




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification <sup>6</sup> : <b>G06F 17/21, 17/30</b></p>	<p><b>A1</b></p>	<p>(11) International Publication Number: <b>WO 95/13585</b>  (43) International Publication Date: 18 May 1995 (18.05.95)</p>
<p>(21) International Application Number: PCT/US94/00049 (22) International Filing Date: 3 January 1994 (03.01.94) (30) Priority Data: 08/151,335 12 November 1993 (12.11.93) US  (71) Applicant: TALIGENT, INC. [US/US]; 10201 N. de Anza Boulevard, Cupertino, CA 95014 (US). (72) Inventors: ANDERSON, David, R.; 10671 W. Estates Drive, Cupertino, CA 95014 (US). PALEVICH, Jack, H.; 1759 Lark Lane, Sunnyvale, CA 94087 (US). SCHAEFFER, Arnold; 5 Skymont Court, Belmont, CA 94002 (US). WATANABE, Ryoji; 22284 Palm Avenue, Cupertino, CA 95014 (US). (74) Agent: STEPHENS, Keith; Taligent, Inc., 10201 N. de Anza Boulevard, Cupertino, CA 95014 (US).</p>		<p>(81) Designated States: AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, ES, FI, GB, HU, JP, KP, KR, KZ, LK, LU, LV, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SK, UA, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p><b>Published</b> <i>With international search report.</i></p>

(54) Title: COMPOUND DOCUMENT FRAMEWORK



(57) Abstract

An object-oriented compound document architecture provides system level support for document processing features. The object-oriented compound document framework supports a variety of document processing functions. The framework provides system level support of collaboration, linking, eternal undo, and content based retrieval, among other things. System level support is provided for document changes, annotation through model and linking, anchors, model hierarchies, enhanced copy and pasting, command objects, and a generic retrieval framework.

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

<b>AT</b>	Austria	<b>GB</b>	United Kingdom	<b>MR</b>	Mauritania
<b>AU</b>	Australia	<b>GE</b>	Georgia	<b>MW</b>	Malawi
<b>BB</b>	Barbados	<b>GN</b>	Guinea	<b>NE</b>	Niger
<b>BE</b>	Belgium	<b>GR</b>	Greece	<b>NL</b>	Netherlands
<b>BF</b>	Burkina Faso	<b>HU</b>	Hungary	<b>NO</b>	Norway
<b>BG</b>	Bulgaria	<b>IE</b>	Ireland	<b>NZ</b>	New Zealand
<b>BJ</b>	Benin	<b>IT</b>	Italy	<b>PL</b>	Poland
<b>BR</b>	Brazil	<b>JP</b>	Japan	<b>PT</b>	Portugal
<b>BY</b>	Belarus	<b>KE</b>	Kenya	<b>RO</b>	Romania
<b>CA</b>	Canada	<b>KG</b>	Kyrgystan	<b>RU</b>	Russian Federation
<b>CF</b>	Central African Republic	<b>KP</b>	Democratic People's Republic of Korea	<b>SD</b>	Sudan
<b>CG</b>	Congo	<b>KR</b>	Republic of Korea	<b>SE</b>	Sweden
<b>CH</b>	Switzerland	<b>KZ</b>	Kazakhstan	<b>SI</b>	Slovenia
<b>CI</b>	Côte d'Ivoire	<b>LI</b>	Liechtenstein	<b>SK</b>	Slovakia
<b>CM</b>	Cameroon	<b>LK</b>	Sri Lanka	<b>SN</b>	Senegal
<b>CN</b>	China	<b>LU</b>	Luxembourg	<b>TD</b>	Chad
<b>CS</b>	Czechoslovakia	<b>LV</b>	Latvia	<b>TG</b>	Togo
<b>CZ</b>	Czech Republic	<b>MC</b>	Monaco	<b>TJ</b>	Tajikistan
<b>DE</b>	Germany	<b>MD</b>	Republic of Moldova	<b>TT</b>	Trinidad and Tobago
<b>DK</b>	Denmark	<b>MG</b>	Madagascar	<b>UA</b>	Ukraine
<b>ES</b>	Spain	<b>ML</b>	Mali	<b>US</b>	United States of America
<b>FI</b>	Finland	<b>MN</b>	Mongolia	<b>UZ</b>	Uzbekistan
<b>FR</b>	France			<b>VN</b>	Viet Nam
<b>GA</b>	Gabon				

## COMPOUND DOCUMENT FRAMEWORK

### *COPYRIGHT NOTIFICATION*

5 Portions of this patent application contain materials that are subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

### *Field of the Invention*

10 The present invention generally relates to computer systems, and more particularly to a method and system for object-oriented compound document processing.

### *Background of the Invention*

15 Document processing has virtually revolutionized the way society generates paper. Typical prior art document processing systems run on top of operating systems, such as DOS or OS/2. More recently, these document processing systems have been designed to run in a Windows environment. Many of these document processing systems are commercially available.

20 While these document processing systems have vastly improved the ability to process documents and text, there is great inconsistency among document processors with respect to the particular methodologies of these processing. The result of these inconsistencies creates problems for both application developers and users of the applications.

25 Application developers must continuously "reinvent the wheel" when creating a new document processor. While operating systems and interface programs provide some tools which may be used, the great majority of the design process for a particular document processor is directed toward creating a group of processing modules which cooperate to allow the user to process documents.  
30 Application developers often design processing modules which have already been developed by another company. This requires great duplication of effort, and requires each developer to deal with the details of how to implement various desired functions.

35 Application users run into other problems. While particular functions may be present in one application, they may be lacking in another. Or a function available in one application may be slightly varied in another, either in use or in performance. For example, a function in application A may require certain user interaction and input to activate the function, while a similar function in

application B may require a slightly varied, or totally different, user interaction and input.

### *Summary of the Invention*

5           It is an object of the present invention to provide a document processing system in which object-oriented frameworks are utilized to implement particular document processing techniques, including an object-oriented compound document system. These and other objects of the present invention are realized by a document framework which supports at the system level a variety of compound  
10 document processing functions. The framework provides system level support of collaboration, linking, eternal undo, and content based retrieval. These and other objects are carried out by system level support of document changes, annotation through model and linking, anchors, model hierarchies, enhanced copy and pasting, command objects, and a generic retrieval framework.

15

### *Brief Description of the Drawings*

Figure 1 is a block diagram of a personal computer system in accordance with a preferred embodiment of the invention;

Figure 2 is a block diagram of a link, anchors, and a model;

20           Figure 3 is a block diagram representing the functions associated with undo;

Figure 4 is a diagram demonstrating the system level indexing and query processing features of the present invention;

Figure 5 is a block diagram of class representation in accordance with the present invention;

25           Figure 6 is a diagram of a typical document which may be made using the present invention;

Figure 7 is a diagram showing the hierarchical structure of the document shown in Figure 6;

Figure 8 is a diagram of the general characteristics of TModel;

30           Figure 9 is a diagram of a notification framework which could be used with the document framework system;

Figure 10 is a diagram of the relationship defining specification classes;

Figure 11 is a diagram of the relationships associated with TModel Command Group;

35           Figure 12 is a diagram demonstrating the processing flow for DoBegin();

Figure 13 is a diagram demonstrating the processing flow for DoRepeat();

Figure 14 is a diagram showing the relationships established for TModelAnchor;

Figure 15 shows the relationships established for TModelLink;

Figure 16 is a diagram depicting the processing of links;

Figure 17 shows the processing of Complete Link;

Figure 18 demonstrates the use of annotations and scripts being linked to a  
5 document;

Figure 19 shows some of the link properties which are available in the  
system;

Figure 20 shows the document framework client/server with respect to  
external documents; and

10 Figure 21 is a block diagram showing the method and system of the present  
invention.

### *Detailed Description of the Invention*

The detailed embodiments of the present invention are disclosed herein. It  
15 should be understood, however, that the disclosed embodiments are merely  
exemplary of the invention, which may be embodied in various forms. Therefore,  
the details disclosed herein are not to be interpreted as limiting, but merely as the  
basis for the claims and as a basis for teaching one skilled in the art how to make  
and/or use the invention.

20 The history of object-oriented programming and the developments of  
frameworks is well-established in the literature. C++ and Smalltalk have been  
well-documented and will not be detailed here. Similarly, characteristics of objects,  
such as encapsulation, polymorphism and inheritance have been discussed at  
length in the literature and patents. For an excellent survey of object oriented  
25 systems, the reader is referred to "Object Oriented Design With Applications" by  
Grady Booch, ISBN 0-8053-0091-0 (1991).

While many object oriented systems are designed to operate on top of a basic  
operating system performing rudimentary input and output, the present system is  
used to provide system level support for particular features.

30 The invention is preferably practiced in the context of an operating system  
resident on a personal computer such as the IBM ® PS/2 ® or Apple ® Macintosh ®  
computer. A representative hardware environment is depicted in Figure 1, which  
illustrates a typical hardware configuration of a computer in accordance with the  
subject invention having a central processing unit 10, such as a conventional  
35 microprocessor, and a number of other units interconnected via a system bus 12.  
The computer shown in Figure 1 includes a Read Only Memory (ROM) 16, a  
Random Access Memory (RAM) 14, an I/O adapter 18 for connecting peripheral  
devices such as disk units 20 and other I/O peripherals represented by 21 to the

system bus 12, a user interface adapter 22 for connecting a keyboard 24, a mouse 32, a speaker 28, a microphone 26, and/or other user interface devices such as a touch screen device (not shown) to the bus 12, a communication adapter 34 for connecting the workstation to a data processing network represented by 23. A display adapter 36  
5 for connecting the bus to a display device 38. The workstation has resident thereon an operating system such as the Apple System 7 ® operating system.

The main goal of the document framework disclosed herein is to raise the base level of applications by enabling several new features at the system level. In addition, the lack of system support for these features limits their implementation.

10 For example, there are applications that allow users to annotate static representations (pictures) of any document, but not the "live" document itself. The content-based retrieval applications have trouble accessing the contents of document because each application has a custom file format. Also, once the application finds a document, it is difficult to do anything with it. There's no  
15 system-level support for opening the document, for example. The document framework also includes a number of higher-level capabilities, such as annotation support, which are built upon these basic services. To the extent possible, the document framework does not specify any policy or user interface decisions. These details will be provided by the particular applications using the document  
20 framework.

### Collaboration

Screen sharing is one popular type of collaboration on the Macintosh, because it is relatively easy to implement and can be put to many uses. Its main  
25 disadvantages are that some applications draw directly to the screen (complicating the implementation) and the large bandwidth required to transmit all drawing operations from one machine to another. Also, it is very restrictive, since it is based on all collaborators viewing the document in exactly the same way.

Screen sharing is one kind of simultaneous, real-time collaboration. The document framework provides support for a different kind of simultaneous, real-  
30 time collaboration. This operates at the level of changes to the document, rather than changes to the screen, which will be more efficient because the amount of data needed to specify a document change is usually less than the amount needed to update the screen.

It is also useful to have asynchronous (i.e., non real-time) collaboration. One  
35 form of this is the ability to annotate a document. The document framework provides low-level support for annotations through its model and linking mechanisms (described below).

### Hypermedia Linking

Figure 2 shows an illustration of the present invention. The blocks represent both the apparatus and the methods involved in the system. As shown in Figure 2, in the document framework, a link 204 is a bi-directional connection between anchors 202 and 206. The meaning of an anchor is application-specific, but in most cases an anchor identifies a sticky selection. An anchor is sticky in that it always refers to the same data regardless of the user's editing changes. For example, if the anchor refers to a word within a text block, it always refers to that word, even if the text around it changes. Anchors are associated with a particular encapsulated block of data (called a model) 200. Each kind of data in the system is represented as a specific subclass of TModel. The abstract base class, TModel, defines generic protocol to enable other models to embed, display, and manipulate this data as a "black box." For example, an application can ask the model to create a presentation (view) of the data. These presentations range from a small thumbnail to a fully editable presentation. There is also protocol for accessing the anchors associated with the model's data and for accessing other models embedded within it. A document's data is represented by a hierarchy of models, with a single model at the root called the "root model."

Once the user creates a link, the user can operate upon it. First, the user can navigate from one end of the link to the other. In general, this involves opening the document containing the target anchor, scrolling the anchor into view, and highlighting the corresponding selection. Applications can change this behavior; for example, navigating to a sound document may simply play the sound without opening the document. It is also possible to transfer data across the link in either direction. The semantics of transferring data is (roughly) equivalent to copying the source data, navigating to the destination anchor, and replacing the existing data with the transferred data. It makes no difference to the document framework whether the data is pushed or pulled through the link (i.e., whether the source or destination initiates the transfer).

The document framework also allows one application to send an arbitrary command across a link. This will allow cooperating applications to implement custom features using the same basic linking mechanism. The straightforward use for the document framework's low-level linking mechanism is to allow users to create links between documents, navigate those links, transfer data across them, etc. This isn't the only use for links, however. Links will also be used to implement other application features. In these features, the fact that links are created and manipulated will be transparent to the user.

For example, the system-wide annotation facility uses links to associate an annotation with the part of the base document to which it refers. The system can position a posted note icon (representing the annotation) near the part of the document to which it refers. In addition, if the annotation contains a suggested  
5 change to the document, it is possible for the author to "accept" the suggestion and have the system automatically transfer data across the link from the annotation to the document.

Another user-transparent use of linking would be to implement a function performed by the Edition Manager of System 7. The user could publish part of a  
10 drawing and subscribe to that data in a word processing document. Internally, the system would create a link between the documents, and perhaps attach an attribute to the link that indicates the nature of the link (e.g., which end is the source of the data). The existence of the link may be transparent to the user.

Changes to the drawing are sent across the link to the word processing  
15 document. The document framework does not restrict which end of the link initiates this transfer. The user can navigate from the destination to the source of the data (from the word processing document to the drawing document). The document framework also supports navigating in the opposite direction because all links are bi-directional.

The document framework's models also improve the way simple copy and  
20 paste works. On the Macintosh, an application handles a paste command in three different ways, depending on the type of the data model in the Clipboard. (1) The document can fully understand the pasted data, and the data becomes a first-class part of the document. (2) The receiving document can display the incoming data  
25 but not manipulate it. A typical example is a picture pasted into a text document. In this case, the incoming data can be displayed but not edited. (3) The incoming data type isn't understood by the receiving document. Here, the paste command can't be completed, and should be disabled.

If the data can't be absorbed, then it can be embedded in the receiving  
30 document in the form of a model. Because models support generic protocol for creating editable presentations of the data, embedded data is not "dead" as is true on the Macintosh. Instead, the user can open up an embedded model and view or edit the data it contains. This capability is similar to that provided by HP's New Wave system or Microsoft's OLE specification. An important difference is that the object-  
35 oriented document framework makes it easier for a developer to implement a new data type. Finally, if an application supports embedded models, then it can paste every type of data. The paste command would never be disabled as long as the Clipboard wasn't empty.



### Eternal Undo

In most Macintosh applications, the undo command is precious, since only the last change can be undone. The document framework uses the same kind of command objects as MacApp, but saves as many command objects as possible.

This decision has several benefits. First, it isn't as important to be choosy about what commands are undoable. For example, in existing Macintosh applications, changes to the selection are not undoable, even though some selections can be difficult to create. In a drawing program, the user can spend much time selecting the right combination of shapes and lose everything with an extra click. If the system supports only one level of undo, then it is unwise to save a selection change if it means forgetting about the last Cut command, for example. With multiple levels of undo, it is feasible to save selection changes.

Another benefit of using command objects is increased reliability. If every command is saved, then it is possible to replay those commands in the event of an application or system crash. The document framework uses concurrency control and recovery classes to save command objects in a robust manner. In the event of a crash, the user should not lose more than the last couple of commands. With multi-level undo comes the added burden of providing a good way to visualize and navigate the list of commands. This is especially true if selection changes are included, because it will be easy for the user to create hundreds of commands.

Figure 3 shows The document framework. linear list 300 of command objects 302, which can be likened to a stack. There are many other ways in which the undo processing could be carried out. This means that undo 306 returns the document to a previous state. It is also contemplated that the user can selectively undo commands 304 (i.e., undo a Cut command but keep all subsequent commands intact). It should be remembered that commands are dependent on one another. A command that copies a shape is dependent on the command that first created the shape. These dependencies would complicate the user interface to undo, as well as the underlying implementation.

A good solution is to integrate the undo and scripting mechanisms, for example to automatically create a script of everything that is done. The user can then edit the script to remove arbitrary commands, rearrange commands, etc. and execute the script. This gives users the maximum flexibility and control.

### Content-based Retrieval

Increasingly, users have more and more information available on their computers. Local hard disks are getting larger, and there are many CD-ROMs

available that contain hundreds of megabytes of data. It is impossible to browse through this data without some assistance from the system. In the Macintosh, the standard tool used to be Find File, which located documents based on their names. System 7 provides a Find command that is integrated within the Finder. Third party developers also provide tools that go beyond Find File and search for documents based on their content, but which aren't well-integrated with the system.

It is important that these retrieval tools be integrated into the system. There's little point in locating documents if the user can't do anything with them. The third party content-based retrieval tools on the Macintosh get no system support in examining the contents of the document, or even opening the document with the appropriate application.

Figure 4 shows the generic retrieval framework of The document framework from a source of information 400. The framework handles both indexing 402 and queries 404. Although many future operating systems will deliver with a default indexing and query package, users will still want the ability to plug in new search packages into the basic framework. The point of designing a framework is so that the background indexing mechanism and query user interface are the same regardless of the underlying retrieval technology.

The framework will provide for automatic, background indexing of documents when they are added to a volume or changed. A retrieval system is worthless if only some of the documents are indexed, and it is unreasonable to place this burden on the user. A generic query front-end will allow a user to install a new retrieval engine and not have to learn a new front-end.

### Classes, The Developer's View

This section contains class and member function descriptions. Certain conventions are followed. All classes have virtual destructors. If the destructor does anything more than release storage, it is discussed; otherwise, nothing is said. Methods inherited from MCollectible (Figure 5, element 500), such as streaming operators, are not discussed here. Many of the classes have getters and setters which simply do field accesses.

### Object Surrogates

Several objects in the document framework provide object surrogates, which act as compact, address space independent references to the real object. This use of surrogates can be described as "Explicit master and surrogates". In a few cases, surrogates and real objects can be used interchangeably, but in most cases the

surrogate serves as a "key" for finding the real object. In all cases the relationship between the surrogate and real object is defined by a common abstract base class. The abstract base class provides the address space independent identification for the object and implements the protocol for comparison between objects (IsEqual). This  
5 allows a surrogate to be compared directly against a real object and used as a key into a set of real objects for lookup. Surrogate objects provide a constructor which will create a surrogate from a real object allowing easy creation of surrogates.

For example, it is necessary for TModelSelections to identify TModel objects in an address space independent manner to allow commands to access data in  
10 multiple collaborator's address spaces. TModel and TModelSurrogate both derive from the common base class TAbstractModel. TAbstractModel provides the address space independent identification of the object and implements the protocol for comparison between objects. TModel provides static member functions for looking-up the real model from a surrogate. When a selection is streamed, it streams only a  
15 surrogate for the model to which it refers. When the command attempts to access the data through the selection, the real model is looked up using the surrogate to provide the command access to the real model data.

### Data Representation

20 Figure 5 shows Representation classes. The abstract base class that encapsulates the data for a particular data type is TModel 506. Derived classes of TModel 506 are the containers for "document" data. They provide type specific protocol for accessing and modifying the data contained in the model. They must also override generic protocol that supports embedding other models and for being  
25 embedded in other models. This mostly involves overriding protocol for creating selections and user interfaces on the data.

TModel objects 506 also provide notification of changes to the contained data to interested objects (typically presentations). Notification could be provided using a standard notification facility of the underlying system, if the system has such  
30 notification available. See the "Data Presentation" section for details on model notification.

The class TModelSurrogate 504 provides a lightweight stand-in or "name" for an actual TModel 506. It does not actually contain the data, only the real model does that, but it does provide protocol for a subset of the behavior supported by TModel.  
35 Specifically, the behavior necessary to appear and operate within the "Workspace". The behavior which TModelSurrogate 504 and TModel 506 share is defined by their common base class TAbstractModel 502.

### Compound Document Structure

A document, such as shown by element 600 in Figure 6, can contain many models which can be of many different classes. The basic structure of the document is a hierarchy of models which reflects the containership hierarchy of the document. A single model exists at the root of the hierarchy and is referred to as the "root model." Each model in the hierarchy may be a container for other models embedded within. A container model refers to its embedded models with model surrogates. C++ pointers must not be used because the embedded models may be filed using separate contexts from the container and are not necessarily in memory with the container. By accessing a model through a surrogate it will automatically be filed in if not already in memory. Optionally, the embedded document can be a directed acyclic graph. Multi-column text can be implemented this way; the individual columns of text were going to be models that all embedded a single shared text-flow model. Also, each model can be locked independently, providing more potential concurrency at the cost of more complexity.

"The Standard Example" shown in Figure 6 illustrates a specific model hierarchy. Each model in the hierarchy which has a presentation included in the document has an associated TUserInterface (which is itself a model containing user interface objects). The TUserInterface for a particular model is managed by that model and is accessible only through that model. Models which are not viewed directly, but rather through another "presentation model", would not have an associated TUserInterface. For example, a TTableModel which was viewed only through a TGraphModel as a scatter plot, would not have a TUserInterface, while the TGraphModel would, since it is providing a presentation to the user.

Figure 7 demonstrates the hierarchical nature of the document of Figure 6. In particular, sections 0-5 represent various types of information which may be embedded within other types of information. For example, Section 1 is the Taligent® text, which is embedded within the Taligent® logo, which is shown by Section 2. The lower right corner of Figure 7 shows the hierarchical structure of the document of Figure 6.

### Data Types

The compound document architecture supports the seamless integration of many data types. These may include, but are not limited to, voice, graphics, bit-mapped images, pictures, tables, video, text, financial data, files, and program components. Virtually any visible or audible data type may be incorporated into the compound document in a seamless manner.

### Managing Embedded Models

It is up to a derived class of TModel to provide an implementation to manage any embedded models it contains. TModel does not provide an implementation for this because it is extremely data dependent as to how the embedded models fit into the container's data structures. The document framework does include a TContainerModel that provides an implementation for simplistic embedding in which the embedded model presentations simply "float" over the containing presentation.

### Copying, Streaming and Filing

Models which are embedded in another model are considered part of the containing model's data for purposes of copying and streaming. When a model is copied, streamed or filed, the embedded models it contains must be handled appropriately. TModel provides default implementations which make the handling of embedded models straightforward for these cases.

#### *Streaming*

The standard streaming operators (operator>>= and operator<<=) provided by TModel stream all data owned by TModel, including anchors and links. These streaming operators also take care of streaming any embedded models contained within the model being streamed. A derived class overrides these operators to stream its own data. If the model contains embedded models, only the surrogates for those models need to be streamed. The inherited streaming operators take care of streaming the real embedded models. For models which file their data by streaming, these methods may be identical to the filing methods. In such cases, it is usually desirable to create private helper functions which can be called by both the streaming operators and the filing methods.

#### *Copying*

When copying an entire model, or copying selected data from the model, the embedded models must be copied by the container. This is easily accomplished by using the "CopyModel" method of the TModelSurrogate for the embedded model. This method copies the surrogate and real model providing a new consistent model/surrogate pair.

#### *Filing*

When filing a model, its embedded models must also be filed. The embedded models must be filed independently from the container to allow customization of

storage for the embedded models. When a model is asked to file its data in FileOutData(), it only files surrogates for any embedded models which it contains. The embedded models are filed automatically by the document framework. The filed surrogates are read back in FileInData(), and may be used to access the real  
5 model. All dirty models in a document are always filed out together. This simplifies the handling of the Roll-Forward log. Models are individually filed in when requested from model surrogates.

When a new model is added to the document it must be given the opportunity to create a new store for itself. The method CreateModelStore() is called  
10 passing in the document root model's store. If this store is acceptable to the model no new store is created. Otherwise, the embedded model will create a new store to use for its storage requirements. The actual creation of the store is done in a lazy fashion, so that models created in memory for temporary use that do not require a store will not create one.

15 When a model is removed from the document it must be deleted from both memory and storage. Deleting a model will delete the model from memory and from store. If the model is to be deleted from memory only the method TModel::DeleteFromMemoryOnly() can be used rather than calling delete.

## 20 Concurrency Control

A single monitor protects all the data in a document. This monitor is entered by creating a TModelEntry object on the stack.

## Attributes

25 TModel provides an extensible interface and implementation for attaching attributes to models. The attributes are inheritable via the model hierarchy. That is, a model can inherit attributes from its ancestors. TModel provides support for looking up attributes on a specific selection in a model and for including or excluding ancestors from the lookup. Attributes are maintained in a simple  
30 attribute group and provides its own inheritance support rather than using the TInheritableAttributeGroup because all models in the hierarchy are not necessarily in memory at one time, and this behavior is not supported by TInheritableAttributeGroup.

## 35 **TAbstractModel**

TAbstractModel is the base class for TModel and TModelSurrogate. It provides the protocol which is common to both of these classes and provides the

common base class between the real and surrogate classes as described in the "Object Surrogates" section of this document. TAbstractModel is an abstract base class. It is only instantiated as part of a derived class. The following methods provide access to the model's storage: 1) SetModelStoreReference, 2) AdoptModelStoreReference, 3) CopyModelStoreReference, 4) GetModelStore. TModel and TModelSurrogate are the only classes which derive from TAbstractModel. TAbstractModel is abstract and therefore is always created and deleted as part of a derived class. Any data created by TAbstractModel as part of its implementation is managed by the class.

10

### **TModelSurrogate**

TModel Surrogate is a lightweight "stand-in" or "name" for a TModel which provides full addressability of the real model in storage. Any number of TModelSurrogates may exist for a single TModel. TModelSurrogate also serves as the address-space independent specifier of a TModel for purposes of collaboration. TModel Surrogate is a concrete class and may be freely instantiated and destroyed. TModelSurrogate provides concrete implementations for all of the abstract functions of TAbstractModel. The following methods support access to and management of the real model represented by this surrogate: 1) GetModel, 2) CopyModel, 3) DeleteModel. The following methods provide access to the model's anchors: 1) CopyAnchor, 2) CreateAnchorIterator. The following methods support presentation of the model represented by the surrogate. These presentations are static. They do not reflect changes to the content of the real model. 1) CreateIcon, 2) CreateThumbnail.

25

### **TModel**

Figure 8 demonstrates TModel 800, which is the container for all document data 802, including anchors and links and serves as the unit of data exchange for cut/copy/paste. It provides protocol that supports embedding 804 the data in other alien data models and/or embedding alien data models within. The base class TModel provides the implementation for managing a model's anchors 806 and links 808. Subclasses provide the protocol and implementation for accessing and managing the type-specific data. Additionally, the subclass must implement the base class protocol for creating a user interface, selections and accessing embedded models. TModel derives from MDelegatingNotifier and therefore can be connected to for notification on changes to the model's data 810. All models in a TDocument share the TModel class notifier and delegate all calls to that notifier. Several member functions support management of a model's user interface model. Subclasses must override CreateUserInterface to create the appropriate user

35

interface for this model. In the future the choice of interface could be determined from a user preference. The member functions are: 1) CreateUserInterface, 2) AdoptUserInterface, 3) GetUserInterface, 4) GetUserInterfaceSurrogate. Numerous member functions also support accessing and modifying a model's set of anchors and links, as well as accessing and modifying attributes associated with the models.

## DATA PRESENTATION

### Notification Classes

Data contained in a TModel is of little interest unless it can be viewed and/or modified by the user. Although the document framework does not directly provide the classes necessary to implement data presentations, the document framework does provide support for managing presentations on models. This section deals primarily with notification, since this is the extent of presentation support provided directly by the document framework. For example, all other presentation support could be provided by an object-oriented User Interface Toolbox.

### User Interface Management

The document framework provides support for managing a model's user interface. Protocol is provided by TModel for creating, storing and retrieving a model's associated user interface. A client of TModel can ask the model for its user interface. In turn the client can ask the user interface for an Embedded, or Window presentation. The document framework provides no support for the contents of a user interface, only for managing its access and storage. See the section, "Data Representation" for a description of this facility of TModel.

### Notification

Figure 9 sets out the notification system for the document framework. The document framework provides change notification to clients of TModels (clients are typically presentation views) on all changes made to the data contained in a model. A presentation can connect to a model for notification on specific changes, or all changes, to a model's data. By performing the proper updates to the presentation when the notifications are received, the presentation remains synchronized with the underlying data. TModel provides notifications for changes to the data it manages, specifically anchors & links, via TLinkNotification 906 and TAnchorNotification 908, and selections, via TSelectionNotification 902. Standard notifications are also provided for Cut/Copy/Paste operations (not explicitly shown). Other notifications relative to the document framework are also contemplated. For example, TDualSelectionNotification provides notification of



dual selections. It is the responsibility of a subclass of TModel to provide notifications specific to that model's data. The notifications must be sufficient for presentations to remain consistent with the underlying data.

5

### Data Specification

In the document framework the class that supports the specification of document data is TModelSelection. It is important to note that model selections are truly specifications of the selected data, and do not actually contain the data. Furthermore, the specification must identify the data in an address independent way, so that selections can be applied in multiple collaborator's address spaces. The document framework manages data at the granularity of whole models and provides default functionality that operates on whole models. TModelSelection is no exception. TModelSelection provides protocol and implementation for specifying the selected model. Typically an implementation of a derived class of TModel will provide a corresponding derived class of TModelSelection which supports specification of the data contained within the model at a finer grain than the whole model. This allows commands to be applied to a subset of the data in a model, rather than the entire model.

20

### Data Exchange

Figure 10 shows the specification classes, which include TModelSelection 1002 and MCollectible 1000. TModelSelection 1002 defines protocol for the exchange of data between selection objects. This protocol is used by a number of standard commands provided by the document framework. These include, Cut, Copy, Paste, Push and Pull commands. TModelSelection 1002 provides the protocol for negotiating data types for exchange between selections. A source selection produces a list of model types in which it can produce the selected data. The destination selection chooses the type of model in which it prefers to receive the data and how it will receive it (absorb or embed). After a type has been selected via the type negotiation process, the source selection is asked for a copy of the data in the chosen type by calling CopyData. The model returned is then passed to the destination selection in AbsorbData or EmbedData depending on how the selection indicated it would receive the data in the type negotiation. The negotiation may occur across document teams when exchanging data between anchors. In this case services of TRemoteDocument are used to access remote information. The entire exchange process described above is typically carried out by a command object such as TPasteCommand or TPushDataCommand.

### **TModelSelection**

The class, TModelSelection 1002, provides most of the protocol that document selections and anchors (persistent selections) are expected to implement. It serves as the base class for all selection objects. TModelSelection objects contain  
5 the protocol for exchanging data between selections using cut, copy, and paste or using push/pull (on anchors). This includes the protocol for type negotiation (what types can I publish this data in, what types can I accept data in) and the protocol for accepting or publishing data in a specified type.

Member functions CreatePreferredTypeList, ChoosePreferredType, and  
10 GetTypeAcceptKind are inherited from MTypeNegotiator and support type negotiation for exchange of data between selections. Member functions are also provided for supporting adding, removing, and creating anchors to the selection. Member functions are also provided to support the exchange of their associated model's specific type of data. The subclass does not need to concern itself with the  
15 copying of anchors and links. This is handled by TModelSelection automatically. Member functions are also provided to override if anchors must be adjusted when imported or export (e.g. anchors in text must be adjusted to be relative to the exported data when exported and to the containing context when imported).

### **Data Modification**

Figure 11 shows the basic relationships necessary for carrying out data modification. TModelCommandGroup 1104, TModelCommand 1102, and TCommand 1100 each cooperate to provide the necessary modifications required  
20 within the system.

### **Modification Classes**

In the document framework the abstract base class that represents a command made by the user is TModelCommand 1102. Subclasses of  
25 TModelCommand 1102 encapsulate the semantics of a user action and can change model based on the user action. Command objects can be "done", "undone", and "redone", and are typically independent of the user interface technique used to create them.

### **TModelCommand 1102**

TModelCommand 1102 is a derived class of the TCommand 1100 class.  
35 TModelCommand 1102 objects encapsulate a user action which changes the model. TModelCommands have protocol for: Doing, undoing and redoing the change to the model, Identifying the collaborator who issued the command, and Doing the change incrementally, as when dragging or typing.

TModelCommand objects 1102 typically operate on a selection which may be part of the command object or more typically the current document selection. The base class, TModelCommand 1102, provides the protocol that all model command objects respond to. Subclasses override the "HandleXXX" methods to provide specific behavior.

Some TModelCommands are intended to be executed incrementally. These commands are used to allow the user to incrementally modify a model. These commands are called repeating commands. One example of a repeating command is a shape dragging command for a drawing program. Even though a tracking command might force the model to go through many intermediate states, it counts as a single command. This means that the entire effect of the command is undone and redone as a single atomic action.

The document framework uses the concept of a command object in its framework for Undo. Command objects are also central to the collaboration features of the document framework.

### Model Based Tracking Details

Figure 12 shows a diagrammatic representation of a possible processing flow for carrying out a part of the tracking operations in the document framework. As indicated by 1200, when the tracker calls DoBegin() the command argument is flattened 1202 and sent to the collaboration server, at 1204. From there it is flattened again and sent to each of the collaborators. At each collaborator the command's HandleDoBegin() method is executed, as indicated at 1206.

Figure 13 shows a diagrammatic representation of a possible processing flow for carrying out another part of the tracking operations in the document framework. When the tracker calls DoRepeat(), as indicated by 1300, the command argument is asked to stream out the delta information 1302. The command delta is sent to the collaboration server, at 1304. From there the delta is sent to each of the collaborators. At each collaborator the delta is streamed into the local copy of the command, as indicated by 1306. Then the command's HandleDoRepeat() method is executed.

When the tracker calls DoEnd() similar operations as those shown with respect to Figure 13 are performed. The command argument is asked to stream out the delta information. The command delta is sent to the collaboration server. From there the delta is sent to each of the collaborators. At each collaborator the delta is streamed into the local copy of the command. Then the command's HandleDoEnd() method is executed. There are two ways an incremental command can finish its extended Do phase. The standard way is for DoEnd() to be called. The

other way is for the collaborator who is doing the tracking to unexpectedly leave the collaboration. In that case the command on the collaboration server has its HandleCollaboratorDied() method called. After the extended Do phase is finished, the command is expected to be in the same state as if its HandleDo() method was called.

**TModelCommandGroup:** TModelCommandGroup is a subclass for TModelCommand which allows complex commands to be easily built from several simple commands. TModelCommandGroup delegates most methods (e.g. HandleDo) to each of the commands in the group. The commands are not truly serialized, in that HandleLocalDo is called for all commands, and then HandleDo is called for all commands. If the HandleLocalDo of a command relies on state set by an earlier command's HandleDo, this will not work.

The following member functions support the management of commands contained in the command group: AdoptFirst, AdoptLast, Orphan, OrphanAll.

#### **Standard Commands**

The following commands are all standard commands provided by the document framework.

**TSelectCommand:** The TSelectCommand should be issued when changing selections. TSelectCommand is an incremental command supporting direct manipulation for selection. This command sets the selection for the collaboration initiating the command and does not affect other collaborator's selections.

**TCutCommand:** The TCutCommand has the local effect of cutting the current selection out of the document. It has the global effect of adding something to the clipboard. The local effect is accomplished by subclassing TReplaceSelection command.

**TCopyCommand:** The TCopyCommand has no local effect. It has the global effect of putting something on the clipboard.

**TPasteCommand:** The TPasteCommand replaces the current selection with the data on the clipboard. This is accomplished by subclassing TReplaceSelection command.

**TReplaceSelectionCommand:** The TReplaceSelectionCommand replaces the data specified by a selection or anchor with data encapsulated in the command object. The command object contains a TModel which is used to replace the selection's data. You will typically never create a TReplaceSelectionCommand yourself. the document framework use this command object in cut, paste, push, pull, etc.

**TNewAnchorCommand:** The TNewAnchorCommand is issued whenever a new anchor is created. A TStartLinkCommand acts as a TNewAnchorCommand after the global effect of TStartLink is done.

**TNewLinkCommand:** The TNewLinkCommand is issued whenever a new link is created. A TCompleteLinkCommand subclasses TModelCommandGroup with TNewLinkCommands embedded in it.

**TStartLinkCommand:** The TStartLinkCommand has the global effect of putting a new anchor on the "link board" and the local effect of adding a new anchor to the document. The local effect is accomplished by subclassing TNewAnchorCommand.

**TCompleteLinkCommand:** The TCompleteLinkCommand has to do a lot of work. It has the (possibly) non-local effect of posting a new link command to another document (the document which issued the start link command). It has the local effect of adding an anchor and a link to the current document. This is accomplished using the appropriate command objects (TNewAnchor and TNewLink) in its implementation.

**TPushDataCommand:** The TPushDataCommand has the (possibly) non-local effect of posting a TPushedData command to the destination anchor. All type negotiation is handled via the selection protocol.

**TPullDataCommand:** The TPullDataCommand command could be called the pull command. It retrieves data from an anchor at the other end of a link.

**TFollowCommand:** The TFollowCommand will "follow" a link. This involves posting a TFollowedCommand to the document containing the other side of the link.

**TFollowedCommand:** The TFollowedCommand is posted to the document containing the destination anchor in a link. The "there" side of the TModelLink embedded in the TFollowedCommand is the destination anchor. The Follow method of the destination anchor will be called. Override this method to implement the proper follow behavior (typically scroll the anchor and its selection into view).

### **Anchors & Links**

Figure 14 demonstrates the relationships between MCollectible 1400, MAbstractModeler 1102, TModelSelection 1404, TModelAnchorSurrogate 1106, and TModelAnchor 1408.

Figure 15 demonstrates the relationships among MCollectible 1500, TAbstractModelLink 1502, TModelLinkSurrogate 1504, and TModelLink 1506.

### **Anchor and Link Classes**

Anchors are typically "sticky" document selections. Sticky means that the data selected by the anchor is constant across editing changes in a model.

Links are connections between two anchors. Operations on links include creating them, removing them, "following" them, pushing data from one sticky selection to the other, or pulling data. To create a link between two anchors, the

user must specify the anchors. The situation is similar to copying and pasting data (the user needs to specify a source and destination), so one way to do this is to maintain a kind of "linkboard," analogous to the clipboard.

5 Figure 16, between 1600 and 1606, shows the Start Link command processing flow. The Start Link command would create an anchor out of the current selection at 1602 and place the anchor on the linkboard at 1604.

10 Figure 17, between 1700 and 1710, shows processing flow of The Complete Link command. The Complete Link Command would also create an anchor at 1702, and then create a link between the new anchor and the one on the linkboard at 1704. It is possible to choose Complete Link at 1706 several times, in order to create several links that share a common anchor, as indicated by the processing flow 1708.

15 Figure 18 demonstrates links in a document 1800 which have been created by an application programmatically. We expect that most of the "interesting" uses of linking will fall into this category. For example, annotations 1802 can be attached to the affected parts of the document with links 1804. Scripts 1806 can also refer to parts of a document with links 1808. It should be noted that the links of Figure 18 are merely representative of software structures and are not intended to convey any particular visual characteristics of links, scripts, annotations, or the document 20 portions linked to. There is only one kind of link in the document framework. It is bi-directional, as indicated by the bidirectional links 1804 and 1808, and supports both navigation (finding the other end of the link) and data transfer.

25 Figure 19 demonstrates some of the characteristics of links. Links also have properties 1900, which applications can use to classify 1902 links and to restrict 1904 how links are used. For example, there could be properties that specify what the user 1908 can do 1906 with a link. It might be desirable to allow certain users only to pull data from a spreadsheet and not navigate to the spreadsheet or push data into it. Links that are used to implement annotations will be identified by a certain 30 property. This property will indicate that the link is part of an annotation, and that the appropriate annotation commands are enabled.

A final example is a link that indicates a master-copy relationship. This would be used for import data. Instead of copying and pasting a static representation of a graph, the user can create a link between the original graph and the copy placed in the word processor document.

35

### **TModelAnchorSurrogate**

TModelAnchorSurrogate, shown in Figure 14 as element 1406, is the surrogate class for TModelAnchor 1408. It provides a subset of the protocol available

from TModelAnchor 1408. Specifically it provides protocol for supporting data exchange between anchors. By using the protocol of the surrogate, the client does not need to distinguish between anchors in the local document or a remote document. TModelAnchor 1408 maintain a list of attributes which describes attributes of the link such as what operations may be performed (e.g. push, pull). Attributes are managed as a simple attribute group. Protocol is provided for adding, removing and looking up attributes. TModelAnchors 1408 are owned by their creator and it is the creator's responsibility to delete them unless ownership has been passed to another object like a TModel. Important member functions are: CreatePreferredTypeList, ChoosePreferredType and CopyData.

### **TModelAnchor**

TModelAnchor, shown in Figure 14 as element 1408, is the base class for all persistent selections which serve as anchors for hyperlinks. The default implementation of TModelAnchor 1408 contains a TModelSelection 1104 to which it delegates all TModelSelection calls. TModelAnchor 1408 maintain a list of attributes which describes attributes of the link such as what operations may be performed (e.g. push, pull). Attributes are managed as a simple attribute group. Protocol is provided for adding, removing and looking up attributes. Important member functions are: SetSelection, Follow, Execute, AddAttribute, scripting, automated testing and AdoptLink.

### **TModelLink**

The TModelLink class, shown in Figure 15 as element 1506, provides the representation for a link. Specialization is accomplished by subclassing TModelAnchor. There are no subclasses of TModelLink 1506. The "here" anchor and the "there" anchor are just names for the two anchors in a link. The "here" anchor in a TModelLink 1506 is typically part of the current document and can be turned into a real anchor using a method of the TModel. The "there" anchor may be part of the current document or could belong to another document.

TModelLinks 1506 also maintain a list of properties which describes attributes of the link such as what operations may be performed (e.g. push, pull). Properties are managed as a dictionary of property names and values. Protocol is provided for adding, removing and looking up properties. Important member functions are: GetHere, GetThere, AddAttribute and RemoveAttribute.

### **TDynamicModelAnchor**

TDynamicModelAnchor is the base class for anchors whose data specification is dynamic rather than static. The anchors are call "dynamic" because they may specify different document data each time they are accessed. Normal anchors always refer to the same data in the document. For example, a dynamic anchor could represent the current selection in the document. When data was pushed from the anchor it would be the data represented by the current selection at that moment.

### **Document Control & Communication**

The classes in this section are included here to give the reader an idea how documents use the process model. The class TDocument is created to start a document in a team. It is responsible for creating the various frameworks and services necessary for local control of a document team. This includes starting servers for collaboration and external document control, and providing control of the document filing operations. The class TRemoteDocument provides a remote interface to documents in other teams. It allows data to be accessed from documents in other teams, and supports control of those documents through commands.

### **Model Filing**

TDocument provides support for initiating filing operations on a document. When filing in a model, there is no real work to be done. TModels are filed in automatically when they are first requested from a surrogate. A model may be deleted from memory at any time by the container if desired. Of course, an exclusive lock on the model must be acquired to delete it, and no pointers to the model should be cached and held beyond the duration of a lock on the model. Model surrogates may cache model pointers, but this will be transparent to clients.

For filing out there is a larger role. In order to support a saveless model and maintain storage consistency, the document framework logs commands on all models in a store to that store. Models cannot be written individually, but rather, all dirty models are always written together. This simplifies handling of the command log, by allowing the entire command log to be removed from each store whenever models are written. The models are written when the document is closed, but also on a periodic basis while the document is open, to reduce the restart time for a document in case of a system failure.



### Inter-document Communication

In the document framework, each document runs in its own team.

Documents must be able to communicate with one another to support many of the standard features of the document framework. For example hyper-text linking allows data to be pushed and pulled across documents, and links to be followed across documents. To accomplish this processing, as shown in Figure 20, the document framework 1700 provides a client/server interface 2002 for external documents 2004. To communicate with an external document, a TRemoteDocument is created from a surrogate for the root model of the document. The TRemoteDocument object may be used to open, send commands, query links, etc.

### Document Startup

A new team must be created for each new document started. Although an "application" can be written which creates a document with a specific root type, this is not necessary. The document framework will use the capabilities provided by the runtime system in TTaskProgram, to create and start the appropriate objects in the document team. A document is started by creating a TRemoteDocument with the surrogate for the root model of the document, and then making a request of the document, typically to open. This will start the document team and open the requested document using the surrogate provided. Other requests may be made of the document which do not require opening a user interface on the document (i.e. copying selected data from an anchor).

### TDocument

The TDocument class acts as the controller for the document data. It creates the frameworks and servers necessary for a document team. It also provides control of the filing process. TDocument is instantiated automatically by TModelApplication when a document is started. Important member functions are: SetSelection, GetSelection.

### TRemoteDocument

This class provides access to a document in another team. This is the only way for two documents to interact for purposes of Cut/Copy, Paste, Linking, etc. because only a single document can operate in a team. It allows the document to be controlled externally. Typical uses include, link navigation & data transfer, opening/closing, etc. The following member function supports doing commands on remote documents (documents in another team): Do.

The following member functions support data exchange with anchors in remote documents: CreatePreferredTypeList, ChoosePreferredType, CopyAnchorData and NotifyLink. The following member functions support

opening and closing user interfaces on a remote document: Open, Close and IsOpen. The following member functions support receiving notification on opening and closing of documents: CreateOpenInterest and CreateCloseInterest.

Figure 21 shows the overall system design of the present system and method.

5 The application layer 2102 interacts with the current 2106 and related 2108 documents in storage 2104 via flow 2112. The application layer 2102 also interacts with the system level document framework 2110 of the present invention via flow 2114, indicated by the circled 1. And finally, the document framework 2110 of the present invention interacts with the current 2106 and related 2108 documents in  
10 storage 2104 via flow 2116, indicated by the circled 2.

The particulars of the document framework 2110 have been discussed in detail above. The document framework 2110 is a collection of objects for Anchor and Link Support 2118, Notification 2120, Collaboration 2122, Embedding 2124, Data Presentation 2126, Data Modification 2128, Multi-Level Undo 2130, Content-Based  
15 Retrieval and Indexing 2132, Compound Document Support 2134, Data Representation 2136, Data Specification 2138, Model Processing 2140 and Document Communication 2142. As previously discussed, while many of the objects find particular inventiveness by being placed at the system level, and providing functions in conjunction with the other objects at the system level, it should also be  
20 understood that certain of the objects may also be considered at other levels of the system.

### StockTicker Example

An example of a real life application may assist in clarifying the practical  
25 application of the subject technology. First, to set the stage, a review of the compound document technology. A model is something that holds data and makes it available to facilities that access, modify, and present the data to users. A presentation is a facility that provides a user with a view onto a model for reading and editing. A clipboard is a special model associated with a user's environment  
30 containing a model that is on the clipboard in the cut/copy-paste operation.

The application implements a stock ticker which emulates the stock tickers found in Merrill Lynch offices around the country. The stock ticker presents current stock prices for various stocks traded on the New York Stock Exchange. It receives as input current stock prices from a database residing ultimately at the stock  
35 exchange in New York City. The communication link is well known in the art and could for example be an asynchronous serial communication link transmitting ASCII characters. The ASCII characters are displayed directly on a display using any number of fonts encoded to synchronize with ASCII 8-bit characters. The display

mechanism will be referred to as a view and could be any number of common window or other display mechanisms that are also well know in the art. A clipboard model contains a model that is a StockTicker instance. The StockTicker instance refers to communication and display information ready for simulating a stock display. A destination model is going to receive the StockTicker instance in a paste operation to show how a preferred embodiment can be used to enable a compound documents with diverse information.

The destination model is part of a document; it could either be the model for the "root" model of the document or it could be an embedded model. A paste command is issued on the destination document. The command includes two specifications: 1) a reference to the source model on the clipboard that is to be copied, and 2) a specification of a location in the destination document that it should be pasted into. The paste command operation involves the following steps. First, a type negotiation interaction between the destination model and source model. The source provides a list of types that it can provide and the destination selects a preferred type. In this case, the destination does not recognize the specific StockTicker data type so it accepts it abstractly as a model and embeds it. The destination side gets a new copy of the StockTicker model, and it is adopted into the destination model. The destination model has been modified, so it sends a change notification message to all presentations opened on it. Each presentation receives the notification and re-synchronizes its state with the new model state that includes the new StockTicker. For instance, the presentation could be done by completely re-generating its state from the new model state. A more optimized model-presentation relationship might involve a more special notification that tells the presentation precisely what changed so that the presentation can update only what changed. The outcome is that presentation queries the StockTicker model to generate a presentation of itself. This is a complex operation that is detailed below.

When the StockTicker model responds by constructing a presentation view instance, the constructor of this StockTicker view subclass performs the following steps. It creates a thread owned by the view subclass and starts it. The thread enters a function that performs the following steps. It prepares windowing and graphic facilities for displaying continuously updating and animating graphics to the StockTicker presentation view. A specific example of this is a "graphics port" facility that is a conduit for graphics drawn by this thread onto the presentation view. Then it goes into a loop which draws an animating presentation of the data and queries the stock data as needed. The properly constructed StockTicker presentation view is then adopted as a subview of the outer presentation, or more

generally, it is wrapped inside an intermediate view that supports facilities for managing embedded entity as a whole.

While the invention has been described in terms of a preferred embodiment in a specific system environment, those skilled in the art recognize that the  
5 invention can be practiced, with modification, in other and different hardware and software environments within the spirit and scope of the appended claims.

**CLAIMS**

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

- 1 1. An apparatus for document processing comprising:
  - 2 (a) processor means;
  - 3 (b) storage means attached to the processor means;
  - 4 (c) display means under control of the processor means;
  - 5 (d) at least one document resident in the storage means for display on the display  
6 means;
  - 7 (e) system level document framework means for processing the displayed  
8 document and documents related to the displayed document, including a  
9 means for receiving at least one document processing request from an  
10 application; and  
11 an object for performing the at least one processing request; and  
12 (f) application means for processing the document and related documents  
13 utilizing the document framework means.
  
- 1 2. The apparatus of claim 1, wherein the document framework means includes  
2 means for linking data.
  
- 1 3. The apparatus of claim 2, wherein the means for linking data includes  
2 means for transferring information across links.
  
- 1 4. The apparatus of claim 3, wherein the means for transferring information  
2 across links includes means for transferring commands across links.
  
- 1 5. The apparatus of claim 3, wherein the means for transferring information  
2 across links includes means for transferring document information across  
3 links.
  
- 1 6. The apparatus of claim 1, wherein the document framework means includes  
2 means for containing document data.
  
- 1 7. The apparatus of claim 6, wherein the means for containing data includes  
2 means for manipulating the data.

-28-

- 1 8. The apparatus of claim 1, wherein the document framework means includes  
2 means for stacking one or more commands; and  
3 means for undoing the commands.
- 1 9. The apparatus of claim 8, wherein the means for undoing commands  
2 includes means for undoing commands in an order opposite from which the  
3 commands were stacked.
- 1 10. The apparatus of claim 8, wherein the means for undoing commands  
2 includes means for selectively undoing commands in an order unrelated to  
3 an order in which the commands were stacked.
- 1 11. The apparatus of claim 1, wherein the document framework means includes  
2 means for managing embedded models.
- 1 12. The apparatus of claim 1, wherein the document framework means includes  
2 means for providing notification.
- 1 13. The apparatus of claim 1, wherein the document framework means includes  
2 means for creating complex command groups.
- 1 14. The apparatus of claim 1, wherein the document framework means includes  
2 hierarchical document support means for embedded data.
- 1 15. The apparatus of claim 14, wherein the hierarchical document support  
2 means includes means for embedding data in a model when data cannot be  
3 absorbed directly by a document.
- 1 16. The apparatus of claim 15, wherein the means for embedding includes means  
2 for overriding protocols associated with the embedded data.
- 1 17. The apparatus of claim 1, wherein the document framework means includes  
2 annotation means for providing additional information with respect to  
3 document data.
- 1 18. The apparatus of claim 1, wherein the document framework means includes  
2 retrieval framework means for providing indexing and query processing.

- 1 19. The apparatus of claim 18, wherein the retrieval framework means includes  
2 background indexing.
- 1 20. The apparatus of claim 1, wherein the document framework means includes  
2 object surrogate means for providing address space independent references to  
3 real objects.
- 1 21. The apparatus of claim 1, wherein the document framework means includes  
2 means for seamlessly integrating audio and visual information in a  
3 compound document.
- 1 22. A method for document processing comprising the steps of:  
2 (a) storing a document;  
3 (b) displaying the document;  
4 (c) providing a system level document framework for processing the document  
5 and related documents;  
6 (d) the processing including the steps of:  
7 (d1) receiving document processing requests from an application; and  
8 (d2) performing the request using system level objects of the document  
9 framework.
- 1 23. The method of claim 22, including the step of linking data.
- 1 24. The method of claim 23, including the step of transferring information across  
2 links.
- 1 25. The method of claim 23, including the step of transferring commands across  
2 links.
- 1 26. The method of claim 23, including the step of transferring document  
2 information across links.
- 1 27. The method of claim 22, including the step of containing document data.
- 1 28. The method of claim 27, including the step of modifying the data.

-30-

- 1 29. The method of claim 22, including the step of stacking one or more  
2 commands; and undoing the commands.
- 1 30. The method of claim 29, including the step of undoing commands in an  
2 order opposite from which the commands were requested.
- 1 31. The method of claim 22, including the step of selectively undoing commands  
2 in an order other than the order in which the commands were requested.
- 1 32. The method of claim 22, including the step of managing embedded models.
- 1 33. The method of claim 22, including the step of providing notification.
- 1 34. The method of claim 22, including the step of creating complex command  
2 groups.
- 1 35. The method of claim 22, including the step of supporting hierarchical  
2 document data including embedded data.
- 1 36. The method of claim 35, including the step of embedding data in a model  
2 when data can not be absorbed directly by a document.
- 1 37. The method of claim 36, including the step of overriding protocols associated  
2 with the embedded data.
- 1 38. The method of claim 22, including the step of annotating for providing  
2 additional information with respect to document data.
- 1 39. The method of claim 22, including the step of indexing and query processing  
2 using a retrieval framework .
- 1 40. The method of claim 39, including the step of background indexing.
- 1 41. The method of claim 22, including the step of providing address space  
2 independent references to real objects using object surrogate means.



- 1 42. The method of claim 21, including the step of seamlessly integrating audio or  
2 visual data into a compound document.
- 1 43. The method of claim 21, including the step of sharing access to a single  
2 document by two or more users utilizing command-based-collaboration.
- 1 44. The method of claim 43, including the step of sharing access to a single  
2 document via a communication link utilizing remote command execution.
- 1 45. A method for enabling an updateable display of data within an existing  
2 document in a storage of a processor with an attached display, comprising the  
3 steps of:  
4 (a) constructing an updateable data model in the storage of the processor;  
5 (b) constructing a source model in the storage of the processor with an instance  
6 of the updateable data model;  
7 (c) constructing a destination model in the storage of the processor;  
8 (d) negotiating an interaction between the destination model and the source  
9 model;  
10 (e) embedding an instance of the updateable data model from the source model  
11 in the destination model;  
12 (f) sending a change notification message to the updateable data model;  
13 (g) re-synchronizing each instance of the updateable data model with a set of  
14 current information from the change notification message; and  
15 (h) displaying the current information on the display in each active instance of  
16 the updateable data model.
- 1 46. A method as recited in claim 45, wherein the updateable data include one or  
2 more stock prices.
- 1 47. A method as recited in claim 45, including the step of dynamically initiating a  
2 task when a new instance of an updateable model is created.
- 1 48. A method as recited in claim 45, including the step of updating the data via a  
2 communication link.




FIG. 1




FIG. 2

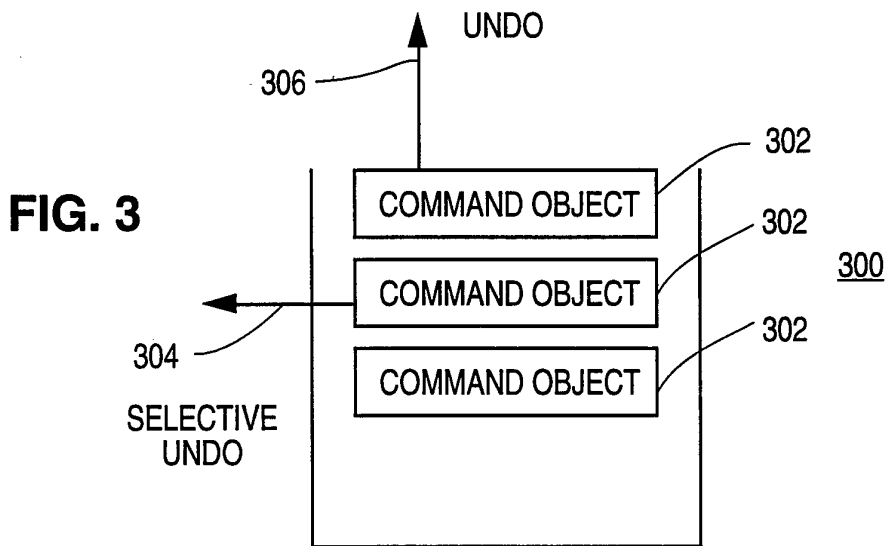


FIG. 3

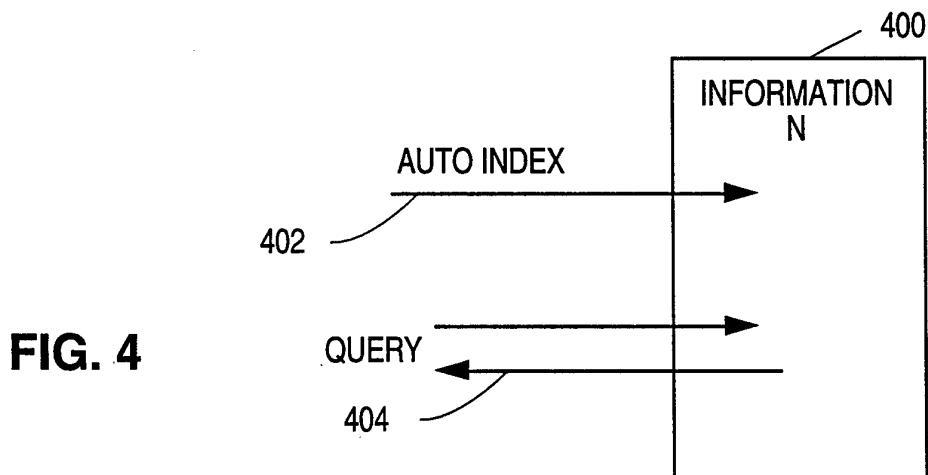




FIG. 4






- 0. DOCUMENT(TPAGESETUPMODEL)
- 1. TALIGENT TEXT (TTEXTMODEL)
- 2. TALIGENTLOGO(TGRAPHICMODEL)
- 3. DOCUMENT BODY (TTEXTMODEL)
- 4. HANDSHAKE ILLUSTRATION (TGRAPHICMODEL)
- 5. TABLE (TTABLEMODEL)



**FIG. 7**



**FIG. 8**

5/13

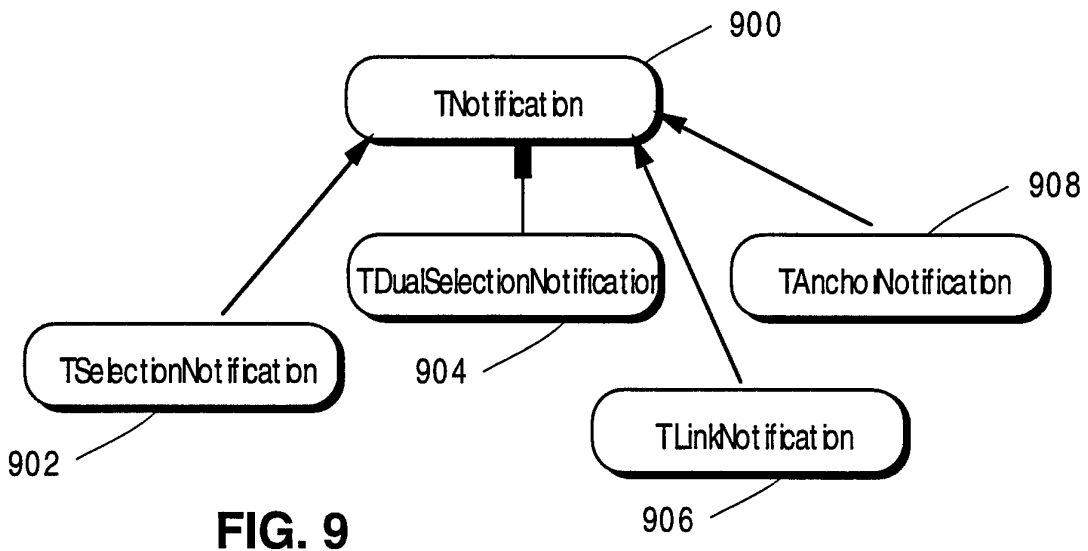


FIG. 9




FIG. 10




FIG. 11

6/13

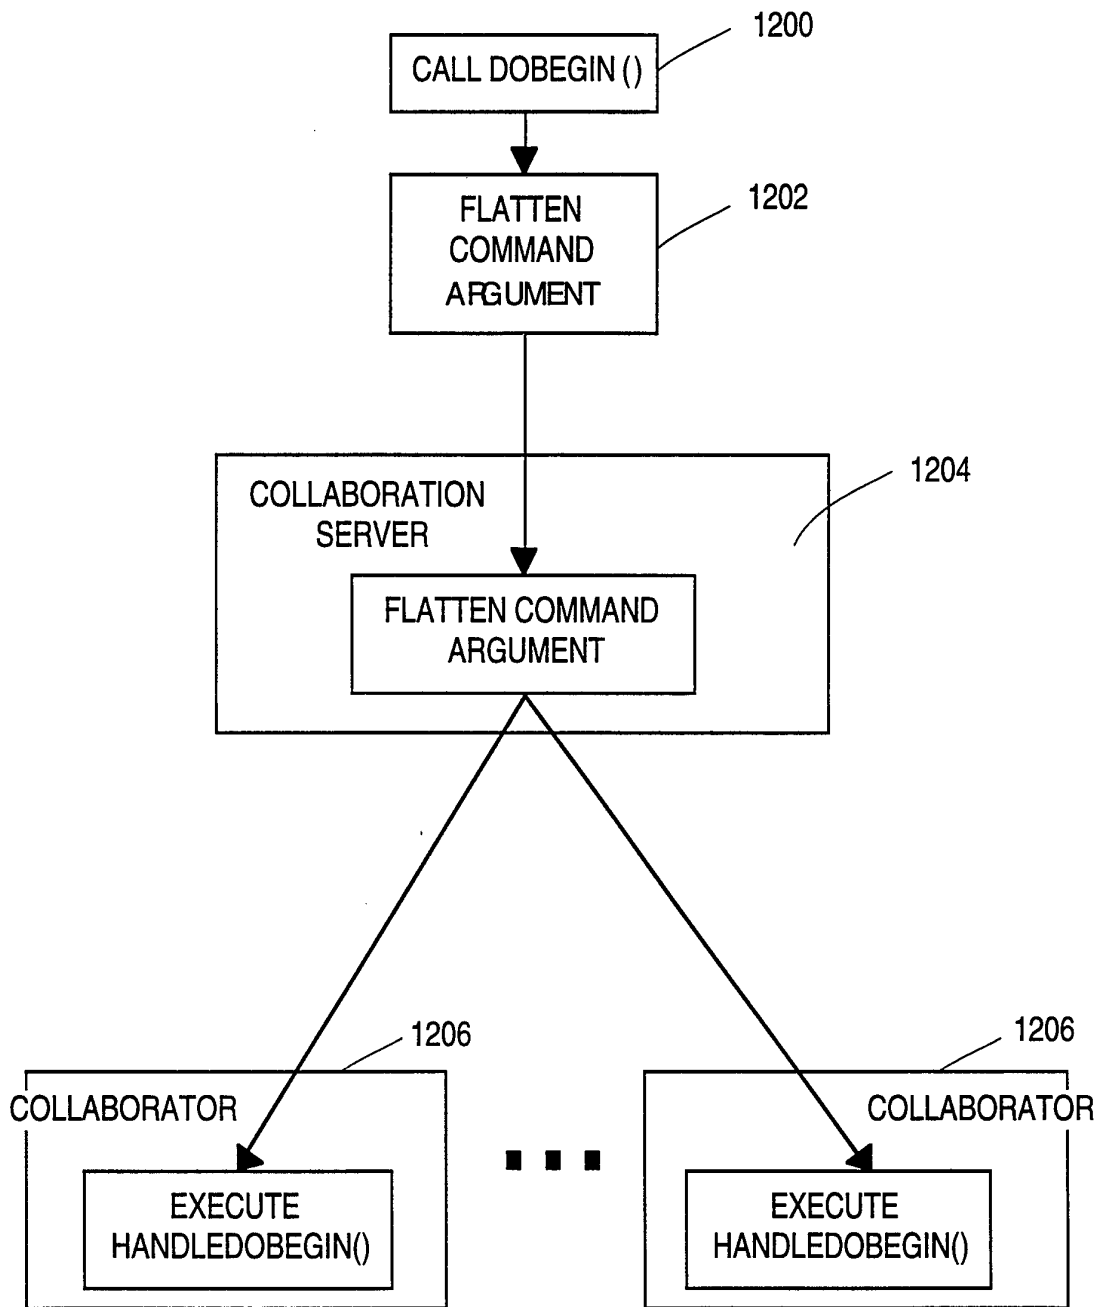


FIG. 12

7/13

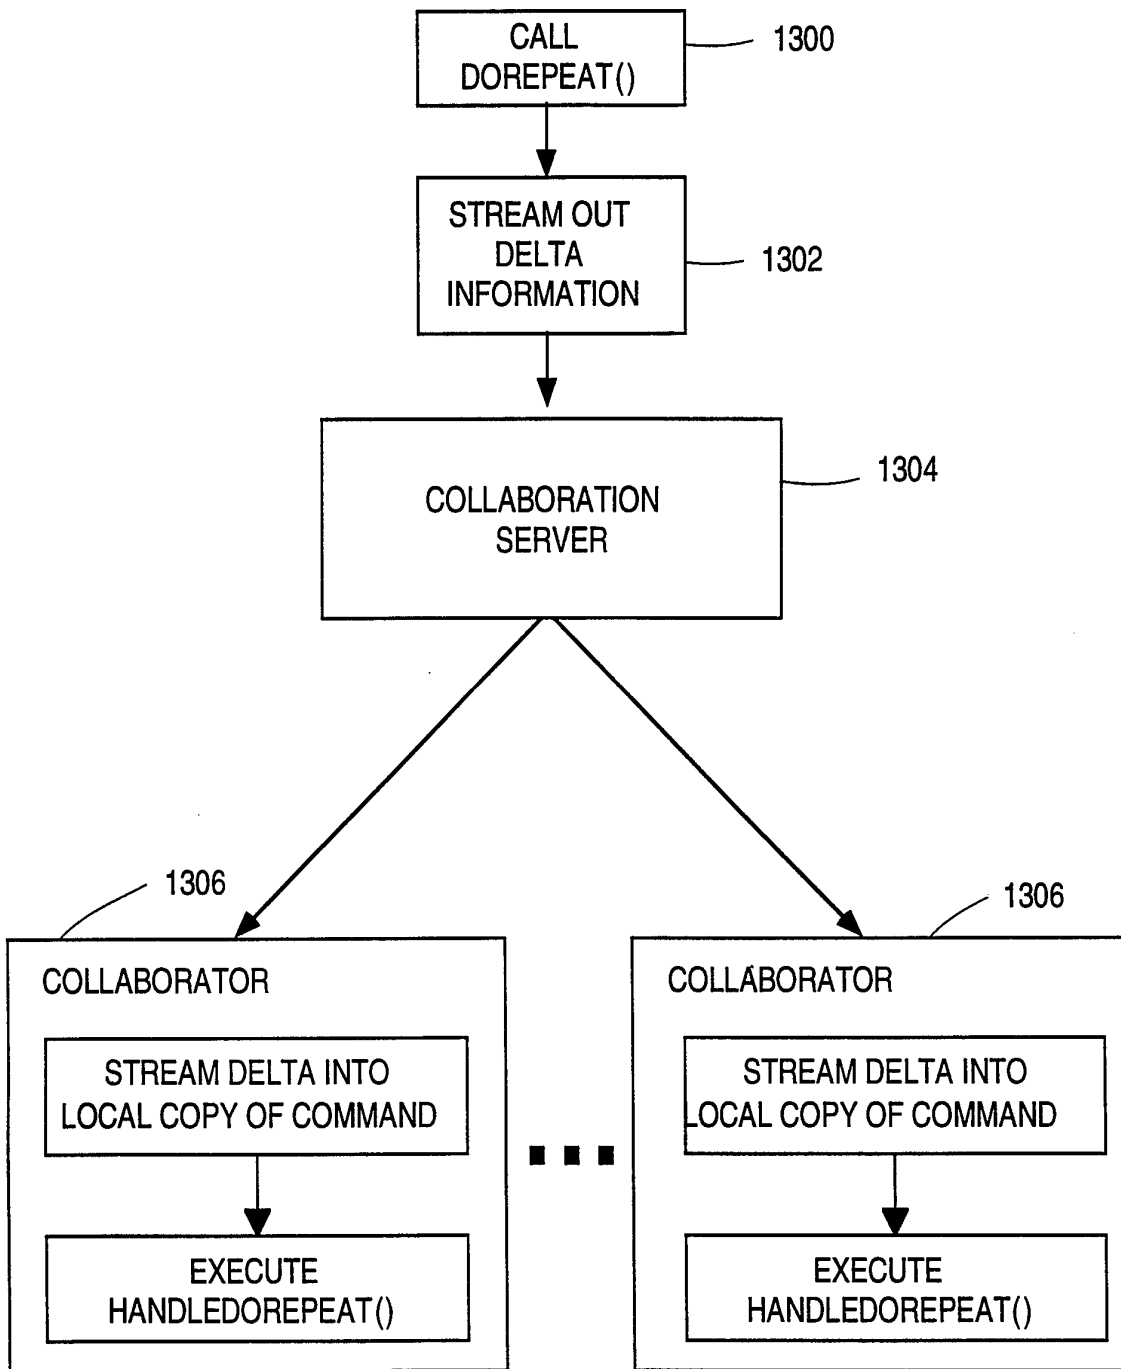


FIG. 13



8/13




FIG. 14




FIG. 15

9/13




FIG. 16

10/13

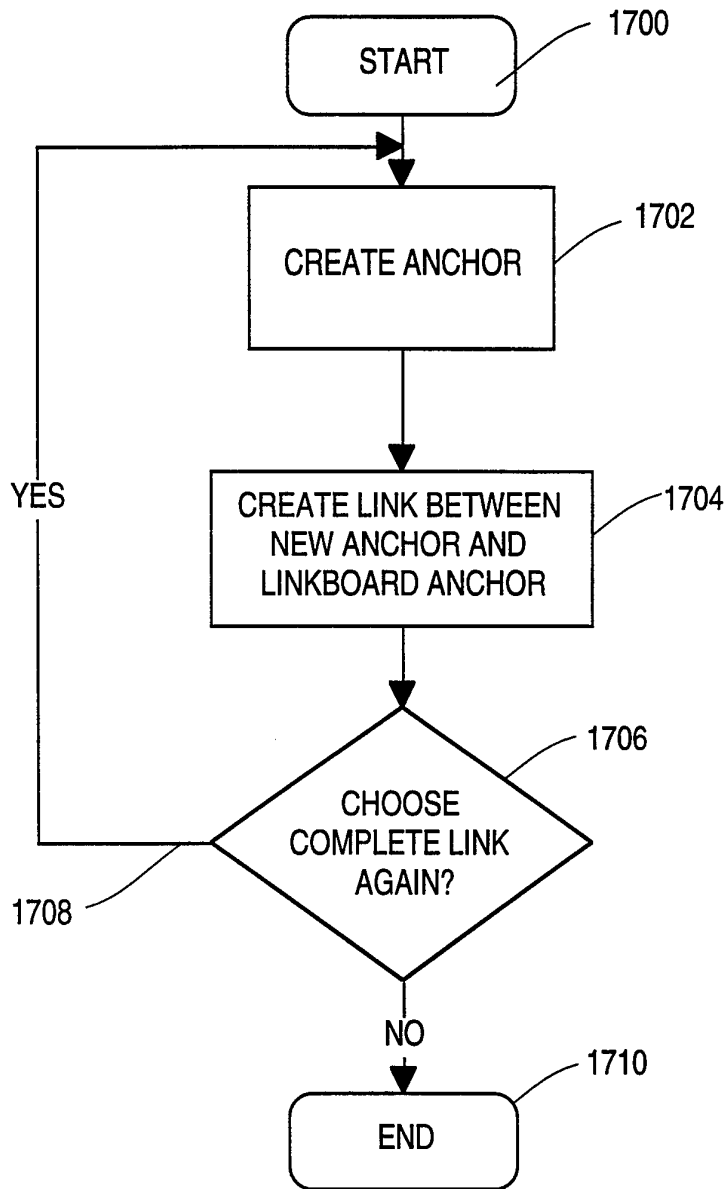


FIG. 17

11/13

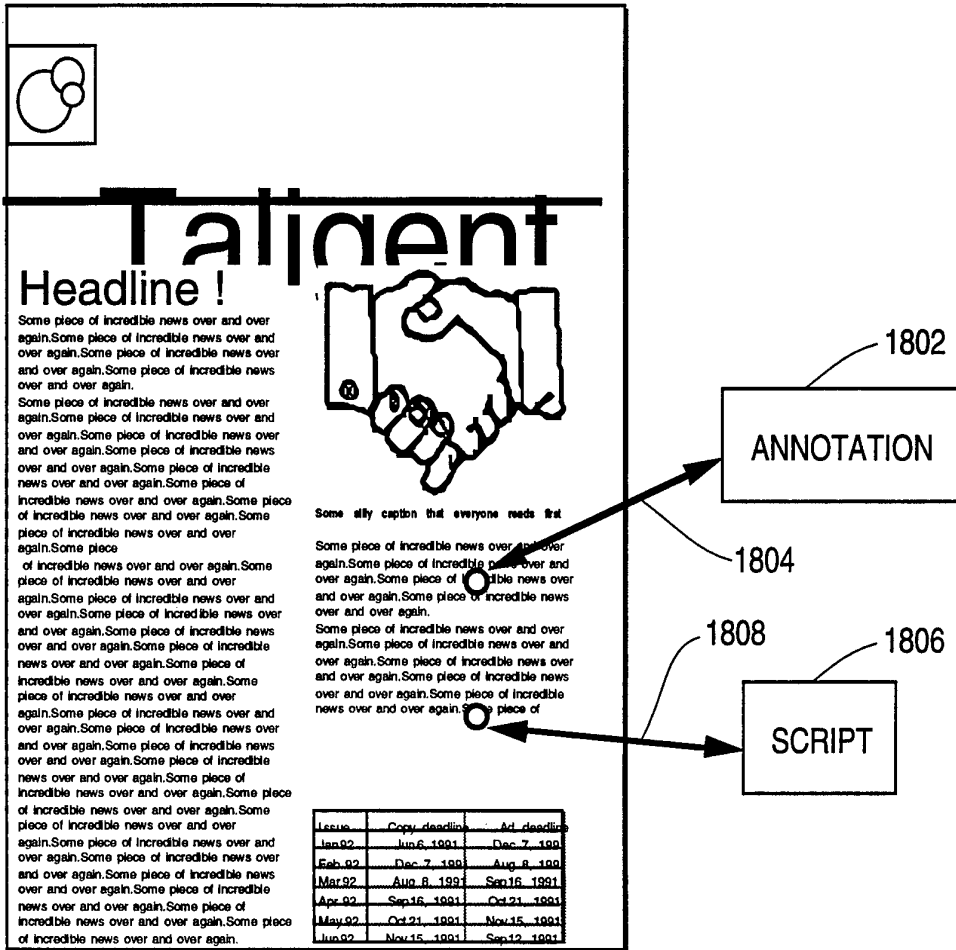


FIG. 18




FIG. 19

12/13

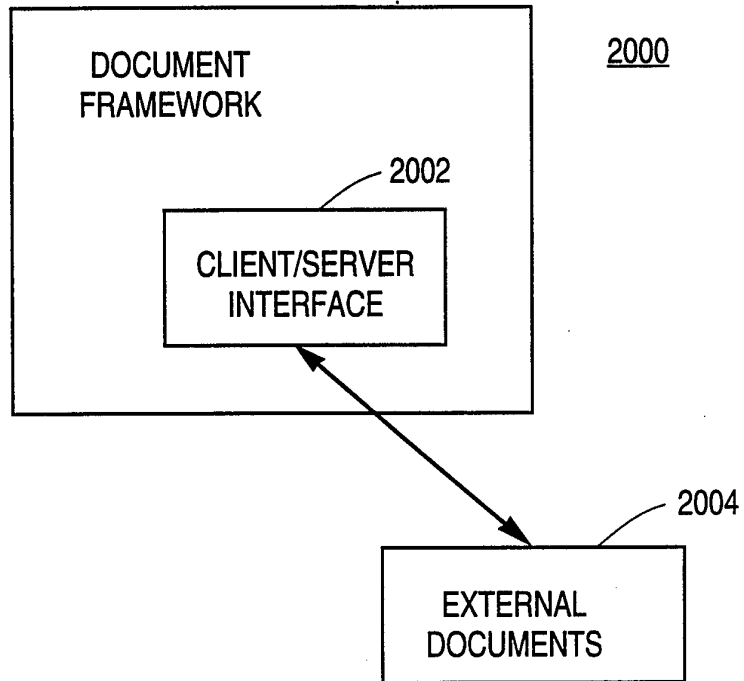


FIG. 20

13/13




FIG. 21

# INTERNATIONAL SEARCH REPORT

International Application No  
**PCT/US 94/00049**

**A. CLASSIFICATION OF SUBJECT MATTER**  
IPC 6 G06F17/21 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>COMPUTER vol. 21, no. 1, January 1988, LONG BEACH US pages 81 - 96 N. YANKELOVICH ET AL. 'Intermedia: The concept and construction of a seamless information environment' see the whole document ---</p>	<p>1-7, 20-28, 41-43,45</p>
X	<p>COMMUNICATIONS OF THE ASSOCIATION FOR COMPUTING MACHINERY vol. 35, no. 1, January 1992, NEW YORK US pages 36 - 51 B.J. HAAN ET AL. 'IRIS hypermedia services' see the whole document ---</p>	<p>1-7, 20-28, 41-43,45</p>

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

\* Special categories of cited documents :

- \*A\* document defining the general state of the art which is not considered to be of particular relevance
- \*E\* earlier document but published on or after the international filing date
- \*L\* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \*O\* document referring to an oral disclosure, use, exhibition or other means
- \*P\* document published prior to the international filing date but later than the priority date claimed

- \*T\* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- \*X\* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- \*Y\* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- \* & \* document member of the same patent family

Date of the actual completion of the international search

9 June 1994

Date of mailing of the international search report

17. 06. 94

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax (+31-70) 340-3016

Authorized officer

Katerbau, R

## INTERNATIONAL SEARCH REPORT

Inter.    nal Application No  
PCT/US 94/00049

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>COMMUNICATIONS OF THE ASSOCIATION FOR COMPUTING MACHINERY vol. 33, no. 3 , March 1990 , NEW YORK US pages 296 - 310 J. NIELSEN 'The art of navigating through hypertext' see the whole document -----</p>	1-7, 20-28, 41-43,45
X	<p>INFORMATIONSTECHNIK IT vol. 32, no. 4 , August 1990 , MUNCHEN BR pages 231 - 240 E. SCHNELL ET AL. 'Konzeptionelle Ansätze für kooperative Applikationen' see the whole document -----</p>	1,22,45
X	<p>PROCEEDINGS OF THE CONFERENCE ON COMPUTER-SUPPORTED COOPERATIVE WORK, CSCW '92 4 November 1992 , TORONTO, CANADA pages 138 - 146 J. HAAKE ET AL. 'Supporting Collaborative Writing of Hyperdocuments in SEPIA' see the whole document -----</p>	1,22,45