missed

transaction deadlines. In fact its performance is close to the best on-line performance that could be achieved using the optimistic locking approach.

3.3. Recovery Issues

Recovery is a complex issue even in traditional databases and is more so in real-time database systems for two reasons. First, the process of recovery can interfere with the processing of ongoing transactions. Specifically, suppose we are recovering from a transaction aborted due to a deadline miss. If locks are used for concurrency control, it is important to release them as soon as possible so that waiting transactions can proceed without delay so as to meet their deadlines. However, it is also necessary to undo the changes done by the transaction to the data if in-place updates are done. But this consumes processing time that can affect the processing of transactions that are not waiting for locks to be released. Whereas optimistic concurrency control techniques or a shadow-pages based recovery strategy can be used to minimize this time, they have several disadvantages including lack of commercial acceptance. Second, unlike traditional databases where permanent data should always reflect a consistent state, in real-time databases, the presence of temporal data, while providing some opportunities for quicker recovery (Vrbsky and Lin, 1988), adds to the complexities of the recovery of transactions. Specifically, if a transaction's deadline expires before it completes the derivation of a data item, then rather than restoring the state of the data to its previous value, it could declare the data to be invalid thereby disallowing other transactions from using the value. The next instance of the transaction, in case the data is updated by a periodic transaction, may produce a valid state.

There is a need for designing new algorithms for logging and recovery in RTDBs because the sequential nature of logging and the lack of time and priority cognizance during recovery are not in tune with the priority oriented and preemptive nature of activities in RTDBs. Motivated by this need, Sivasankaran et al. (2000) present SPLIT, a partitioned logging and recovery algorithm for RTDBs, and (algorithms for recovery using non-volatile high speed store (ARUN)), a suite of logging and recovery algorithms for RTDBs. In order to improve performance, ARUN makes use of the solid state disc (SSD) technology that is referred to as non-volatile high speed store (NVHSS).

The logging and recovery algorithms are based on dividing data into a set of equivalence classes derived from a taxonomy of data and transaction attributes. For example, data can be broadly classified into hot and cold depending on the type of the data (like critical, temporal, etc.), and type of the transactions (like high priority, low priority etc.), that accesses the data. The logging and recovery algorithms assume that the classes of data are disjoint and that the recovery of one class can be done without having to recover the other classes, and if one class is recovered, then the transactions accessing that class can proceed without having to wait for the system to recover the other classes. For instance, network management systems consist of different kinds of data such as the performance, traffic, configuration, fault and security management data. The fault management data, which could potentially get updated frequently, is very critical to the operation that recovering it immediately after a crash can improve the performance of the

system. A recent work towards achieving bounded and predictable recovery using real-time logging based on these ideas is presented in Shu et al. (2004).

3.4. *Managing I/O and Buffers*

While the scheduling of CPU and data resources has been studied fairly extensively in the real-time database literature, studies of scheduling approaches for dealing with other resources, such as disc I/O, and buffers has begun only recently. In this section we review some recent work in this area and discuss some of the problems that remain.

I/O scheduling is an important area for real-time systems given the large difference in speeds between CPU and discs and the resultant impact of I/O devices' responsiveness on performance. Since the traditional disc scheduling algorithms attempt to minimize average I/O delays, just like traditional CPU scheduling algorithms aim to minimize average processing delays, time-cognizant I/O scheduling approaches are needed.

It must be recognized that what is important is the meeting of transaction deadlines and not the individual deadlines that may be attached to I/O requests. Assume that we model a transaction execution as a sequence of (disc I/O, computation) pairs culminating in a set of disc I/O's, the latter arising from writes to log and to the changed pages. Suppose we assign (intermediate) deadlines to the I/O requests of a transaction given the transaction's deadline. One of the interesting questions with regard to disc I/O scheduling is: How does one derive the deadline for an I/O request from that of the requesting transaction? First of all, it must be recognized that depending on how these I/O deadlines are set, deadlines associated with I/O requests may be soft since even if a particular I/O deadline is missed, the transaction may still complete by its deadline. This is the case if I/O deadlines are set such that the overall laxity (i.e., the difference between the time available before the deadline and the total computation time) of a transaction is uniformly divided among the computations and the I/O. On the other hand, assume that an intermediate deadline is equal to the latest completion time (i.e., the time an I/O must complete assuming that subsequent computations and I/O are executed without delay). This is the less preferred method since we now have a firm deadline associated with I/O requests—if an I/O deadline is missed, there is no way for the transaction to complete by its deadline and so the requesting transaction must be aborted.

Work on I/O scheduling includes Carey et al. (1989), Abbott and Garcia-Molina (1990) and Chen et al. (1991). The priority driven algorithm described in Carey et al. (1989) is a variant of the traditional SCAN algorithm which works on the elevator principle to minimize disc arm movement. Without specifying how priorities are assigned to individual I/O requests, (Carey et al., 1989) proposes a variant in which the SCAN algorithm is applied to each priority level. Requests at lower priority are serviced only after those at higher priority are served. Thus, if after servicing a request, one or more higher priority requests are found waiting, the disc arm moves towards the highest priority request that is closest to the current disc arm position. In the case of requests arising from transactions with deadlines, priority assignment could be based on the deadline assigned to the I/O request.

Another variant of SCAN, one which directly takes I/O deadlines into account is FD-SCAN Abbott and Garcia-Molina (1990). In this algorithm, given the current position of the disc arm, the disc arm moves towards the request with the earliest deadline that can be serviced in time. Requests that lie in that direction are serviced and after each service it is checked whether (1) a request with an even earlier deadline has arrived and (2) the deadline of the original result cannot be met. In either case, the direction of disc arm movement may change.

Clearly, both these protocols involve checks after each request is served and so incur substantial run-time overheads. The protocols described in Chen et al. (1991) are aimed at avoiding the impact of these checks on I/O performance. Specifically, the protocols perform the necessary computations while I/O is being performed. In the algorithm shortest-seek and earliest deadline by ordering (SSEDO), the need to give higher priority to requests with earlier deadlines is met while reducing the overall seek times. The latter is accomplished by giving a high priority to requests which may have large deadlines but are very close to the current position of the disc arm. A variant of SSEDO is SSEDV which works with specific deadline values, rather than deadline orderings. Chen et al. (1991) shows how both the algorithms can be implemented so as to perform disk scheduling while service is in progress and shows that the algorithms have better performance than the other variants of the SCAN algorithms.

Another resource for which contention can arise is the database buffer. What we have here is a conflict over buffer slots—akin to conflicts that occur over a time slot, in the case of a CPU. Thus, similar issues arise here also. Specifically, how to allocate buffer slots to transactions and which slots to replace when a need arises are some of the issues. Consider buffer replacement: in case there is a need to replace an existing buffer slot to make room for a new entry, the replacement policy may have an impact on performance, especially if the slot being replaced is used by an uncommitted transaction. Studies discussed in (Carey et al., 1989) show that transaction priorities must be considered in buffer management.

4. Satisfying QoS/QoD Requirements

In many important real-time applications including e-commerce, agile manufacturing, sensor networks, traffic control, target tracking, and network management, transactions should be processed within their deadlines, using fresh (temporally consistent) sensor data that reflect the current real-world status. Meeting both requirements of timeliness and data freshness is challenging. Generally, transaction execution time and data access pattern are not known *a priori*, but could vary dynamically. For example, transactions in stock trading may read varying sets of stock prices, and perform different arithmetic/logical operations to maximize the profit considering the current market status. Further, transaction timeliness and data freshness can often pose conflicting requirements. By preferring user requests to sensor updates, the deadline miss ratio is improved; however, the data freshness might be reduced. Alternatively, the freshness increases if updates receive a higher priority. In this section, we first review previous work in quality of

service (QoS) management in RTSBs, discuss the performance metrics for QoS management, and then present several approaches based on feedback control in RTDBs.

4.1. *Earlier Work*

Despite the abundance of the QoS research, QoS-related work is relatively scarce in database systems. Priority adaptation query resource scheduling (PAQRS) provided timeliness differentiation of query processing in a memory-constrained environment (Pang et al., 1995). From the observation that the performance of queries can vary significantly depending on the available memory, per-class query response time was differentiated by an appropriate memory management and scheduling. Given enough memory, queries can read the operand relations at once to produce the result immediately. If less memory is allocated, they have to use temporary files to save the intermediate results, therefore, the query processing may slow down. In this way, query deadline miss ratios were differentiated between the classes. However, no data freshness issues were considered, and the performance could easily fluctuate under the workload changes.

An on-line update scheduling policy has been proposed in the context of the web server (Labrinidis and Roussopoulos, 2001b). The performance of a web server can be improved by caching dynamically generated data at the web server and the back-end database continuously updates them. Given the views to materialize, the proposed update scheduling policy can significantly improve the data freshness compared to FIFO scheduling. A complementary problem of view materialization in web service is discussed in Labrinidis and Roussopoulos (2001a). While the paper considers the trade-offs between response time and data freshness, it provides neither miss ratio nor data freshness guarantees.

Stanford real-time information processor (STRIP) addressed the problem of balancing between the freshness and transaction timeliness (Adelberg et al., 1995). To study the trade-off between freshness and timeliness, four scheduling algorithms were introduced to schedule updates and transactions, and the performance was compared. In their later work, a similar trade-off problem was studied for derived data (Adelberg et al., 1996). Ahmed and Vrbsky (2000) proposed a new approach to maintain the temporal consistency of derived data. Different from STRIP, an update of a derived data object is explicitly associated with a certain timing constraint, and is triggered by the database system only if the timing constraint could be met. By a simulation study, the relative performance improvement was shown compared to the forced delay scheme of STRIP. None of the two approaches considers dynamic adaptations of update policy, and no performance guarantee is provided.

The correctness of answers to database queries can be traded off to enhance the timeliness. A query processor, called APPROXIMATE (Vrbsky, 1993), can provide approximate answers depending on the availability of data or time. An imprecise computation technique (milestone approach (Lin et al., 1987) is applied by APPROXIMATE. In the milestone approach, the accuracy of the intermediate result increases monotonically as the computation progresses. Therefore, the correctness of answers to the query could monotonically increase as the query processing progresses. A

relational database system called CASE-DB (Ozsoyoglu and Hou, 1990) can produce approximate answers to queries within certain deadlines. Approximate answers are provided processing a segment of the database by sampling, and the correctness of answers can improve as more data are processed. Before beginning each data processing, CASE-DB determines if the segment processing can be finished in time. In replicated databases, consistency can be traded off for shorter response time. For example, epsilon serializability (Pu and Leff, 1991) allows a query processing despite the concurrent updates. Notably, the deviation of the answer to the query can be bounded, different from a similar approach called quasi serializability (Du and Elmagarmid, 1989). An adaptable security manager is proposed in Son et al. (2000), in which the database security level can be temporarily degraded to enhance timeliness. These performance trade-off schemes lack a systematic QoS management architecture and none of them consider providing guarantees for both miss ratio and data freshness.

4.2. QoS Metrics for RTDBs

Before we discuss QoS management and performance metrics for real-time data services, it is important to note that almost all the work on QoS in RTDBs is not aiming to provide hard guarantees. Instead, the main focus is on soft real-time applications, in which infrequent deadline misses/temporal consistency violations can be tolerated, if neither of them exceeds the limits specified in the QoS requirements.

Two major performance metrics usually considered for QoS specification are the following:

- *User transaction deadline miss ratio*: In a QoS specification, a database administrator can specify the target deadline miss ratio that can be tolerated for a specific real-time application.
- *Data freshness*: We categorize data freshness into database freshness and perceived freshness. Database freshness is the ratio of fresh data to the entire temporal data in a database. Perceived freshness is the ratio of fresh data accessed to the total data accessed by timely transactions. To measure the current perceived freshness, we exclusively consider timely transactions. This is because tardy transactions, which missed their deadlines, add no value to the system. Note that only the perceived freshness can be specified in the QoS requirement. A QoS-sensitive approach provides the perceived freshness guarantee while managing the database freshness internally (hidden to the users) (Kang et al., 2002).

Long-term performance metrics such as ones listed above are not sufficient for performance specification of dynamic systems, in which the system performance can be widely time-varying. For that reason, transient performance metrics such as overshoot and settling time, adopted from control theory, have been considered as performance metrics for real-time data services (Franklin et al., 1994; Phillips and Nagle, 1995; Lu et

al., 2000). Similar transient performance metrics are proposed in (Rosu et al., 1997) to capture the responsiveness of adaptive resource allocation in real-time systems.

- *Overshoot* is the worst-case system performance in the transient system state (e.g., the highest miss ratio in the transient state). Overshoot is an important metric because a high transient miss ratio can cause serious implications in several applications such as robots and e-commerce.
- *Settling time* is the time for the transient overshoot to decay and reach the steady state performance. The settling time represents how fast the system can recover from overload. This metric has also been called reaction time or recovery time.

4.3. Feedback Control-Based QoS Management

As one of the first efforts to address QoS management in RTDBs, Kang et al. (2002, 2004) developed a novel QoS management architecture called QMF (a QoS management architecture for deadline Miss ratio and data Freshness). Figure 3 shows the QMF architecture. In QMF, a formal control theoretic approach is applied to compute the required workload adjustment considering the miss ratio error, that is, the difference between the desired miss ratio and the currently measured miss ratio. Feedback control is used since it is very effective to support the desired performance when the system model includes uncertainties (Lu et al., 2001; Philips and Nagle, 1995). At each sampling instant, the feedback-based miss ratio controller measures the miss ratio and computes the

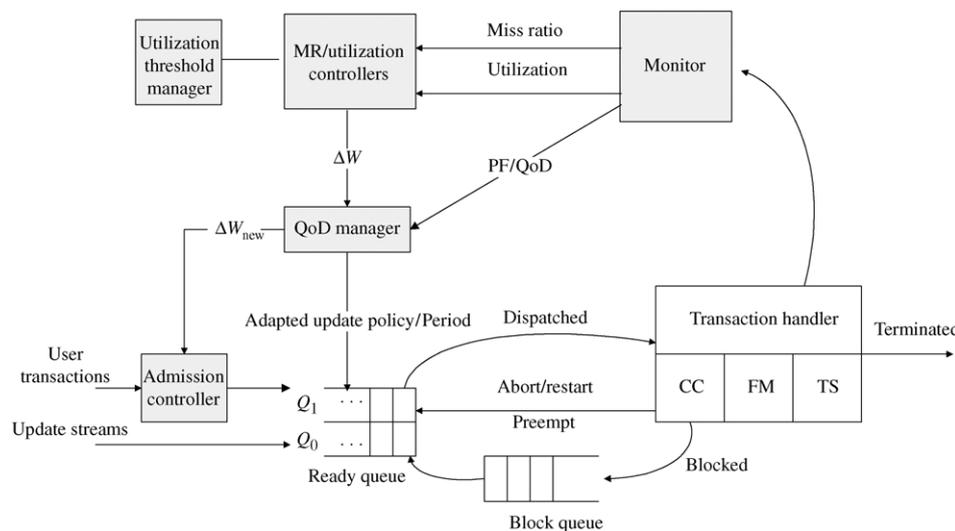


Figure 3. RTDBS architecture for QoS management using feedback control.

miss ratio error, and the controller computes the control signal, that is, the required workload adjustment to react to the error.

According to the required workload adjustment computed in the feedback loop, flexible freshness management and admission control are applied to support the desired miss ratio without affecting the freshness. Using the notion of perceived freshness, an adaptive update policy was developed to maintain the freshness in a cost-effective manner. Initially, all data are updated immediately when their new sensor readings arrive. Under overload conditions, some sensor data can be updated with reduced update rates, if necessary, to improve the miss ratio (as long as the target perceived freshness is supported). This flexible approach contrasts to the existing database update policy, commonly accepted in the real-time database research such as (Adelberg et al., 1995; Kim et al., 2002), which is fixed and not adaptable regardless of the current system status.

To prevent potential deadline misses or stale data accesses due to the delay for on-demand updates, Kang et al. (2004) developed an alternative approach in which all data are updated immediately. In that approach, the notions of QoD and flexible validity intervals are used to manage the freshness. When overloaded, update periods of some sensor data can be relaxed within the specified range of QoD to reduce the update workload, if necessary. However, sensor data are always maintained fresh in terms of flexible validity intervals. Therefore, the age of sensor data is always bounded. According to the performance study, QMF shows satisfying stringent QoS requirements for a wide range of workloads and access patterns. QMF achieves a near zero miss ratio and perfect freshness even given dynamic workloads and data access patterns.

4.4. QoS Management Using Imprecision

Imprecise computation techniques (Liu et al., 1994) provide means for achieving graceful degradation during transient overloads by trading off resource needs for the quality of a requested service. Imprecise computation and feedback control scheduling have been used for QoS management of RTDBs (Amirijoo et al., 2003a–c). In this approach the notion of imprecise computation is applied on transactions as well as data, that is, data objects are allowed to deviate, to a certain degree, from their corresponding values in the external environment. Although a RTDB models an external environment that changes continuously, the values of data objects that are slightly different in age or in precision are often interchangeable as consistent read data for user transactions. The time it takes to update a data object alone introduces a time delay which means that the value of the data object cannot be the same as the corresponding real-world value at all times. This means that during transient overloads the system skips update transactions with values similar to the values stored in the database. To measure data imprecision the notion of data error, denoted de_i , is used which gives an indication of how much the value of a data object d_i stored in the RTDB deviates from the corresponding real-world value given by the latest arrived transaction updating d_i . A user transaction T_j is discarded if the data error of the data object d_i to be updated by T_j is less or equal to the maximum data error mde (i.e., $de_i \leq mde$). If mde increases, more update transactions are discarded, degrading the

precision of the data. Imprecision at transaction level can be expressed in terms of certainty, accuracy, and specificity (Zilberstein and Russell, 1996). Certainty refers to the probability of a result to be correct, accuracy refers to the degree of accuracy of a value returned by an algorithm (typically through a bound on the difference from the exact solution), and specificity reflects the level of detail of the result. The imprecision of the result of a user transactions T_i , denoted the transaction error te_i , increases as the resource available for the transaction decreases.

The database operator expresses QoS requirements, in terms mde and average te_i of terminated transactions (denoted ate), by specifying not only the desired steady-state performance, but also the transient-state performance describing the worst-case system performance and system adaptability in the face of unexpected failures or load variation.

The general outline of the QoS management using imprecision is the following: Admitted transactions are placed in the ready queue. Periodically, ate is monitored and fed into the QoS controller, which compares the desired ate with the actual ate to get the current performance error. Based on this the controller computes a change, denoted δl , to the total estimated requested load. Based on δl , the (QoD) manager changes the total estimated requested load by adapting mde. The precision controller discards an update transaction writing to a data object d_i having an error less or equal to the maximum data error allowed, that is, $de_i \leq mde$. The portion of δl not accommodated by the QoD manager, denoted δl_{new} , is returned to the admission controller (AC), which enforces the remaining load adjustment.

4.5. QoS Management for Distributed RTDBS

Providing QoS guarantees for data services in a distributed environment is a challenging task. The presence of multiple sites in distributed environments raises issues that are not present in centralized systems. The transaction workloads in distributed real-time databases may not be balanced and the transaction access patterns may be time-varying and skewed. The work in Wei et al. (2003) describes an algorithm that provides QoS guarantees for data services in distributed real-time databases with full replication of temporal data. The algorithm consists of feedback-based local controllers and heuristic global load balancers (GLB) working at each site. The local controller controls the admission process of incoming transactions. The global load balancers collect the performance data from different nodes and balance the system-wide workload.

At each node, there are a local miss ratio controller and a local utilization controller. The local miss ratio controller takes the miss ratios from the latest sampling period, compares them with the QoS specification and computes the local miss ratio control signal, which is used to adjust the target utilization at the next sampling period. In order to prevent under-utilization, a utilization feedback loop is added. It is used to avoid a trivial solution, in which all the miss ratio requirements are satisfied due to under-utilization. At each sampling period, the local utilization controller compares the utilization of the last period with the preset utilization threshold and generates the local utilization control signal. The local miss ratio controller reduces the incoming transaction

workload when the QoS specification is violated; the utilization increases the incoming transaction workload when the system is under-utilized.

To balance the workload between the nodes and thus provide distributed QoS management, decentralized GLB are used. GLBs sit at each node in the system, collaborating with each other for load balancing. At each sampling period, nodes in the system exchange their system performance data. The GLB at each node compares the system condition of different nodes. If a node is overloaded while some other nodes in the system are not, in the next period, the GLB at the overloaded node will send some transactions to the nodes that are not overloaded.

A more recent work extends the algorithm so that it scales to large distributed systems (Wei et al., 2004). In that work, partial data replication and efficient replica management algorithms are studied because full replication is inefficient or impossible in large distributed systems.

5. New and Emerging Applications

Much of the research described in the previous sections has dealt with the application of real-time techniques and theory to databases. Recently, there has been a trend towards applying the results of this core RTDB research to other, related applications. The key ingredients of these applications are the requirements for real-time response to requests for real-time data. Sensor networks are a natural application for real-time data services because their main purpose is to provide sensed data to some requesting entity, very often with real-time constraints on the data and on the requests. A similar application that has recently benefited from early RTDB research is real-time mobile databases in which some or all of the nodes in the database can change location and are connected via a wireless network. Both of these applications bring new constraints to be considered. They both must deal with the unpredictability and lower bandwidth of wireless networks. They also both may need to consider power issues in some or all of the devices used in the system. A third application that has recently become a subject of research is real-time web-based data services. These applications include programmed stock trading and information services. The research involved in providing real-time data from such services pulls from various areas already discussed in this paper. A related area of research deals with dissemination of streaming real-time data. A key goal in these applications is delivering temporally valid real-time data to requestors within specified timing constraints. This section will discuss how each of these new application areas have provided new challenges to consider along with the RTDB issues discussed earlier. It will also describe how these applications have leveraged the existing research and developed new solutions to tackle some of these challenges.

5.1. Sensor Network Applications

Sensor networks are large-scale wireless networks that consist of numerous sensor and actuator nodes used to monitor and interact with physical environments (Estrin et al.,

1999; Hill et al., 2000). From one perspective sensor networks are similar to distributed database systems. They store environmental data on distributed nodes and respond to aperiodic and long-lived periodic queries (Bonnet et al., 2000; Jaikao et al., 2000; Madden and Franklin, 2002). Data interest can be pre-registered to the sensor network so that the corresponding data is collected and transmitted only when needed. These specified interests are similar to views in traditional databases because they filter the data according to the application's data semantics and shield the overwhelming volume of raw data from applications (Bonnet et al., 2001; Shen et al., 2001).

Sensor networks have inherent real-time properties. The environment that sensor networks interact with is usually dynamic and volatile. The sensor data usually has an absolute validity interval of time after which the data values may not be consistent with the real environment. Transmitting and processing "stale" data wastes communication resources and can result in wrong decisions based on the reported out-of-date data. Besides data freshness, often the data must also be sent to the destination by a deadline. To date, not much research has been performed on real-time data services in sensor networks.

Despite their similarity to conventional distributed real-time data-bases, sensor networks differ in the following important ways. First, individual sensors are small in size and have limited computing resources, while they must also operate for long periods of time in an unattended fashion. This makes power conservation an important concern in prolonging the lifetime of the system. In current sensor networks, the major source of power consumption is communication. To reduce unnecessary data transmission from each node, data collection and transmission in sensor networks are always initiated by subscriptions or queries. Second, any individual sensor is not reliable. Sensors can be damaged or die after consuming the energy in the battery. The wireless communication medium is also unreliable. Packets can collide or be lost. Because of these issues we must build trust on a group of sensor nodes instead of any single node. Previous research emphasizes reliable transmission of important data or control packets at the lower levels, but less emphasis is on the reliability on data semantics at the higher level (Ratnasamy et al., 2002). Third, the large amount of sensed data produced in sensor networks necessitates in-network processing. If all raw data is sent to base stations for further processing, the volume and burstiness of the traffic may cause many collisions and contribute to significant power loss. To minimize unnecessary data transmission, intermediate nodes or nearby nodes work together to filter and aggregate data before the data arrives at the destination. Fourth, sensor networks can interact with the environment by both sensing and actuating. When certain conditions are met, actuators can initiate an action on the environment. Since such actions are difficult to undo, reducing false alarms is crucial in certain applications.

There are many ongoing data service middleware research projects for the sensor network applications, including Cougar, Rutgers Dataman, SINA, SCADDS, Smart-msgs, and some virtual-machine-like designs (Cougar Project; Dataman Project; SCADDS; Smart-msgs; Bonnet et al., 2001; Feng et al., 2002; Mattern et al., 2002; Shen et al., 2001). COUGER and SINA are two typical data-centric middleware designs which have goals that are similar to our design goal of providing data services. In COUGER, sensor data is viewed as tables and query execution plans are developed and possibly optimized

in th middleware. DSWare project (Li et al., 2004) is more tailored to sensor networks, including supporting group-based decision, reliable data-centric storage, and implementing other approaches to improve the performance of real-time execution, reliability of aggregated results and reduction of communication. SINA is a cluster-based middleware design which focusses on the cooperation among sensors to conduct a task. Its extensive SQL-like primitives can be used to issue queries in sensor networks. However, it does not provide schemes to hide the faulty nature of both sensor operations and wireless communication. In SINA it is the application layer that must provide robustness and reliability for data services. The real-time scheduling component and built-in real-time features of other service components make DSWare more suitable than SINA for real-time applications in wireless sensor networks.

Multisensor data fusion research focusses on solutions that fuse data from multiple sensors to provide more accurate estimation of the environment (Jayasimha et al., 1991; Qi et al., 2001). In mobile-agent-based data fusion approaches, software that aggregates sensor information are packed and dispatched as mobile agents to ‘hot’ areas (e.g., the area where an event occurred) and work independently there. The software migrates among sensors in a cluster, collects observations, then infers the real situation (Qi et al., 2001). This group-based approach makes use of consensus among a number of nearby sensors of the same type to increase the reliability of a single observation. The mobile-agent-based approach, however, leverages on the migration traffic of mobile agents and their appropriate processing at each sensor node in its routes. For instance, if a node in the route inserts wrong data or refuses to forward the mobile agents, the aggregation and subsequent analysis are untrustful.

A fuzzy modeling approach is sometimes used for data fusion in sensor networks. It is used to model the uncertainty in sensor failures and faulty observations (Samarasooriya and Varshney, 2000). This approach is useful in modeling the sensor error rates due to equipment wear and aggregating local decisions from multiple sensors that measure the same type of data. Some optimal decision schemes focus on the fusion of asynchronously arriving decisions (Chang and Kam, 1994; Samarasooriya and Varshney, 1996). E. Bosse et al. (2000) presented a modeling and simulation approach for a real-time algorithm in multi-source data fusion systems. These data fusion schemes are suitable for increasing the accuracy of decisions, but require extensive computing resources. Dempster–Shafer evidential theory is also applied to incorporate uncertainty into decisions in some sensor fusion research (Murphy, 1999). This scheme uses belief and plausibility functions to describe the reliability feature of each source and uses a normalized Dempster’s combination rule to integrate decisions from different sources. The confidence function in DSWare is similar to Dempster–Shafer method except that it places the evidence in both temporal and spatial spectrums to take the real-time validity intervals of data and possible contexts into consideration.

5.2. *Mobile Real-Time Databases*

Recent advances in wireless communications technology have made mobile information services a valuable mechanism for applications to share and distribute data. Applications

like real-time traffic information systems, and mobile internet stock trading systems have requirements for timely transactions and valid data. However, the resource constraints of mobile systems make it difficult to meet timing constraints. These constraints include low bandwidth, unreliable wireless networks and frequent disconnections. The research described in this section attempts to apply RTDB solutions and solve the resource constraint issues for mobile real-time database applications.

The work in Kayan and Ulusoy (1999) describes a mobile real-time database system model that consists of a number of fixed and mobile hosts. Each mobile host has a fixed host that acts as a mobile support station (MSS) and coordinates the operations of the transactions submitted by that mobile host. A mobile transaction in this model can execute by the submitting mobile host, and/or at one or more fixed hosts, and has a deadline by which it must be completed. The mobile transaction can be executed either in a single request message to the fixed network, or as a sequence of messages for read and write operations to the fixed network, thus placing the majority of the execution of the transaction either on the fixed host (former case) or on the mobile host (latter case). The authors performed various simulations to test these execution strategies, as well as other chosen techniques. The results of these simulations indicate that as the number of hosts increases, the fraction of transactions that satisfy their deadlines decreases. Further, transactions that are executed on the fixed host, consistently perform better than those that are executed on the mobile host.

The work in Lam et al. (2000), extends a RTDB concurrency control technique to handle the specific characteristics of a mobile RTDB. The protocols in this work address two main issues: the cost of resolving data conflicts, and the impact of the mobile network. The system model consists of both a mobile network and a reliable wired network. The concurrency control mechanisms described in this paper are based on the HP-2PL protocol (Abbott and Garcia-Molina, 1992). The distributed high priority two phase locking protocol (DHP-2PL) uses a transaction restart mechanism to resolve data conflicts. However, it attempts to minimize the restarting of transactions that access data on mobile hosts by allowing a committing transaction to hold a lock until it has finished the commit procedure. Priority inversion is minimized by increasing the priority of the committing transaction higher than all executing transactions. To handle the problem of disconnection of a mobile transaction, a cautious waiting scheme is used in which a higher priority transaction is restarted by a lower priority transaction if the higher priority transaction is suspected of having been disconnected.

The idea of transaction shipping is presented in Lam et al. (2000) to reduce communication costs of executing mobile transactions. Recall that in Kayan and Ulusoy (1999), it was shown that transactions that are executed on the fixed host perform better than those executed on the mobile host. Transaction shipping sends the entire transaction to the fixed database server. The main issue of concern is how to identify the execution path and the required data objects of a transaction ahead of time. A pre-analysis approach is suggested in which the set of operations of the transaction is identified, and the precedence relationships of the operations are determined. Given this information, the database server can execute the transaction without the need to communicate with the mobile host on which it originated.

In Lam et al. (2001), a multi-version data model is adopted to allow for more flexibility in resolving data conflicts. This provides a looser correctness criterion based on relative consistency, and therefore increases data availability in the mobile system. In this model, new updates of data items are generated by mobile hosts, and a specific fixed host maintains a backup copy of, not only the current version, but also of previous copies of the data item. An image transaction is generated at a fixed host, called the anchor host, when a mobile transaction is received. Its job is to get its required data items to the anchor host in time for the mobile transaction to access it. Techniques similar to those in Lam et al. (2000) are used for pre-analysis. When the mobile transaction is executed, the relative consistency of the data items is checked to ensure that they all reflect the same temporal state of the data. After the mobile transaction accesses a version of a particular data item, all other data items are checked for relative consistency. If at any point, such a version cannot be found, the transaction is restarted, and the next latest version of its first accessed data item will be selected.

An enhanced multi-version data model is presented in Leung et al. (2002) for data broadcasts in time constrained mobile systems. The broadcast cycle is made smaller by relaxing the requirement to broadcast every data item in every cycle. For each version of a data item in the broadcast cycle, this technique keeps track of the cycle number when the version was created, and the cycle number of the latest version of the data item. This allows for the distance between versions to be recorded. This work also proposes an on-demand broadcast for minimizing missing deadlines. After the broadcast of all data items that are scheduled, the system examines the on-demand buffer, containing any request that occurred during the broadcast cycle. These requests are sorted by deadline, and processed in this order. The amount of time spent processing broadcast data versus on-demand data is a tunable parameter in the system.

As mentioned above, mobile real-time database systems must contend with power constraints on the mobile devices that generate and access data. The work in Gruenwald and Banik (2001) presents transaction management that reduces missed deadlines while balancing the amount of energy consumed by mobile hosts in the network. The system model in this work is a mobile ad-hoc network (MANET), in which only mobile hosts exist, not fixed hosts. Each mobile host can be in one of three energy modes: active (executing as usual), doze (CPU runs at lower rate and can receive messages), and sleep (CPU and communication are suspended). Small mobile hosts (SMH) have highly constrained memory, storage and power capabilities. Large mobile hosts (LMH) have more resources than small mobile hosts. Each LMH maintains a global schema that describes the status of the other LMHs. Each LMH broadcasts its ID, position and energy level to all other hosts. When an SMH starts a transaction, it is sent in its entirety to a LMH to be processed. To schedule the transactions, an energy-efficient modification of Least Slack Time scheduling is proposed that considers the problem of disconnection. Highest priority is given to the transaction with the smallest slack time. In case of equal slack time, higher priority is given to the transaction whose requestor has the lower energy level.

RTMonitor (Lam et al., 2002) is a real-time data monitoring system that uses mobile agents to adaptively monitor real-time data based on the urgency of the requests and on system workload. The target application for this system is traffic navigation. The agents

us Adaptive PUSH or PULL (APoP) to maintain the temporal consistency of the data. To manage the traffic system, local traffic graphs connect to neighboring local graphs through external nodes. The system consists of two types of agents: stationary and mobile. Stationary agents manage the traffic graphs. The mobile agents serve the client navigation requests. In the traffic system, traffic sensors collect traffic data for road segments and send the data whenever there is a significant change in a value. In the APoP approach used here, the monitoring period is defined based on the urgency of the request, the system status, and the dynamic properties of the traffic data. When a request is received, a GraphAgent is created to find the shortest path. This is accomplished by computing a path on the local graph of the originating region, a virtual path to the destination region, and a local path from an external node to the destination in the destination region. To keep the calculated path up to date, the system initially uses a PUSH mode for data updates. When the workload at a server becomes heavy, some less critical requests will be switched to a PULL mode. The QueryAgent for a client generates a pull based on the slack time and the dynamic properties of the traffic data.

Uncorrected Proof

5.3. Web-based Real-Time Data Services

Recently, the world wide web has offered a new venue for providing real-time data services. Many applications, such as program stock trading, require information services that can provide real-time data to widely distributed users. The main challenge in these applications is to provide timely access to fresh data, in the face of the highly dynamic environment of the web. Due to the unpredictable nature of web-based applications, most of the research in this area has been focussed on providing QoS management for real-time data services.

In Sebastine et al. (2001), a web-based content distribution service for industrial applications is presented. The service allows for remote monitoring of industrial devices that may be geographically distributed. It is based on an active web caching architecture that provides on-demand replication of dynamically changing web content, like industrial process state. The active web cache is made up of a standard web proxy cache that handles all static data, and an active server that handles one or more types of dynamic content. The approach in Sebastine et al. (2001) is to trade-off data temporal consistency with timeliness, while keeping within specified inconsistency bounds. This trade-off is made through the formulation of a QoS optimization problem and a heuristic for QoS management. $\min C_i$ and $\max C_i$ are the minimum and maximum allowed temporal consistency of a data item as defined by the publisher of the data. The consistency level of a data item is defined as:

$$QoS_i = (\max C_i - \text{Current}C_i) / (\max C_i - \min C_i) \quad (1)$$

and is used in the heuristic to manage QoS. The heuristic provides a mechanism for degrading the consistency level when the CPU utilization is above a given threshold.

A QoS management framework for web-based real-time data services, based on feedback control theory is presented in Son and Kang (2002). In this approach, a designer specifies desired behavior using performance metrics such as deadline miss ratio, data freshness, highest miss ratio and settling time, which are mapped to dynamic responses of control systems. The designer then determines a dynamic model of the real-time data services for performance control in order to define the relationship between the control input and the system performance. Using existing mathematical techniques of feedback control theory, the designer creates a feedback controller that provides analytic guarantees on the desired behavior. Using this technique, this work provides a solution to maintaining acceptable freshness in lazy replication. Each replicated data object has a primary copy that is updated periodically. Secondary copies of this data are created on sites where it is needed. These copies are updated by either a push or a pull model. At first, all secondary copies are updated at the same rate as the primary copy. When an overload occurs on a node, the replication can be adjusted by either performing on-demand copies, or delaying updates for some copies.

The work in Amirijoo et al. (2003a) handles QoS management in web-based data services by allowing imprecise computation. It describes two algorithms for feedback control with imprecision, which were described previously in Section 4.4.

5.4. Dissemination of Streaming/Dynamic Web Data

Recent studies have shown that an increasing fraction of the data on the world wide web is time-varying (i.e., changes frequently). Examples of such data include sports information, news, and financial information such as stock prices. The coherency requirements associated with a data item depends on the nature of the item and user tolerances. To illustrate, a user may be willing to receive sports and news information that may be out-of-sync by a few minutes with respect to the server, but may desire to have stronger coherency requirements for data items such as stock prices. A user who is interested in changes of more than a dollar for a particular stock price need not be notified of smaller intermediate changes.

An important issue in the dissemination of such time-varying web data is the maintenance of temporal coherency. In the case of servers adhering to the HTTP protocol, clients need to frequently pull the data based on the dynamics of the data and a user's coherency requirements. In contrast, servers that possess push capability maintain state information pertaining to clients and push only those changes that are of interest to a user. These two canonical techniques have complementary properties with respect to the level of temporal coherency maintained, communication overheads, state space overheads, and loss of coherency due to (server) failures.

Specifically a pull-based approach does not offer high fidelity when the data changes rapidly or when the coherency requirements are stringent. Moreover, the pull-based approach imposes a large communication overhead (in terms of the number of messages exchanged) when the number of clients is large. A push-based algorithm can offer high fidelity for rapidly changing data and/or stringent coherency requirements. However, it incurs a significant computational and state-space overhead resulting from a large number

of open push connections. Moreover, the approach is less resilient to failures due to its stateful nature. These properties indicate that a push-based approach is suitable when a client requires its coherency requirements to be satisfied with a high fidelity, or when the communication overheads are the bottleneck. A pull-based approach is better suited to less frequently changing data or for less stringent coherency requirements, and when resilience to failures is important.

In Bhide et al. (2002) combinations of push- and pull-based techniques are proposed to achieve the best features of both approaches. The combined techniques tailor the dissemination of data from servers to clients based on (i) the capabilities and load at servers and proxies, and (ii) clients' coherency requirements.

Time-varying web data is frequently used for online decision making, and in many scenarios, such decision making involves multiple time-varying data items, possibly from multiple independent sources. Examples include a user tracking a portfolio of stocks and making decisions based on the net value of the portfolio. Observe that computing the overall value of the portfolio at any instant requires up-to-date values of each stock in the portfolio. In some scenarios, users might be holding these stocks in different (brokerage) accounts and might be using a third-party aggregator, such as yodlee.com, to provide them a unified view of all their investments. The third-party aggregator assembles a unified view of the portfolio by periodically gathering and refreshing information from each individual account (source), all of which may be independent of one another. If the user is interested in certain events (such as a notification every time the value of the portfolio increases or decreases by \$1000), then the aggregator needs to periodically and intelligently refresh the value of each stock price to determine when these user-specified thresholds are reached. More formally, the aggregator runs a continuous query over the portfolio by intelligently refreshing the components of the portfolio from their sources in order to determine when user-specified thresholds are exceeded. We refer to such queries as continuous threshold queries (CTQ). Another example of CTQ would be an intersection of a busy thoroughfare where an increase in the traffic beyond a pre-specified limit triggers flow control measures. Much of the Web infrastructure assumes that it is difficult to cache time-varying data at a proxy, and consequently, requests for these data items as well as queries on these data items are handled directly by the server (data source). Such an approach has two important limitations. First, since continuous threshold queries are compute-intensive (due to the need to continuously evaluate the query condition), the server could become a bottleneck, and thereby, limit the scalability of the system. Second, a pure server-based approach makes CTQs on data items from multiple independent sources impossible (e.g., the server-based approach is unfeasible in the case where a user specifies a query on stocks held in multiple independent accounts). Executing queries at a proxy eliminates both restrictions: the approach is scalable, since computations are offloaded from the server to proxies, and queries on data items from multiple sources become feasible. In addition, a proxy-based approach improves user response times. However, such an approach raises new research challenges. The key challenge is to ensure that the results of the query are no different from the case where the query is handled by the server (i.e., correctness of the results should be ensured). To do so, the proxy should ensure that cached values of time-varying data items are temporally coherent with the source.

6. Challenges and Research Agenda

The field of real-time database research has evolved a great deal over the relatively short time of its existence. In the early days of the 1980s, much of the research concentrated on examining how to add real-time support to traditional databases. Much of this work involved developing new scheduling and concurrency control techniques that extended existing techniques to allow RTDBs to provide timely transactions and temporally valid data. This has included techniques to relax traditional ACID properties of databases, and managing QoS provided by the database in order to allow for timing constraints to be met. More recently, the focus of RTDB research has been on applying the lessons learned from the earlier research towards providing real-time functionality to emerging applications that have data delivery and data management needs. In the most recent Real-Time and Embedded Applications Symposium (RTAS, 2004), there was no session devoted to RTDBs, as there has been in the past. However, papers in real-time data dissemination, real-time data distribution, and real-time sensor networks, were presented in various other sessions. On the other hand, there was a session on RTDBs at the EuroMicro Conference on Real-Time Systems (ECRTS, 2004) with papers on QoS issues in derived data and on embedded RTDBs in automobiles. This is an indication of the evolution of RTDB research, and the widening of the areas that are covered in the field.

The future of RTDB research will depend on this continuing evolution. Research in this field should continue to follow cutting edge applications and apply existing techniques to them, as well as developing new ones when needed. It is also important that researchers keep up with new technology in traditional database research. As new approaches to solving database issues are developed, it is important to investigate how such innovations can apply to real-time applications. For example, a survey of papers in the current Conference on Very Large Databases (VLDB, 2004) indicates that there is substantial research in XML, XQuery, database privacy, OLAP, data mining, data warehousing, query optimization and databases for the scientific grid. Research in RTDBs must keep up with the new technology in these areas and continue to ask the question, "Will this database research help solve real problems for RTDBs?" We can also ask the question, "Does it make sense to approach some of these issues from a real-time perspective?" For example, what does it mean to apply real-time technology to XML and XQuery? Also, are there useful applications for real-time data mining and data warehousing?

On a more practical note, it is important that RTDB research be widely applicable. Many academic papers mention various applications like programmed stock trading, medical patient monitoring, etc., to which their technologies will apply. However, there are very few real-time databases in use in such applications. In order for this research to continue, it must be shown to be practical and useful in real applications. There are currently several commercial solutions that provide some RTDB support (Lockheed Martin Corporation, 1998; TimesTen; eXtremeDB; Empress). However, none of these employs the latest RTDB technology. Most are main memory database systems that provide "fast" response times. In order to further advance the state of the art research into deployed systems, it may be necessary for RTDB researchers to team with commercial vendors to develop systems that can be used by applications that require the

kinds of technology that a RTDB can provide. Standards efforts may also be useful towards bringing the RTDB technology to the marketplace.

References

- Abbott, R., and Garcia-Molina, H. 1992. Scheduling real-time transactions: a performance evaluation. *ACM Transactions on Database Systems* 17(3): 513–560.
- Abbott, R., and Garcia-Molina, H. 1990. Scheduling I/O requests with deadlines: a performance evaluation. *Proceedings of the Real-Time Systems Symposium*.
- Adelberg, B., Garcia-Molina, H., and Kao, B. 1995. Applying update streams in a soft real-time database system. In *Proceedings of the 1995 ACM SIGMOD*, pp. 245–256.
- Adelberg, B., Kao, B., and Garcia-Molina, H. 1996. Database support for efficiently maintaining derived data. In *ETDB*.
- Ahmed, Q. N., and Vrbsky, S. V. 2000. Triggered updates for temporal consistency in real-time databases. *Real-Time Systems Journal* 19: 209–243.
- Amirijoo, M., Hansson, J., and Son, S. H. 2003a. Algorithms for managing QoS for real-time data services using imprecise computation. In *Proceedings of the Conference on Real-Time and Embedded Computing Systems and Applications (RTCSA)*.
- Amirijoo, M., Hansson, J., and Son, S. H. 2003b. Error-driven QoS management in imprecise real-time databases. In *Proceedings of the Euromicro Conference on Real-Time Systems (ECRTS)*.
- Amirijoo, M., Hansson, J., and Son, S. H. 2003c. Specification and management of QoS in imprecise real-time databases. In *Proceedings of the International Database Engineering and Applications Symposium (IDEAS)*.
- Bell, D., and LaPadula, L. 1976. Secure computer systems: unified exposition and multics interpretation. Technical Report MTR-2997, The Mitre Corp., March.
- Bestavros, A., and Braoudakis, S. 1994. Timeliness via speculation for Real-time databases. In *Proceedings of 15th Real-Time Systems Symposium*.
- Blaze, M., Feigenbaum, J., and Lacy, J. 1996. Decentralized trust management. In *IEEE Symposium on Security and Privacy*. Oakland, CA, pp. 164–173.
- Bhattacharya, S., Kim, H., Prabh, S., and Abdelzaher, T. 2003. Energy-conserving data placement and asynchronous multicast in wireless sensor networks. In *Proceedings of the 1st International Conference on Mobile Systems, Applications, and Services*. San Francisco, CA.
- Bhide, M., Deolasse, P., Katker, A., Panchbudhe, A., Ramamritham, K., and Shenoy, P. 2002. Adaptive push-pull: Disseminating dynamic web data. *IEEE Transactions Computers (special issue on Quality of Service)* 51: 652–668.
- Blum, B., Nagaraddi, P., Wood, A., Abdelzaher, T., Son, S. H., and Stankovic, J. A. 2003. An entity maintenance and connection service for sensor networks. In *Proceedings of the 1st International Conference on Mobile Systems, Applications, and Services*. San Francisco, CA.
- Bonnet, P., Gehrke, J., and Seshadri, P. 2000. Querying the physical world. *IEEE Personal Communication Magazine* 7: 10–15.
- Bonnet, P., Gehrke, J., and Seshadri, P. 2001. Towards sensor database systems. In *Proceedings of the 2nd International Conference on Mobile Data Management*. Hong Kong.
- Bosse, E., Roy, J., and Paradis, S. 2000. Modeling and simulation in support of design of a data fusion system. *Information Fusion* 1: 77–87.
- Buchmann, A. P., McCarthy, D. R., Chu, M., and Dayal, U. 1989. Time-critical database scheduling: a framework for integrating real-time scheduling and concurrency control. In *Proceedings of the Conference on Data Engineering*.
- Carey, M. J., Jauhari, R., and Livny, M. 1989. Priority in DBMS resource scheduling. In *Proceedings of the 15th VLDB Conference*, pp. 397–410.
- Chang, W., and Kam, M. 1994. Asynchronous distributed detection. In *IEEE Transactions on Aerospace Electronic Systems*, pp. 818–826.

- Chen, S., Stankovic, J., Kurose, J., and Towsley, D. 1991. Performance evaluation of two new disc scheduling algorithms for real-time systems. *Real-Time Systems*.
- Chen, D., and Mok, A. K. 1999. SRDE-application of data similarity to process control. In *The 20th IEEE Real-Time Systems Symposium*. Phoenix, Arizona, December.
- Cougar Project. www.cs.cornell.edu/database/cougar.
- Dataman Project. www.cs.rutgers.edu/dataman.
- Davidson, S., Lee, I., and Wolfe, V. 1989. A protocol for timed atomic commitment. In *Proceedings of 9th International Conference on Distributed Computing Systems*.
- Du, W., and Elmagarmid, A. 1989. Quasi serializability: A correctness criterion for global concurrency control in interbase. In *Proceedings of the 15th International Conf. on Very Large Data Bases*.
- Empress. Real-time data acquisition system: Empress real-time data collector. Whitepaper. www2.empress.com.
- Estrin, D., Govindan, R., Heidemann, J., and Kumar, S. 1999. Next century challenges: scalable coordination in sensor networks. In *Proceedings of the 5th Annual International Conference on Mobile Computing and Networks*. Seattle, WA.
- eXtremeDB. Real-time Databases for Embedded Systems. Whitepaper. www.mobject.com.
- Feng, J., Koushanfar, F., and Potkonjak, M. 2002. System-architectures for sensor networks: issues, alternatives, and directions. In *Proceedings of the 20th International Conference on Computer Design*. Freiburg, Germany.
- Franklin, G. F., Powell, J. D., and Emami-Naeini, A. 1994. *Feedback Control of Dynamic Systems*, 3rd edn. Addison Wesley.
- George, B., and Haritsa, J. 2000. Secure Buffering in firm real-time database systems. *Journal on Very Large Data Bases (VLDB)* 8(3): 178–198.
- Graham, M. C. 1991. Issues in real-time data management. CMU/SEL-91-TR-17.
- Greenberg, I. et al., 1993. The Secure Alpha study—Final summary report. Technical Report, CS Lab, SRI International.
- Griffeth, N., and Weinrib, A. 1990. Scalability of a real-time distributed resource counter. In *Proceedings of the Real-Time Systems Symposium*. Orlando, Florida.
- Gruenwald, L., and Banik, S. 2001. A power-aware technique to manage real-time database transactions in mobile ad-hoc networks. In *Proceedings of the 12th International Conference on Database and Expert Systems Applications (DEXA'01)*. Munich, Germany.
- Haritsa, J. R., Carey, M. J., and Livny, M. 1990a. On being optimistic about real-time constraints. In *Proceedings of ACM PODS*.
- Haritsa, J. R., Carey, M. J., and Livny, M. 1990b. Dynamic real-time optimistic concurrency control. In *Proceedings of the Real-Time Systems Symposium*.
- Haritsa, J. R., Carey, M. J., and Livny, M. 1991. Earliest deadline scheduling for real-time database systems. In *Proceedings of the Real-Time Systems Symposium*.
- Haritsa, J.R., Carey, M. J., and Livny, M. 1992. Data access scheduling in firm real-time database systems. *The Journal of Real-Time Systems* 4(3): 203–241.
- Haritsa, J., Ramamritham, K., and Gupta, R. 2000. The PROMPT real-time commit protocol. *IEEE Transactions on Parallel and Distributed Systems* 11(2): 160–181.
- He, T., Stankovic, J. A., Lu, C., and Abdelzaher, T. 2003. SPEED: a stateless protocol for real-time communication in ad hoc sensor networks. In *Proceedings of the 23rd International Conference on Distributed Computing Systems*. Providence, RI.
- Hill, J., Szewczyk, R., Woo, A., Hollar, S., Culler, D. E., and Pister, K. S. J. 2000. System architecture directions for networked sensors. In *Architectural Support for Programming Languages and Operating Systems*, pp. 93–104.
- Hou, W., Ozsoyoglu, G., and Taneja, B. K. 1989. Processing aggregate relational queries with hard time constraints. In *Proceedings of the ACM SIGMOD International Conference on the Management of Data*.
- Huang, J., Stankovic, J. A., Towsley, D., and Ramamritham, K. 1989. Experimental evaluation of real-time transaction processing. In *Proceedings of the Real-Time Systems Symposium*.
- Huang, J., Stankovic, J. A., Ramamritham, K., and Towsley, D. 1991. Experimental evaluation of real-time optimistic concurrency control schemes. In *Proceedings of the Conference on Very Large Data Bases*.
- Huang, J., Stankovic, J. A., Ramamritham, K., Towsley, D., and Purimetla, B. 1992. On using priority inheritance in real-time databases. *Special Issue of Real-Time Systems Journal* 4(3).

- Jaikao, C., Srisathapornphat, C., and Shen, C.-C. 2000. Querying and tasking in sensor networks. In *Proceedings of SPIE's 14th Annual International Symposium on Aerospace/Defense Sensing, Simulation, and Control (Digitization of the Battlespace V)*. Orlando, FL.
- Jayasimha, D., Ivengar, S., and Kashyap, R. 1991. Information integration and synchronization in distributed sensor networks. *IEEE Transactions on Systems, Man and Cybernetics* 21(5): 1032–1043.
- Kang, K. D., Son, S. H., and Stankovic, J. A. 2004. Managing deadline miss ratio and sensor data freshness in real databases. *IEEE Transactions on Knowledge and Data Engineering* 16(7).
- Kang, K. D., Son, S. H., Stankovic, J. A., and Abdelzaher, T. F. 2002. A QoS-sensitive approach for timeliness and freshness guarantees in real-time databases. In *the 14th Euromicro Conference on Real-Time Systems*.
- Kayan, E., and Ulusoy, O. 1999. An evaluation of real-time transaction management issues in mobile database systems. *The Computer Journal (Special Issue on Mobile Computing)* 42(6).
- Kim, S., Son, S. H., Stankovic, J. A., Li, S., and Choi, Y. 2003. SAFE: A Data Dissemination Protocol for Periodic Updates in Sensor Networks. In *IEEE Workshop on Data Distribution for Real-Time Systems (DDRTS)*. Providence, RI.
- Kim, S., Son, S., and Stankovic, J. 2002. Performance evaluation on a real-time database. In *IEEE Real-Time Technology and Applications Symposium*.
- Kim, W., and Srivastava, J. 1991. Enhancing real-time DBMS performance with multi-version data and priority-based disc scheduling. In *Proceedings of the Real-Time Systems Symposium*, pp. 222–231.
- Korth, H. F., Soparkar, and Silberschatz, A. 1990. Triggered real-time databases with consistency constraints. In *Proceedings of the Conference on Very Large Data Bases*.
- Kuo, T., and Mok, A. K. 1992. Real-time data semantics and similarity-based concurrency control. *IEEE 13th Real-Time Systems Symposium*.
- Kuo, T., and Mok, A. K. 1993. SSP: a semantics-based protocol for real-time data access. *IEEE 14th Real-Time Systems Symposium*.
- Kuo, T., and Mok, A. K. 2000. Real-time data semantics and similarity-based concurrency control. *IEEE Transactions on Computers* 49(11): 1241–1254.
- Ho, S., Kuo, T., and Mok, A. K. 1997. Similarity-based load adjustment for static real-time transaction systems. In *Proceedings of the 18th Real-Time Systems Symposium*.
- Labrinidis, A., and Roussopoulos, N. 2001a. Adaptive WebView materialization. In *Proceedings of the Fourth International Workshop on the Web and Databases, held in conjunction with ACM SIGMOD*.
- Labrinidis, A., and Roussopoulos, N. 2001b. Update propagation strategies for improving the quality of data on the web. In *Proceedings of the 27th International Conference on Very Large Data Bases (VLDB'01)*. Rome, Italy.
- Lam, K. 1994. Concurrency control in distributed real-time database systems. Ph.D. thesis, City University of Hong Kong.
- Lam, K.-Y., Kwan, A., and Ramamritham, K. 2002. RTMonitor: Real-time data monitoring using mobile agent technologies. In *Proceedings of the 28th VLDB Conference*. Hong Kong, China.
- Lam, K.-Y., Kuo, T.-W., Tsang, W.-H. G. 2000. Law concurrency control in mobile distributed real-time database systems. *Information Systems* 25(4): 261–322.
- Lam, K.-Y., Li, G., Kuo, T.-W. 2001. A multi-version data model for executing real-time transactions in a mobile environment. In *Proceedings of the 2nd ACM International Workshop on Data Engineering for Wireless and Mobile Access*.
- Lampson, B. 1973. A note on the confinement problem. *Communications of the ACM* 16(10): 613–615.
- Leung, H.-W., Yuen, J., Lam, K.-Y., and Chan, E. 2002. Enhanced multi-version data broadcast schemes for time-constrained mobile computing systems. In *Proceedings of the 13th International Workshop on Database and Expert Systems Applications (DEXA'02)*.
- Li, S., Lin, Y., Son, S. H., Stankovic, J., and Wei, Y. 2004. Event detection services using data service middleware in distributed sensor networks. *Telecommunication Systems* (26): 351–368.
- Lin, K. J., Natarajan, S., and Liu, J. W. S. 1987. Imprecise results: utilizing partial computations in real-time systems. In *Proceedings of the Real-Time System Symposium*.
- Lin, Y., and Son, S. H., 1990. Concurrency control in real-time databases by dynamic adjustment of serialization order. In *Proceedings of the Real-Time Systems Symposium*.

- Liu, C. L., and Layland, J. 1973. Scheduling algorithms for multiprogramming in a hard real-time environment. *Journal of the ACM* 20(1).
- Liu, J., Lin, K., Shih, W., Yu, A., Chung, J., and Zhao, W. 1991. Algorithms for scheduling imprecise computation. *IEEE Computer* 24(5).
- Liu, J. W. S., Shih, W.-K., Lin, K.-J., Bettati, R., and Chung, J.-Y. 1994. Imprecise computations. *Proceedings of the IEEE*. Vol. 82.
- Locke, D. 1997. Real-time databases: real-world requirements. In Azer Bestavros, Kwei-Jay Lin, and Sang, H. Son (eds), *Real-Time Database Systems: Issues and Applications*. Kluwer Academic Publishers, pp. 83–91.
- Lockheed Martin Corporation. 1998. <http://www.lmco.com>.
- Lu, C., Abdelzaber, T., Stankovic, J., and Son, S. 2001. A feedback control approach for guaranteeing relative delays in web servers. In *Proceedings of the 7th IEEE Real-Time Technology and Applications Symposium (RTAS 01)*. Taipei, Taiwan, pp. 51–62.
- Lu, C., Blum, B., Abdelzaber, T., Stankovic, J. A., and He, T. 2002. RAP: A real-time communication architecture for large-scale wireless sensor networks. In *Proceedings of the 8th IEEE Real-Time Technology and Applications Symposium*. San Jose, CA.
- Lu, C., Stankovic, J., Abdelzaber, T., Tao, G., Son, S. H., and Marley, M. 2000. Performance specifications and metrics for adaptive real-time systems. In *Real-Time Systems Symposium*. Orlando, Florida.
- Lu, C., Stankovic, J. A., Tao, G., and Son, S. H. 2001. Feedback control real-time scheduling: framework, modeling and algorithms. *Journal of Real-Time Systems* (Special Issue on Control-Theoretical Approaches to Real-Time Computing) (to appear).
- Lu, C., Stankovic, J. A., Tao, G., and Son, S. H. 2002. Feedback control real-time scheduling: Framework, modeling and algorithms. *Journal of Real-time Systems* (Special Issue on Control-Theoretical Approaches to Real-Time Computing) 23(1/2).
- Madden, S., Franklin, M., Hellerstein, J., and Hong, W. 2002. TAG: A tiny aggregation service for ad-hoc sensor networks. In *Proceedings of the 18th International Conference on Data Engineering*. San Jose, CA.
- Madden, S., and Franklin, M. J. 2002. Fjording the stream: an architecture for queries over streaming sensor data. In *Proceedings of the 18th International Conference on Data Engineering*. San Jose, CA.
- Mattern, F., Romer, K., and Kasten, O. 2002. Middleware challenges for wireless sensor networks. *ACM SIGMOBILE Mobile Computing and Communication Review (MC2R)*.
- Mostowitz, I., Greenwald, S., and Kang, M. 1996. An analysis of timed Z-channel. *IEEE Symposium on Security and Privacy*. Oakland, CA, pp. 2–11.
- Murphy, P. R. 1999. Dempster–Shafer theory for sensor fusion in autonomous mobile robots. *IEEE Transactions on Robotics and Automation* 14(2): 197–206.
- O’Neil, P., Ramamritham, K., and Pu, C. 1992. Towards predictable transaction executions in real-time database systems. Technical Report 92-35, University of Massachusetts.
- Ozsoyoglu, G., and Hou, W.-C. 1990. Research in time- and error-constrained database query processing. In *Proceedings of the Workshop on Real-Time Operating Systems and Software*.
- Pang, H., Carey, M. J., and Livny, M. 1995. Multiclass query scheduling in real-time database systems. *IEEE Transactions on Knowledge and Data Engineering* 7(4).
- Parekh, S., Gandhi, N., Hellerstein, J., Tilbury, D., Jayram, T., and Bigus, J. 2002. Using control theory to achieve service level objectives in performance management. *Journal of Real-time Systems* (Special Issue on Control-Theoretical Approaches to Real-Time Computing) 23(1/2).
- Park, C., Park, S., and Son, S. H. 2002. Multiversion locking protocol with freezing for secure real-time database systems. *IEEE Transactions on Knowledge and Data Engineering* 14(5): 1141–1154.
- Phillips, C. L., and Nagle, H. T. 1995. *Digital Control System Analysis and Design*. 3rd edn. Prentice Hall.
- Pu, C., and Leff, A. 1991. Replica control in distributed systems: as asynchronous approach. In *Proceedings of ACM SIGMOD International Conference on Management of Data*.
- Qi, H., Wang, X., Iyengar, S. S., and Chakrabarty, K. 2001. Multisensor data fusion in distributed sensor networks using mobile agents. In *Proceedings of 5th International Conference on Information Fusion*. Annapolis, MD.
- Ramamritham, K. 1993. Real-time databases. *Distributed and Parallel Databases* 1: 199–226.
- Ramamritham, K., 1996. Where do deadlines come from and where do they go? *Journal of Database Management* Spring.

- Rastogi, R., Seshadri, S., Bohannon, J., Leinbaugh, D., Silberschatz, A., and Sudarshan, S., 2000. Improving predictability of transaction execution times in real-time databases. *Real-Time Systems Journal* 19(3): 205–208.
- Ratnasamy, S., Estrin, D., Govindan, R., Karp, B., Shenker, S., Yin, L., and Yu, F. 2002. Data-centric storage in sensornets. In *Proceedings of the 1st Workshop on Sensor Networks and Applications*. Atlantic, GA.
- Rosu, D., Schwan, K., Yalmanchili, S., and Jha, R. 1997. On adaptive resource allocation for complex real-time applications. In *Proceedings of the IEEE Real-Time Systems Symposium*.
- Samarasooriya, V. N. S., and Varshney, P. K. 1996. A sequential approach to asynchronous decision fusion. *Optical Engineering* 35(3): 625–633.
- Samarasooriya, V. N. S., and Varshney, P. K. 2000. A fuzzy modeling approach to decision fusion under uncertainty. *Fuzzy Sets and Systems* 114(1): 59–69.
- SCADDS: Scalable coordination architectures for deeply distributed systems. www.isi.edu/scadds.
- Sebastine, S., Kang, K.-D., Abdelzaher, T., and Son, S. H. 2001. A scalable web-based real-time information distribution service for industrial applications. In *Proceedings of the 27th Annual Conference of IEEE Industrial Electronics Society*. Denver, CO.
- Sha, L., Rajkumar, R., and Lehoczky, J. P. 1988. Concurrency control for distributed real-time databases. *ACM SIGMOD Record*.
- Sha, L., Rajkumar, R., and Lehoczky, J. 1990. Priority inheritance protocols: an approach to real-time synchronization. *IEEE Transactions on Computers* 39(9): 1175–1185.
- Shen, C.-C., Srisathapornphat, C., and Jaikao, C. 2001. Sensor information networking architecture and applications. *IEEE Personal Communication Magazine* 8(4): 52–59.
- Sivasankaran, R., Ramamritham, K., and Stankovic, J. A. 2000. System failure and recovery. In Lam, K.-Y., and Kuo, T.-W. (eds), *Real-Time Database Systems: Architecture and Techniques*. Kluwer Academic Publishers, pp. 109–124.
- Sivasankaran, R. M., Stankovic, J. A., Towsley, D., Purimetla, B., and Ramamritham, K. 1996. Priority assignment in real-time active databases. *The International Journal on Very Large Data Bases* 5(1).
- Smart Messages Project. discolab.rutgers.edu/sm.
- Shu, L., Son, S. H., and Stankovic, J. 2004. Achieving bounded and predictable recovery using real-time logging. *The Computer Journal* 47(4): 373–394.
- Son, S. H., David, R., and Chaney, C. 1997. Design and analysis of an adaptive policy for secure real-time locking protocol. *Journal of Information Sciences* 99: 101–135.
- Son, S. H., Chaney, C., and Thomlinson, N. 1998. Partial security policies to support timeliness in secure real-time databases. In *Proceedings of the IEEE Symposium on Security and Privacy*. Oakland, CA, pp. 136–147.
- Son, S. H., and Kang, K.-D. 2002. QoS management in web-based real-time data services. In *Proceedings of the 4th IEEE Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems (WECWIS'02)*. Newport Beach, CA.
- Son, S. H., Lin, Y., and Cook, R. P. 1991. Concurrency control in real-time database systems. In Andre van Tilborg, and Gary, Koob (eds), *Foundations of Real-Time Computing: Scheduling and Resource Management*. Kluwer Academic Publishers, pp. 185–202.
- Son, S. 1987. Using replication for high performance database support in distributed Real-Time systems. In *Proceedings of the 8th IEEE Real-Time Systems Symposium*. p. 7986.
- Son, S., and Kouloumbis, S., 1993. A real-time synchronization scheme for replicated data in distributed database systems. *Information Systems* 18(6).
- Son, S. H., Mukkamala, R., and David, R. 2000. Integrating security and real-time requirements using covert channel capacity. *IEEE Transactions on Knowledge and Data Engineering* 12(6): 865–879.
- Son, S., and Zhang, F. 1995. Real-Time replication control for distributed database systems: algorithms and their performance. In *Proceedings of the 4th International Conference on Database Systems for Advanced Applications*. Singapore.
- Son, S. H., Zimmerman, R., and Hansson, J. 2000. An adaptable security manager for real-time transactions. In *Proceeding of the Euromicro Conference on Real-Time Systems*. Stockholm, Sweden, pp. 63–70.
- Song, X., and Liu, J. W. S. 1992. How well can data temporal consistency be maintained. In *Proceedings of the IEEE Symposium on Computer-Aided Control Systems Design*.

- Song, X., and Liu, J. W. S. 1995. Maintaining temporal consistency: pessimistic vs. optimistic concurrency control. *IEEE Transactions on Knowledge and Data Engineering* 7(5): 786–796.
- Soparkar, N. et al. 1992. Adaptive commitment for real-time distributed transactions. Technical Report TR9215, Department of Computer Science, University of Texas, Austin.
- Stankovic, J., He, T., Abdelzaher, T., Marley, M., Tao, G., Son, S., and Lu, C. 2001. Feedback control scheduling in distributed real-time systems. In *Proceedings of the 22nd IEEE Real-Time Systems Symposium (RTSS 01)*. London, UK, pp. 59–72.
- Stankovic, J. A., Ramamritham, K., and Towsley, D. 1991. Scheduling in real-time transaction systems. In Andre, van Tilborg, and Gary Koob (eds), *Foundations of Real-Time Computing: Scheduling and Resource Management*. Kluwer Academic Publishers, pp. 157–184.
- Stankovic, J. A., Son, S., and Hansson, J. 1999. Misconceptions about real-time databases. *IEEE Computer* 32(6): 29–36.
- Steere, D., Goel, A., Gruenberg, J., McNamee, D., Pu, C., and Walpole, J. 1999. A feedback-driven proportion allocator for real-rate scheduling. In *Proceeding of the 3rd Symposium on Operating Systems Design and Implementation (OSDI 99)*. Berkeley, CA, February 22–25, pp. 145–158.
- Thuraisingham, B., and Ford, W. 1995. Security constraint processing in a multilevel secure distributed database management system. *IEEE Transaction on Knowledge and Data Engineering* 7(2).
- TimesTen. TimesTen Real-Time Event Processing System: Product and Technology Overview. Whitepaper. www.timesten.com.
- Timmerman, B. 1997. A security model for dynamic adaptive traffic masking. In *Proceedings of the New Security Paradigms Workshop*, pp. 1–25.
- Ting, T. C. How Secure is Secure?. *IFIP WG 11.3 Working conference on Database Security*. Opening Remarks.
- Tremblay, M. R., and Cutkosky, M. R. 1995. Using sensor fusion and contextual information to perform event detection during a phase-based manipulation task. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. Pittsburgh, PA.
- Ulusoy, O. 1994. Processing real-time transactions in a replicated database system. *International Journal of Distributed and Parallel Databases* 2(4).
- Vrbsky, S. 1993. *APPROXIMATE: a query processor that produces monotonically improving approximate answers*. Ph.D. thesis, University of Illinois at Urbana-Champaign.
- Vrbsky, S. V., and Lin, K. J. 1988. Recovering imprecise computations with real-time constraints. In *Proceedings of the Seventh Symposium on Reliable Distributed Systems*, pp. 185–193.
- Ahmed, and Vrbsky, S. 2000. Triggered updates for temporal consistency in real-time databases. *Real-Time Systems-Journal* 19(3): 205–208.
- Wei, Y., Son, S., Stankovic, J., and Kang, K. D. 2003. QoS management in replicated real-time databases. In *Proceedings of the 24th IEEE Real-Time Systems Symposium (RTSS 2003)*.
- Wei, Y., Aslinger, A., Son, S., and Stankovic, J. 2004. ORDER: a dynamic replication algorithm for periodic transactions in distributed real-time databases. In *10th International Conference on Real-Time and Embedded Computing Systems and Applications (RTCSA 2004)*.
- Wray, J. 1991. An analysis of covert timing channels. In *Proceedings of the IEEE Symposium on Security and Privacy*. Oakland, CA, pp. 2–7.
- Xiong, M., Sivasankaran, R. M., Stankovic, J., Ramamritham, K., and Towsley, D. 1996. Scheduling transactions with temporal constraints: exploiting data semantics. In *Proceedings of the IEEE 17th Real-Time Systems Symposium*, pp. 240–251.
- Xiong, M., and Ramamritham, K. 1999. Deriving deadlines and periods for real-time update transactions. In *Proceedings of the 20th IEEE Real-Time Systems Symposium*. Phoenix, Arizona.
- Xiong, M., Ramamritham, K., Haritsa, J., and Stankovic, J. A. 2002a. *Mirror: a state-conscious concurrency control protocol for replicated real-time databases information systems*. *An International Journal* 27(4): 277–297.
- Xiong, M., Ramamritham, K., Stankovic, J. A., Towsley, D., and Sivasankaran, R. 2002b. Scheduling transactions with temporal constraints exploiting data semantics. *IEEE Transactions on Knowledge and Data Engineering* 14(5): 1155–1166.
- Xiong, M., and Ramamritham, K. 2004. Deriving deadlines and periods for real Time update transactions. *IEEE Transactions on Computers* 53(5): 567–583.

- Yan, S., Wang, S., and Dou, Z. 2001. An energy model in fire detection and integrated analysis on false alarms. In *Proceedings of the 12th International Conference on Automatic Fire Detection*. Gaithersburg, MD.
- Yoon, Y. 1994. Transaction scheduling and commit processing for real-time distributed database systems. Ph.D. thesis, Korea Advanced Institute of Science and Technology.
- Zeng, X., Bagrodia, R., and Gerla, M. 1998. GloMoSim: a library for parallel simulation of large-scale wireless networks. In *Proceedings of the 12th Workshop on Parallel and Distributed Simulation*. Alberta, Canada.
- Zhou, H., and Jahanian, F. 1998. Real-time primary-backup (RTPB) replication with temporal consistency guarantees. In *Proceedings of ICDCS*.
- Zilberstein, S., and Russell, S. J. 1996. Optimal composition of real-time systems. *Artificial Intelligence* 82(1–2): 181–213.
- Zurko, M., and Simon, R. 1996. User-centered security. In *Proceeding of the New Security Paradigms Workshop*. Lake Arrowhead, CA, pp. 27–33.

Uncorrected Proof