Applications for the Clustering, SAN and Fibre Channel Lab

Sharad Mehrotra, Alex Veidenbaum, and Nalini Venkatasubramanian
424 Computer Science
University of California at Irvine
Irvine, CA 92612
tel: 949 824 5975
fax :949 824 4056
email: (sharad,alexv,nalini)@ics.uci.edu

November 19, 1999

Introduction

The primary research thrusts of the faculty involved with the clustering laboratory include:

- **Data Management and Data Mining**

  Support for complex multidimensional datasets, management of highly dynamic data sets, spatio-temporal database management, management of uncertainty in database systems, machine learning and pattern recognition over large data sets, visualization of dynamic data sets.

- **Clustering for Numerical and Scientific Computing**

  Scalability, memory hierarchy performance, parallel I/O, compilation, application issues, portable libraries including message passing.

- **Distributed Systems**

  Network centric computing, availability and fault tolerance in distributed and parallel systems, replication and caching in distributed environments, open and heterogeneous systems environments, Quality of Service enabled resource management, composable middleware for distributed systems, concurrent object based computing, ubiquitous and mobile computing environments, security protocols, systems and performance analysis.

- **Multimedia Management**

  Analysis and interpretation of multimedia data, representation and storage of multimedia information, multimedia information retrieval and database management, communication and transmission of multimedia objects, multimedia in mobile environments, quality of service specification and enforcement in multimedia applications.

1 Cluster Enabled Applications

We describe two specific applications designed to test the limits of scalability and availability in the design and management of cluster based environments. The number of concurrent users in both applications are likely to scale to large numbers; one of our goals is to test system performance under these extreme conditions. Multimedia applications, such as video-on-demand specify QoS requirements that define the extent to which application requirements such as timeliness can be violated. We intend to test the impact of different failure modes on application QoS for applications executing in the cluster framework.

1.1 WEBMARS – Multimedia Search Engine for the World Wide Web

MARS is a integrated multimedia information retrieval and database management system being developed at the University of California, Irvine with the focus of providing seamless integration of multimedia information in database management systems. Our objectives are to develop efficient techniques to support content-based similarity retrieval over multimedia data (specially images and video) and explore mechanisms to integrate such techniques as access methods into database management systems. As part of the MARS project, we have built a MARS retrieval system that provides content-based retrieval over multimedia objects consisting of image properties along with textual descriptions. The existing system consists of (1) a MARS server that supports similarity queries over multimedia objects. Similarity is computed based on both textual and visual properties, (2) a client that supports insertion and deletion of new/old objects into the database, and (3) a web-based client that allows users to construct queries (using both example images and text), display answers, and to refine queries by providing relevance feedback.

The objective of WEBMARS is to build a multimedia web search engine that uses the MARS server to provide retrieval of matching HTML pages based both on their textual and visual content. The proposed search engine will provide a vehicle to explore effectiveness of combining the text-based and image-based retrieval paradigms into a single search framework for the Web environment.

Many existing search engines (both commercial systems and research prototypes) have explored similarity retrieval over visual objects and textual data. However, to the best of our knowledge, these efforts have concentrated on exploring these paradigms independently. Often, user's information need is best expressed using multiple modalities. As an example, consider a user, who while browsing the web, comes across an image of a dress with a particular pattern that she likes in an online fashion magazine. The user might then be interested in retrieving matching dresses (based on visual properties) available at a local store in her size at a given price range (based on textual properties). It is difficult to express the user's information need using only visual properties or the textual properties in this case. Another advantage of supporting an integrated retrieval framework is that often, information in different modalities can complement each other improving the effectiveness of retrieval. As an example of the potential improvements, consider that the text surrounding the image in the fashion magazine in the previous example contained a description of the style of the dress. Again, utilizing this textual description along with the visual properties in evaluating the query will improve search performance.

WebMARS will support techniques to formulate and execute integrated queries over multiple modalities (images and text) using the MARS server. Another novelty of the search engine developed are techniques to support query refinement using relevance feedback. We have, in the past, developed many relevance feedback techniques for multimedia retrieval and explored efficient mechanisms to implement refined queries. These techniques have been built into the MARS server and will be available to the search engine.

Developing a search engine to support an integrated retrieval framework requires research in diverse interrelated fields. The first issue is the design of an interface that allows users to express their queries using different modalities. Another challenge is to extract and represent semantically meaning features from both the textual descriptions and the images from the multimedia objects (viz., HTML pages). Finally, a model to represent content of HTML pages, retrieval models to match a multi-modality query with the content description, weighting techniques to give appropriate importance to different modalities in the query, relevance feedback techniques to support effective refinement of multi-modality queries, and efficient techniques to evaluate such queries need to be developed. Many of these research issues are being explored as part of the MARS project.

We are currently building an experimental version of WEBMARS that allows us to explore the above mentioned research issues. WEBMARS will also provide us with a tool to experimentally explore and quantify the improvements in retrieval effectiveness that result by integrating multiple modalities in a single retrieval framework. The search engine has the following components:

- **A Web Crawler:** that navigates the Web to populate the MARS database.
- **An HTML processor:** that parses the HTML files to extract (1) links to be followed by the crawler; (2) text to be indexed by the text indexer into the database; (3) images linked to the HTML page to be indexed by the image indexer.
- **Text Indexer:** that extracts meaningful text from HTML pages which capture its content. The text indexer will also construct a summary of the text to aid browsing of the collection. While text

summarization is a rich and complex field in information retrieval, we will use well established simple techniques in our initial development (extracting first lines from paragraphs, spotting discriminating words).

- **Image Indexer:** this component will extract meaningful image properties (e.g., colors, shape, layout of colors and textures). Current MARS image analyzer will be used for this purpose. However, the existing analyzer, which only deals with global image properties, will be extended to extract and represent local properties of objects in an image. Extracting local properties requires a difficult problem of object segmentation. We will bypass this problem by providing an interactive tool which allows a human-indexer to draw an outline of salient objects in an image. Once the object has been outlined, the objects can be segmented and represented based on their local properties such as color, texture, shape, etc. The system will allow for asynchronous extraction of local properties of objects. That is, when an image first arrives into the system, it will be represented using its global properties which can be automatically extracted. As the collection builds, a human-indexer will outline salient objects in images contained in the collection. The retrieval system will be flexible in the sense that local properties of objects will be used if and when they are available. Else, retrieval will be based only on global properties.

- **Retrieval Subsystem:** The retrieval subsystem including a multi-modality retrieval model, techniques for query refinement, and efficient query processing are already being built in the context of the MARS system; the same techniques will be used within WEBMARS. We are currently extending the retrieval subsystem to adapt to the representation and and retrieval of HTML documents. For example, the current multimedia object model supported by the MARS server does not handle local properties of objects in an image. Both the model and the code are being extended to enable retrieval based on local properties.

- **User Interface:** The user interface developed will not only allow users to pose integrated queries over multiple modalities but will also provide a powerful model of navigating the Web which combines both browsing and retrieval into a single framework. A user, while browsing the Web, would be able to mark web pages he or she finds interesting. Later a user can construct a query based on components (or parts) of these Web pages that are relevant to his/her information need. The system will support a flexible interface that allows a user to input and construct queries at multiple levels – a user can specify keywords, identify a (set of) interesting images, choose specific property of images, or specific objects within an image, choose a HTML page for its text content, or choose a specific part of text when constructing the query. The query expressed will be mapped to an internal representation supported by the MARS server and the answers returned by the server will be displayed to the user along with the relevance scores. A user, if not satisfied with the answers, may provide a feedback to the system based on which the server will adjust the query (models for adjusting queries based on the feedback are already built into MARS server). A user may continue with multiple iterations of feedback until his/her information need is satisfied.

1.2 Multimedia Distance Learning Toolkit

The distributed multimedia distance learning toolkit will utilize a cluster computing environment to provide Quality of Service enabled multimedia delivery to multiple clients at remote locations simultaneously.

The toolkit will contain the following components

- An Integrated Multimedia-on-Demand System for Lectures
- A Multimedia Conferencing Component for Interactive Learning

The toolkit utilizes a QoS-enabled customizable middleware framework, **CompOSE|Q**, that provides scalable and composable distributed resource management services We are developing application-aware resource management techniques in the middleware layer that will ensure the enforcement of end-to-end QoS to applications as the system scales in magnitude.

The integrated multimedia-on-demand (MOD) system for distance learning exploits the utilization of multiple media types to facilitate online classrooms. Instructional media include online lecture notes, video and audio capturing the content of the lecture and whiteboards (via pen-based tablets). To enable selective

learning, a navigation and indexing tool is integrated into the framework to allow the student to selectively view a portion of the lecture. Such systems have 2 basic modes of operation:

- Live mode: Live capture and delivery of multimedia content
- Stored mode: On-demand retrieval and delivery of stored media objects

Traditional multimedia-on-demand systems contain limited capabilities for interaction; multimedia conferencing applications permit a higher degree of interactivity. However, this implies that the resource requirements are symmetric since all endpoints function as both senders and receivers. We are currently developing a testbed framework for multimedia conferencing that is capable of ensuring specified QoS guarantees - both in transmission and presentation. We are currently enhancing off-the-shelf conferencing tools such as Microsoft NetMeeting with QoS capabilities and providing a translation mechanism into low-level QoS parameters in the middleware.

While there exist commercial VOD and multimedia conferencing systems that operate in restricted environments, delivering reliable QoS for the different media types over a wide-area network can be an issue as the number of requests in the system grows. Intelligent middleware that are aware of client heterogeneity and resource constraints in the network and servers are needed in order to ensure high throughput. CompOSE|Q contains within it a number of policies and mechanisms for QoS-based resource management. These techniques must also provide decisions about the degree of network awareness that applications and middleware must possess to ensure performance under varying network and system conditions.

Addressing Scalability:

The application environment developed can be used to address two kinds of scalability (a) Scalability in system configuration (size of network) (b) Scalability in system load (number of concurrent requests). We are studying technologies for the development of scalable and customizable directory services to be used in the development of services such as dynamic resource discovery and QoS-aware resource allocation. System Load scalability can be achieved by QoS-aware load management techniques implemented in the application framework. We will study mechanisms for adaptive scheduling of requests to servers, replication and placement of data objects in the distributed environment, as well as composite scheduling and routing of multimedia requests to replicated services over a wide-area network.

Addressing Availability:

In order to support varying levels of availability for simultaneously executing requests, we are currently working on a application interface that will allow a user to specify varying degrees of availability at a high level. Specification of fault tolerance as a QoS parameter and translation of high level availability specifications to concrete parameters that drive algorithms is an area of active research. We are currently exploring the translation of availability-QoS metrics such as MTBF, MTTR to concrete parameters such as replication degree values and replica distribution strategies that can be used to drive the the proposed fault tolerant data placement algorithms. Based on the level of application fault tolerance desired, we are implementing specific replication degree enforcement (RDE) techniques that to ensure appropriate satisfaction of replication degree while ensuring QoS to the application. By integrating replication degree enforcement with intelligent replica placement for each video object, we can tolerate multiple node failures in the system. We are currently conducting performance evaluations of the various RDE and object placement algorithms to determine request admission ratios, request completion ratios and recovery characteristics under a variety of failure modes.