

# AIDA-based Communication for Distributed Time-critical Applications

AZER BESTAVROS  
(best@cs.bu.edu)

Computer Science Department  
Boston University  
Boston, MA 02215

January 20, 1993

## Abstract

In this paper, we propose a superior technique for distributed time-critical communication using AIDA (Adaptive Information Dispersal Algorithm), a novel elaboration on Michael O. Rabin's IDA [Rabi89]. We show that using a small *dynamically-controlled* amount of redundancy, stringent timing constraints imposed on periodic as well as sporadic communication requests in a distributed real-time system can be fulfilled up to any degree of confidence. AIDA is *randomized* in the sense that it does not guarantee the fulfillment of hard time constraints. Instead, it guarantees a lower bound on the *probability* of fulfilling such constraints. We contrast AIDA with traditional communication scheduling techniques used in conjunction with time-critical application in general, and distributed multimedia systems in particular. The suitability of AIDA-based bandwidth allocation for a variety of time-critical application is established and plans for future research experiments are mentioned.

DO NOT DISTRIBUTE WITHOUT AUTHOR'S PERMISSION

# 1 Introduction

The successful execution of a time-critical task running in a distributed environment often requires that a set of communication transactions<sup>1</sup> be successfully completed before a set deadline. In order to guarantee such conditions, an accurate knowledge of the delays introduced by the communication network is often required. For communication scheduling purposes (hereinafter referred to as *bandwidth allocation*), such knowledge can be acquired either statically or dynamically. Using static techniques, worst-case delays are determined and accounted for *a priori* when scheduling the communication transactions. Alternatively, using dynamic techniques the average (or maximum) delays experienced through a communication network can be measured and used as an *estimate* for use with future communication transactions. Static communication scheduling (using a priori knowledge about the communication network delays) can be safely and efficiently used in systems with predetermined communication patterns (*e.g.* , broadcasting) and systems with predetermined computation requirements (*e.g.* , periodic tasks). For systems with unpredictable communication patterns or systems with sporadic computation requirements, dynamic communication scheduling becomes necessary.

Distributed multimedia applications represent an important class of time-critical applications, for which bandwidth allocation is crucial.<sup>2</sup> A distributed multimedia presentation uses audio, video, text, and graphics from local and remote sources. Multimedia data such as audio and video require special considerations when supported by a computer system. They have well-defined (*natural*) presentation timing constraints that must be satisfied by the system during presentation (or *playout*). Examples of natural timing constraints include the maximum tolerable delay in playing out a voice (or video) packet, beyond which dropping the packet might be necessary to avoid disturbing the continuity (smoothness) of the presentation. Other data types (*e.g.* , text and graphics) do not have natural timing constraints, but can be subject to *synthetic* constraints, such as those arising from synchronization requirements (a piece of text or graphics might be required to appear when a particular sequence of video frames is displayed).

Several techniques have been suggested in the literature for dynamic bandwidth allocation. Most of these techniques rely on the use of feedback from the communication network to establish a performance model that can be used in conjunction with a scheduling algorithm to allocate/reserve the communication bandwidth needed for the successful execution of time-critical tasks [Laza90, Litt92a].

While bandwidth allocation is an important consideration in the design of distributed time-critical systems (such as multimedia), it is not the only one. Issues of reliability, availability, and fault-tolerance are equally (if not more) important. The most common technique used to tackle these issues is *replication*. For example, in distributed database applications [Bern87, Elma89],

---

<sup>1</sup>For example fetching a number of pages from a set of remote sites in a distributed shared memory system.

<sup>2</sup>Multimedia applications have received a lot of attention lately among members of the real-time community. Indeed, during the 1992 IEEE RTOS workshop in Pittsburgh, multimedia has been singled out as one of the important applications to be tackled in the nineties [Toku92].

several copies of a particular data object might be kept in a number of different sites so that the failure (whether intermittent or permanent [Rand78]) of any proper subset of these sites would not render the data object unavailable. For distributed applications operating under strict time constraints, replication alone might not be sufficient. In particular, failures should not be allowed to increase the retrieval delay for data objects (at least not considerably). In this respect, techniques that rely on watchdog timers and/or retransmit protocols may not be adequate. Instead, techniques that use redundant communication (*e.g.*, requesting the same data object from a set of failure-independent sites/paths) might be necessary. This, however, might have an adverse impact on the overall performance of the system due to the added “redundant” communication traffic.

In this paper we propose AIDA, a novel technique for dynamic bandwidth allocation, which makes use of minimal, controlled redundancy to guarantee timeliness and fault-tolerance up to any degree of confidence. Our technique is an elaboration on the Information Dispersal Algorithm of Michael O. Rabin [Rabi89], which we have previously shown to be a sound mechanism that considerably improves the performance of I/O systems and parallel/distributed storage devices [Best89a, Best89b].

## 2 Real-time Bandwidth Allocation: Related work

A real-time communication system must manage the communication of time-dependent data to provide timely and predictable data delivery; it provides performance guarantees for the delivery of data from source to destination. These guarantees can be absolute through deterministic scheduling and resource allocation, or probabilistic through the use of statistical approaches. In this section, we review a few of the representative techniques currently being used/investigated for real-time bandwidth allocation in multimedia applications.

One way to schedule data transmission is to maintain statistics characterizing each of the communication resources (channels) in the system. Whenever the channel characteristics of the network change, the server responsible for delivering the time-critical data can adjust accordingly to maintain predictable service. This can be achieved by decreasing the demand on the network. For example, when a network becomes congested and the percentage of late data elements (missed deadlines) increases, dropping the demand on the network helps clear the congestion [Gilg91]. This effectively allows data elements scheduled for transmission to traverse the communication network and reach their destination *on time* rather than be lost due to lateness.

Another mechanism to deal with the adverse effect of network congestion is to distinguish between the various communication requirements. This was proposed in the Asynchronous Time-sharing System (ATS) [Laza90], in which data traffic is divided into four classes. A control class  $C$  has the highest priority; it delineates a class of communication where data loss or unpredictable communication delays cannot be tolerated. Class I is next on the priority scale; it delineates a class of communication where data loss cannot be tolerated, but a user-specified maximum end-to-end communication delay is allowed. Class II has a set maximum percent of lost packets and a maximum count of consecutive packets lost. Finally, class III has zero loss and no maximum end-to-end

delay for communication that is not subject to time constraints. A similar treatment of the different communication requirements imposed on a distributed system is under investigation at the University of California at Berkeley, where an experimental RAID-II network file server that treats requests with a low end-to-end delay requirement differently from requests with a high bandwidth requirement is being implemented [Lee92].

The network protocol presented in [Ferr90, Ferr91] handles performance requirements in a different manner. When a connection is requested, the user provides the network manager with maximum end-to-end delay, maximum packet size, maximum packet loss rate, minimum packet inter-arrival time, and maximum jitter, where jitter is defined as the difference in the delays experienced between two packets on the same connection. Three types of channels can be requested: *deterministic*, *statistical*, and *best-effort*. For deterministic channels, the communication delay is guaranteed to fall below a given time bound. For statistical channels, the probability that the delay is less than a given time  $D$  is kept greater than or equal to a requested factor  $q$ . This can be thought of as establishing a *confidence interval* about the expected delay rather than a *deterministic bound* on that delay. Best-effort channels provide no guarantees for the percentage of messages reaching their destination on-time; they merely attempt to make the best use of the available bandwidth.<sup>3</sup>

Statistical approaches to overcoming delay and bandwidth limitations are attractive because they provide application programs with a flexible framework, in which a continuum of communication priorities can be easily expressed as confidence intervals. In particular, we argue that the distinction between deterministic, statistical, and best-effort channels in the protocol proposed in [Ferr90, Ferr91] is artificial. Deterministic and best-effort channels can be thought of as *special* statistical channels, for which the confidence interval (determined by  $q$ ) describing the communication delay is taken to its limits.<sup>4</sup> Therefore, in this paper (without loss of generality), we consider only statistical channels.

Current techniques for statistical bandwidth allocation [Litt92b] rely on choosing an end-to-end time delay  $T$  per packet that is larger than the delay expected to be experienced by a percentage  $P$  of the retrieved packets. This time  $T$  is used as an estimate for the time it will take to retrieve packets from a given source. Figure 1 illustrates a typical relationship between  $T$  and  $P$ . While such a delay function can accurately represent delay characteristics over a given period of time, network loading does change with time, possibly making the delay distribution (and thus the delay function) outdated. One way to accommodate this dynamic behavior is to monitor the delays experienced by retrieved packets and adjust the delay function accordingly. In [Gibb92], a mechanism called *Limited A Priori* (LAP) scheduling is proposed, in which adjustments to the delay function are made either periodically or whenever sudden changes in network traffic are detected. Using second and higher order moments, linear and quadratic extrapolation of the network delay characteristics can be more accurately predicted. This research work, however, is yet to be pursued.<sup>5</sup>

---

<sup>3</sup>Deterministic channels are necessary for computations with *hard* time constraints, whereas statistical channels are appropriate for computations with *soft* time constraints. Best-effort channels are adequate for computations with no time constraints.

<sup>4</sup>For deterministic channels,  $q = 1$ . For best-effort channels,  $q = 0$ .

<sup>5</sup>For more information, please contact the author.




Figure 1: A typical end-to-end delay characteristic function.

All of the bandwidth allocation mechanisms described so far (with the exception of RAID-II) assume a single source of data for a given transaction. In a truly distributed environment, this is not likely to be the case. The storage of a single object might span a number of nodes, either because of fault-tolerance requirements<sup>6</sup> or else to accommodate data placement constraints. Although it is possible to extend the aforementioned bandwidth allocation mechanisms to deal with data distributed over a number of nodes, the performance of these protocols deteriorates significantly. The mechanism we are proposing in this paper is inherently distributed and, in that respect, is far more superior.

### 3 Information Dispersal and Retrieval Using IDA

In this section we overview the original Information Dispersal Algorithm (IDA). We refer the reader to the original paper on IDA [Rabi89] for a more thorough presentation.

Let  $F$  represent the original data object (hereinafter referred to as the *file*) in question. Furthermore, let's assume that the storage of file  $F$  is to be distributed over  $N$  sites. Using the IDA algorithm, the file  $F$  will be processed to obtain  $N$  distinct pieces in such a way that recombining *any*  $m$  of these pieces,  $m \leq N$ , is sufficient to retrieve  $F$ . The process of processing  $F$  and distributing it over  $N$  sites is called the *dispersal* of  $F$ , whereas the process of retrieving  $F$  by collecting  $m$  of its pieces is called the *reconstruction* of  $F$ .

Figure 2 illustrates the dispersal and reconstruction of an object using IDA. The dispersal and reconstruction operations are simple linear transformations using *irreducible polynomial arithmetic*.<sup>7</sup> Both the dispersal and reconstruction of information using IDA can be performed in

<sup>6</sup>For example, striping data for a video presentation over  $N$  nodes would increase the *availability* of the system by allowing a graceful degradation of the quality of the presentation by  $1/N\%$ , should any of the  $N$  nodes fail.

<sup>7</sup>For a concrete implementation and for examples, the reader is referred to our previous work on SETH [Best90]

real-time. This was demonstrated in [Best90], where we presented an architecture and a CMOS implementation of a VLSI chip<sup>8</sup> that implements IDA.

Let  $|F|$  be the size of the file  $F$ . The IDA approach *inflates*  $F$  by a factor of  $\frac{N}{m}$ . In particular, the size of each one of the dispersed pieces of  $F$  would be  $\frac{|F|}{m}$ . This added redundancy makes the system capable of tolerating up to  $N - m$  faults without any effect on timeliness. More importantly (as we will demonstrate shortly), this added redundancy will boost the performance of the information retrieval process *significantly*.




Figure 2: Dispersal and reconstruction of information using IDA.

Several *redundancy-injecting* protocols have been suggested in the literature to deal with fault-tolerance issues. In most of these protocols, redundancy is injected in the form of parity blocks, which are only used for error detection and/or correction purposes [Gibs88]. The IDA approach is radically different in that redundancy is added *uniformly*; there is simply *no* distinction between data and parity. It is this feature that makes it possible for IDA to be used not only to boost communication fault-tolerance, but also to improve bandwidth allocation and utilization.

An important aspect of IDA is that, unlike other redundancy-injecting protocols [Schu89, Lee92], the amount of redundancy to be used with a given object, or in a given communication session, does *not* have to be constant. In particular, as we will describe later, our AIDA-based bandwidth allocation strategy *controls* the amount of redundancy to be used with a particular object in a particular communication session so as to reflect the *priority* and/or the *urgency* of the transaction at hand. By increasing the redundancy allocated for a given communication session, the expected retrieval delay can be reduced, thus increasing the chances of meeting the possibly tight time constraint imposed on the transaction.

and IDA-based RAID I/O systems [Best91].

<sup>8</sup>The chip (called SETH) has been fabricated by MOSIS and tested in the VLSI lab of Harvard University, Cambridge, MA. The performance of the chip was measured to be about 1 megabyte per second. By using proper pipelining and more elaborate designs, this figure can be boosted significantly.

## 4 Performance Characteristics

Let  $X$  be a data object dispersed using IDA into  $N$  pieces, each residing in a different site. Let  $m$  be the minimum number of pieces needed to reconstruct  $X$ . Obviously, in order to retrieve  $X$ , at least  $m$  of the  $N$  sites must be consulted. It is possible, however, to consult more than  $m$  of these sites. Let  $n$  (where  $m \leq n \leq N$ ) denote the total number of sites *consulted* for the retrieval of  $X$ . In this section, we derive an expression for the expected communication delay for accessing such an object. Later, we will use this result to establish the merits of our proposed AIDA-based bandwidth allocation protocol.

$$\text{Prob}(t \geq z) = \text{Prob}(\text{Response time of at least } (n - m) \text{ of the sites } \geq z) \quad (1)$$

$$= \sum_{r=n-m+1}^n \binom{n}{r} P^r (1 - P)^{n-r} \quad (2)$$

where  $P$  is the probability that the response time of a single site will be  $z$  or more.  $P$  can be estimated using delay characteristic functions such as the one illustrated in figure 1.

### 4.1 Approximation using a uniform distribution delay model

As a first and safe approximation, we will assume that the delays experienced through the communication network are *uniformly distributed* random variables with lower and upper bounds ( $D_{\min}$  and  $D_{\max}$ ) as illustrated in figure 3. We denote by  $P_u$  the value of  $P$  (in equation 2) under the uniform distribution assumption.

$$P_u = \begin{cases} 1 & \text{if } 0 < z < D_{\min} \\ 1 - \frac{z - D_{\min}}{D_{\max} - D_{\min}} & \text{if } D_{\min} \leq z \leq D_{\max} \\ 0 & \text{if } D_{\max} < z < \infty \end{cases}$$




Figure 3: End-to-end delay characteristic function under the uniform delay assumption.

The random variable  $t$  (in equation 2) is simply the  $(n - m + 1)^{\text{th}}$  largest of these  $n$  uniformly distributed independent random variables. It can be shown that  $t$  follows the *beta probability law* and that the mean and standard deviation for  $t$  are given by:<sup>9</sup>

$$\tau_u = D_{\min} + \frac{m}{n+1}(D_{\max} - D_{\min}) \quad (3)$$

$$\sigma_u = \sqrt{\frac{m(n-m+1)}{(n+1)^2(n+2)}}(D_{\max} - D_{\min}) \quad (4)$$

## 4.2 Approximation using an exponential distribution delay model

We denote by  $P_e$  the value of  $P$  (in equation 2) under the exponential distribution delay model.

$$P_e = \begin{cases} 1 & \text{if } 0 < z < D_{\min} \\ 1 - \lambda \cdot e^{-\lambda(z - D_{\min})} & \text{if } D_{\min} < z < \infty \end{cases}$$




Figure 4: End-to-end delay characteristic function under the exponential delay assumption.

Let  $\tau_e$  denote the average delay experienced using an IDA-based strategy under an exponential distribution delay model with parameter  $\lambda$  (see figure 4). To compute  $\tau_e$ , we need to compute the *mean* value of the random variable  $t$ . This can be done as follows:

---

<sup>9</sup>Derivation is omitted for space limitations. For a reference, refer to [Lars82].

$$\begin{aligned}
\tau_e &= \text{Expected delay under an exponential distribution assumption} \\
&= D_{\min} + E(t - D_{\min}) \\
&= D_{\min} + \int_{D_{\min}}^{\infty} \text{Prob}(t \geq z) \cdot dz \\
&= D_{\min} + \sum_{r=n-m+1}^n \binom{n}{r} \int_{D_{\min}}^{\infty} P^r (1-P)^{n-r} \cdot dz \\
\tau_e &= D_{\min} + \sum_{r=n-m+1}^n \binom{n}{r} \int_0^1 P^r (1-P)^{n-r} \cdot \frac{1}{\lambda P} \cdot dP \\
&= D_{\min} + \frac{1}{\lambda} \sum_{r=n-m+1}^n \binom{n}{r} \frac{\Gamma(r)\Gamma(n-r+1)}{\Gamma(n+1)} \\
&= D_{\min} + \frac{1}{\lambda} \sum_{r=n-m+1}^n \frac{\Gamma(n+1)}{\Gamma(r+1)\Gamma(n-r+1)} \frac{\Gamma(r)\Gamma(n-r+1)}{\Gamma(n+1)} \\
\tau_e &= D_{\min} + \frac{1}{\lambda} \sum_{r=n-m+1}^n \frac{1}{r} \tag{5}
\end{aligned}$$

Unless stated otherwise, the remainder of this paper assumes an exponential delay model.

### 4.3 Effect of distribution and redundancy on delay characteristics

There are a number of interesting observations to be made from the delay analysis of the previous section. By varying the values of  $n$  and  $m$ , the negative effect of data distribution and the positive effect of data redundancy on the delay characteristics can be demonstrated. The following cases can be readily examined:

- a.  $\underline{n = m = 1}$ : This is the case when the object  $X$  is not distributed. The expected delay reduces to  $\frac{1}{2}(D_{\min} + D_{\max})$  under the uniform delay model and reduces to  $D_{\min} + \frac{1}{\lambda}$  under the exponential delay model. This corresponds to the average delay for one transmission.
- b.  $\underline{n > m = 1}$ : This is the case when the object  $X$  is replicated over  $n$  sites. For  $n \gg 1$ , the expected delay approaches  $D_{\min}$ , which is the minimum delay for one transmission under both the uniform and exponential delay models.
- c.  $\underline{n = m > 1}$ : This is the case when the object  $X$  is distributed without any added redundancy. For  $n \gg 1$ , the expected delay approaches  $D_{\max}$ , which is the maximum delay for one transmission under the uniform delay model. Under the exponential delay model, the expected delay approaches  $D_{\min} + \frac{1}{\lambda} \ln(n)$ , making the communication delay logarithmically proportional to the distribution level.

IDA-based communication attempts at striking a balance between the above three extreme setups. Figure 5 illustrates the improvement (speedup) in communication delay that can be achieved through the use of even remarkably small levels of redundancy.<sup>10</sup> For example, at a 20% redundancy level ( $\frac{1}{5}$  of the communicated data is redundant), IDA cuts the expected delay through a communication network by almost 50% (a 2-fold speedup) for an object distributed over 8 sites. This gain is even larger for objects distributed over a larger number of sites. If the level of redundancy is increased further, the gain is substantial. For example, IDA can deliver a 5-fold speedup in communication with the redundancy level set at 50% for an object that is distributed over 32 sites.

For the same amount of redundancy, other protocols (such as replication) yield minuscule speedups compared to IDA. For example, if an object is replicated once ( $\frac{1}{2}$  of the communicated data is redundant) and each of the two replica is distributed over 16 sites (for a total of 32-site distribution), then it can be shown that under the exponential delay model, the achievable speedup will be less than 1.1-fold. Under the same conditions, IDA delivers over 5-fold speedups.

## 5 Fault-tolerance and Security Characteristics

The storage and transmission of data in a distributed system raises significant security and reliability problems. In particular, data might be lost due to hardware failures, it might be accidentally (or even maliciously) garbled or destroyed, and it might be read and interpreted by unauthorized tapping of communication links. In this section, we examine the fault-tolerance and security characteristics of our proposed bandwidth allocation strategy.

### 5.1 Effect of distribution and redundancy on fault-tolerance

The usual technique employed to deal with communication failures is to retransmit on errors (or timeouts). For time-critical applications, this detect-then-recover approach might not be feasible due to the time constraints imposed on the system. Instead, masking techniques are employed. In particular, error-correcting codes are used to tolerate communication failures, whereas replication and/or n-modular redundancy (NMR) techniques are used to protect against site failures [Rand78]. The main drawback of these techniques is their excessive use of redundancy, which might adversely affect performance. For example, to mask one site failure an approach relying on replication will require that a particular object be retrieved from two different sites, thus doubling the network traffic. The blowup is even larger when error-correction for a relatively small number of communication-induced errors is taken into account.

The AIDA-based protocol we are proposing in this paper is a failure-masking protocol that is provably optimal in its use of redundancy. The main reason for AIDA's superiority is that it does not distinguish between communication failures and site failures, thus making the best use of allotted redundancy in the system. To tolerate up to  $r$  simultaneous failures – be it site failures or communication failures – AIDA requires that the total number of sites from which the dispersed

---

<sup>10</sup>These results were obtained under an exponential delay model, but can be easily reproduced for any other model.




Figure 5: Expected AIDA speedups – Only the random part of the delay is considered.

object  $X$  will be requested exceeds the minimum number of data pieces needed to reconstruct  $X$  by  $r$ . Thus, a total of  $n = m + r$  sites is needed for every  $m$  pieces of data, a redundancy of  $100(n - m)/m$  percent. For example, if an object  $X$  is to be dispersed over  $n = 12$  sites and coverage for up to 3 failures is required, then using AIDA, the total redundancy would be 25% (a blowup of 33%). To provide the same coverage using replication, the total redundancy would soar to 75% (a blowup of 300%).

## 5.2 Effect of distribution and redundancy on security

A common technique to insure communication security is to store and communicate information using some form of encryption, where only authorized users are enabled to decrypt the information through the use of appropriate *secret keys* [Sham79]. The proven difficulty of decrypting the information without knowing the secret key guarantees a high level of *security*. The main disadvantage of this technique is that the information (although encrypted) is available in one site – whether stored in or communicated through that site – for long periods of time. This might make it possible for adversaries to break the secret key of the encryption.

The AIDA-based protocol we are proposing in this paper guarantees the security of the communicated information by making it unavailable as a whole in any one particular site. As a matter of fact, it is hard to get any clue about the original information unless at least  $m$  pieces from the dispersed file are collected. This makes the task of the adversaries more difficult, since they have to control  $m$  of the sites and not only one. Even if this happens, it is provably very difficult to reconstruct the original file unless the secret key is known.

## 6 AIDA-based Bandwidth Allocation

In this section, we highlight the features of AIDA that enable it to deal effectively with deadline and priority issues in time-critical systems.

### 6.1 Using redundancy to control communication delays

Let the retrieval of an object  $X$  be subject to a soft time-constraint that requires  $X$  to be fetched within  $T_{\max}^X$  units of time. According to equation 5 the expected delay in retrieving  $X$  decreases *predictably* as  $n - m$  increases. Incorporating the time constraint in equation 5, we can solve for  $n$  as follows.

$$\begin{aligned} T_{\max}^X &\geq \tau_e \\ &\geq D_{\min} + \frac{1}{\lambda} \sum_{r=n-m+1}^n \frac{1}{r} \\ \lambda(T_{\max}^X - D_{\min}) &\geq \sum_{r=n-m+1}^n \frac{1}{r} \end{aligned}$$

Using the lower bound  $[\ln(n) - \ln(n - m)]$  to approximate the value of  $\sum_{r=n-m+1}^n \left(\frac{1}{r}\right)$ , we get:

$$\lambda(T_{\max}^X - D_{\min}) > \ln(n) - \ln(n - m)$$

Solving the above inequality for a lower bound on  $n$  we get:

$$n > \frac{m}{1 - e^{-\lambda(T_{\max}^X - D_{\min})}} \quad (6)$$

In order to compute the appropriate value of  $n$  using equation 6, it is necessary to evaluate *dynamically* the values of  $D_{\min}$  and  $\lambda$ . This can be done using statistical techniques similar to the those described in [Gibb92].

## 6.2 Priority-based rationing of redundant bandwidth

Equation 6 establishes a lower bound on  $n$  that guarantees an *expected* communication delay, not an *actual* communication delay. In other words, while it is very possible for the actual communication delay to be less than the desired expected delay (thus satisfying the imposed time constraint), it is very possible as well for the actual communication delay to exceed the desired expected delay (thus resulting in a violation of the imposed time constraint). This *randomness* factor can be accounted for and controlled by using second order moments (*e.g.*, Standard Deviation) to build a *confidence interval* about the actual communication delay. One way of building such a confidence interval is to set the value of  $n$  so as to make  $T_{\max}^X$ , the available slack for completing the communication session, greater than or equal to  $\tau_e + \alpha\sigma_e$  (rather than simply  $\tau_e$ ).

$$n > \frac{m}{1 - e^{-\lambda(T_{\max}^X - D_{\min} - \alpha\sigma_e)}} \quad (7)$$

The value of  $n$  in the above equation defines a confidence interval that corresponds to a specific probability of meeting the time constraint imposed on the communication session. This probability can be made arbitrarily high by increasing the value of  $\alpha$ . This, however, is not without cost. In a distributed real-time system, the total communication bandwidth is *finite*, and increasing the amount of *redundant* information flowing in the system might adversely affect the end-to-end delay characteristics that we were aiming to improve in the first place!

One way of solving the aforementioned problem is to set the value of  $\alpha$  in such a way that the total available bandwidth in the system is *rationed* among the different communication sessions in a way that reflects the *priority* assigned to these sessions. In other words, the value of  $\alpha$  for a particular task is related to its priority *and* the priority of all the other tasks sharing the available bandwidth in the system.

It is important to notice that using AIDA, the priority of the transaction (how critical it is to the mission of the system) and the urgency of the transaction (how tight its time constraint

is) are *both* taken into account when the value of  $n$  is determined. This stands in sharp contrast with protocols that deal only with either the priority of the transaction *or* its urgency, making it necessary for applications to express (artificially) one of these attributes using the other.

## 7 Conclusion

Communication timeliness, fault-tolerance, and security are at the heart of many time-critical distributed computing systems. This has been aptly expressed in a recent broad agency announcement by the US Army Research Laboratory: “On the battlefield, *efficient*, *reliable*, and *timely* communications over a variety of radio links is vital to mission success” [ARL92].

AIDA is a novel bandwidth allocation strategy suitable for distributed fault-tolerant time-critical systems. In AIDA redundancy is used to tackle several crucial problems. In particular, redundancy is used to *tolerate* failures, to *increase* the likelihood of meeting tight time-constraints, and to *ration* (based on task priorities) the limited bandwidth in the system. AIDA is a *randomized* protocol in the sense that it does not guarantee the fulfillment of hard time constraints. Instead, it guarantees a lower bound on the *probability* of fulfilling such constraints. In this paper we have presented AIDA’s potentials and established its superiority with respect to existing protocols. The implementation of an AIDA-based network file-server that would act as an interface between the application programs and the communication network is our next goal.

Several variants of the basic ideas presented in this paper are under investigation. For example, we are evaluating a number of possible mechanisms (both centralized and distributed) to be used for the selection of the  $n$  out of  $N$  sites to be consulted for an object retrieval. While the correctness and efficacy of AIDA are not dependent on such mechanisms, its performance might benefit greatly. Similar performance gains can be achieved by classifying communication requests as was done in [Laza90, Lee92]. In this paper there is an inherent assumption that communication requests are all homogeneous – for example, they are all requests for equally-sized data granules (*e.g.*, pages of memory). If such an assumption is not preserved, then it might be useful to classify communication requests based on their requirements and maintain different estimates about the network delay characteristics for each one of these classes. Such a treatment is likely to reduce the uncertainty associated with communication delays, thus providing for a more efficient allocation of bandwidth.

Another area of particular interest is the susceptibility of AIDA to sudden changes in the communication delay characteristics, whether as a result of an influx of sporadic communication requests, or as a result of a sudden decrease in bandwidth due to failures. Simulations that would relate such susceptibilities to system design parameters (for example frequency and sample size to be used for estimating the *actual* delay characteristics), are underway.

In this paper we focussed on information retrieval. Issues pertaining to information *update* were not tackled. These issues are particularly important in distributed time-critical systems to ensure data consistency and recency. In particular, it is of utmost importance to investigate the interaction between AIDA and other consistency-preserving protocols such as distributed shared memory protocols [Teva87], caching protocols [Arch86], and non-coherent memory protocols [Hedd93].

## References

- [Arch86] J. Archibald and J-L. Baer. “Cache coherence protocols: Evaluation using a multiprocessor simulation model.” *ACM Transactions on Computer Systems*, 4(4):273–298, November 1986.
- [ARL92] ARL. “United States Army Research Laboratory Broad Agency Announcement.” . Commerce Business Daily, October 1992. 37(2), Page 8.
- [Bern87] A. Bernstein, A. Philip, V. Hadzilacos, and N. Goodman. *Concurrency Control And Recovery In Database Systems*. Addison-Wesley, 1987.
- [Best89a] Azer Bestavros. “IDA-based disk arrays.” Technical Memorandum 45312-890707-01TM, AT&T, Bell Laboratories, Department 45312, Holmdel, NJ, July 1989.
- [Best89b] Azer Bestavros, Danny Chen, and Wing Wong. “The reliability and performance of parallel disks.” Technical Memorandum 45312-891206-01TM, AT&T, Bell Laboratories, Department 45312, Holmdel, NJ, December 1989.
- [Best90] Azer Bestavros. “SETH: A VLSI chip for the real-time information dispersal and retrieval for security and fault-tolerance.” In *Proceedings of ICPP'90, The 1990 International Conference on Parallel Processing*, Chicago, Illinois, August 1990.
- [Best91] Azer Bestavros. “IDA disk arrays.” In *Proceedings of the First International Conference on Parallel and Distributed Information Systems*, Miami Beach, Florida, December 1991.
- [Elma89] Ramez Elmasri and Shamkant Navathe. *Fundamentals of Database Systems*. The Benjamin/Cummings Publishing Company Inc., 1989.
- [Ferr90] D. Ferrari and D.C. Verma. “A scheme for real-time channel establishment in wide-area networks.” *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.
- [Ferr91] D. Ferrari. “Design and application of a delay jitter control scheme for packet-switching internetworks.” In *Proceedings of the second International Conference on Network and Operating System Support for Digital Audio and Video*, Heidelberg, Germany, November 1991.
- [Gibb92] John F. Gibbon, Azer Bestavros, and Tom Little. “Limited a priori scheduling for distributed multimedia systems.” . Work in progress, November 1992.
- [Gibs88] Garth Gibson, Lisa Hellerstein, Richard Karp, Randy Katz, and David Patterson. “Coding techniques for handling failures in large disk arrays.” Technical Report UCB/CSD 88/477, Computer Science Division, University of California, July 1988.
- [Gilg91] M. Gilge and R. Gussella. “Motion video coding for packet-switching networks – an integrated approach.” In *Proceedings of the SPIE Conference on Visual Communications and Image Processing*, Boston, MA, September 1991.
- [Hedd93] Abdelsalam Heddaya and Himanshu S. Sinha. “An overview of MERMERA: a system and formalism for non-coherent distributed parallel memory.” In *Proceedings of the 26th Hawaii International Conference on System Sciences*, January 1993.
- [Lars82] Harold Larson. *Probability theory and statistical inference, Third Edition*. John Wiley & Sons, 1982.
- [Laza90] Aurel A. Lazar, Adam Temple, and Rafael Gidron. “An architecture for integrated networks that guarantees quality of service.” *International Journal of Digital and Analog Cabled Systems*, 3(2), 1990.
- [Lee92] Edward Lee, Peter Chen, John Hartman, Ann Drapeau, Ethan Miller, Randy Katz, Garth Gibson, and David Patterson. “RAID-II: a scalable storage architecture for high-bandwidth network file service.” Technical Report CSD 92/672, University of California at Berkeley, Spring 1992.

- [Litt92a] T.D.C. Little and A. Ghafoor. “Scheduling of bandwidth-constrained multimedia traffic.” *Computer Communications (Special Issue on Multimedia Communications)*, 15(6):381–387, July/August 1992.
- [Litt92b] T.D.C. Little and J.F. Gibbon. “Management of time-dependent multimedia data.” In *Proceedings of the SPIE Symposium OE/FIBERS’92: Enabling technologies for multimedia, multiservice networks*, Boston, MA, September 1992.
- [Rabi89] Michael O. Rabin. “Efficient dispersal of information for security, load balancing and fault tolerance.” *Journal of the Association for Computing Machinery*, 36(2):335–348, April 1989.
- [Rand78] B. Randell, P. Lee, and P. Treleaven. “Reliability issues in computing system design.” *ACM Computing Surveys*, 10:84–98, June 1978.
- [Schu89] Martin Schulze, Garth Gibson, Randy Katz, and David Patterson. “How reliable is a RAID?.” In *Proceedings of COMPCON-89, the Thirty-fourth IEEE Computer Society International Conference*, March 1989.
- [Sham79] A. Shamir. “How to share a secret?.” *Communication of the ACM*, 22:612–613, November 1979.
- [Teva87] Avadis Tevanian and (et al). “A Unix interface for shared memory and memory mapped files under Mach.” Technical report, Carnegie-Mellon University, Department of Computer Science, February 1987.
- [Toku92] Hideyuki Tokuda. “Summary of the recommendations of the 1992 IEEE RTOS workshop.” . Pittsburgh, PA, May 1992.