# 12

# The Architecture and Implementation of a Distributed Hypermedia Storage System*

Douglas E. Shackelford
John B. Smith
F. Donelson Smith
University of North Carolina, Chapel Hill

## INTRODUCTION AND MOTIVATION

Future hypermedia systems will integrate diverse information resources, systems, and technologies. They will be based on modular architectures (e.g., Thompson, 1990) that separate orthogonal concerns into plug-compatible components such as change management, query and content search, notification, application-specific concurrency control, computational semantics, and window conferencing. Some of these components, such as change management, may be highly dependent on the semantics of a particular domain, whereas others will provide general support for all applications.

> The key point ... is that it is modular and open. This modularity is based on the observations that the functions the modules perform are independent of each other, that is *orthogonality implies modularity*. (Thompson, 1990, p. 234)

Orthogonality implies modularity; *modularity implies choice*. The importance of this observation is that every service has a cost associated with it. For example, transactions may be the appropriate concurrency mechanism for one application, while imposing prohibitively high overhead on another.

Ideally, one should be able to use a service when it is needed without having to pay for it when it is not.

In this chapter, we describe the architecture and implementation of our Distributed Graph Storage (DGS) system. We have designed it in a way that supports modular expansion to add services such as those enumerated earlier. A fundamental requirement has been that the basic hypermedia services for data storage and access should be inexpensive, efficient, and scalable. This is particularly important because the performance of these basic services is an upper bound on the performance of the system as a whole.

The DGS has been developed as a part of a larger program of research that focuses on the process of collaboration and on technology to support that process. We are concerned with the intellectual collaboration that is required for designing software systems or other similar tasks in which groups of people work together to build large, complex structures of ideas. The work of such groups—either directly or indirectly—is concerned with producing some tangible artifact. For software systems, the artifact may include concept papers, architecture, or specification documents, programs, diagrams, reference and user manuals, as well as administrative documents. A subtle but important point is that we view a group's tangible creations as parts of a single artifact.

Our research in the UNC Collaboratory project studies how groups merge their ideas and their efforts to build an artifact, and we are developing a computer system (called ABC for Artifact-Based Collaboration; Smith & Smith, 1991) to support that process. ABC has six key components (Jeffay, Lin, Menges, Smith, & Smith, 1992): the Distributed Graph Storage system, a set of graph browsers, a set of data application programs, a shared window conferencing facility, real-time video and audio, and a set of protocol tools for studying group behaviors and strategies.

## REQUIREMENTS FOR THE DISTRIBUTED STORAGE SERVICE

In this section we give a brief summary of key requirements that have shaped our storage service design.

*Permanent (persistent) storage*: obvious but fundamental.

*Sharing with protection*: because the artifact effectively constitutes the group's collective memory, it must be sharable by all. There are, however, requirements for mechanisms to authorize or deny access to selected elements of the artifact by individuals or subgroups.

*Concurrent access*: because collaborators must work together, it is often necessary for more than one user to read or modify some part of the arti-

fact at the same time. Data consistency semantics in these cases should be easily understood and provide minimal barriers to users' access to the artifact.

*Responsive performance*: sufficient to support interactive browsing of the artifact, is required.

*Scalable*: we are concerned about scale in two respects: the number of users in a group (and consequent size and complexity of the artifact), and the geographic dispersion of group members. To be scalable, it must be possible to distribute the system over available processing and network resources and to add resources incrementally as necessary. Performance (responsiveness) as perceived by users must not degrade significantly as the system grows in scale,

*Available*: if data becomes unavailable because of system faults, users may be severely impacted. The system must, therefore, be designed to tolerate most common faults and continue to provide access to most or all elements of the artifact. Replication of data and processing capacity is required to achieve high availability.

*User and artifact mobility*: users will need to change locations and system administrators will need to move data or processing resources to balance loads and capacity. The system must support this mobility in a way that is transparent to users and application programs. There should be no location dependencies inherent in the storage system.

*Private data*: these are created by individuals for their own use. Examples include personal notes, annotations on documents, and correspondence. Users must be able to create and protect such data and still establish relationships among them and the public artifact.

*Support for applications*: many applications used by a group are likely to be existing tools such as editors, drawing packages, compilers, and utilities, which use a conventional file model for persistent storage. The system should make it possible to use such tools on node data-content with no changes.

## DATA MODEL CONCEPTS

### Attributes and Content

The most basic element of the data model is the *node*, which usually contains the expression of a single thought or idea. Structural and semantic relationships between nodes are represented explicitly as *links* between nodes.[1]

---

[1]Links to links are prohibited.

The data model provides two mechanisms for storing information within a node: node attributes and node content. *Attributes* are typed, named variables for storing fine-grained information (approximately 1–100 bytes). Some attributes (such as creation time and size) are maintained automatically by the system. There may also be an unlimited number of application-defined attributes.

In comparison to attributes, node *content is* designed to reference larger amounts of information. This content can take one of two forms:

1. a stream of bytes (accessed using a file metaphor)
2. a composite object (accessed using a graph metaphor)

Applications control whether the content of a particular node is of Type 1 or of Type 2. Because Type I content obeys the standard file metaphor, it can be used to store the same types of information as files (e.g., text, bitmaps, line drawings, digitized audio and video, spreadsheets, and other binary data). Applications that can read and write conventional files can read and write Type I content with no changes. Type I content is stored with the node that contains it.

When a node has Type 2 content, then the content is stored *separately* as a composite object called a subgraph.[2] A *subgraph is* defined as a subset of the nodes and links in the artifact that is consistent with graph-theoretic constraints. For example, all subgraphs satisfy the condition that if a link belongs to a subgraph, then so do the link's source node and target node. Nodes and links may belong to multiple subgraphs at the same time, but every node and link must belong to at least one subgraph. Our data model also provides *strongly typed* subgraphs (e.g., trees and lists) that are guaranteed to be consistent with their type.

Links can have both attributes and content associated with them. Moreover, the data model defines two classes of links: structural and hyperstructural. Structural links (S-links) are used to store the essential structure of an artifact. By contrast, hyperstructural links (HS-links) are lighter weight objects that represent relationships that cut across the basic structure (see Fig. 12.1). Subgraphs containing only structural links are called S-subgraphs; those containing hyperstructural links are called HS-subgraphs.

## Using the Data Model to Organize Information

The data model encourages users to compose a large artifact from small subgraphs using subgraph content. This organization can improve human comprehension of the artifact and increase the potential for concurrent

---

[2]Hereafter, Type 1 content will be referred to as *file* content and Type 2 content will be called *subgraph* content.
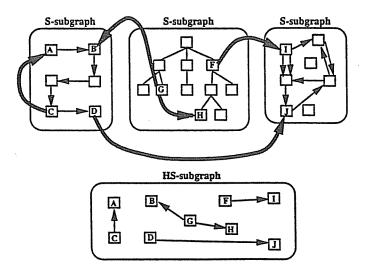
FIG. 12.1. Examples of hyperstructural linking.

access to individual components. The best way to understand these mechanisms is by example.

Figure 12.2 illustrates one way to organize the public and private materials associated with a large research project (node content is indicated by dashed lines). One can observe that Fig. 12.2 subsumes the organization of data in a conventional file system while providing additional mechanisms for storing metainformation about files (in attributes) and for representing semantic and structural relationships between files (in links).

Subgraph SG 9 in Fig. 12.2 is the top-level subgraph of a document. A useful exercise is to compare this graph structure with the way that the conference paper would be stored in a conventional file system. The most striking difference is the number and size of the nodes that compose the document. Whereas a conventional document would normally be stored in a single file or a small number of files, the DGS data model encourages a user to divide documents into many smaller nodes and subgraphs. This maximizes the benefits of hyperstructural linking because each node expresses a single concept or idea. By dividing a document into different subgraphs, collaborators may be able to structure their materials for easier concurrent access.

## Fine-Grained Linking Using Anchors

Although nodes are finer grained than traditional files, there are still times when one would like to reference information at an even finer level. For example, an application might want to create a link that points to a specific
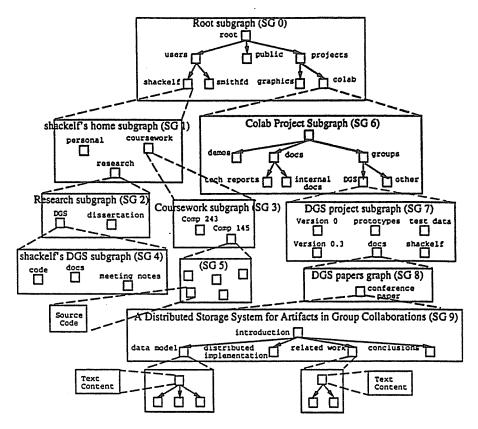
FIG. 12.2.   Organizing the public and private pieces of an artifact.

word within a node, rather than to the node itself. To achieve fine-grained linking like this, the data model provides the concept of an anchor within a node. An anchor identifies part of a node's content, such as a function declaration in a program module, a definition in a glossary, or an element of a line drawing. An anchor can be used to focus an HS-link onto a specific place within the content of a node. When an HS-link is paired with one or more anchors in its source or target nodes, it is called an *anchored HS-link*. The relationship between anchors and HS-links is many-to-many.

## Common Attributes and Graph Attributes

Some attributes are called *common attributes* because their values are independent of the context from which they are accessed. All objects—nodes, links, and subgraphs—can have common attributes. In addition, nodes and links can have context-sensitive attributes whose value may be different

depending on the context from which they are accessed. This second type of attribute is called a *graph attribute* because a subgraph provides the context.

## DESIGN AND IMPLEMENTATION

### System Architecture

As shown in Fig. 12.3, the DGS has a layered architecture that can be configured in a number of different ways. The *Application Layer* contains the user interface and other code that is application-specific. The top layer of the DGS is the *Application Programming Interface (API)* which exports a graph-oriented data model to applications. An overview of this data model was presented in a previous section. Most of the DGS is implemented in the bottom two layers: the Graph-Cache Manager (GCM) and the Storage Layer. The GCM implements the data model and performs local caching; the *Storage Layer* is responsible for permanently storing results.

Because the API isolates the application from the rest of the DGS, application code is portable across different implementations of the bottom two layers. We currently support two different implementations of the storage layer and two different methods for connecting the API with the GCM. This
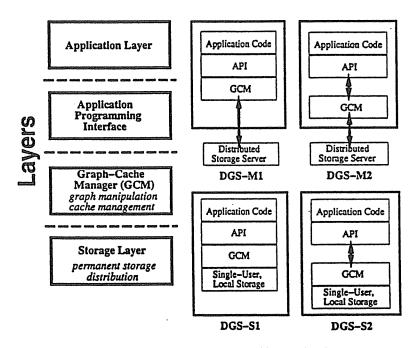


FIG. 12.3. Four implementations of the DGS layered architecture.

yields the four implementations that are shown in Fig. 12.3. In DGS-M2, the application and the GCM run in different processes on the same machine; the Storage Layer is implemented as a multiuser, distributed storage server. DGS-M I is the same except that the GCM is linked with the application to become a single process. The advantage of this design is better local response time due to reduced Interprocess Communication (IPC). A disadvantage is that it increases the size of application executables. DGS-S I and DGS-S2 follow a similar pattern except that the distributed storage server is replaced by a single-user, nondistributed storage layer.

## The Object-Oriented API

The API for the DGS is a C++ class library (for a complete description, see Shackelford, 1993). Figure 12.4 shows the major classes in the inheritance hierarchy. The class *Object* defines operations that are common to all objects such as the functions for manipulating the common attributes of an object. Subclasses inherit the API of their parent class and extend the inherited API with more specialized functions.

All node, link, and subgraph objects are identified by an object identifier (OID) that is universal and unique. Once an object is created by the DGS, its OID is never changed and the value is never reused even if the object is deleted. To applications, an OID is an "opaque" (uninterpreted) key that can be used to retrieve the corresponding object. However, we discourage application programmers from making direct reference to OIDs. Most operations can be performed without even knowing that OIDs exist.

## Concurrent Access to Objects

Because the DGS data model is object-oriented, the objects of the data model—nodes, links, and subgraphs—exist as distinct entities within the storage system. Before a user's application can access the data within a particular object (see Fig. 12.5), the application must explicitly open the object using its *Open()* function.
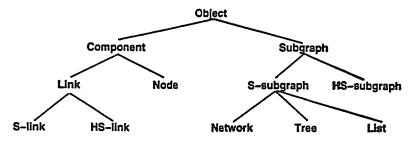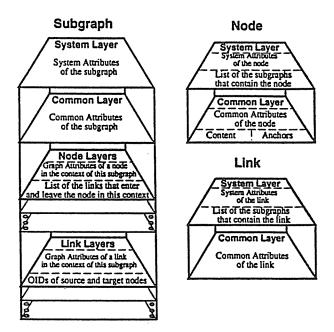
FIG. 12.4.   API class hierarchy.

FIG. 12.5.   Information stored in nodes, links, and subgraphs.


*Open()* will fail if the user lacks the proper access authorizations and if the request is in conflict with other requests in progress. Conflict can occur when different users try to access the same object concurrently. To specify allowable concurrent accesses, the API defines three access modes for nodes, links, and subgraphs: DGS-READ, DGS-WRITE, and DGS-READ–NO-ANCHOR. Applications must specify one of these modes as a parameter to *Open()* DGS-READ allows operations that do not change subgraph membership, linking information, or attribute or content values. In the case of nodes, DGS–READ also allows anchor creation and deletion, but only when the application has read–write authorization on the HS-link that is being anchored. DGS-READ–NO-ANCHOR is defined only for nodes and allows all operations of DGS-READ except anchor creation and deletion. DGS-WRITE allows all operations.

The following rules govern concurrent access to an object:

- For links and subgraphs, multiple opens with DGS-READ access and a single open with DGS-WRITE access are allowed concurrently (as is the weaker case of multiple DGS-READ opens alone).
- For nodes, multiple opens with DGS-READ–NO-ANCHOR access and a single open with DGS-WRITE access are allowed concurrently (as is the weaker case of multiple DGS-READ and/or DGS-READ–NO-ANCHOR opens alone).

Thus, for nodes the design supports multiple nonannotating readers and a single writer OR multiple annotating readers. A consequence of this is that a writer is blocked from accessing a node that is being annotated by a reader and vice versa. Changes to an object are not visible to any applications with overlapping opens of the object until it is closed by the writer and then only to applications that open it after the close completes.

## Access Control for Objects

Groups can control access to parts of the artifact by specifying access authorizations for node, link, and subgraph objects. Authorizations are expressed in an access control list that is stored with each object. An *access control list* maps names of users or groups of users to categories of operations that they are allowed to perform on the associated object. Two categories of authorizations are defined: access and administer. *Access authorizations* give users permission to access the data associated with a particular object. *Administer authorizations* give users permission to perform operations such as changing the object's access control list. Although the API does not define an explicit annotate permission, a similar effect can be accomplished by restricting the access authorizations associated with HS-subgraphs.

## Distributed Implementation

In this section we discuss the distributed implementations (DGS-M1 and DGS-M2 in Fig. 12.3) with emphasis on key design decisions.

Given an artifact composed from small elements and user access via interactive browsers, we believe many characteristics and access patterns of objects will strongly resemble those observed in distributed file systems supporting software teams using workstations (Baker, Hartman, Kupfer, Shirriff, & Ousterhout, 1991; Kistler & Satyanarayanan, 1991). Our design is based on the notion that a scalable implementation can be achieved by applying design principles such as local caching, bulk-data transfer, and minimal client–server interactions pioneered in high-performance, scalable file systems like AFS (Howard et al., 1988), Sprite (Nelson, Welch, & Ousterhout, 1988), and Coda (Kistler & Satyanarayanan, 1991). We also model our approaches to data consistency, concurrency semantics, and replication after these distributed file systems. This provides a sufficient level of function to users without requiring the full complexity of mechanisms (e.g., distributed transactions) used in database systems.

The basic structure of the system is shown in Fig. 12.6. A browser or application process acts on behalf of a user to read and modify objects. Each user's workstation runs a single Graph-Cache Manager (GCM) process that
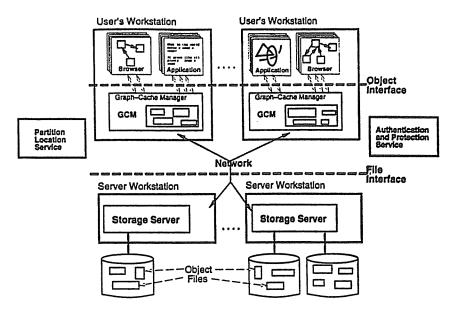
FIG. 12.6.   DGS system structure.

services all applications running on that machine. Application requests are directed over local interprocess communication facilities to the GCM. The GCM maintains a local copy of node, link, and subgraph objects used by application processes and is responsible for implementing all operations on objects in the data model except for anchor table merging. The GCM is also responsible for maintaining the consistency of typed S-subgraphs. It is important to note that this design distributes the processing for all complex object operations to the users' workstations and thus minimizes the processing demands on shared (server) resources.

When an application opens an object, the GCM, in turn, opens the object at the storage server and retrieves it using a whole-file transfer. The received object is converted from its representation in a file to an object representation designed for fast access in memory. As the application makes requests, the GCM performs those operations on the copy in its local cache. Write operations are reflected in the storage server only when the GCM closes the object and returns the modified file representation to the storage server. Each file retrieved from the storage server contains either a whole node (including data content, if present), a whole subgraph, or a group of links. An important performance optimization is that context-dependent attributes (graph attributes) and link information for all nodes in a subgraph are stored in one subgraph file. Thus, all of the data needed by a browser to

display a subgraph is available from a single request (open) to the storage server. The structure of each type of file is shown in Fig. 12.7. Nodes and subgraphs are stored individually, whereas links are grouped according to the subgraph in which they were created.

The file-oriented interface to the storage server is designed to isolate it as much as possible from the representation and semantics of objects. The primary responsibility of the storage server, therefore, is to store and control access to files indexed by an object's OID. Storage servers are also responsible for maintaining access control lists, enforcing access authorizations, enforcing concurrency semantics, creating unique OIDs and anchor IDs, and merging anchor table information created by concurrent readers of the same node. The storage server must perform several checks before completing an open request. First, it must determine whether the user who is running the application has the correct authorizations to open the object in the requested access mode. Then, the storage server must determine whether the requested access mode is in conflict with any overlapping opens for the same object. An open request will fail if the user lacks proper access authorization or if the open conflicts with other opens in progress.

Each GCM may need to communicate with multiple storage servers, including servers that provide protection services and mappings from an OID to the host system that is the custodian for that object. Object location is based on dividing the artifact store into nonoverlapping collections of nodes, links, and subgraphs called partitions. Each partition is associated with real storage devices. Partitions form boundaries for administrative controls such as space quotas, load balancing among servers, and replication of data. The partition number of an object is embedded in its OID but this substructure is never made visible outside the storage service. An object must (logically) remain in the same partition for its entire lifetime because its OID cannot be changed.

We distinguish the partition number of an object from its absolute physical location(s) and, by introducing a level of indirection (a partition directo-
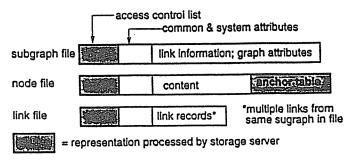


FIG. 12.7. Structure of object files.

ry), it is possible to change the physical location of an object while preserving its OID and, therefore, all its link and composition relationships with other objects (see Fig. 12.8). Partition–location servers maintain a mapping of logical partitions to host(s) running server processes for that partition. The GCM extracts the partition number from the OID of the object and uses the partition location service to find the host running a storage server process maintaining a directory for that partition (the GCM can also cache the partition location information for use in references to other objects). We expect that in most cases one storage server maintains both the partition directory and data storage for an object. Despite their importance, partitions are invisible to users. Only system administrators and system programmers need to understand partitions. An RPC interface to the storage servers is provided for administrative processes to use in creating new partitions, moving objects from one physical partition to another, and performing backup and recovery operations.

Storage servers are responsible for managing partitions on disk, replicating partitions for availability and fault tolerance in case of media or process failures, and for recovering from most failures. The key to our implementation of fault tolerance is the ISIS system developed by Ken Birman and his colleagues at Cornell University (Joseph & Birman, 1986). In particular, we use ISIS process groups to maintain replicated copies of physical partitions and to provide the location independence of logical partitions. Each logical partition corresponds to an ISIS process group.

Performance and scalability are two key requirements for the system. To evaluate the current implementation with respect to these requirements, we
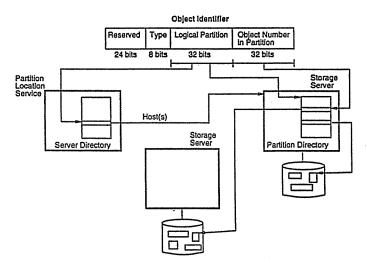


FIG. 12.8.   OID and object location.

have begun a series of benchmark experiments similar to those used to evaluate performance and scalability of distributed file systems such as AFS (Howard et al., 1988) and Sprite (Nelson et al., 1988). We have created several benchmark programs designed to stress different aspects of the system. The most interesting of these is a "synthetic browser" program that mimics the requests that result when users search for information in an artifact stored in the system. Load on the storage service is generated by running copies of the synthetic browser on several workstations. This program has parameters that can be used to produce a wide range of browsing behaviors. In our first experiments we are using parameter values that represent the observed behavior of human subjects in a series of experiments we conducted to understand how people would use a hypertext system for problem solving (Smith, 1992). With these values, each instance of the program running on one workstation generates a load on the server corresponding to approximately 10 users working with interactive browsing applications. We have also written an "artifact generator" program that, based on a number of input parameters, creates a structure of subgraphs, nodes, and links to serve as data for the browsing benchmark.

The results of our initial measurements have been very encouraging. The configuration for these measurements consisted of one storage server running on a DECstation 5000/25c and up to seven workstations (DECstation 5000/120s) each running a copy of the synthetic browser program. All workstations were connected by a single ethernet segment. The most significant results are:

- CPU utilization on the server is at most 0.5–1.0% per active user.
- Server response times to requests from 50 users increased by less than 20% over response times to requests from 10 users.

The results show that one server can support at least 50 users. More extensive benchmark experiments are underway to validate this conclusion for a variety of configurations.

We are currently using the DGS for developing browsers and other collaboration support tools. We continue to make enhancements (mostly for operations and administration) and plan to have a version suitable for distribution to other groups by Fall 1993.

## COMPARISON WITH RELATED WORK

In this section, we compare our design with several hypertext systems that have significant capability for supporting collaborating groups (i.e., Intermedia; Haan, Kahn, Riley, Coombs, & Meyrowitz, 1992; Yankelovich et al.,

1988), HyperBase/CHS (Schütt & Haake, 1993; Schütt & Streitz, 1990), Augment (Engelbart, 1984), Telesophy (Caplinger, 1987; Schatz, 1987), KMS (Akscyn, McCracken, & Yoder, 1988), and HAM (Campbell & Goodman, 1988; Delisle & Schwartz, 1987). These systems differ widely on factors such as the data model supported, scalability, concurrent reader–writer semantics, and protection.

DGS, HyperBase/CHS, and Dexter (Halasz & Schwartz, 1990) support rich data models that include aggregates (named groups of objects), aggregates of aggregates, and aggregates as endpoints of links. Intermedia, HAM, and Augment do not use aggregates in composition or linking. Telesophy's data model has aggregates but does not give first-class status to links. HB I (Schnase, Leggett, & Hicks, 1991) and Trellis (Stotts & Furuta, 1989) provide strong support for computation within hypertext but do not have aggregates. The DGS data model benefits from the graph-theoretical metaphor on which it is based and is the only system to provide strongly typed aggregate objects.

Other areas in which these systems differ substantially are in the semantics of concurrent reading and writing and in the access protection mechanisms (see Table 12.1). These systems also differ in their capability to scale up to large numbers of users (and objects) while preserving the illusion of location transparency. Both Telesophy and the DGS have made scalability a central issue in their designs. However, the DGS provides more flexibility in its data model and stronger consistency semantics.

## SUMMARY AND CONCLUSIONS

Collaborative groups face many problems, but one of the hardest and most important is to meld their thinking into a conceptual structure that has integrity as a whole and that is coherent, consistent, and correct. Seeing that construct as a single, integrated artifact can help. But groups must also be able to view specific parts of the artifact in order to understand and manage it. Our design was guided by these requirements, along with others discussed earlier. The graph-based data model permits us to both partition the artifact and to compose those pieces to build larger components and the whole. The distributed architecture, in turn, permits us to build a system that can scale up in terms of the size of the artifact, the number of users, and their geographic distances from one another.

We observe that most of the academic research in hypermedia is not based on the sort of modular architecture that was described at the beginning of this chapter. Although many communities view hypermedia as an "interesting" application, we take the perspective (also expressed in Schnase et al., 1991) that hypermedia has a broader role to play. In our opin-

**TABLE 12.1**
Concurrent Reader–Writer Semantics and Object Protect

| Hypermedia System | Concurrent Reader–Writer Semantics | Protection of Objects |
|---|---|---|
| Augment | Can have multiple readers of documents that have been submitted to the Journal system | Objects in the Journal are read-only. Access to Journal entries can be restricted at submission time |
| HAM | Could not be determined | Access Control Lists (Optional): access, annotate, update, and destroy permissions |
| HyperBase/CHS | Activity markers are provided to warn applications of concurrent activity, but these markers are advisory I nature. All applications are notified when data is changed, so that they can update their view (if desired). | Access control will be based on user roles such as "manager" and "secretary" (not yet implemented). |
| Intermedia | Supports multiple users reading and annotating, and a single writer. First user to write an object locks out other potential writers. | Provides read, write, and annotate permissions that can be granted to users and groups of users. |
| KMS | Uses an optimistic concurrency method. When a writer attempts to save a node, he/she may be denied because someone else has concurrently written to the same node. In this case, the human user must manually merge the two conflicting versions. | Owner can protect a frame from modification or read access. In addition, an intermediate form allows users to add annotation items, but not to modify existing items. |
| Telesophy | Supports multiple concurrent readers and writers. When writes overlap, the last writer completely overwrites the work of others. | could not be determined |
| DGS | Supports multiple nonannotating readers and a single writer OR multiple annotating readers. Applications must declare their intent at the time that they open an object. Intent can be one of: read and annotate; read-only; read/write and annotate. | Access Control Lists: access (read or read/write) and administer permissions. Rather than associate a single annotate permission with a node, the DGS provides a more flexible mechanism of associating annotate permission with the HS-subgraphs that contain the node. Thus, a user might be allowed to annotate a node within his personal context at the same time that he is denied the ability to annotate the node in a public context. |

ion, hypermedia is not just an application, but is a new paradigm for the way we work and collaborate with each other. As such, it will be an essential component of the next generation of operating system support. Our experiences with DGS strongly indicate that it is possible to achieve the richer functions needed for hypermedia storage with cost, performance, and scalability comparable to the best conventional distributed file systems (e.g., AFS).

As we look to the future, additional issues we will explore pertain to wide-area network access, dynamic change notification, graph traversal, and support of a richer set of graph and set operations and queries. Many of these extensions lend themselves to the sort of modular approach that is suggested in the Strawman Reference Model (Thompson, 1990).

## ACKNOWLEDGMENTS

## REFERENCES

Akscyn, R. M., D. L. McCracken, and E. A. Yoder (1988, July). KMS: A distributed hypermedia system for managing knowledge in organizations. *Communications of the ACM, 31*(7), 820–835.

Baker, M. G., J. H. Hartman, M. D. Kupfer, K. W. Shirriff, and J. K. Ousterhout (1991, October). Measurements of a distributed file system. Operating Systems Review, Special Issue: *Proceedings of the 13th ACM Symposium on Operating Systems Principles* (Pacific Grove, CA), 25(5), 198–212.

Campbell, B. and J. M. Goodman (1988). HAM: A general purpose hypertext abstract machine. *Communications of the ACM, 31*(7), 856–861.

Caplinger, M. (1987, October). An information system based on distributed objects. In *OOPSLA '87 Proceedings*, pp. 126–137.

Delisle, N. M. and M. D. Schwartz (1987, April). Contexts: A partitioning concept for hypertext. *ACM Transactions on Office Information Systems, 5*(2), 168–186.

Engelbart, D. C. (1984, February). Authorship provisions in AUGMENT. In *Proceedings of the 1984 COMPCON Conference*, San Francisco, CA, pp. 465–472.

Haan, B. J., P. Kahn, V. A. Riley, J. H. Coombs, and N. K. Meyrowitz (1992, January). IRIS hypermedia services. *Communications of the ACM, 35*(1), 36–51.

Halasz, F. and M. Schwartz (1990). The Dexter hypertext reference model. In *Proceedings of the NIST Hypertext Standardization Workshop* (Gaithersburg, MD), pp. 1–39.

Howard, J. H., M. L. Kazar, S. G. Menees, D. A. Nichols, M. Satyanarayanan, R. N. Sidebotham, and M. J. West (1988, February). Scale and performance in a distributed file system. *ACM Transactions on Computer Systems, 6*(1), 51–81.

Jeffay, K., J. K. Lin, J. Menges, F. D. Smith, and J. B. Smith (1992). Architecture of the artifact-based collaboration system matrix. In *Proceedings of ACM CSCW '92 Conference on Computer-Supported Cooperative Work, CSCW Architectures*, pp. 195–202.

Joseph, T. A. and K. P. Birman (1986, February). Low cost management of replicated data in fault-tolerant distributed systems. *ACM Transactions on Computer Systems, 4*(1), 54–70.

Kistler, J. J. and M. Satyanarayanan (1991, October). Disconnected operation in the Coda file system. Operating Systems Review, Special Issue: *Proceedings of the 13th ACM Symposium on Operating Systems Principles* (Pacific Grove, CA), *25*(5), 213–225.

Nelson, M. N., B. B. Welch, and J. K. Ousterhout (1988, February). Caching in the Sprite network file system. *ACM Transactions on Computer Systems, 6*(1), 134–154.

Schatz, B. R. (1987). Telesophy: A system for manipulating the knowledge of a community. In *Proceedings of Globecom'87* (New York), pp. 1181–1186. ACM.

Schnase, J. L., J. J. Leggett, and D. L. Hicks (1991, October). *HBI: Initial design and implementation of a hyperbase management system* (Tech. Rep. TAMU-HRL 91-003). Hypertext Research Lab, Texas A&M University.

Schütt, H. and J. M. Haake (1993, March). Server support for cooperative hypermedia systems. In *Hypermedia'93*, Zurich.

Schütt, H. A. and N. A. Streitz (1990). Hyperbase: A hypermedia engine based on a relational database management system. In *Proceedings of the ECHT'90 European Conference on Hypertext, Databases, Indices and Normative Knowledge*, pp. 95–108.

Shackelford, D. E. (1993, January). *The Distributed Graph Storage System: A users manual for application programmers* (Tech. Rep. TR93-003). Department of Computer Science, The University of North Carolina at Chapel Hill.

Smith, D. K. (1992). Hypermedia vs. paper: User strategies in browsing SNA materials (Tech. Rep. TR92-036). Department of Computer Science, The University of North Carolina at Chapel Hill.

Smith, J. B. and F. D. Smith (1991). ABC: A hypermedia system for artifact-based collaboration. In *Proceedings of ACM Hypertext'91, Construction and Authoring*, pp. 179–192.

Stotts, P. D. and R. Furuta (1989). Petri-net-based hypertext: Document structure with browsing semantics. *ACM Transactions on Information Systems, 7*(1), 3–29.

Thompson, C. W. (1990, January). Strawman reference model for hypermedia systems. In *Proceedings of the NIST Hypertext Standardization Workshop* (Gaithersburg, MD), pp. 223–246.

Yankelovich, N. et al. (1988, January). Intermedia: The concept and the construction of a seamless information environment. *IEEE Computer, 21*(1), 81–96.