



(19) **United States**

(12) **Patent Application Publication**
Atkinson et al.

(10) **Pub. No.: US 2007/0061816 A1**

(43) **Pub. Date: Mar. 15, 2007**

(54) **METHOD AND SYSTEM FOR NAMING AND BINDING OBJECTS**

(75) Inventors: **Robert G. Atkinson**, Woodinville, WA (US); **Antony S. Williams**, Redmond, WA (US); **Edward K. Jung**, Bellevue, WA (US)

Correspondence Address:
KLARQUIST SPARKMAN LLP
121 S.W. SALMON STREET
SUITE 1600
PORTLAND, OR 97204 (US)

(73) Assignee: **Microsoft Corporation**, Redmond, WA

(21) Appl. No.: **11/481,154**

(22) Filed: **Jul. 3, 2006**

Related U.S. Application Data

(60) Division of application No. 10/283,627, filed on Oct. 29, 2002, now Pat. No. 7,150,024, which is a continuation of application No. 09/867,853, filed on May 29, 2001, now Pat. No. 6,519,764, which is a continuation of application No. 08/916,999, filed on Aug. 20, 1997, now Pat. No. 6,263,379, which is a continuation of application No. 08/467,917, filed on Jun. 6, 1995, now Pat. No. 5,740,439, which is a division

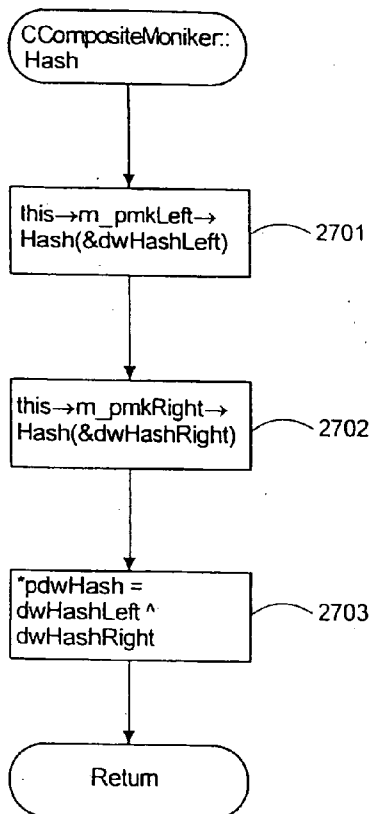
of application No. 08/088,724, filed on Jul. 6, 1993, now Pat. No. 5,581,760, which is a continuation-in-part of application No. 07/909,983, filed on Jul. 6, 1992, now abandoned.

Publication Classification

(51) **Int. Cl.**
G06F 9/46 (2006.01)
(52) **U.S. Cl.** **719/316**

(57) **ABSTRACT**

A method and system for referring to and binding to objects using a moniker object is provided. In preferred embodiments a moniker object contains information to identify linked source data and provides methods through which a program can bind to the linked source data. A binding method is provided that returns an instance of an interface through which the linked source data can be accessed. In a preferred embodiment, the moniker object provides other methods including a reducing method that returns a more efficient representation of the moniker object; equality and hash methods for comparing moniker objects; and inverse, common prefix, and relative-path-to methods for comparing and locating moniker objects from other moniker objects. Several implementations of a moniker object are provided. Each implementation is a moniker class and has a class identifier that identifies code to manage the moniker class.






FIG. 1

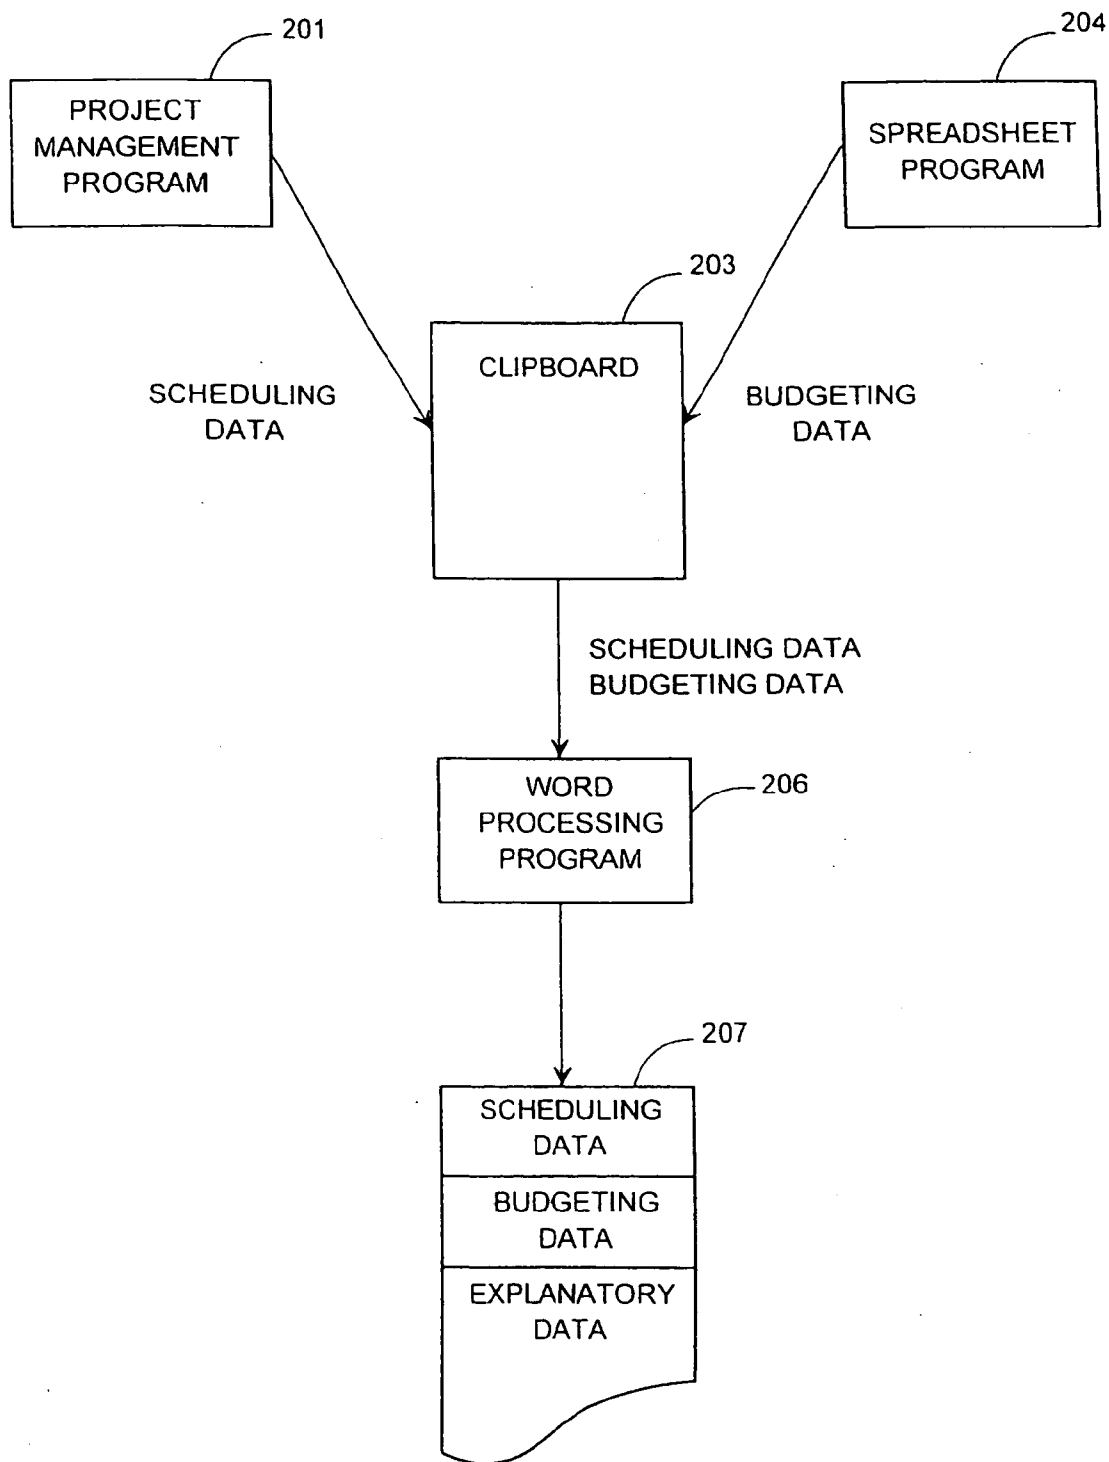


FIG. 2

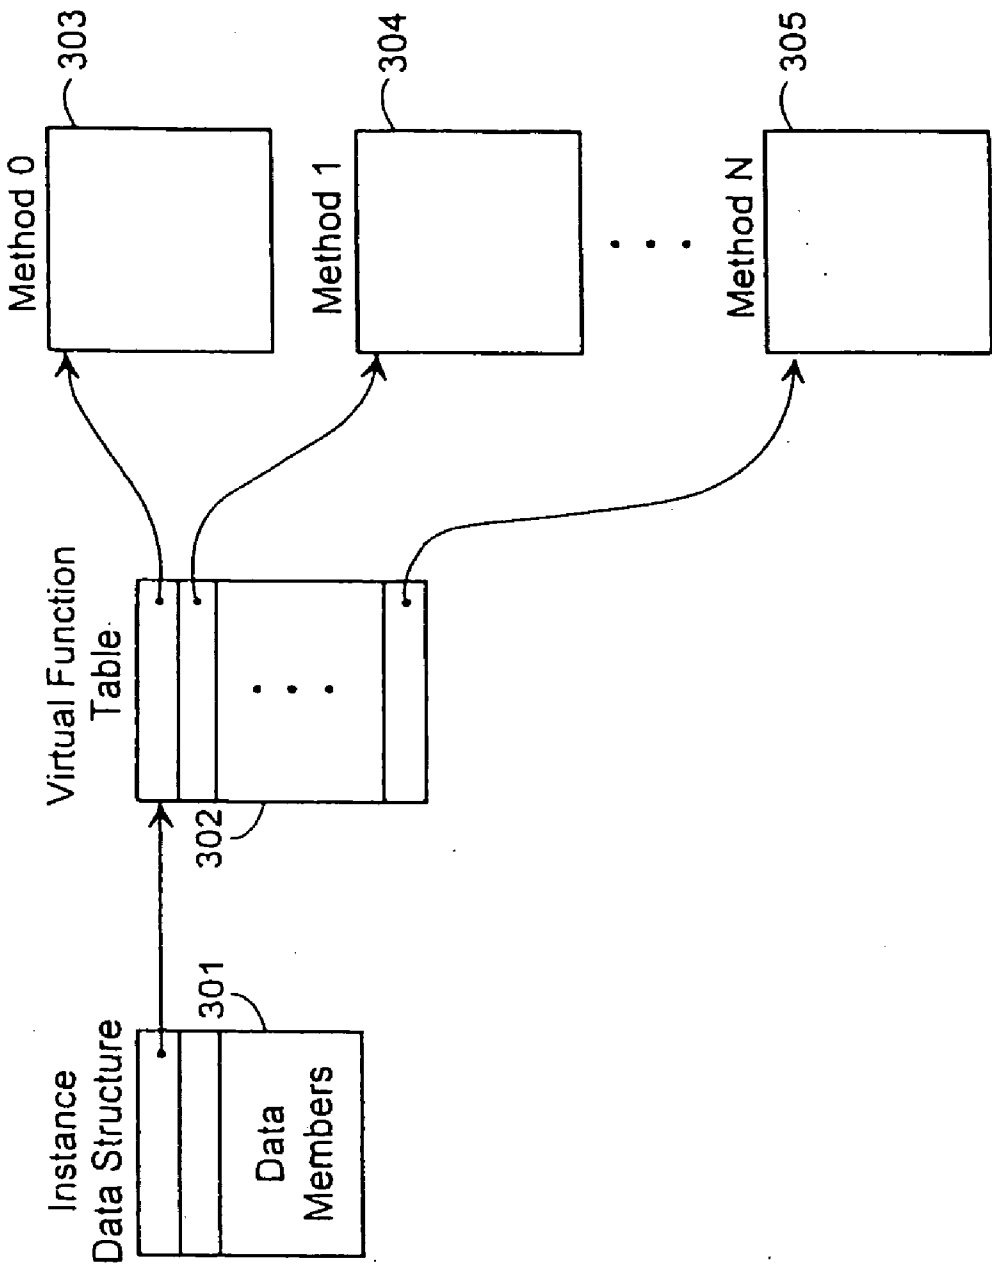


FIG. 3




FIG. 4




FIG. 5




FIG. 6

Suffix Class		ClassFactory Server		
CHT	CLSID_Chart	Chart.DLL	Chart.Exe	704B
		FileMkr.DLL		

Suffix Class	
CHT	CLSID_Chart






FIG. 7

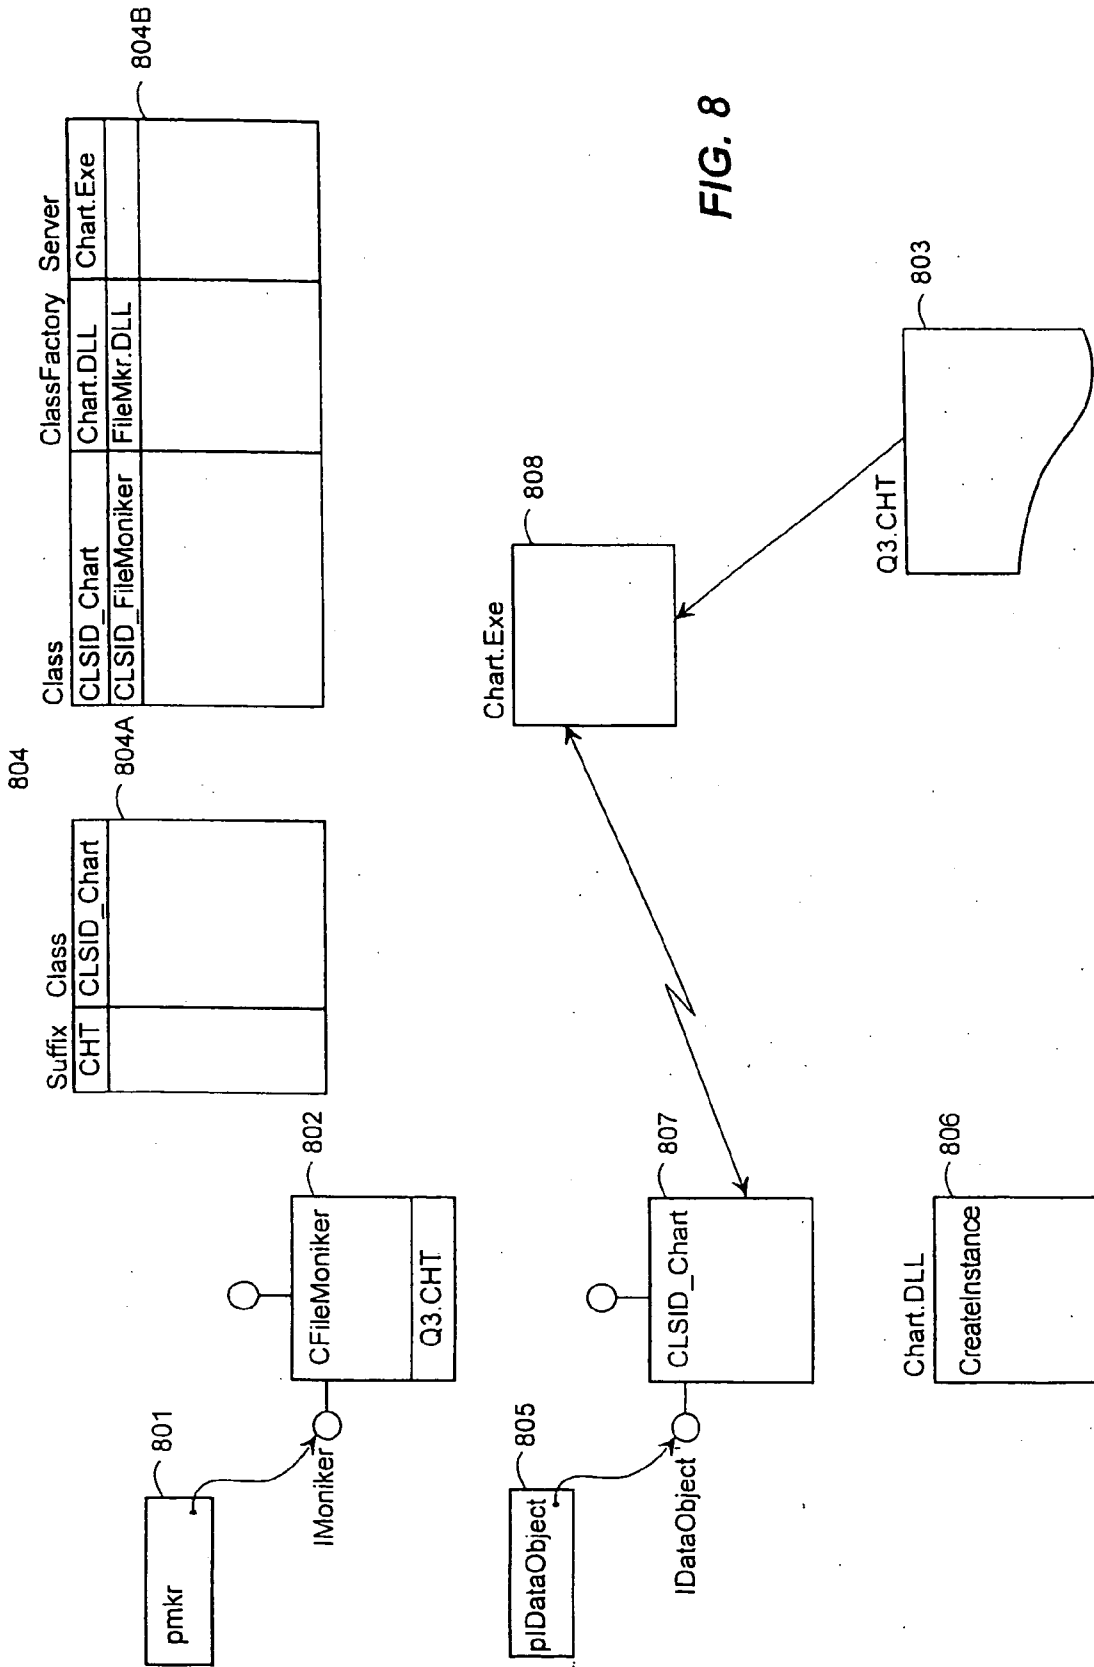


FIG. 8




FIG. 9




FIG. 10B




FIG. 10A

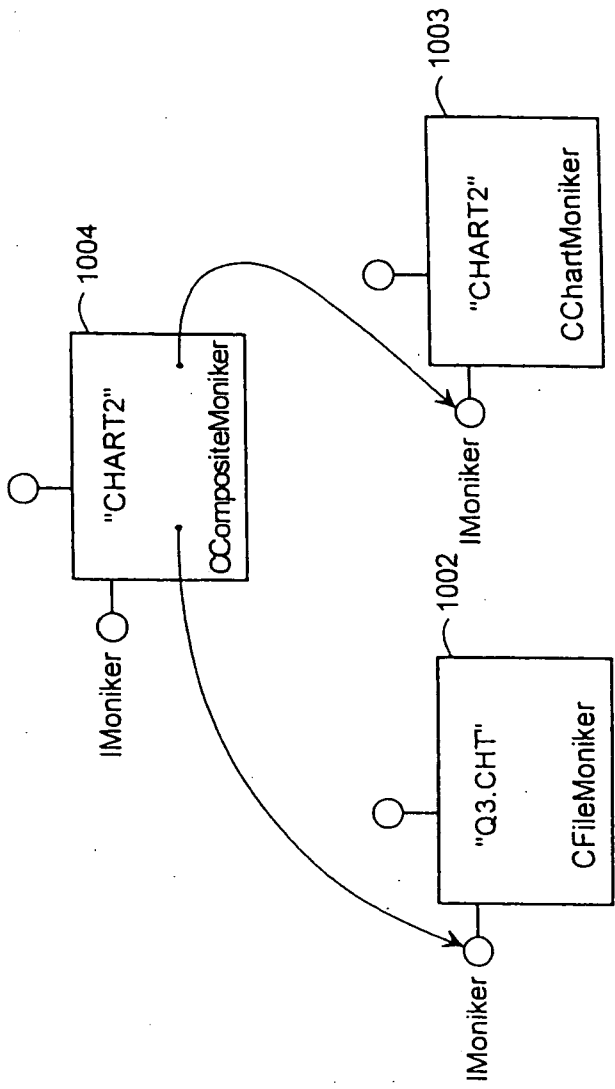


FIG. 10C

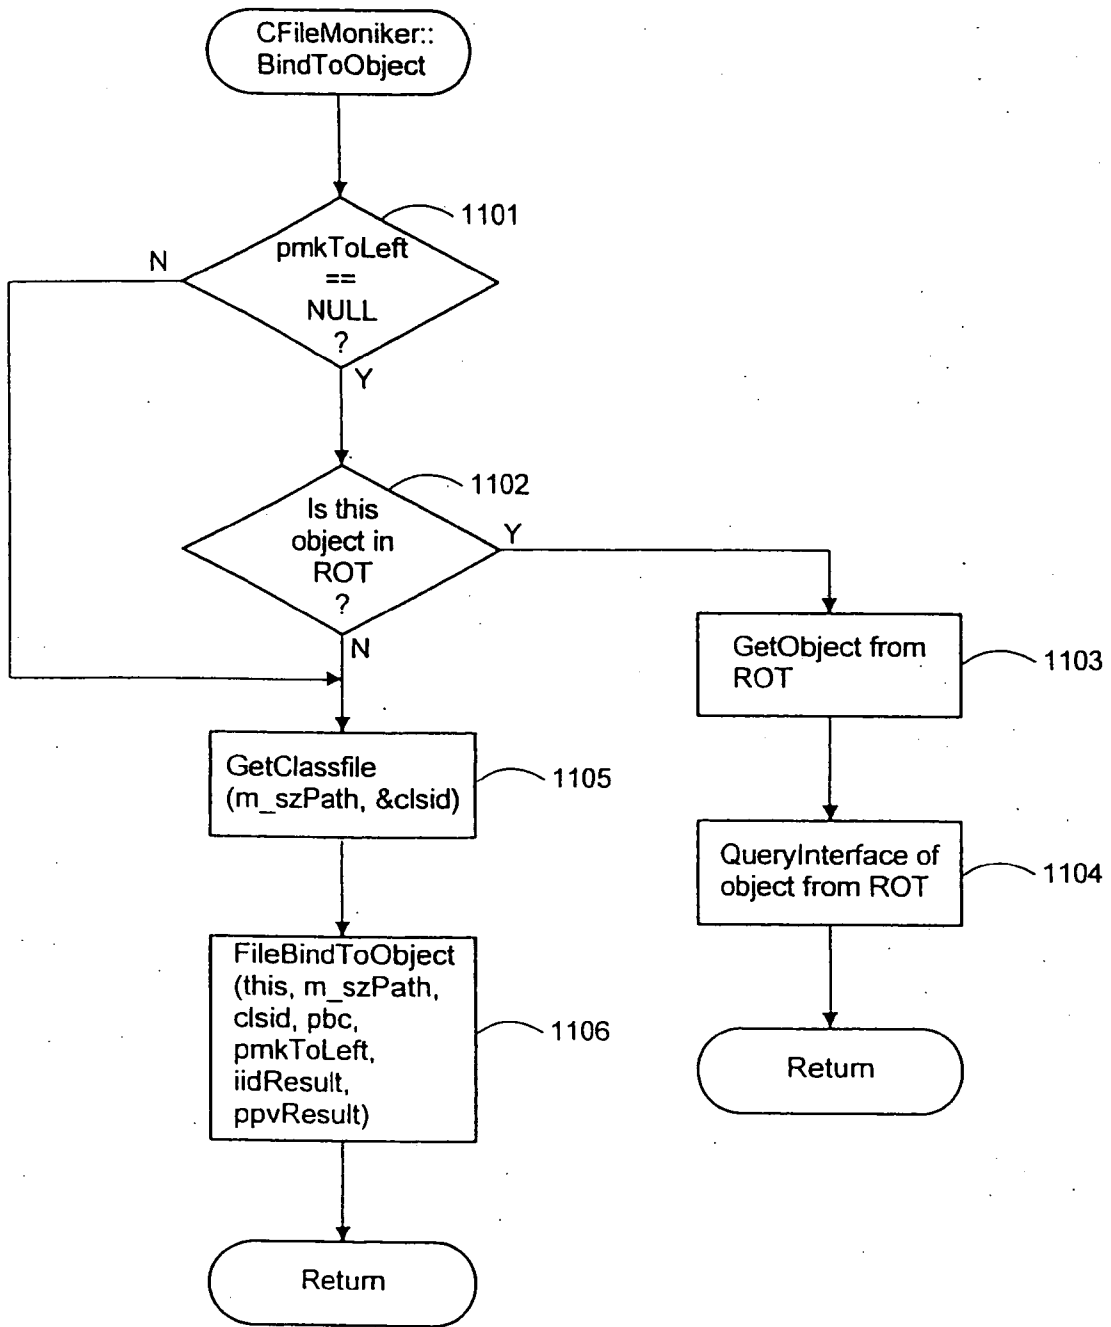


FIG. 11

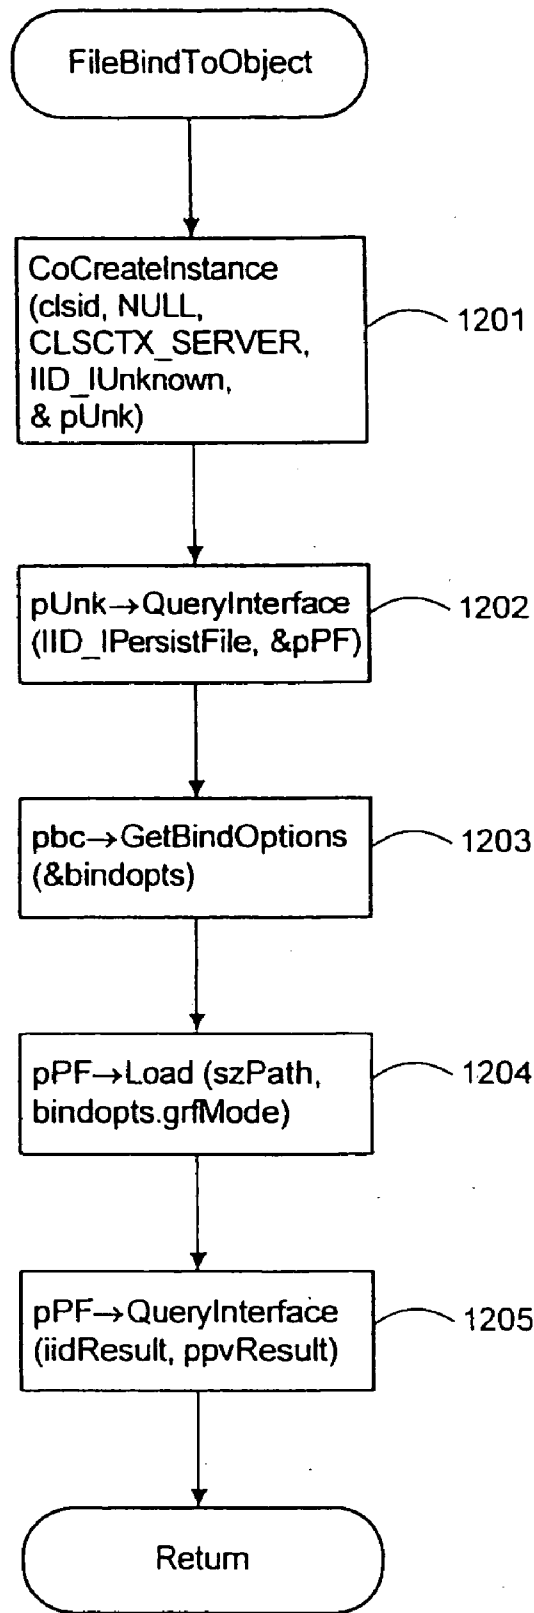


FIG. 12

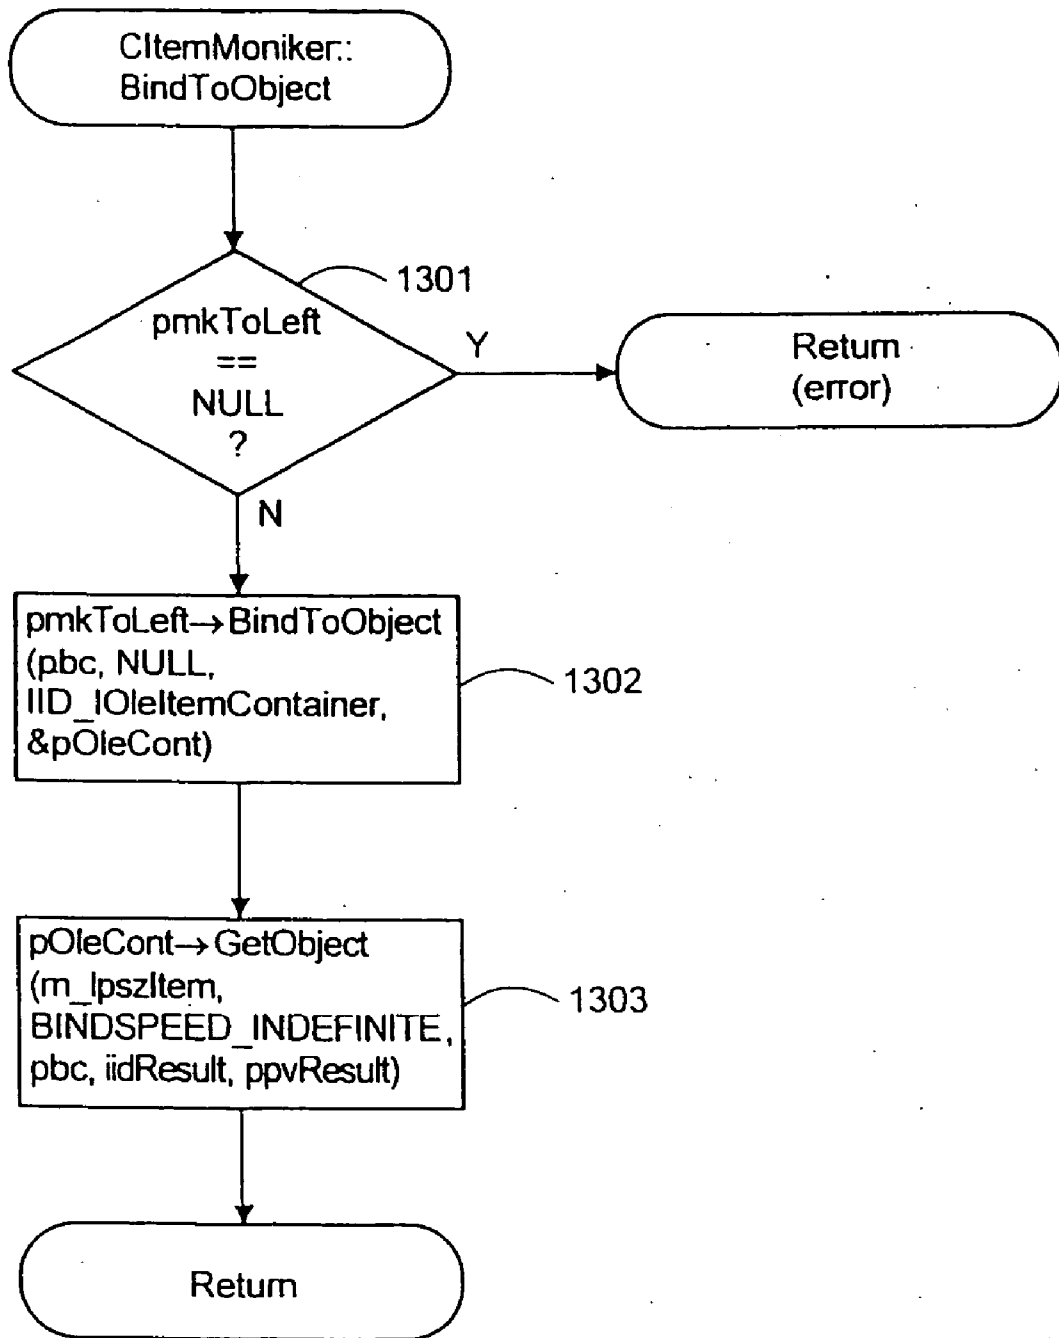


FIG. 13

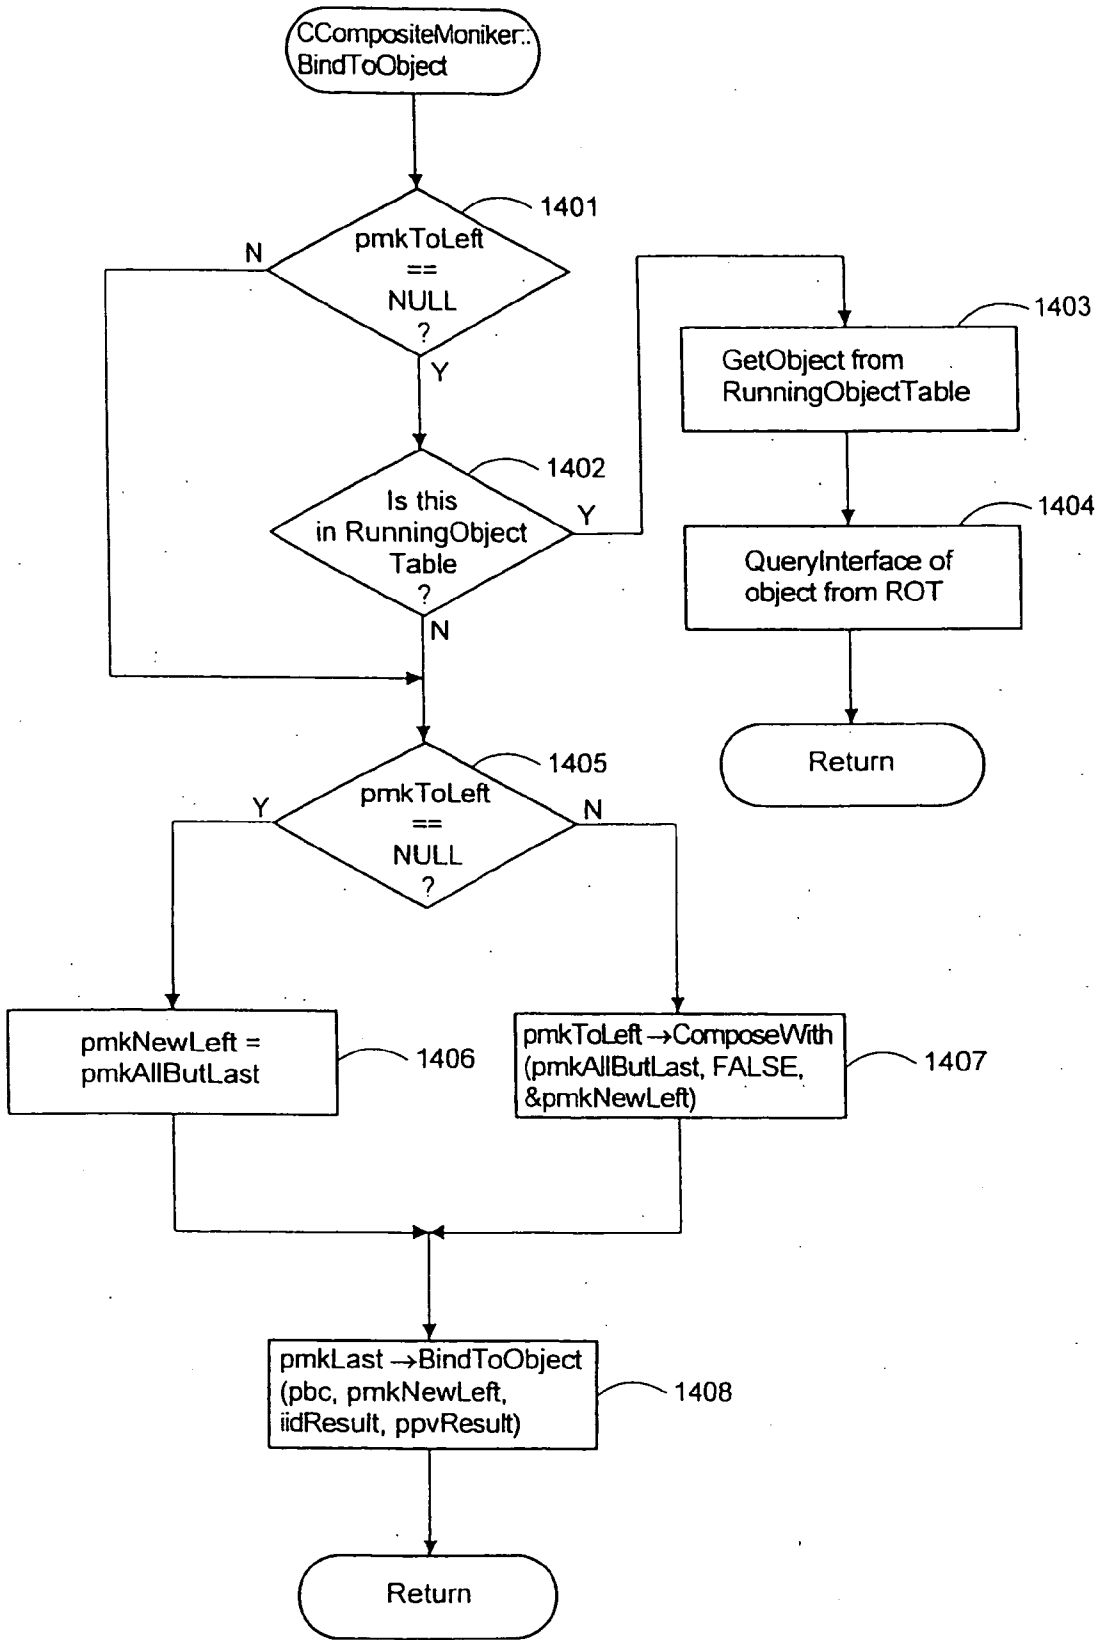


FIG. 14




FIG. 15B

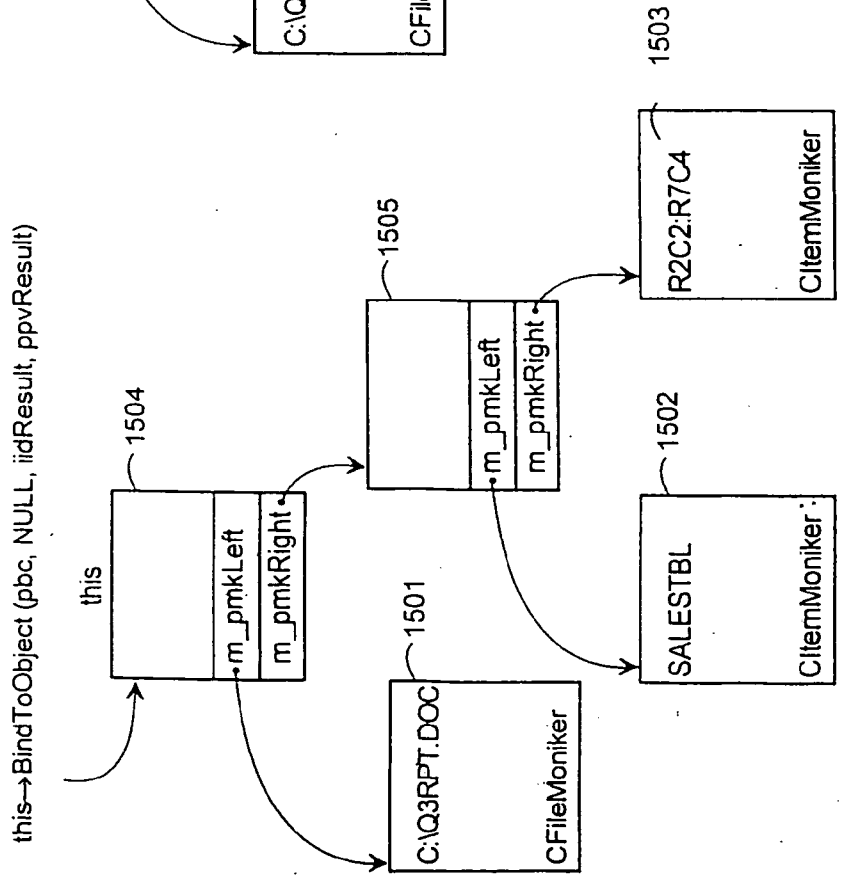


FIG. 15A




FIG. 15E

FIG. 15C




FIG. 15D




FIG. 15F




FIG. 15G

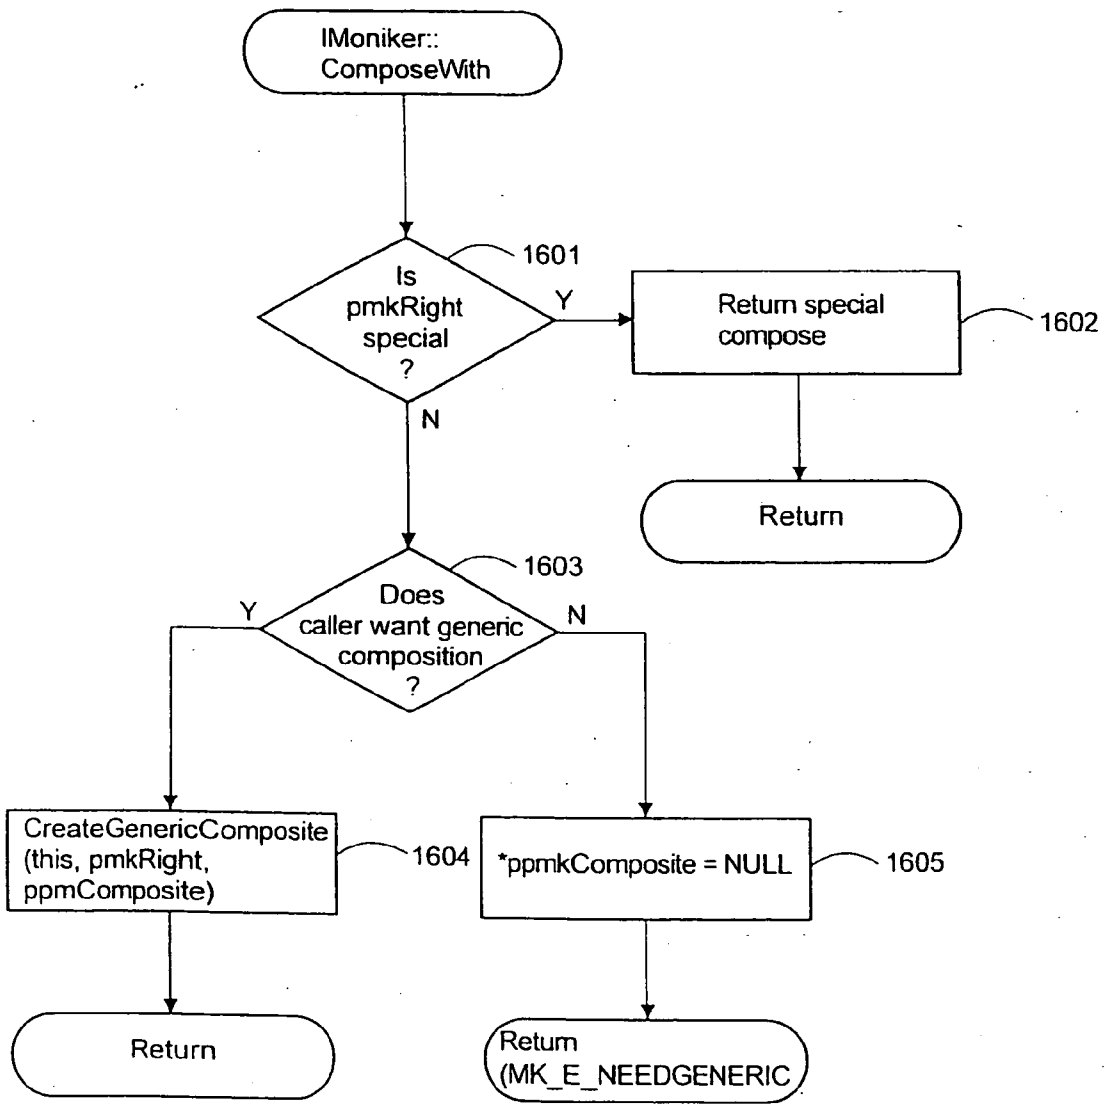


FIG. 16

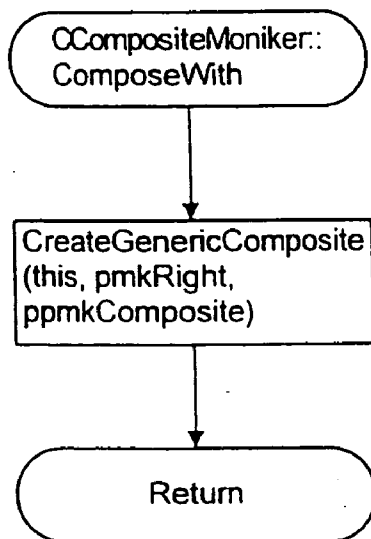


FIG. 17

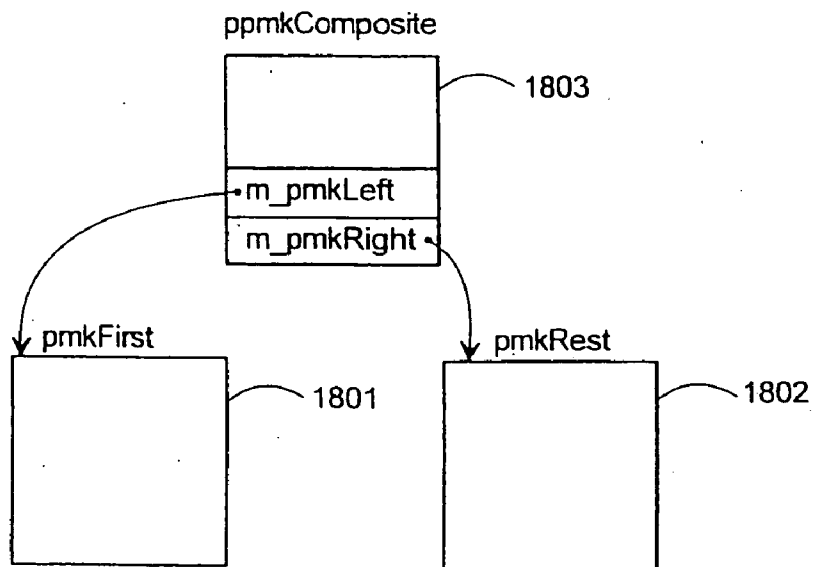


FIG. 18




FIG. 19B




FIG. 19C




FIG. 19A




FIG. 20B




FIG. 20C




FIG. 20A




FIG. 21A




FIG. 21C

FIG. 21B




FIG. 22

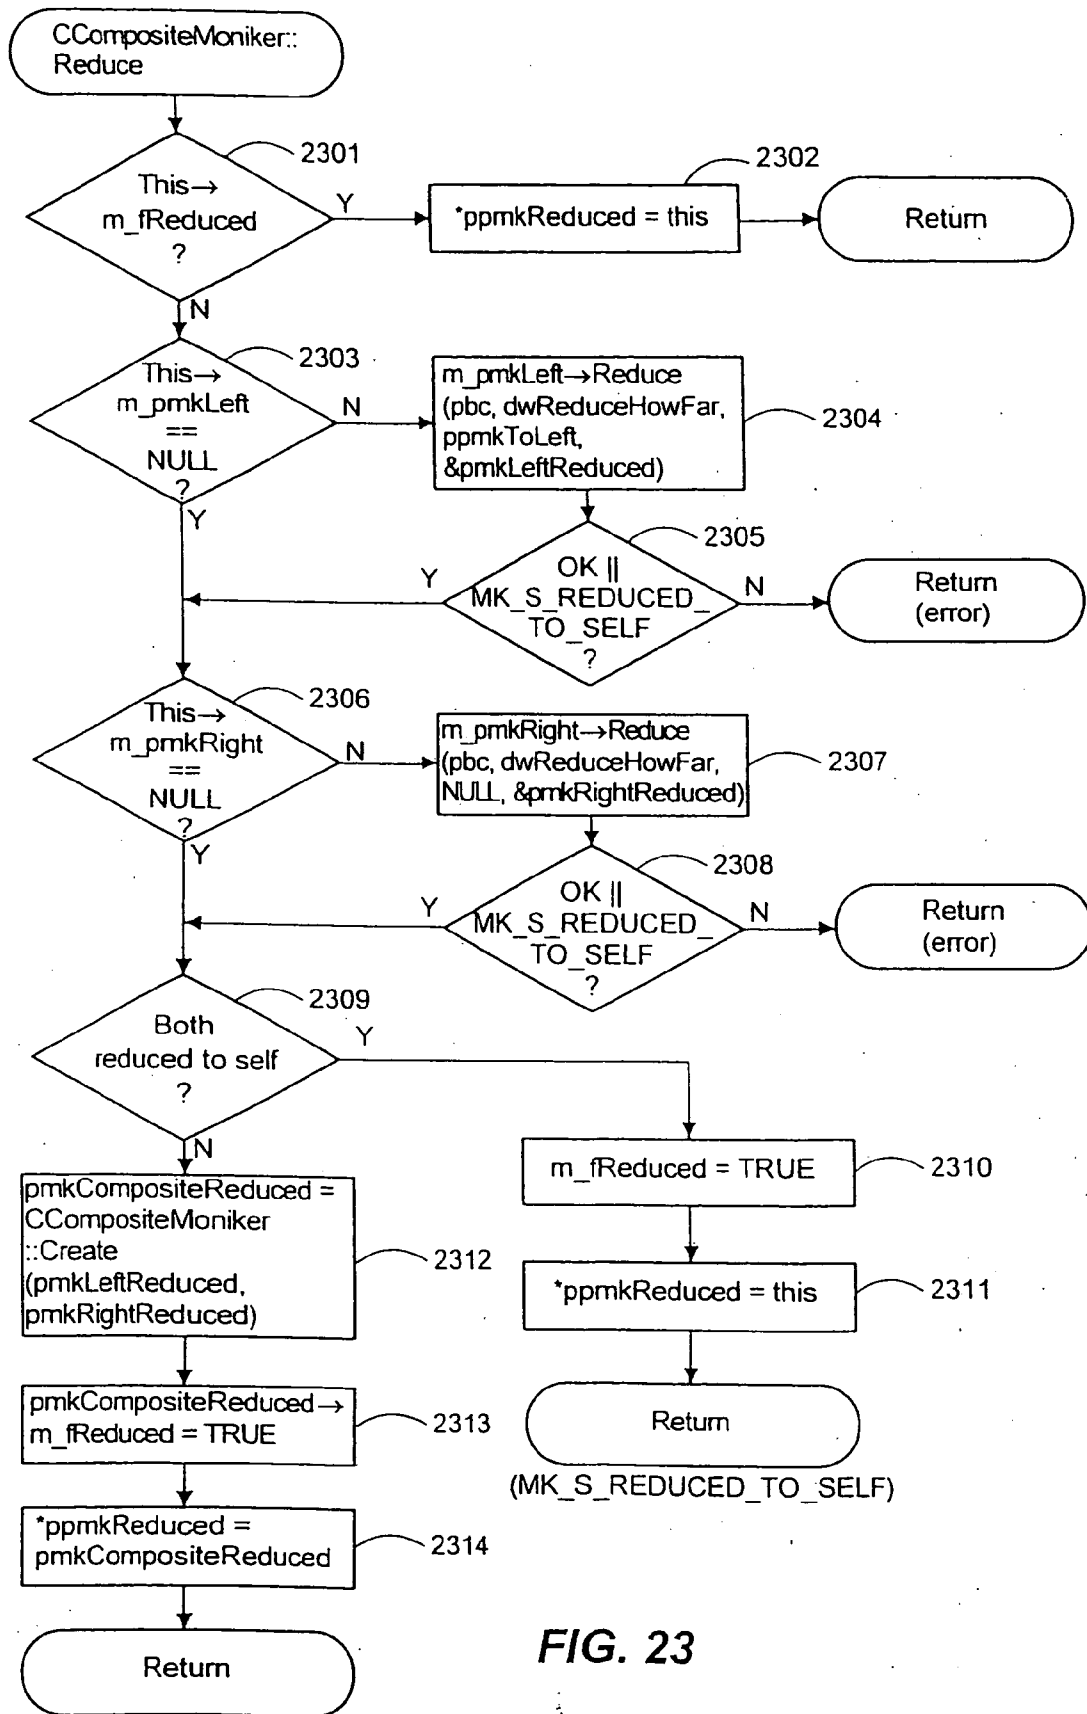


FIG. 23

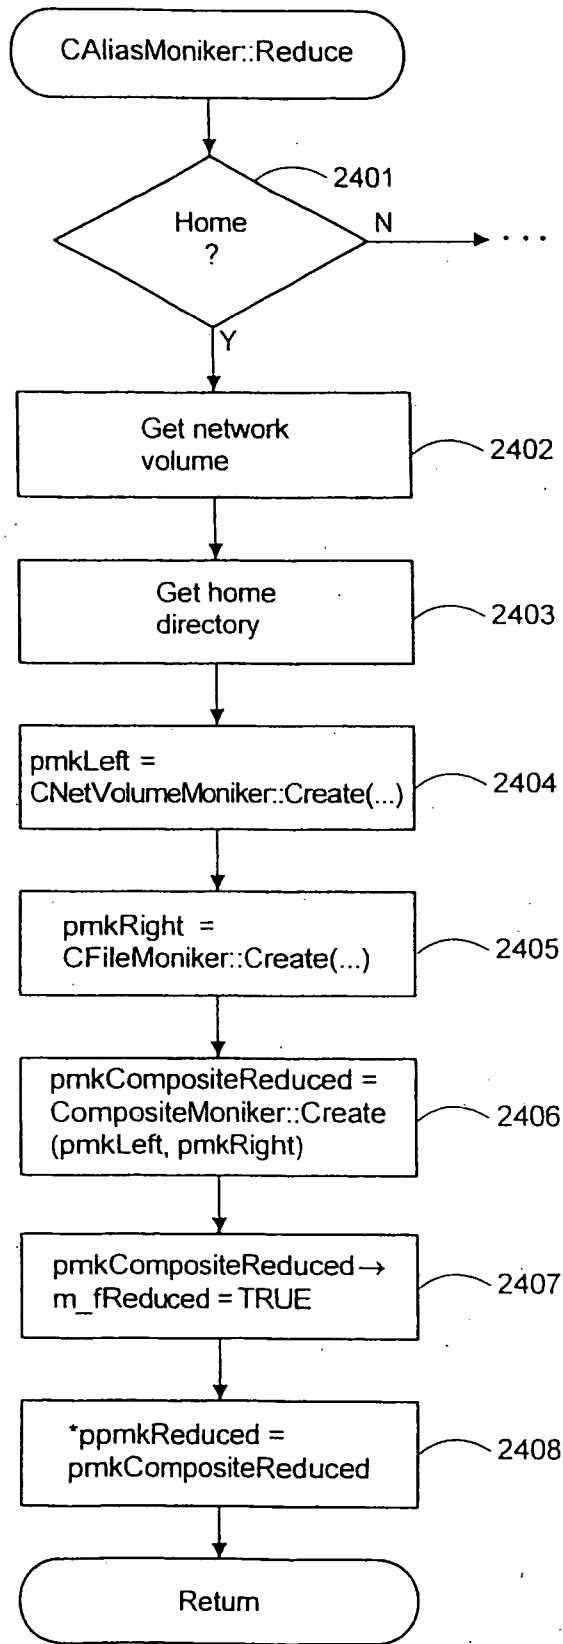


FIG. 24

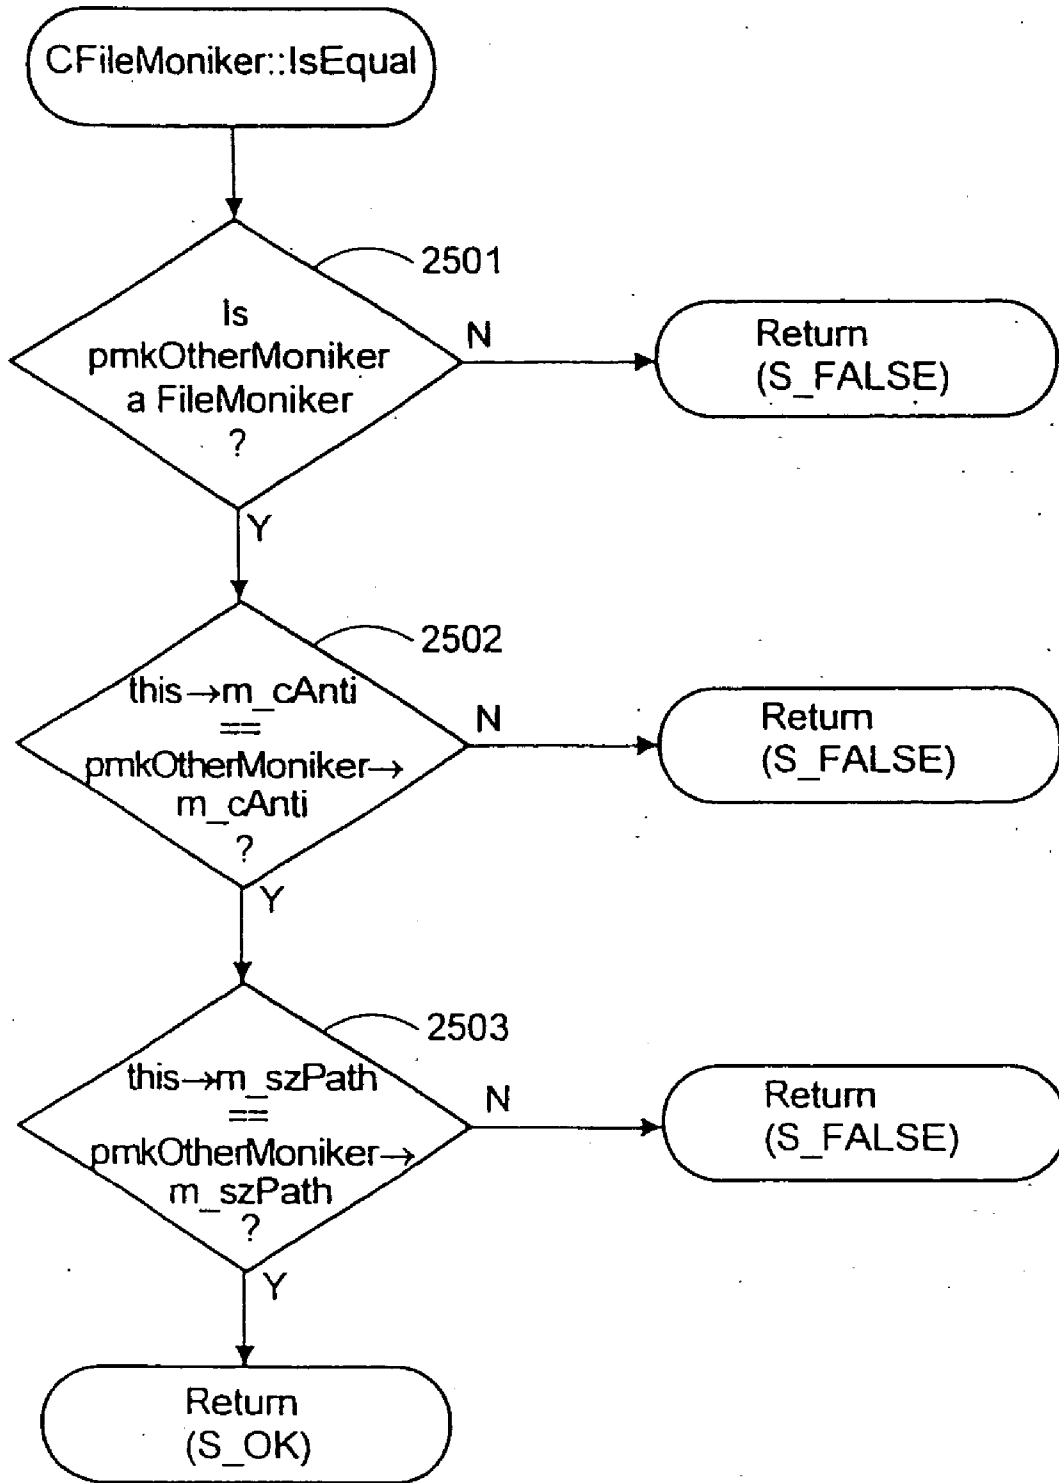


FIG. 25

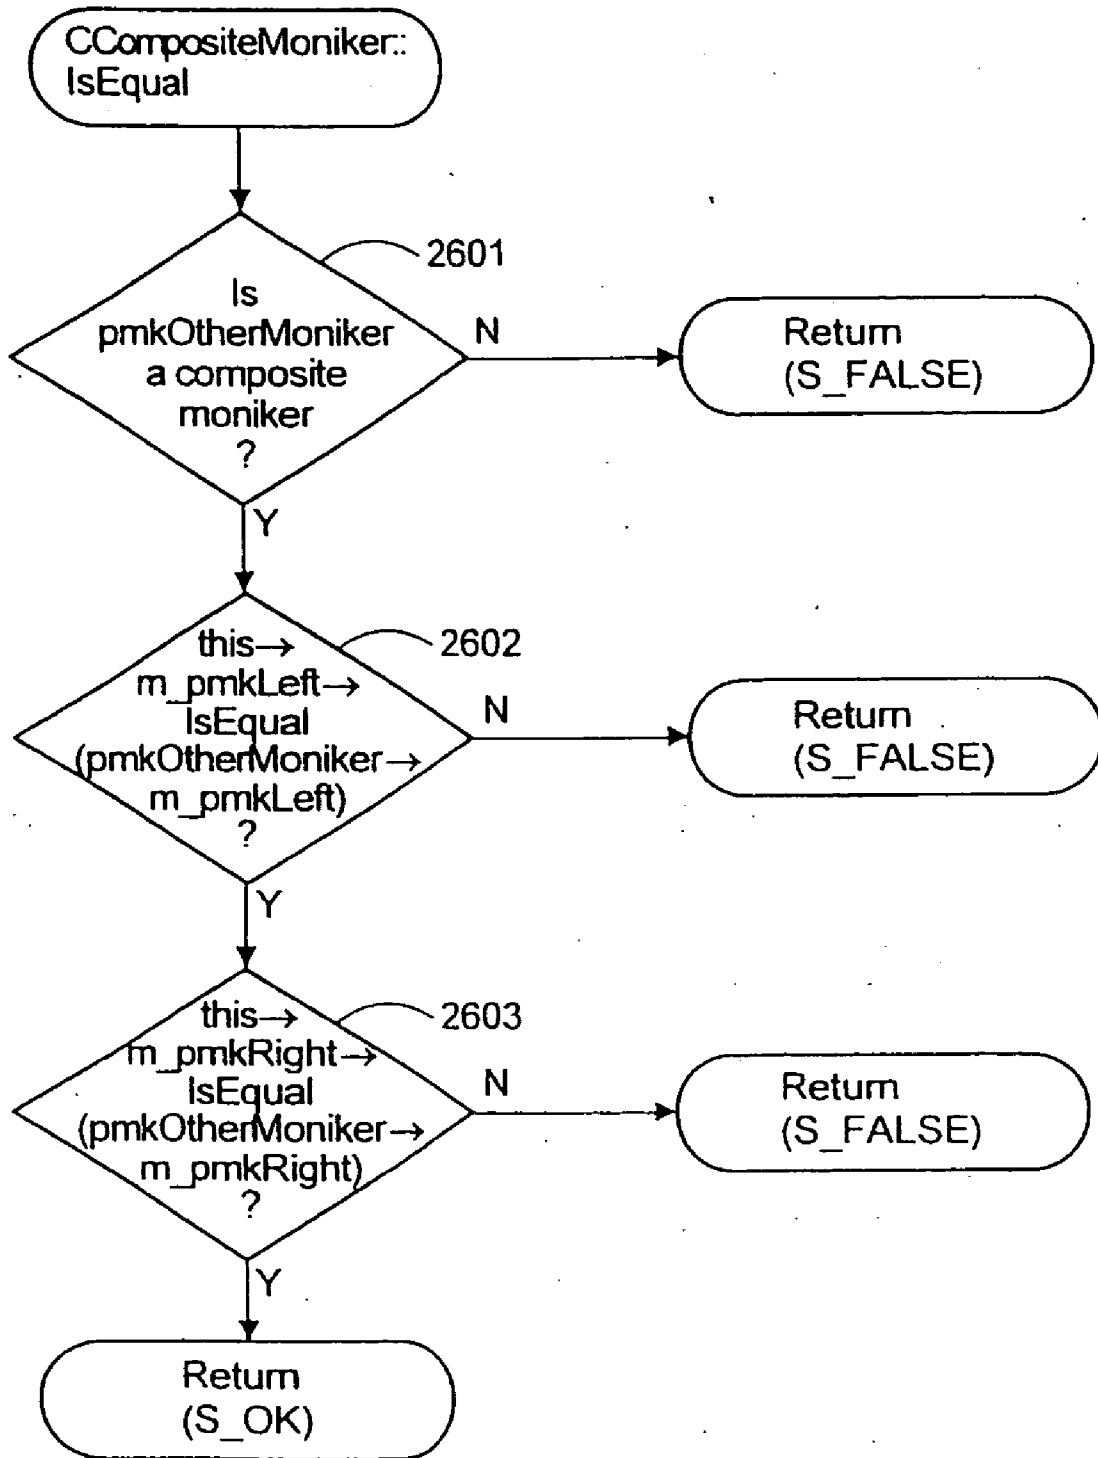


FIG. 26

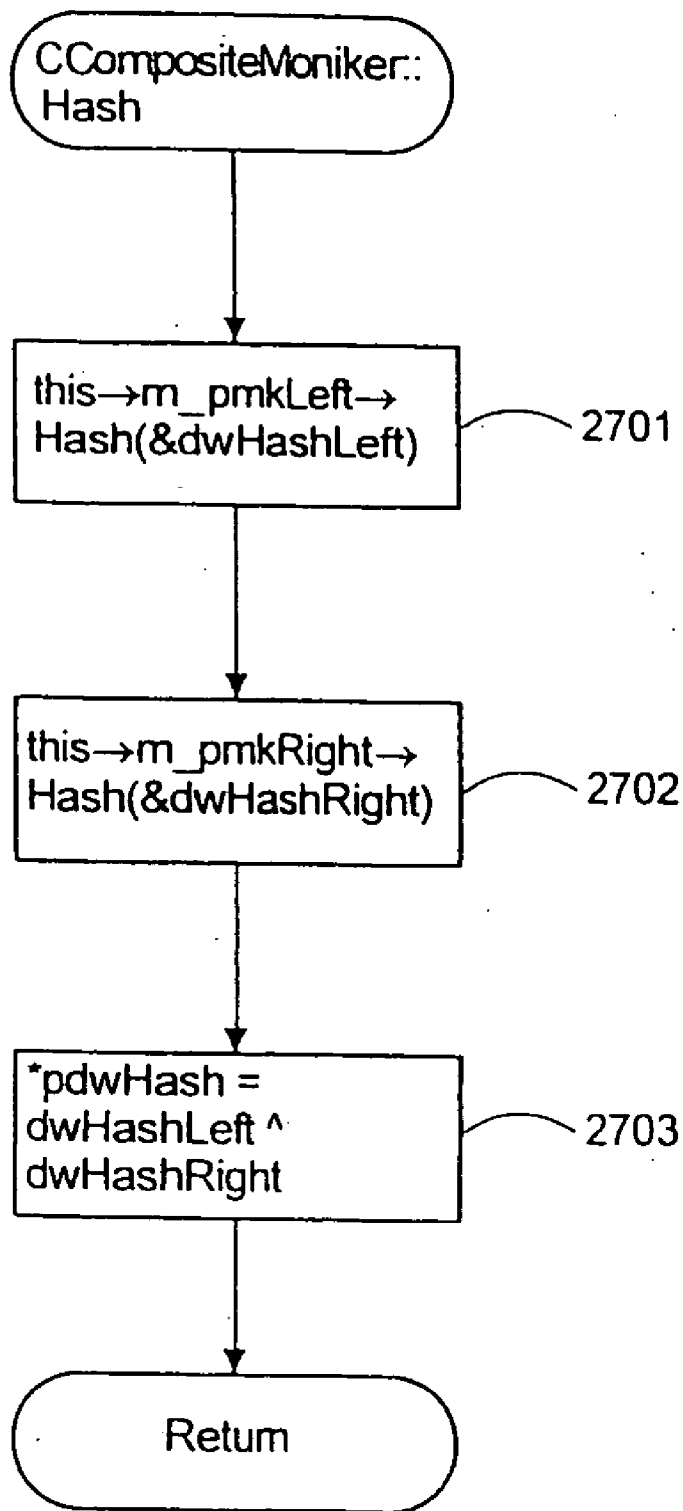


FIG. 27




FIG. 28A




FIG. 28B




FIG. 28C

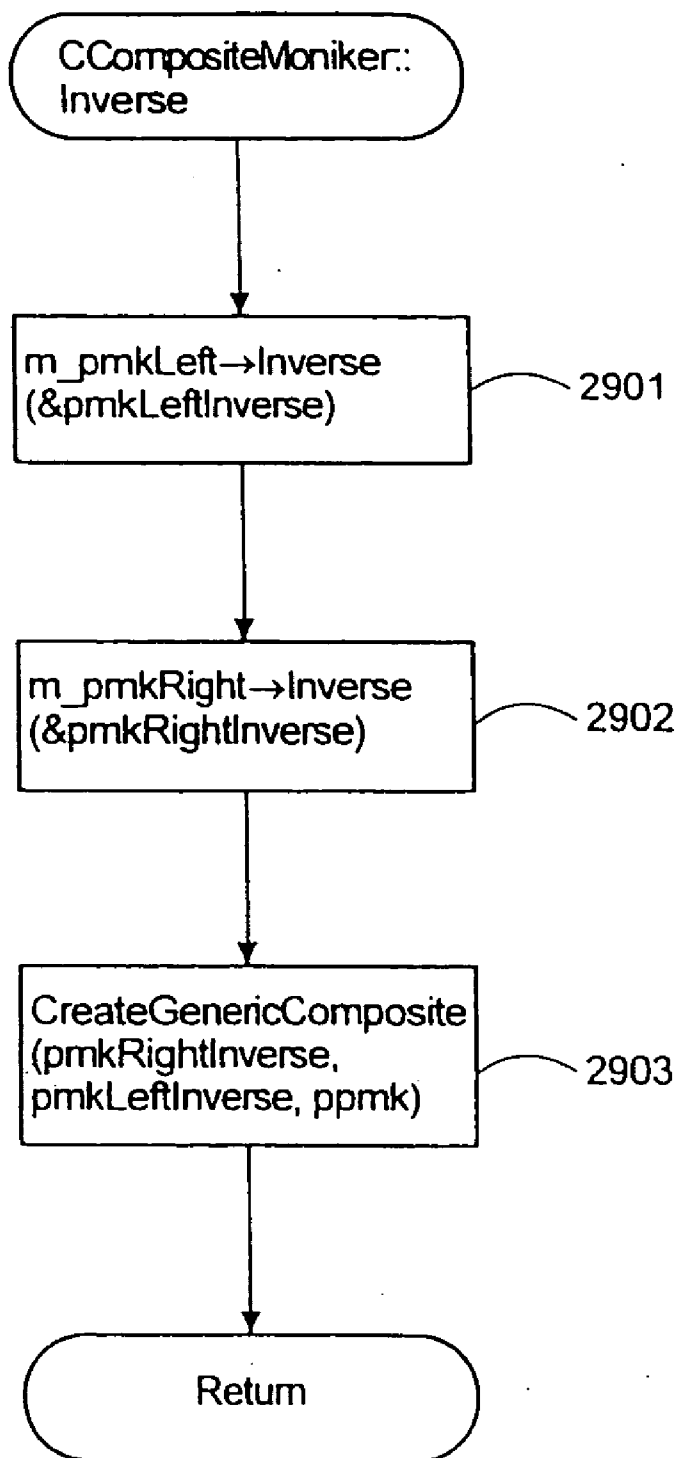


FIG. 29




FIG. 30

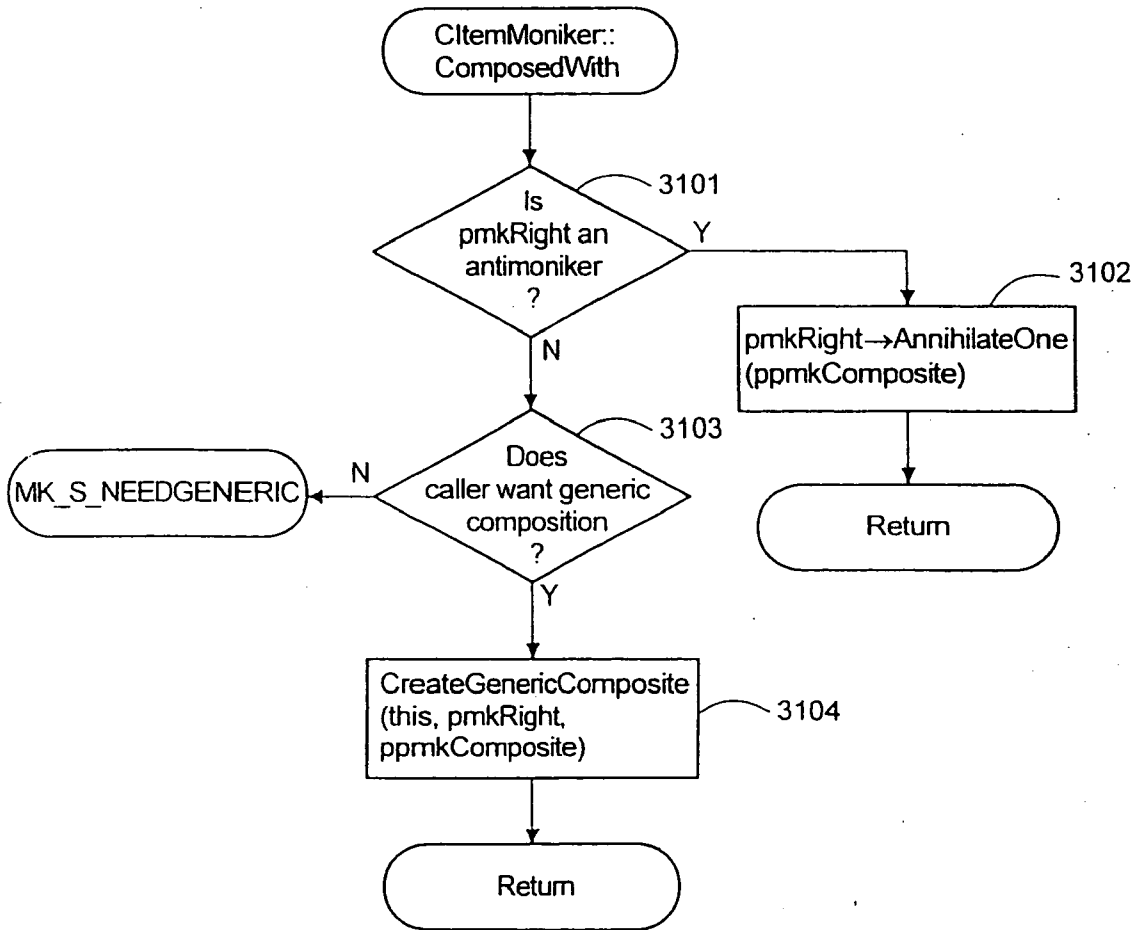


FIG. 31

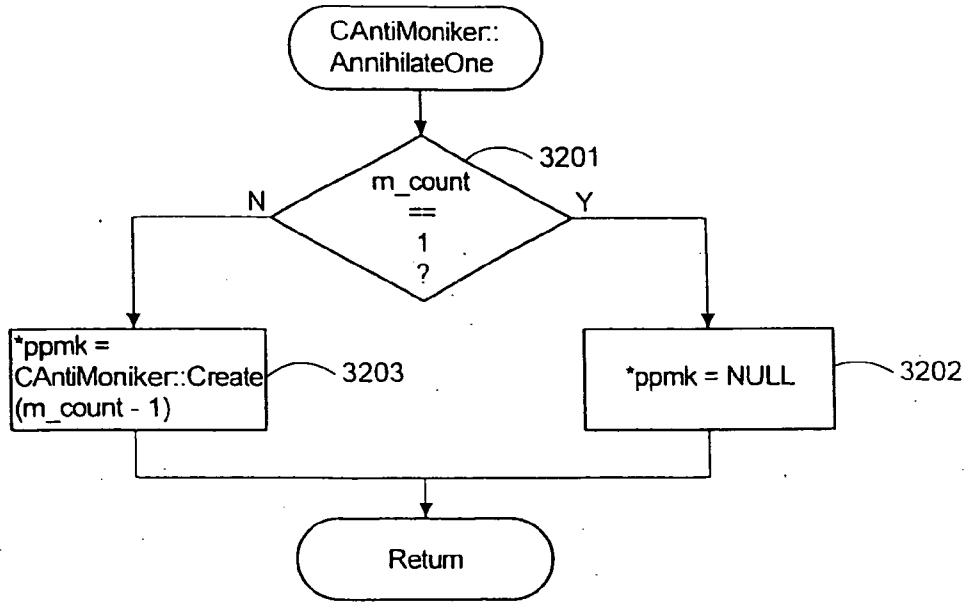


FIG. 32

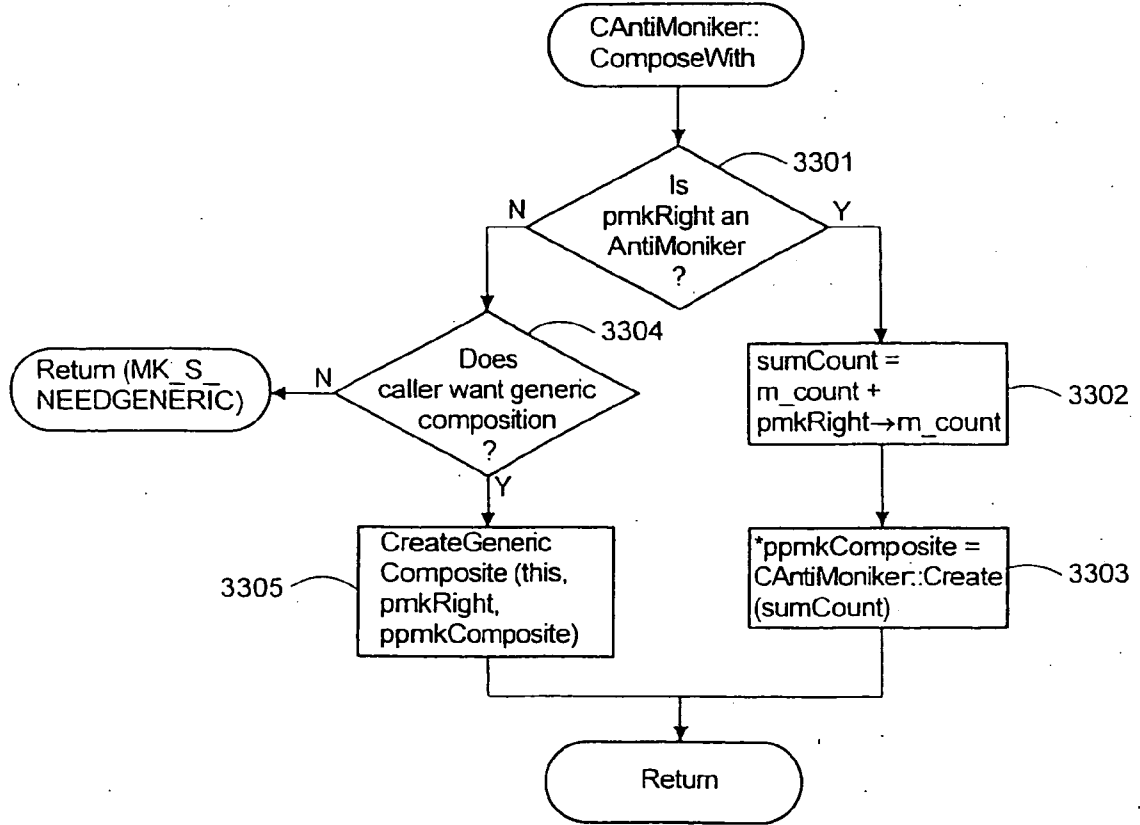


FIG. 33




FIG. 34

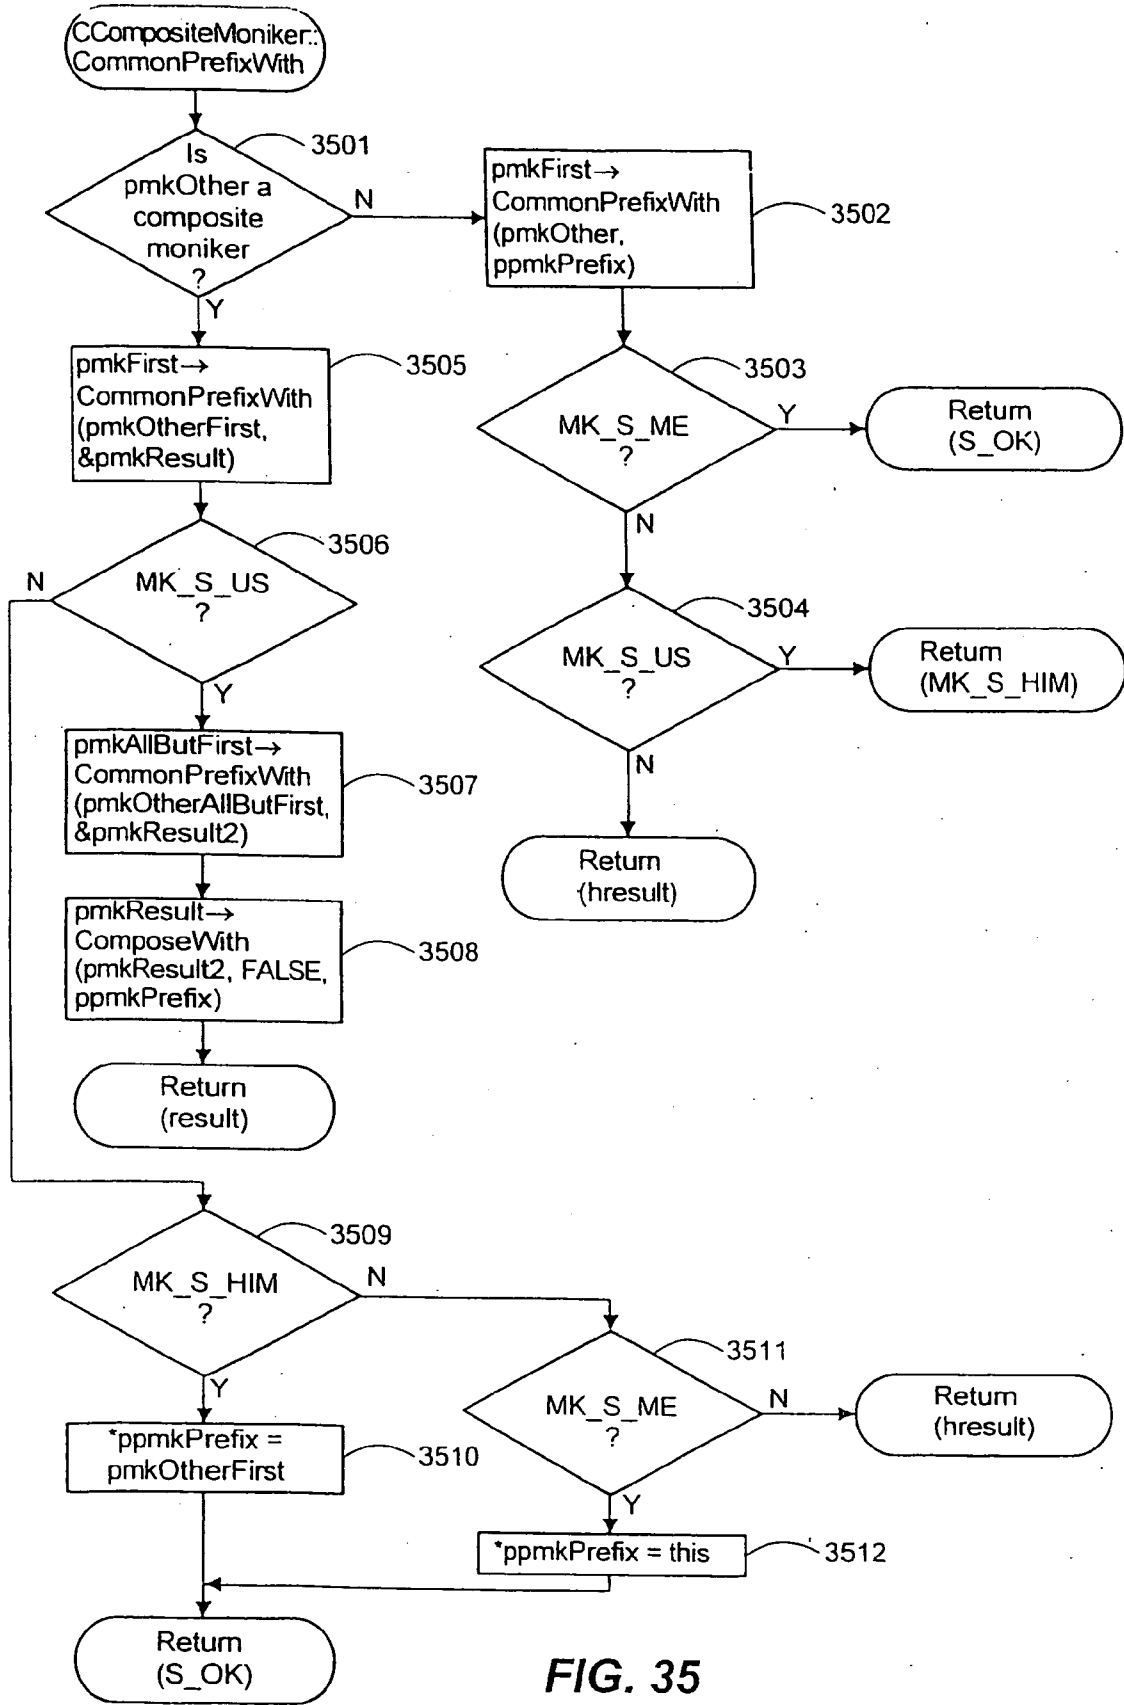


FIG. 35

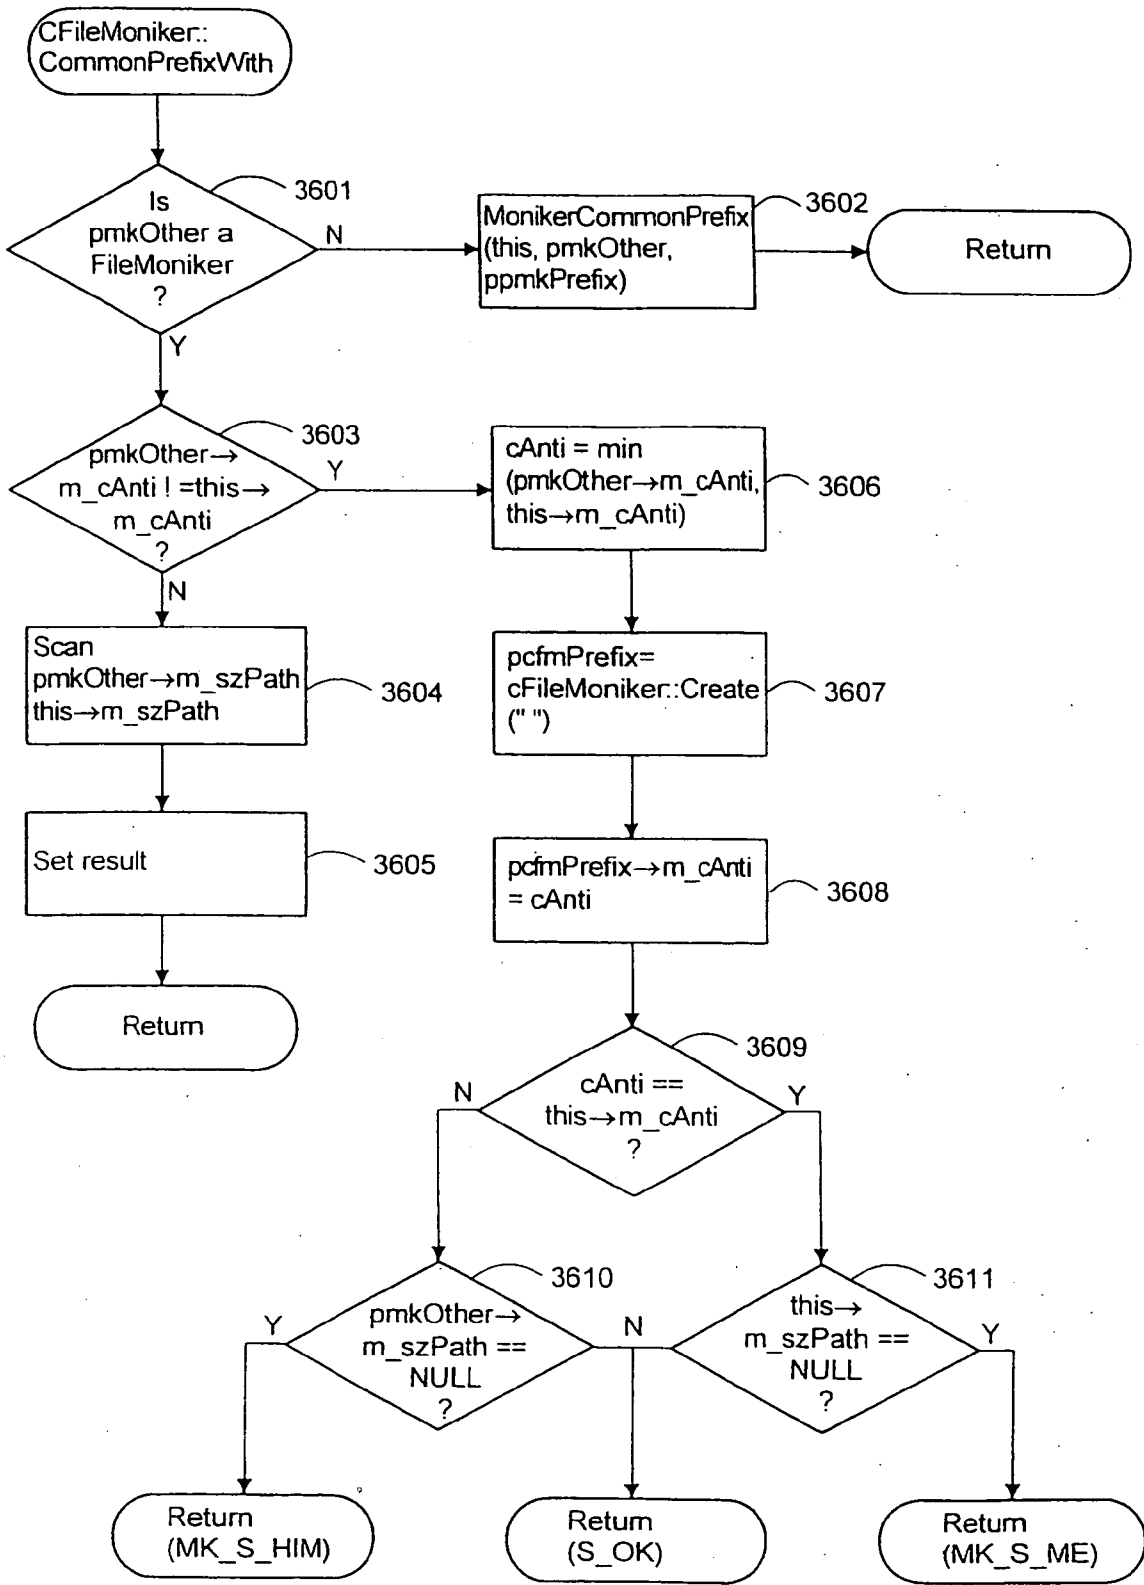


FIG. 36

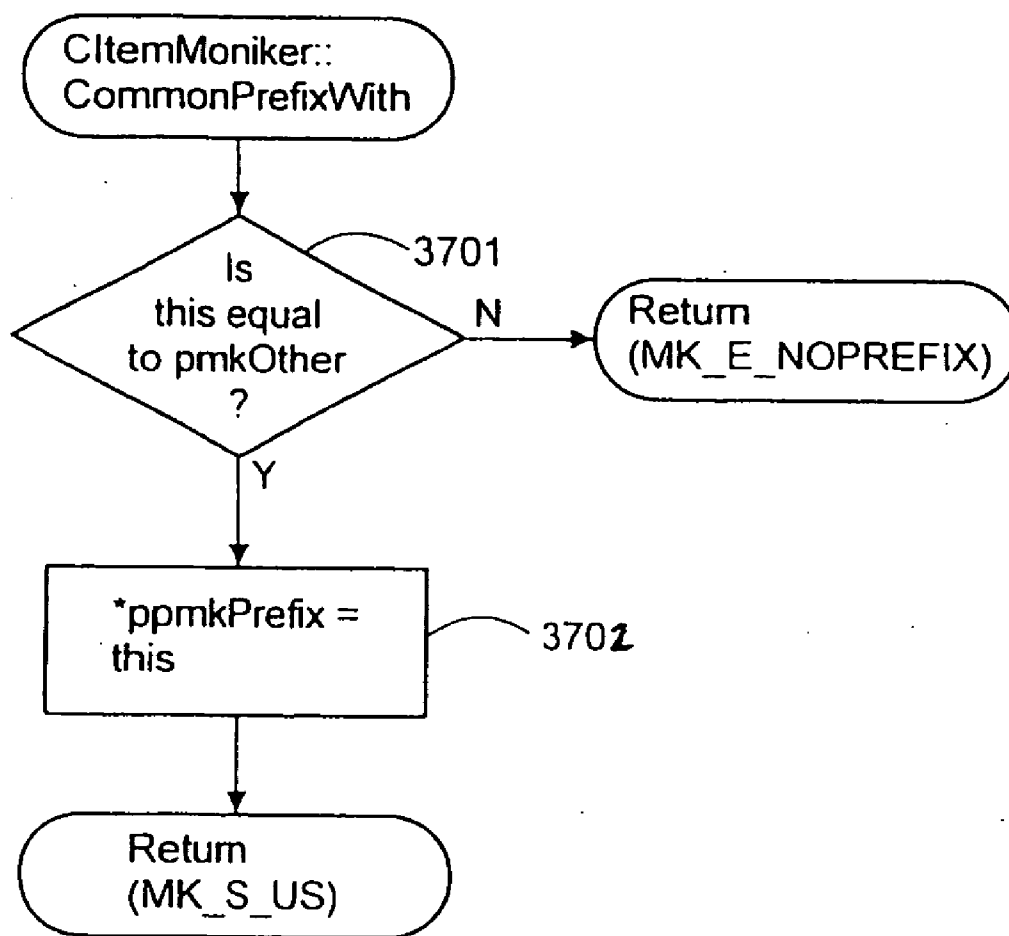


FIG. 37

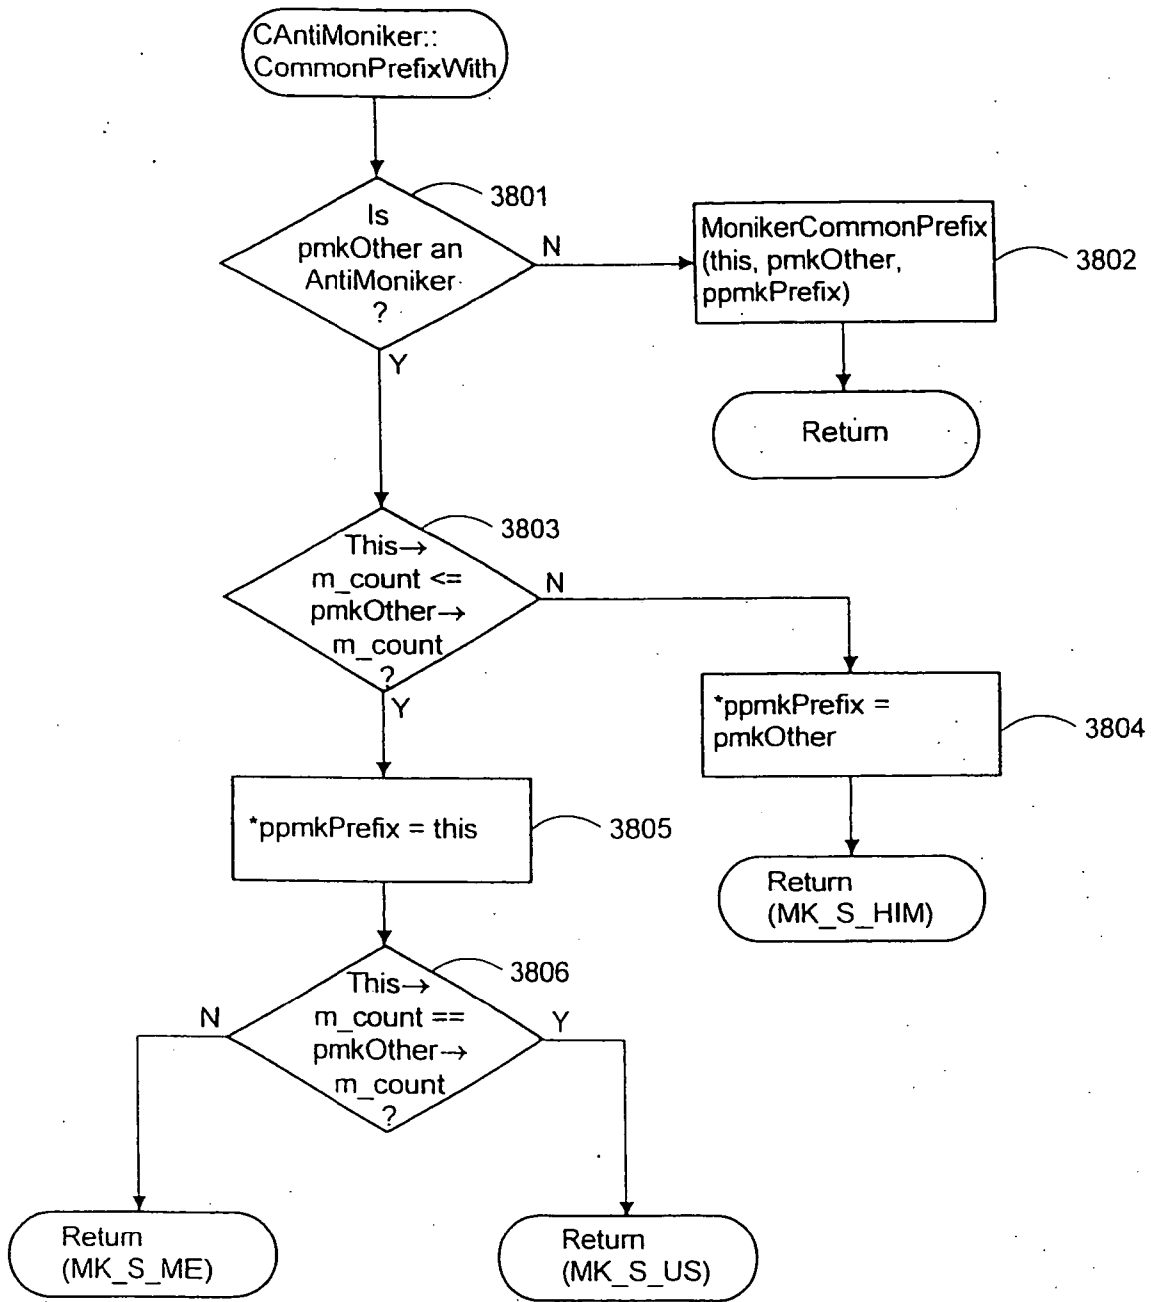


FIG. 38

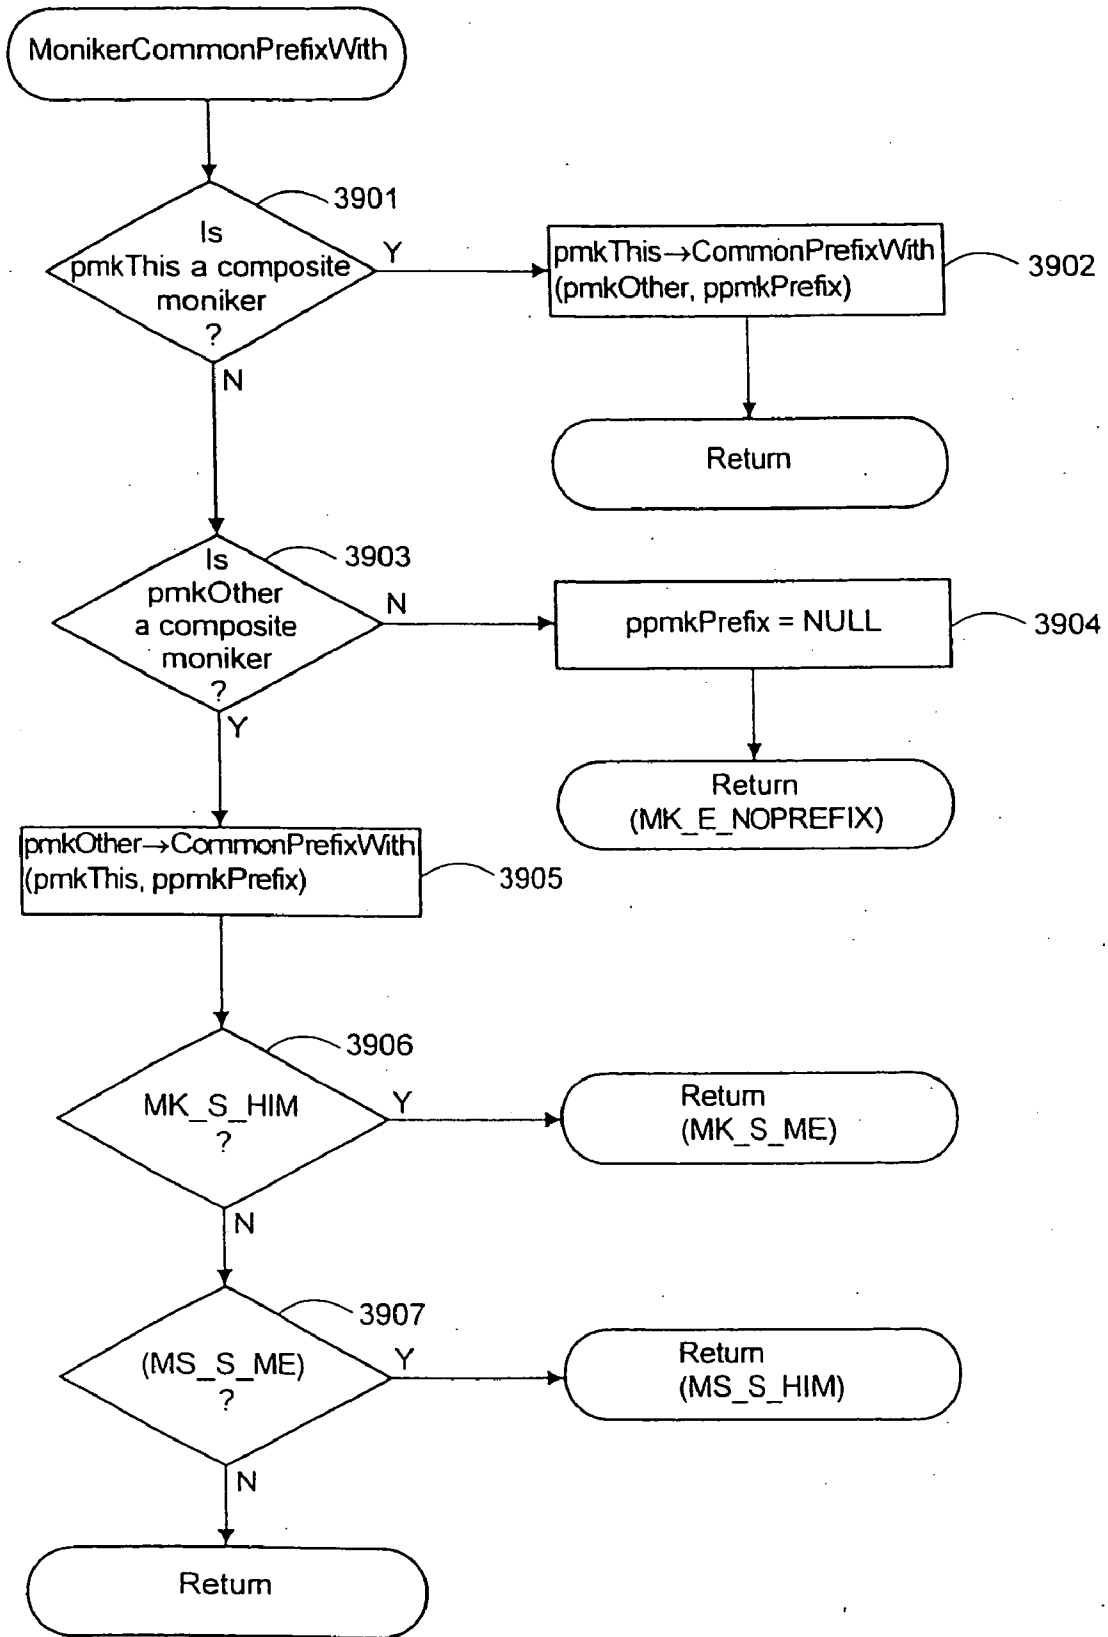


FIG. 39




FIG. 40





FIG. 41A



FileMoniker(c:\reports\monthly\June.doc) ◦ ((AntiMoniker(Anti = 1) ◦ FileMoniker(data\June.XLS)) ◦ ItemMoniker(R4C7:R8C12)) =
 FileMoniker(c:\reports\monthly\data\June.XLS) ◦ ItemMoniker(R4C7:R8C12)

FIG. 41B

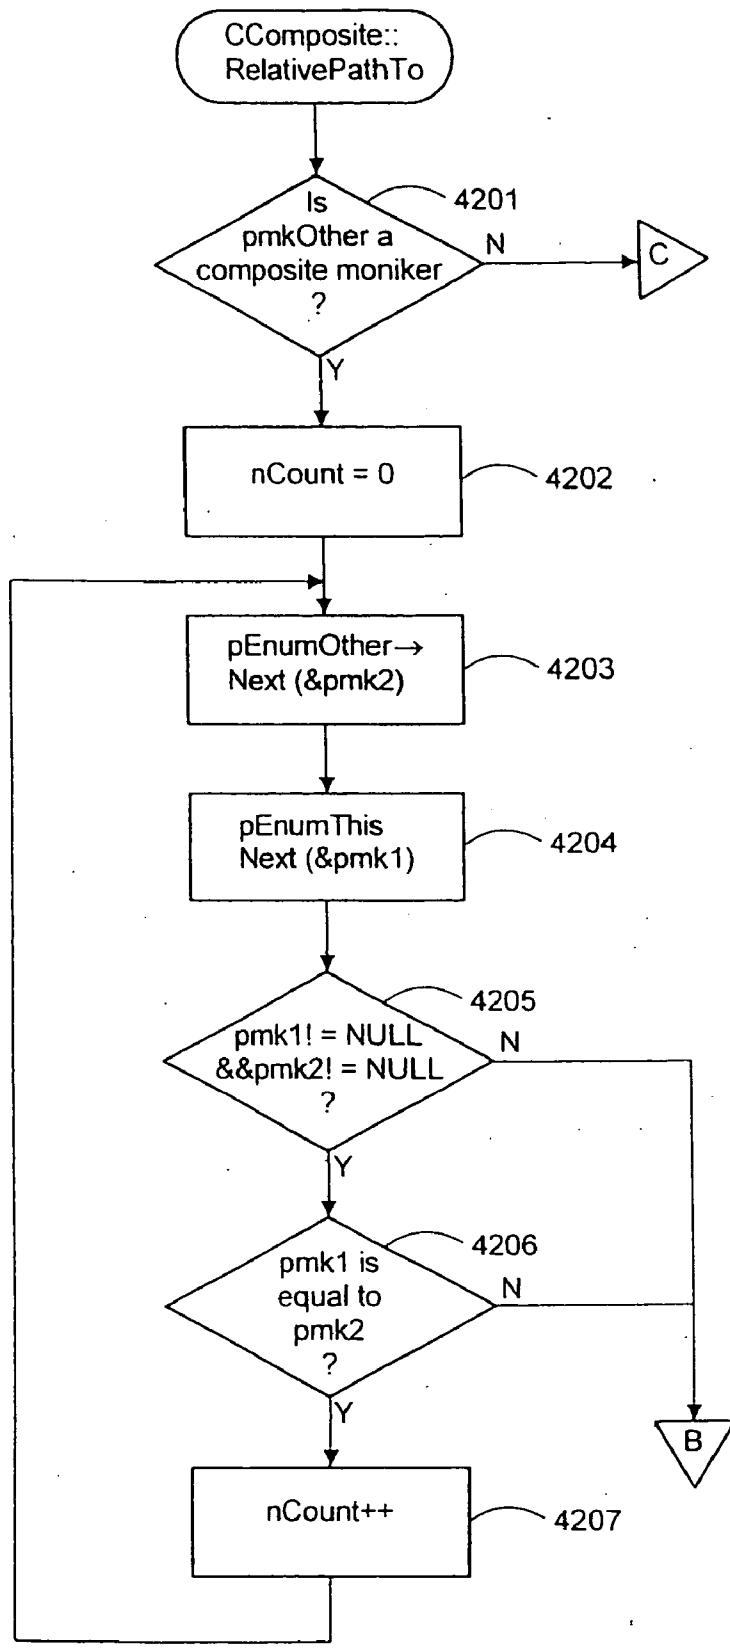


FIG. 42A

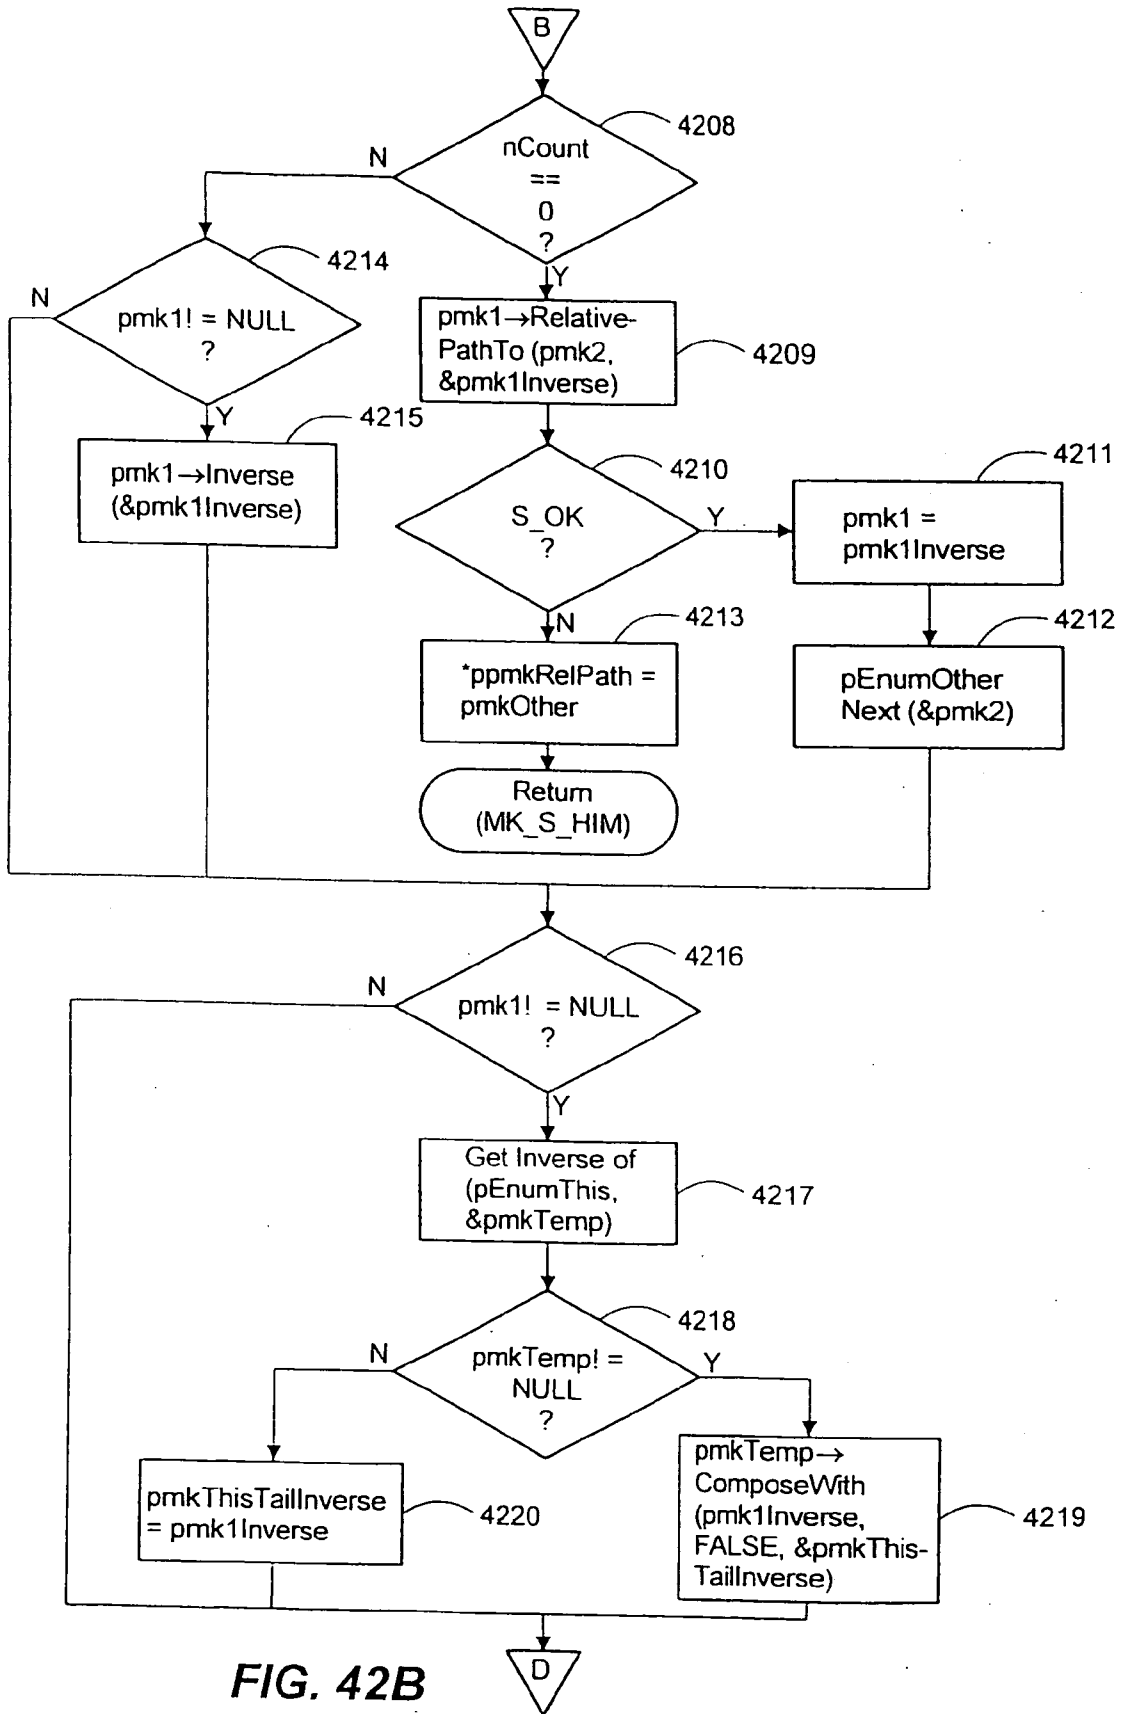


FIG. 42B

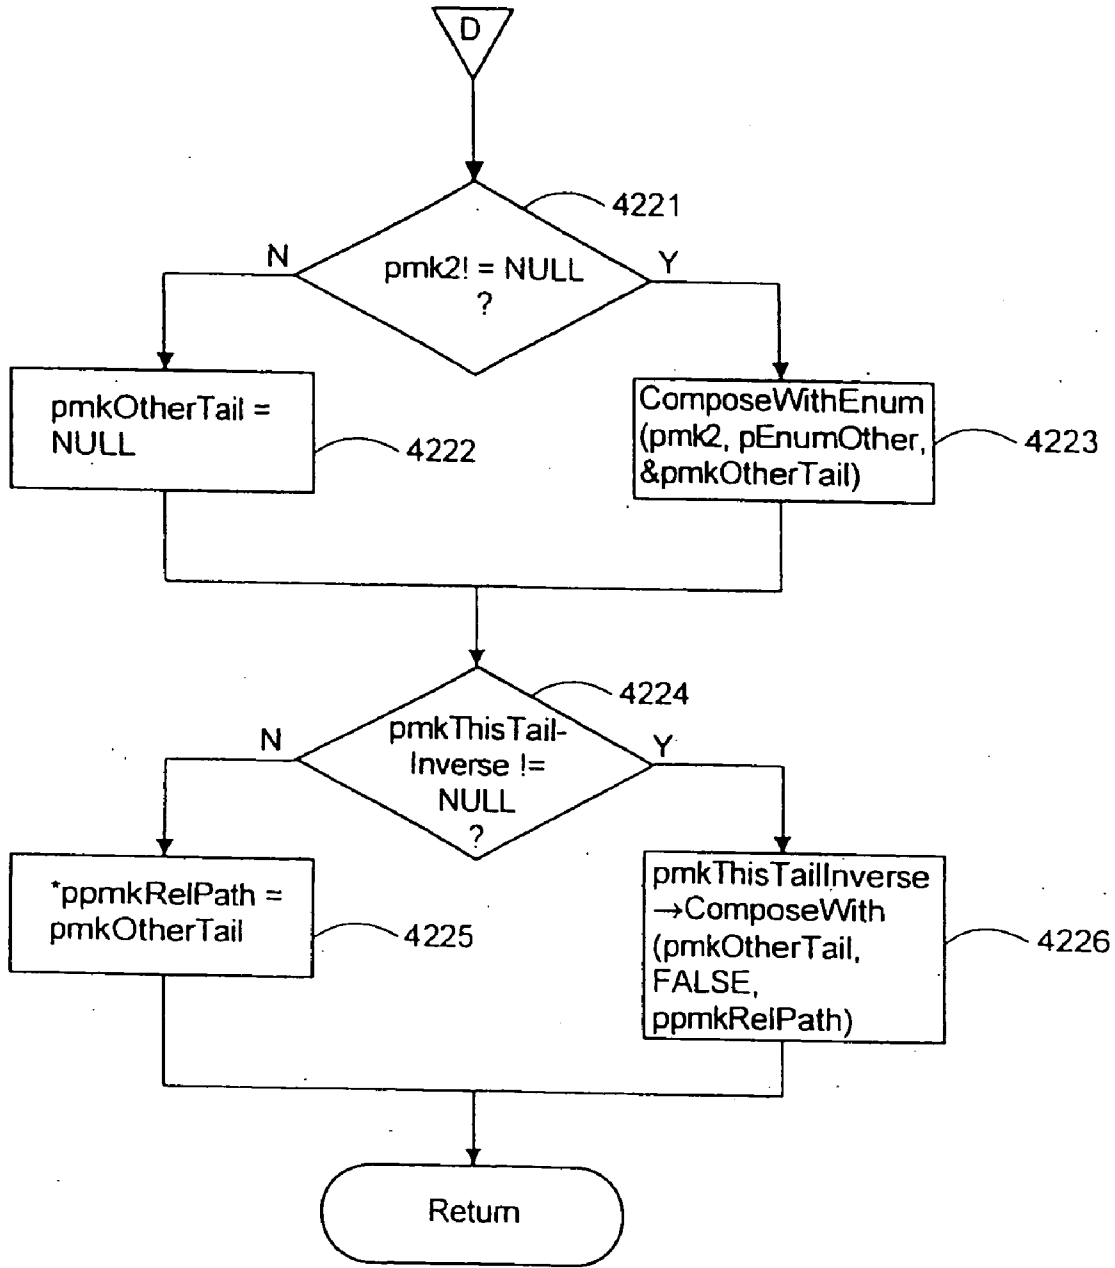


FIG. 42C

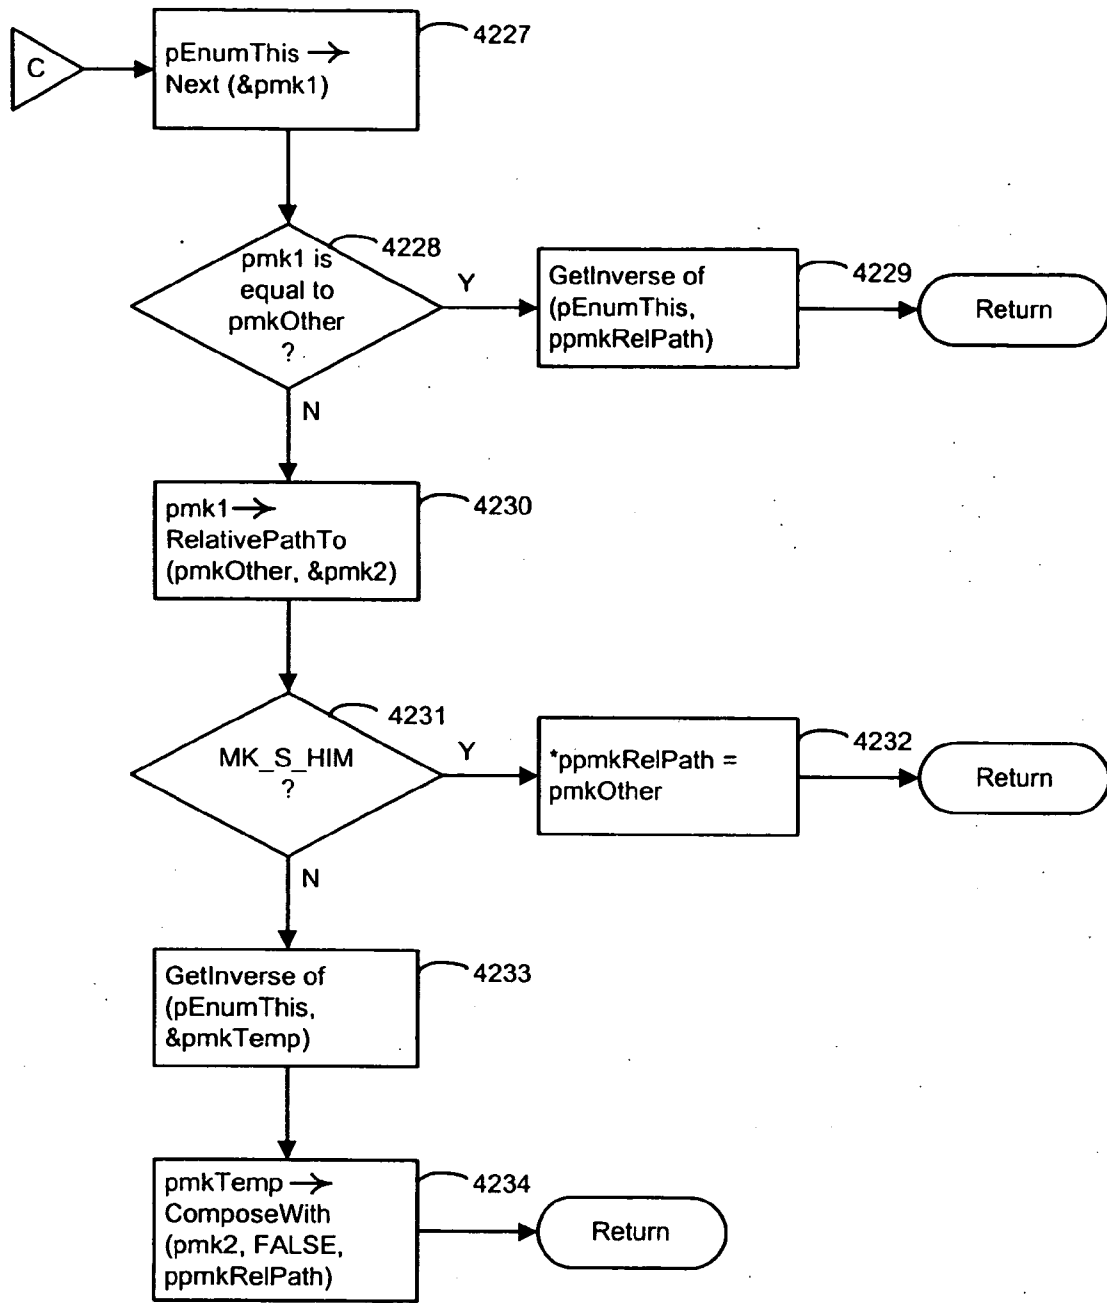


FIG. 42D

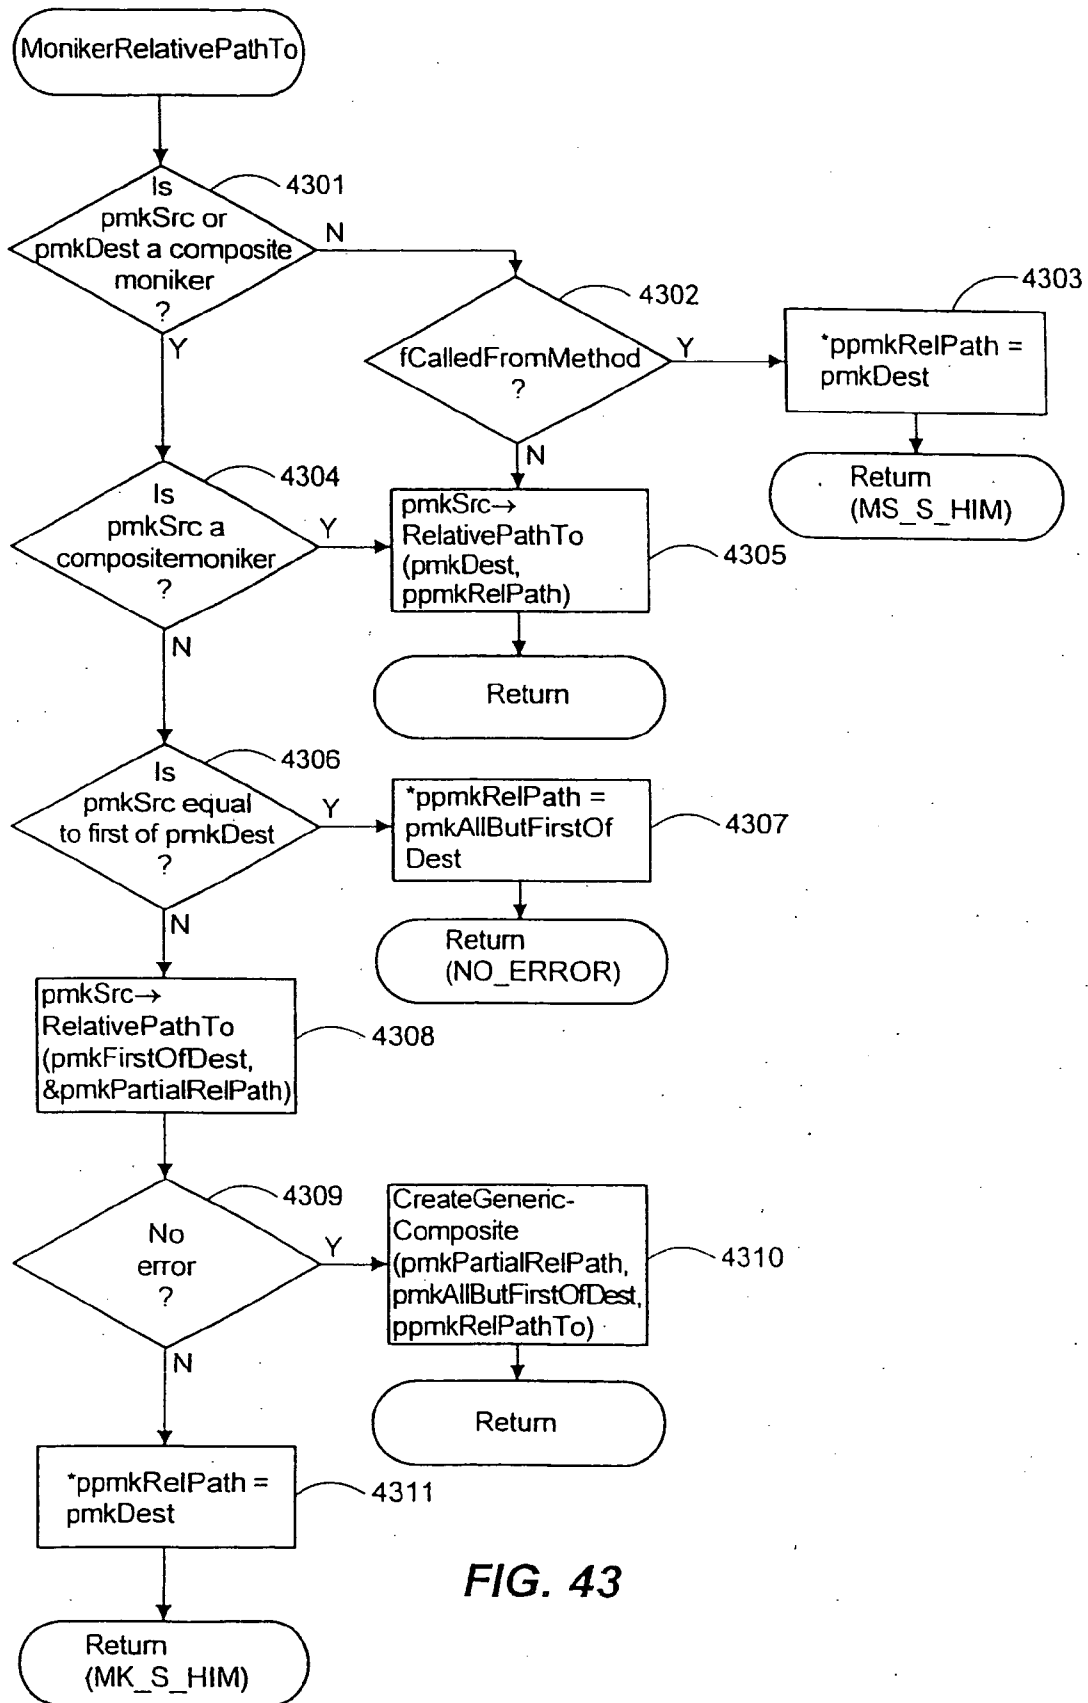


FIG. 43

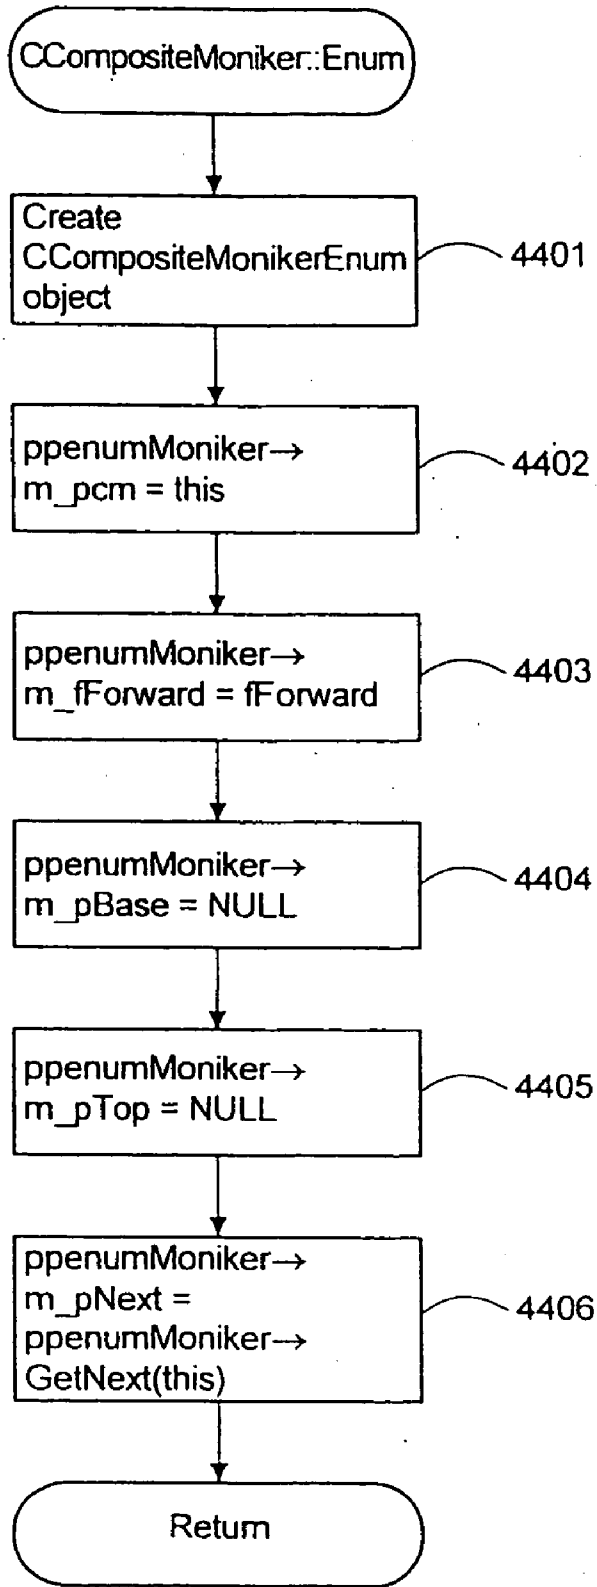


FIG. 44

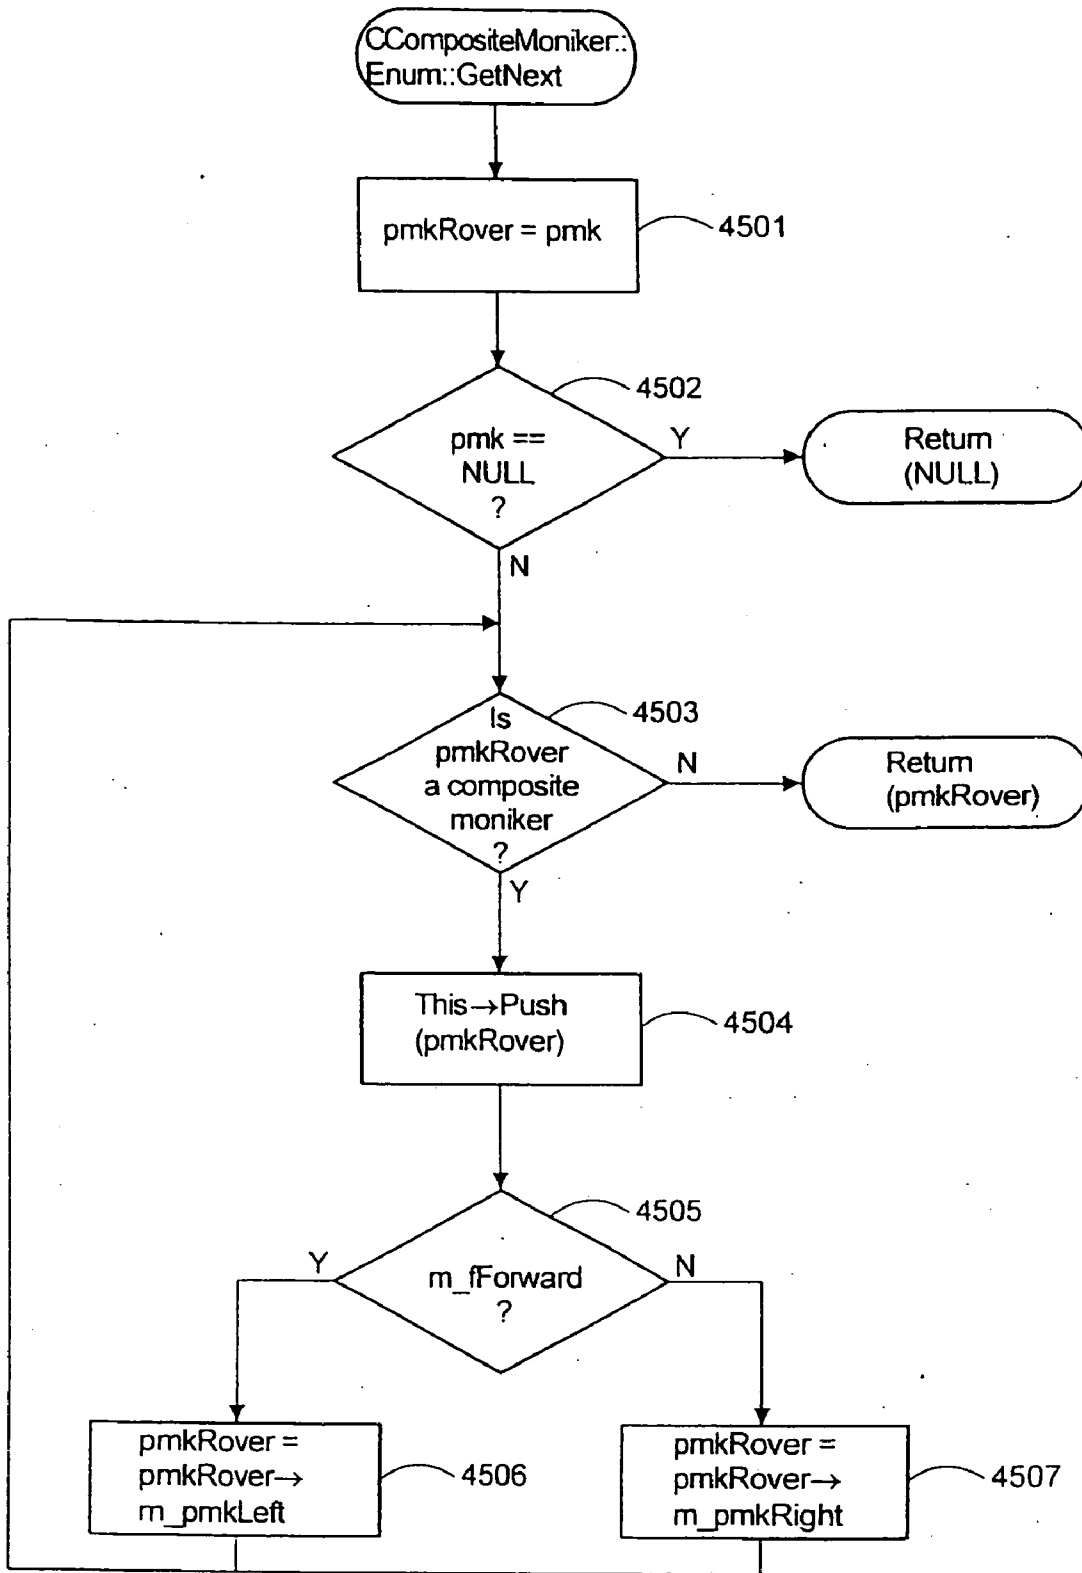


FIG. 45

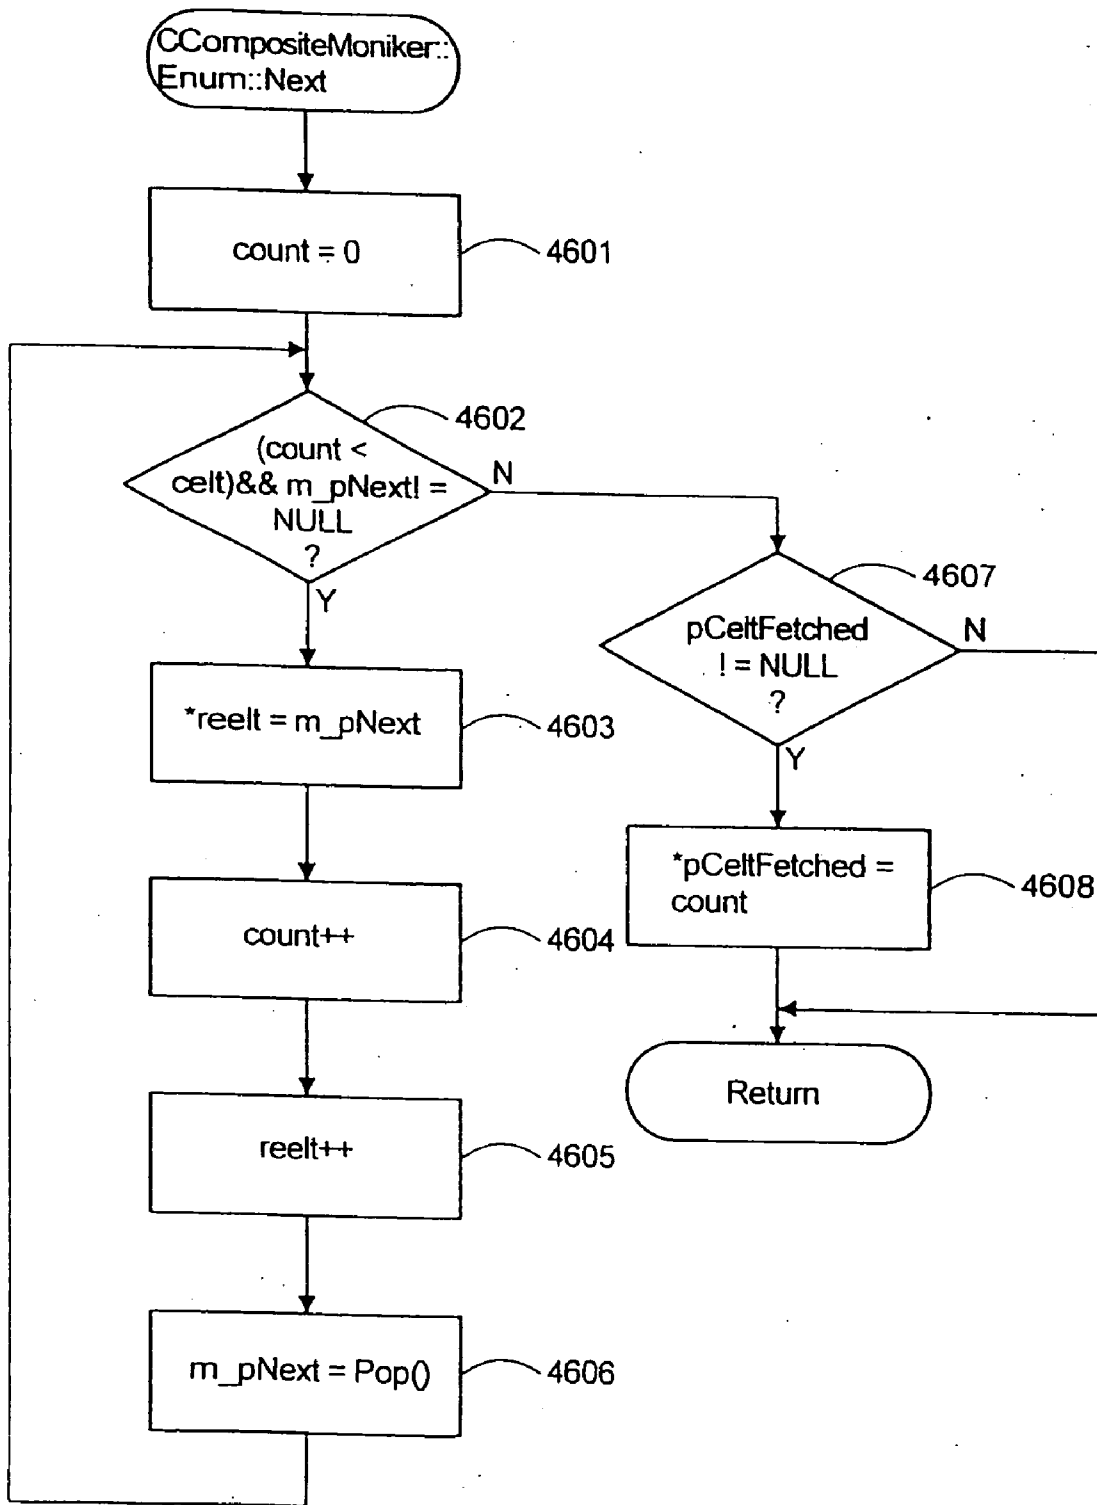


FIG. 46




FIG. 47

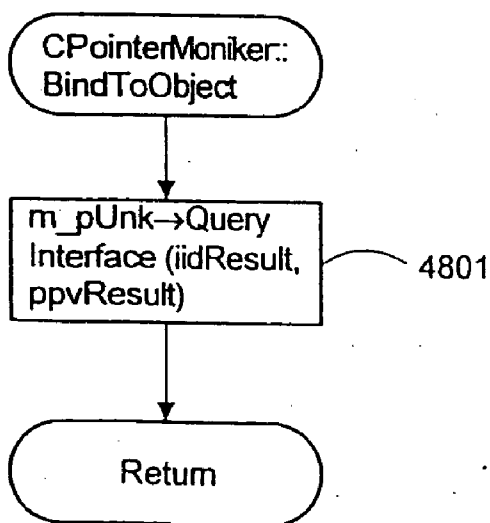


FIG. 48




FIG. 49

METHOD AND SYSTEM FOR NAMING AND BINDING OBJECTS

[0001] This application is a continuation-in-part of U.S. Ser. No. 07/909,983, filed Jul. 6, 1992, now pending.

TECHNICAL FIELD

[0002] This invention relates generally to a computer method and system for referencing objects and, more specifically, to a method and system for naming objects and binding to objects

BACKGROUND OF THE INVENTION

[0003] Current document processing computer systems allow a user to prepare compound documents. A compound document is a document that contains information in various formats. For example, a compound document may contain data in text format, chart format, numerical format, etc. FIG. 1 is an example of a compound document. In this example, the compound document 101 is generated as a report for a certain manufacturing project. The compound document 101 contains scheduling data 102, which is presented in chart format; budgeting data 103, which is presented in spreadsheet format; and explanatory data 104, which is presented in text format. In typical prior systems, a user generates the scheduling data 102 using a project management computer program and the budgeting data 103 using a spreadsheet computer program. After this data has been generated, the user creates the compound document 101, enters the explanatory data 104, and incorporates the scheduling data 102 and budgeting data 103 using a word processing computer program.

[0004] FIG. 2 shows a method for incorporating the scheduling data, budgeting data, and explanatory data into the compound document. A user generates scheduling data using the project management program 201 and then stores the data in the clipboard 203. The user also generates budgeting data using the spreadsheet program 204 and then stores the data in the clipboard 203. The clipboard 203 is an area of storage (disk or memory) that is typically accessible by any program and is used to transfer data between programs. The project management program 201 and the spreadsheet program 204 typically store the data into the clipboard in a presentation format. A presentation format is a format in which the data is easily displayed on an output device. For example, the presentation format may be a bitmap that can be displayed with a standard bitmap block transfer operation (BitBlt). The storing of data into a clipboard is referred to as "copying" to the clipboard.

[0005] After data has been copied to the clipboard 203, the user starts up the word processing program 206 to create the compound document 101. The user enters the explanatory data 104 and specifies the locations in the compound document 101 to which the scheduling data and budgeting data that are in the clipboard 203 are to be copied. The copying of data from a clipboard to a document is referred to as "pasting" from the clipboard. The word processing program 206 then copies the scheduling data 102 and the budgeting data 103 from the clipboard 203 into the compound document 101 at the specified locations. Data that is copied from the clipboard into a compound document is referred to as "embedded" data. The word processing program 206 treats the embedded data as simple bitmaps that it displays with a

BitBlt operation when rendering the compound document 101 on an output device. In some prior systems, a clipboard may only be able to store data for one copy command at a time. In such a system, the scheduling data can be copied to the clipboard and then pasted into the compound document. Then, the budgeting data can be copied to the clipboard and then pasted into the compound document.

[0006] Since word processors typically process only text data, users of the word processing program can move or delete embedded data, but cannot modify embedded data, unless the data is in text format. Thus, if a user wants to modify, for example, the budgeting data 103 that is in the compound document 101, the user must start up the spreadsheet program 204, load in the budgeting data 103 from a file, make the modifications, copy the modifications to the clipboard 203, start up the word processing program 206, load in the compound document 101, and paste the modified clipboard data into the compound document 101.

[0007] Some prior systems store links to the data to be included in the compound document rather than actually embedding the data. When a word processing program pastes the data from a clipboard into a compound document, a link is stored in the compound document. The link points to the data (typically residing in a file) to be included. These prior systems typically provide links to data in a format that the word processing program recognizes or treats as presentation format. For example, when the word processing program 206 is directed by a user to paste the scheduling data and budgeting data into the compound document by linking, rather than embedding, the names of files in which the scheduling data and budgeting data reside in presentation format are inserted into the document. Several compound documents can contain links to the same data to allow one copy of the data to be shared by several compound documents.

[0008] A link is conceptually a path name to the data. Some prior systems store two-level links. A two-level link identifies both a file and an area within the file. For example, the two-level link "\BUDGET.XLS\R2C2:R7C4" identifies a spreadsheet file "BUDGET.XLS" and the range of cells "R2C2:R7C4." The use of two-level links limits the source of the links to data that is nested one level within a file. If a file contains multiple spreadsheets, then a two-level link could identify the file and a spreadsheet, but could not identify a range within the spreadsheet. It would be desirable to have a method and system of supporting links to an arbitrary level.

[0009] Since the present invention is described below using object-oriented programming, an overview of well-known object-oriented programming techniques is provided. Two common characteristics of object-oriented programming languages are support for data encapsulation and data type inheritance. Data encapsulation refers to the binding of functions and data. Inheritance refers to the ability to declare a data type in terms of other data types.

[0010] In the C++ language, object-oriented techniques are supported through the use of classes. A class is a user-defined type. A class declaration describes the data members and function members of the class. For example, the following declaration defines data members and a function member of a class named CIRCLE.

```

class CIRCLE
{ public:
  int x, y;
  int radius;
  void draw( );
};

```

Variables x and y specify the center location of a circle and variable radius specifies the radius of the circle. These variables are referred to as data members of the class CIRCLE. The function draw is a user-defined function that draws the circle of the specified radius at the specified location. The function draw is referred to as a function member of class CIRCLE. The data members and function members of a class are bound together in that the function operates on an instance of the class. An instance of a class is also called an object of the class.

[0011] In the syntax of C++, the following statement declares the objects a and b to be of type class CIRCLE.

[0012] CIRCLE a, b;

This declaration causes the allocation of memory for the objects a and b. The following statements assign data to the data members of objects a and b.

[0013] a.x=2;

[0014] a.y=2;

[0015] a.radius=1;

[0016] b.x=4;

[0017] b.y=5;

[0018] b.radius=2;

The following statements are used to draw the circles defined by objects a and b.

[0019] a.draw();

[0020] b.draw();

[0021] A derived class is a class that inherits the characteristics—data members and function members—of its base classes. For example, the following derived class CIRCLE_FILL inherits the characteristics of the base class CIRCLE.

```

class CIRCLE_FILL : CIRCLE
{ public:
  int pattern;
  void fill( );
};

```

This declaration specifies that class CIRCLE_FILL includes all the data and function members that are in class CIRCLE in addition to those data and function members introduced in the declaration of class CIRCLE_FILL, that is, data member pattern and function member fill. In this example, class CIRCLE_FILL has data members x, y, radius, and pattern and function members draw and fill. Class CIRCLE_FILL is said to “inherit” the characteristics of class CIRCLE. A class that inherits the characteristics of another class is a derived

class (e.g., CIRCLE_FILL). A class that does not inherit the characteristics of another class is a primary (root) class (e.g., CIRCLE). A class whose characteristics are inherited by another class is a base class (e.g., CIRCLE is a base class of CIRCLE_FILL). A derived class may inherit the characteristics of several classes, that is, a derived class may have several base classes. This is referred to as multiple inheritance.

[0022] A derived class may specify that a base class is to be inherited virtually. Virtual inheritance of a base class means that only one instance of the virtual base class exists in the derived class. For example, the following is an example of a derived class with two nonvirtual base classes.

[0023] class CIRCLE_1 : CIRCLE { . . . };

[0024] class CIRCLE_2 : CIRCLE { . . . };

[0025] class PATTERN : CIRCLE_1, CIRCLE_2{ . . . };

In this declaration class PATTERN inherits class CIRCLE twice nonvirtually through classes CIRCLE_1 and CIRCLE_2. There are two instances of class CIRCLE in class PATTERN.

[0026] The following is an example of a derived class with two virtual base classes.

[0027] class CIRCLE_1 : virtual CIRCLE { . . . };

[0028] class CIRCLE_2 : virtual CIRCLE { . . . };

[0029] class PATTERN: CIRCLE_1, CIRCLE_2{ . . . };

[0030] The derived class PATTERN inherits class CIRCLE twice virtually through classes CIRCLE_1 and CIRCLE_2. Since the class CIRCLE is virtually inherited twice, there is only one object of class CIRCLE in the derived class PATTERN. One skilled in the art would appreciate virtual inheritance can be very useful when the class derivation is more complex.

[0031] A class may also specify whether its function members are virtual. Declaring that a function member is virtual means that the function can be overridden by a function of the same name and type in a derived class. In the following example, the function draw is declared to be virtual in classes CIRCLE and CIRCLE_FILL.

```

class CIRCLE
{ public:
  int x, y;
  int radius;
  virtual void draw( );
};
class CIRCLE_FILL : CIRCLE
{ public:
  int pattern;
  virtual void draw( );
};

```

[0032] The C++ language provides a pointer data type. A pointer holds values that are addresses of objects in memory. Through a pointer, an object can be referenced. The following statement declares variable c_ptr to be a pointer on an

object of type class CIRCLE and sets variable c_ptr to hold the address of object c.

[0033] CIRCLE *c_ptr;

[0034] c_ptr=&c;

Continuing with the example, the following statement declares object a to be of type class CIRCLE and object b to be of type class CIRCLE_FILL.

[0035] CIRCLE a;

[0036] CIRCLE_FILL b;

The following statement refers to the function draw as defined in class CIRCLE.

[0037] a.draw();

Whereas, the following statement refers to the function draw defined in class CIRCLE_FILL.

[0038] b.draw();

[0039] Moreover, the following statements type cast object b to an object of type class CIRCLE and invoke the function draw that is defined in class CIRCLE_FILL.

```
CIRCLE *c_ptr;
c_ptr = &b;
c_ptr->draw();           // CIRCLE_FILL::draw()
```

Thus, the virtual function that is called is function CIRCLE_FILL::draw.

[0040] FIG. 3 is a block diagram illustrating typical data structures used to represent an object. An object is composed of instance data (data members) and member functions, which implement the behavior of the object. The data structures used to represent an object comprise instance data structure 301, virtual function table 302, and the function members 303, 304, 305. The instance data structure 301 contains a pointer to the virtual function table 302 and contains data members. The virtual function table 302 contains an entry for each virtual function member defined for the object. Each entry contains a reference to the code that implements the corresponding function member. The layout of this sample object conforms to the model defined in U.S. patent application Ser. No. 07/682,537, entitled "A Method for Implementing Virtual Functions and Virtual Bases in a Compiler for an Object Oriented Programming Language," which is hereby incorporated by reference. In the following, an object will be described as an instance of a class as defined by the C++ programming language. One skilled in the art would appreciate that objects can be defined using other programming languages.

[0041] An advantage of using object-oriented techniques is that these techniques can be used to facilitate the sharing of objects. In particular, object-oriented techniques facilitate the creation of compound documents. A compound document (as described above) is a document that contains objects generated by various computer programs. (Typically, only the data members of the object and the class type are stored in a compound document.) For example, a word processing document that contains a spreadsheet object generated by a spreadsheet program is a compound docu-

ment. A word processing program allows a user to embed a spreadsheet object (e.g., a cell) within a word processing document. To allow this embedding, the word processing program is compiled using the class definition of the object to be embedded to access function members of the embedded object. Thus, the word processing program would need to be compiled using the class definition of each class of objects that can be embedded in a word processing document. To embed an object of a new class into a word processing document, the word processing program would need to be recompiled with the new class definition. Thus, only objects of classes selected by the developer of the word processing program can be embedded. Furthermore, new classes can only be supported with a new release of the word processing program.

[0042] To allow objects of an arbitrary class to be embedded into compound documents, interfaces are defined through which an object can be accessed without the need for the word processing program to have access to the class definitions at compile time. An abstract class is a class in which a virtual function member has no implementation (pure). An interface is an abstract class with no data members and whose virtual functions are all pure.

[0043] The following class definition is an example definition of an interface. In this example, for simplicity of explanation, rather than allowing any class of object to be embedded in its documents, a word processing program allows spreadsheet objects to be embedded. Any spreadsheet object that provides this interface can be embedded, regardless of how the object is implemented. Moreover, any spreadsheet object, whether implemented before or after the word processing program is compiled, can be embedded.

```
class ISpreadSheet
{
    virtual void File( ) = 0;
    virtual void Edit( ) = 0;
    virtual void Formula( ) = 0;
    virtual void Format( ) = 0;
    virtual void GetCell( string RC, cell *pCell) = 0;
    virtual void Data( ) = 0;
}

```

The developer of a spreadsheet program would need to provide an implementation of the interface to allow the spreadsheet objects to be embedded in a word processing document. When the word processing program embeds a spreadsheet object, the program needs access to the code that implements the interface for the spreadsheet object. To access the code, each implementation is given a unique class identifier. For example, a spreadsheet object developed by Microsoft Corporation may have a class identifier of "MSSpreadsheet," while a spreadsheet object developed by another corporation may have a class identifier of "LTSSpreadsheet." A persistent registry in each computer system is maintained that maps each class identifier to the code that implements the class. Typically, when a spreadsheet program is installed on a computer system, the persistent registry is updated to reflect the availability of that class of spreadsheet objects. So long as a spreadsheet developer implements each function member defined by the interface and the persistent registry is maintained, the word processing program can embed the developer's spreadsheet objects into a word processing document.

[0044] Various spreadsheet developers may wish, however, to implement only certain function members. For example, a spreadsheet developer may not want to implement database support, but may want to support all other function members. To allow a spreadsheet developer to support only some of the function members, while still allowing the objects to be embedded, multiple interfaces for spreadsheet object are defined. For example, the interfaces IDatabase and IBasic may be defined for a spreadsheet object as follows.

```

class IBasic
{ virtual void File() = 0;
  virtual void Edit() = 0;
  virtual void Formula() = 0;
  virtual void Format() = 0;
  virtual void GetCell (string RC, cell *pCell) = 0;
}
class IDatabase
{ virtual void Data() = 0;
}
    
```

Each spreadsheet developer would implement the IBasic interface and, optionally, the IDatabase interface.

[0045] At run time, the word processing program would need to determine whether a spreadsheet object to be embedded supports the IDatabase interface. To make this determination, another interface is defined (that every spreadsheet object implements) with a function member that indicates which interfaces are implemented for the object. This interface is named IUnknown (and referred to as the unknown interface or the object management interface) and is defined as follows.

```

class IUnknown
{ virtual HRESULT QueryInterface (REFIID iid, void **ppv) = 0;
  virtual ULONG AddRef() = 0;
  virtual ULONG Release () = 0;
}
    
```

[0046] The IUnknown interface defines the function member (method) QueryInterface. The method QueryInterface is passed an interface identifier (e.g., "IDatabase") in parameter iid (of type REFIID) and returns a pointer to the implementation of the identified interface for the object for which the method is invoked in parameter ppv. If the object does not support the interface, then the method returns a false. (The type HRESULT indicates a predefined status, and the type ULONG indicates an unsigned long integer.)

CODE TABLE 1

```

HRESULT XX::QueryInterface(REFIID iid, void **ppv)
{ ret = TRUE;
  switch (iid)
  { case IID_IBasic:
    *ppv = *pIBasic;
    break;
    case IID_IDatabase:
    *ppv = *pIDatabase;
    break;
  }
}
    
```

CODE TABLE 1-continued

```

case IID_IUnknown:
  *ppv = this;
  break;
default:
  ret = FALSE;
}
if (ret == TRUE){AddRef( );};
return ret;
}
    
```

[0047] Code Table 1 contains C++ pseudocode for a typical implementation of the method QueryInterface for class XX, which inherits the class IUnknown. If the spreadsheet object supports the IDatabase interface, then the method QueryInterface includes the appropriate case label within the switch statement. The variables pIbasic and pIDatabase point to a pointer to the virtual function tables of the Ibasic and IDatabase interfaces, respectively. The method QueryInterface invokes the method AddRef (described below) to increment a reference count for the object of class XX when a pointer to an interface is returned.

CODE TABLE 2

```

void XX::AddRef() {refcount++;}
void XX::Release() {if (--refcount==0) delete this;}
    
```

[0048] The interface IUnknown also defines the methods AddRef and Release, which are used to implement reference counting. Whenever a new reference to an interface is created, the method AddRef is invoked to increment a reference count of the object. Whenever a reference is no longer needed, the method Release is invoked to decrement the reference count of the object and, when the reference count goes to zero, to deallocate the object. Code Table 2 contains C++ pseudocode for a typical implementation of the methods AddRef and Release for class XX, which inherits the class IUnknown.

[0049] The IDatabase interface and IBasic interface inherit the IUnknown interface. The following definitions illustrate the use of the IUnknown interface.

```

class IDatabase : public IUnknown
{ public:
  virtual void Data() = 0;
}
class IBasic : public IUnknown
{ public:
  virtual void File() = 0;
  virtual void Edit() = 0;
  virtual void Formula() = 0;
  virtual void Format() = 0;
  virtual void GetCell (string RC, cell *pCell) = 0;
}
    
```

[0050] FIG. 4 is a block diagram illustrating a sample data structure of a spreadsheet object. The spreadsheet object comprises object data structure 401, IBasic interface data structure 403, IDatabase interface data structure 404, the virtual function tables 402, 405, 406 and methods 407 through 421. The object data structure 401 contains a pointer

to the virtual function table 402 and pointers to the IBasic and IDatabase interface. Each entry in the virtual function table 402 contains a pointer to a method of the IUnknown interface. The IBasic interface data structure 403 contains a pointer to the virtual function table 405. Each entry in the virtual function table 405 contains a pointer to a method of the IBasic interface. The IDatabase interface data structure 404 contains a pointer to the virtual function table 406. Each entry in the virtual function table 406 contains a pointer to a method of the IDatabase interface. Since the IBasic and IDatabase interfaces inherit the IUnknown interface, each virtual function table 405 and 406 contains a pointer to the methods QueryInterface, AddRef, and Release. In the following, an object data structure is represented by the shape 422 labeled with the interfaces through which the object may be accessed.

[0051] The following pseudocode illustrates how a word processing program determines whether a spreadsheet object supports the IDatabase interface.

```

if (pIBasic->QueryInterface("IDatabase", &pIDatabase) == S_OK)
  \* IDatabase supported
else
  \* IDatabase not supported

```

The pointer pIBasic is a pointer to the IBasic interface of the object. If the object supports the IDatabase interface, the method QueryInterface sets the pointer pIDatabase to point to the IDatabase data structure and returns the value S_OK.

[0052] Normally, an object can be instantiated (an instance of the object created in memory) by a variable declaration or by the "new" operator. However, both techniques of instantiation need the class definition at compile time. A different technique is needed to allow a word processing program to instantiate a spreadsheet object at run time. One technique provides a global function CreateInstanceXX, which is defined in the following.

```

[0053] static void CreateInstanceXX (REFIID iid, void
**ppv)=0;

```

The method CreateInstanceXX (known as a class factory) instantiates an object of class XX and returns a pointer ppv to the interface of the object designated by parameter iid.

SUMMARY OF THE INVENTION

[0054] It is an object of the present invention to provide a method and system for generating links to source data incorporated within a compound document.

[0055] It is another object of the present invention for binding links to source data.

[0056] It is another object of the present invention for interfacing with these links in a manner that is independent of the underlying source data.

[0057] It is another object of the present invention for linking to data nested to an arbitrary level within a compound document.

[0058] These and other object, which will become apparent as the invention is more fully described below, are provided by a method and system for naming and binding

data objects. In a preferred embodiment, a link to an object incorporated is stored as a moniker. A moniker is an identifier object that encapsulates the information needed to access the incorporated data and provides methods which bind to the incorporated data.

BRIEF DESCRIPTION OF THE DRAWINGS

[0059] FIG. 1 is a block diagram of an example of a compound document.

[0060] FIG. 2 is a block diagram illustrating scheduling data, budgeting data, and explanatory data.

[0061] FIG. 3 is a block diagram illustrating typical data structures used to represent an object.

[0062] FIG. 4 is a block diagram illustrating a sample data structure of a spreadsheet object.

[0063] FIG. 5 is a block diagram showing a sample compound document.

[0064] FIGS. 6, 7, and 8 are block diagrams illustrating the use of a moniker by a word processing program.

[0065] FIG. 9 is a block diagram of a generic composite moniker.

[0066] FIGS. 10A, 10B, and 10C are block diagrams illustrating moniker composition.

[0067] FIG. 11 is a flow diagram of the method BindToObject of the class CFileMoniker.

[0068] FIG. 12 is a flow diagram of the function FileBindToObject.

[0069] FIG. 13 is a flow diagram of the method BindToObject of the class CItemMoniker.

[0070] FIG. 14 is a flow diagram of the method BindToObject of the class CCompositeMoniker.

[0071] FIGS. 15A through 15G are block diagrams illustrating the binding to an object identified by a generic composite moniker.

[0072] FIG. 16 is a flow diagram illustrating the overall behavior of implementations of the method ComposeWith.

[0073] FIG. 17 is a flow diagram of the method ComposeWith of the class CCompositeMoniker.

[0074] FIGS. 18, 19A, 19B, 19C, 20A, 20B, 20C, 21A, 21B, and 21C are block diagrams illustrating sample generic composite monikers.

[0075] FIG. 22 is a block diagram illustrating moniker reduction.

[0076] FIG. 23 is a flow diagram of the method Reduce of the class CCompositeMoniker.

[0077] FIG. 24 is a flow diagram of the method Reduce of the sample class CAliasMoniker.

[0078] FIG. 25 is a flow diagram of the method IsEqual of the class CFileMoniker.

[0079] FIG. 26 is a flow diagram of the method IsEqual of the class CCompositeMoniker.

[0080] FIG. 27 is a flow diagram of the method Hash of the class CCompositeMoniker.

[0081] FIGS. 28A, 28B, and 28C are block diagrams illustrating composition with inverse monikers.

[0082] FIG. 29 is a flow diagram of the method Inverse of the class CCompositeMoniker.

[0083] FIG. 30 is a flow diagram of the method Inverse of the class CItemMoniker.

[0084] FIG. 31 is a flow diagram of the method ComposeWith of the class CItemMoniker.

[0085] FIG. 32 is a flow diagram of the method AnnihilateOne of the class CAntiMoniker.

[0086] FIG. 33 is flow diagram of the method ComposeWith of the class CAntiMoniker.

[0087] FIG. 34 is a block diagram illustrating a common prefix of generic composite monikers.

[0088] FIG. 35 is a flow diagram of the method CommonPrefixWith of the class CCompositeMoniker.

[0089] FIG. 36 is a flow diagram of the method CommonPrefixWith of the class CFileMoniker.

[0090] FIG. 37 is a flow diagram of the method CommonPrefixWith of the class CItemMoniker.

[0091] FIG. 38 is a flow diagram of the method CommonPrefixWith of the class CAntiMoniker.

[0092] FIG. 39 is a flow diagram of the function MonikerCommonPrefixWith.

[0093] FIG. 40 is a block diagram illustrating a relative path to moniker.

[0094] FIGS. 41A and 41B are block diagrams illustrating a usage of the method RelativePathTo.

[0095] FIGS. 42A, 42B, and 42C comprise a flow diagram of the method RelativePathTo of the class CCompositeMoniker.

[0096] FIG. 43 is a flow diagram of the function MonikerRelativePathTo.

[0097] FIG. 44 is a flow diagram of the method Enum of the class CCompositeMoniker.

[0098] FIG. 45 is a flow diagram of the method GetNext of the class CCompositeMonikerEnum.

[0099] FIG. 46 is a flow diagram of the method Next of the class CCompositeMonikerEnum.

[0100] FIG. 47 is a flow diagram of the method Create of the class CPointerMoniker.

[0101] FIG. 48 is a flow diagram of the method BindToObject of the class CPointerMoniker.

[0102] FIG. 49 is a block diagram illustrating a pointer moniker.

DETAILED DESCRIPTION OF THE INVENTION

[0103] The present invention provides a computer implemented method and system for naming and binding to linked data. In a preferred embodiment, a compound document that incorporates linked data stores a persistent data handle, called a "moniker," which is a reference to the link source. A moniker is an identifier object that contains information to identify the linked data and provides methods through which a program can bind to the linked data. A binding method

returns an instance of an interface through which the linked data can be accessed. A moniker may link to data that is itself embedded data within another compound document. For example, a moniker may link to a range of cells within a spreadsheet table that is contained in a word processing document. A moniker may link to data at any level within a compound document. During execution of the binding method, several applications may be invoked to locate the link data. For example, to bind to the range of cells within a spreadsheet table that is within a word processing document, the word processing program may be invoked to locate the embedded spreadsheet table and the spreadsheet program may be invoked to bind to the range of cells. The present invention defines an interface through which a moniker is accessed. A moniker can identify source data that is stored persistently or non-persistently.

[0104] In a preferred embodiment, monikers can be composed to form a composite moniker. A composite moniker is conceptually a path to a source object that is identified by the concatenation of the monikers. For example, if a moniker specifying a certain path (e.g., "c:\reports") is composed with a moniker specifying a certain file name (e.g., "Q3.doc") then the result is the complete path name to the file (e.g., "c:\reports\Q3.doc"). Each composite moniker comprises a plurality of component monikers. The present invention provides a method and system for decomposing a composite moniker. In a preferred embodiment, each moniker provides a method that is used to retrieve each component moniker.

[0105] In a preferred embodiment, a moniker provides a reducing method which returns another moniker that is a more efficient representation of a moniker to the same source object. The reducing method may interpret a macro script that identifies the source object. Alternatively, the reducing method may evaluate a query request that identifies the source object.

[0106] In a preferred embodiment, a moniker provides an equality method and a hash method. The equality method determines whether two monikers identify the same source object. The hash method provides a hash value for a moniker. The equality method and hash method are used to implement hash tables indexed by monikers.

[0107] In a preferred embodiment, a moniker provides an inverse method that generates another moniker that is the inverse of the moniker. When a moniker is composed with its inverse, the result is NULL. The inverse moniker is said to annihilate the moniker. An inverse moniker may be used, for example, to remove portions of a path and is analogous to the "." functionality of traditional file systems.

[0108] In a preferred embodiment, a moniker provides a common prefix with method and a relative path to method. The common prefix with method determines the common prefix portion of two monikers. For example, if one moniker identifies the object "c:\reports\Q3.doc" and another moniker identifies the object "c:\reports\data\Q3.xls" the common prefix is "c:\reports". The relative path to method generates relative path to moniker that when composed with one moniker results in specified moniker. For example, the moniker specifying the path that is the inverse of a moniker identifying object "Q3.doc" composed with a moniker specifying the path "data\Q3.xls" is a relative path to moniker from the moniker "c:\reports\Q3.doc" to the moniker

"c:\reports\data\Q3.xls". Relative paths to monikers are preferably used when identifying objects by relative paths from another object.

[0109] In a preferred embodiment, the present invention provides several implementation monikers including a file moniker, an item moniker, a generic composite moniker, a pointer moniker, and an anti moniker. Each implementation is referred to as a moniker class and has a class identifier. A file moniker provides a moniker that conceptually is a path name in a file system. An item moniker provides a moniker that conceptually identifies a portion of an object. A generic composite moniker provides a mechanism for composing monikers with arbitrary implementations. For example, a file moniker can be composed with an item moniker to specify a portion of a file. A generic composite moniker is preferably created by the composing method of the file moniker. A pointer moniker is a moniker that wraps an instantiated source object in a moniker. A pointer moniker contains a pointer to the instantiated source object and when a pointer moniker is bound, it returns the pointer. An anti moniker is a moniker that is the inverse of other monikers. When a moniker is composed with an anti moniker, the result is NULL. If a generic composite moniker is composed with an anti moniker, the result is a moniker comprising all but the last component moniker. The anti moniker annihilates the last component moniker of a generic composite moniker.

[0110] In a preferred embodiment of the present invention, an application program that creates a compound document controls the manipulation of linked or embedded data generated by another application. In object-oriented parlance, this data is referred to as an object. (The reference Budd, T., "An Introduction to Object-Oriented Programming," Addison-Wesley Publishing Co., Inc., 1991, provides an introduction to object-oriented concepts and terminology.) An object that is either linked or embedded into a compound document is "contained" within the document. Also, a compound document is referred to as a "container" object and the objects contained within a compound document are referred to as "containee" objects. Referring to FIGS. 1 and 2, the scheduling data 102 and budgeting data 103 are containee objects and the compound document 101 is a container object. The user can indicate to the word processor that the user wants to edit a containee object, such as the budgeting data 103. When the user indicates that the budgeting data 103 is to be edited, the word processing program determines which application should be used to edit the budgeting data (e.g., the spreadsheet program) and launches (starts up) that application. The user can then manipulate the budgeting data using the launched application, and changes are reflected in the compound document. The same procedure is used whether the budgeting data is stored as an embedded or linked object.

[0111] FIG. 5 is a block diagram showing a sample compound document. The weekly project report 501 is the same compound document of FIG. 1. The executive summary report 503, contains a budgeting chart 505 that is linked to the weekly project 501. The weekly project 501 contains an embedded spreadsheet 502. The embedded spreadsheet 502 was created by the spreadsheet program 204 in FIG. 2. The data for this spreadsheet, the budget for the project, is stored within the storage the weekly project report 501 because it is an embedded object. The executive summary document 503 is a compound document which contains native text 504

and a contained object, the budget chart 505. The budget chart 505 is linked to the data contained within in the spreadsheet 502 which is embedded in the compound document 501.

[0112] In a preferred embodiment, application programs ("applications") cooperate using object linking and embedding facilities to create and manipulate compound documents. An application that creates a compound document is referred to as a client application, and an application that creates and manipulates containee objects are referred to as server applications. An application can behave as both a client server. Referring to FIG. 2, the project management program 201 and the spreadsheet program 204 are server applications, and the word processing program 206 is a client application. A client application is responsible for selection of the various objects within the container object and for invoking the proper server application to manipulate the selected containee object. A server application is responsible for manipulating the contents of the containee objects.

[0113] In a preferred embodiment, applications are provided with an implementation-independent Application Programming Interface (API) that provides the object linking and embedding functionality. Appendix A contains a detailed description of a preferred object linking and embedding system. The API is a set of functions that are invoked by client and server applications. These functions manage, among other things, the setup and initialization necessary for client applications to send and receive messages and data to and from server applications. The API provides functions to invoke the correct server application to act upon a particular containee object and to manipulate containee objects.

[0114] In addition, the object linking and embedding API defines "interfaces" through which client applications can communicate with their contained objects. An interface is a set of methods which abide by certain input, output, and behavior rules. If a contained object supports a particular interface, the client application can invoke the methods of that interface to effect the defined behavior. In a preferred embodiment, the client application is not allowed direct access to the object data; it manipulates the object using the supported interfaces. A client application is bound to a contained object through a pointer to an interface. The client application accesses the object by invoking the methods of the interface. To access the object data, the methods may send messages to the server application requesting the specified access. In a preferred embodiment, messages are sent between clients and servers using interprocess communications mechanisms provided by the underlying operating system.

[0115] An example will help illustrate the relationship between a client process and a server process. Referring again to FIG. 1, if a user wants to edit the budgeting data 103 of the compound document 101, then the following sequence of events occurs. First, the user starts up the word processor program, which is dynamically linked to the object linking and embedding API. Second, the user opens the compound document for editing. Third, the user selects the budgeting data, which is a containee object, and indicates that the selected object is to be edited. Fourth, the client application invokes a client API routine for performing an action on an object passing the routine a handle (which uniquely identifies the selected object) to the object

and an indicator that the action is edit. Fifth, the client API routine determines that the spreadsheet program provides the actions for the budgeting data. Sixth, the client API code starts up the spreadsheet program as a server process, if it is not already started. Seventh, the word processor application sends a message to the spreadsheet program that it should edit the budgeting data. Eighth, the server API code receives the request to edit and invokes a routine in the spreadsheet program for editing the data. When editing is complete, the spreadsheet routine returns to the server API code. The server API code sends a message to the word processor application to indicate that editing is complete. The client API code receives the message and returns from its invocation. Upon return from the invocation, the word processor application knows that the editing is complete.

[0116] In addition to the client and server API, the object linking and embedding facilities of the present invention provide information to client and server applications through a persistent global “registry.” This registry is a database of information such as (1) for each type of object, the server application that implements the object type, (2) the actions that the each server application provides to client applications, (3) where the executable files for each server application are located, and (4) whether each server application has an associated object handler. An object handler is a collection of functions in a dynamic link library. An object handler can be used to provide certain functions without launching the server.

[0117] FIGS. 6, 7, and 8 are block diagrams illustrating the use of a moniker by a word processing program. In FIG. 6, the document “RPT.DOC” 601 contains a link 602 to the chart file 603. The link 602 is a moniker that is persistently stored in the document 601. The persistent storage of the moniker includes the class identifier “CLSID_FileMoniker” and the name of the chart file 603 (“Q3.CHT”). When the word processing program displays the chart of chart file 603, it first instantiates a moniker object of type CLSID_FileMoniker, requests the moniker to load its persistent data (e.g., “Q3.CHT”), and then requests the moniker to bind to the file indicated by the loaded data. FIG. 7 is a block diagram illustrating the instantiation of a moniker object. The word processing program first reads in the class identifier of link 602. To determine how to instantiate an object of that class, the program accesses the global registry 704. The global registry includes a mapping from class identifiers to the location of a class factory to create an instance of that class. For example, table 704B indicates that the class factory for the moniker class identified by CLSID_FileMoniker is contained in the dynamic link library named “FileMkr.DLL.” The program links to the class factory code within the dynamic link library 706. The program invokes the function CreateInstance to create an instance of an object of class CFileMoniker 702 (a file moniker). The program then requests the IPersistStream interface of the file moniker. The IPersistStream interface (described in detail in Appendix A) provides methods through which the internal state (e.g., “Q3.CHT”) of a moniker can be saved to persistent storage and then loaded into memory from the persistent storage. Using the methods of the IPersistStream interface, the program loads the moniker internal state persistently stored in link 602 into the file moniker 702. The program then requests the IMoniker interface to the file moniker 702. The pointer to the IMoniker interface is stored in pointer 701. FIG. 8 is a block diagram illustrating the binding of the file

moniker 702 to the chart file 603. When a binding method of the file moniker 702 is invoked, the method determines the class identifier for the file identified by the file moniker by accessing the global registry table 704A. The class identifier for files with suffix “CHT” is CLSID_Chart. The program then retrieves the class factory for the class CLSID_Chart from the registry table 704B. The program links to the dynamic link library “CHART.DLL” 806. The program then invokes the CreateInstance method within the dynamic link library 806, which creates an instance of a chart object 807 and returns a pointer to the IDataObject interface. The IDataObject interface (described in detail in Appendix A) provides methods to pass data to and from an object (e.g., methods GetData and SetData). Through the chart object 807, the program can access the chart file 603 through the chart server 808.

[0118] In a preferred embodiment, a moniker is an object that supports the IMoniker interface of Code Table 3. The IMoniker interface inherits the IPersistStream interface; thus, monikers can be saved to and loaded from streams. The persistent form of a moniker contains the class identifier (CLSID) of its implementation which is used during the loading process, and new classes of monikers can be created transparently to clients.

[0119] The IMoniker interface provides for binding to the object to which it points, which is supported by the method BindToObject. This method takes as a parameter the interface identifier by which the caller wishes to talk to the object, runs whatever algorithm is necessary in order to locate the object, then returns a pointer of that interface type to the caller. Each moniker class can store arbitrary data its persistent representation, and can run arbitrary code at binding time.

[0120] If there is an identifiable piece of persistent storage in which the object referenced by the moniker is stored, then the method BindToStorage can be used to access it. Many objects have such identifiable storage (e.g., a file), but some, such as the objects which are the ranges in a spreadsheet do not.

[0121] In a preferred embodiment, a particular moniker class is designed to be one step along the path (a component) to a data source. These components can be composed together to form a moniker which represents the complete path to the data source. For example, the moniker stored inside the chart of FIG. 5 might be a generic composite moniker formed from three component as illustrated in FIG. 9. This composite is itself a moniker; it is a moniker which is a sequenced collection of other composite monikers. The composition is generic in that it has no knowledge of the component monikers involved other than that they are monikers.

CODE TABLE 3

class IMoniker : IPersistStream {			
virtual	HRESULT	BindToObject(pbc, pmkToLeft,	iidResult, ppvResult) = 0;
virtual	HRESULT	BindToStorage(pbc, pmkToLeft,	iid, ppvObj) = 0;
virtual	HRESULT	Reduce(pbc, dwReduceHowFar,	ppmkToLeft, ppmkReduced) = 0;
virtual	HRESULT	ComposeWith(pmkRight,	fOnlyIfNotGeneric, ppmkComposite)

CODE TABLE 3-continued

virtual	HRESULT	Enum(fForward, ppenmMoniker) = 0;
virtual	HRESULT	IsEqual(pmkOtherMoniker) = 0;
virtual	HRESULT	Hash(pdwHash) = 0;
virtual	HRESULT	IsRunning(pbc, pmkToLeft, pmkNewlyRunning) = 0;
virtual	HRESULT	GetTimeOfLastChange(pbc, pmkToLeft, pfiletime) = 0;
virtual	HRESULT	Inverse(ppmk) = 0;
virtual	HRESULT	CommonPrefixWith(pmkOther, ppmkPrefix) = 0;
virtual	HRESULT	RelativePathTo(pmkOther, ppmkRelPath);
virtual	HRESULT	GetDisplayName(pbc, pmkToLeft, lppszDisplayName) = 0;
virtual	HRESULT	ParseDisplayName(pbc, pmkToLeft, lpaszDisplayName, pcchEaten, ppmkOut) = 0;
virtual	HRESULT	IsSystemMoniker(pdwMksys);
		};

[0122] The example of FIGS. 6, 7, and 8 illustrate the use of a moniker that identifies a file. The present invention allows moniker to be combined (composed) to an arbitrary level. FIGS. 10A, 10B, and 10C illustrate moniker composition. For example, if the chart file 603 contained multiple charts, it would be useful to designate a specific chart to be the source of a link. In one embodiment of the present invention, a moniker class named "CChartMoniker" could be implemented by the developer of the chart program. A chart moniker 1001 would contain a name of a chart file ("Q3.CHT") and an indication of a chart within the file ("CHART2"). The methods of the class CChartMoniker would have a behavior similar to that provided by the class CFileMoniker plus behavior needed to bind to the identified chart. As described above, the present invention allows two monikers to be composed to form a third moniker. By composing monikers, a developer can use an implementation developed by someone else. For example, the developer of the chart program could define and implement the class CChartMoniker to contain only an indication of a chart within a file. The class CChartMoniker can be developed assuming that an instance of chart moniker 1003 will be composed with a file moniker (e.g., file moniker 1002). In a preferred embodiment, to facilitate the composing of monikers, a moniker of class CCompositeMoniker is defined and implemented. The class CCompositeMoniker encapsulates any two monikers into single generic composite moniker. The generic composite moniker 1004 encapsulates the file moniker 1002 and the chart moniker 1003. A link to a chart is stored as a generic composite moniker which encapsulates a file moniker and a chart moniker. The client of the link need only know that the moniker supports the IMoniker interface.

[0123] In the following, each method of the IMoniker interface is defined. In addition, several implementations of various methods are described. In particular, implementations of methods of the classes CFileMoniker, CCompositeMoniker, and CItemMoniker are described. The class CFileMoniker (a file moniker) is a moniker class that identifies a path name in a file system. When a file moniker is bound to, it determines the class of the file by using the persistent global registry, ensures that the appropriate class server is running, and then requests the server to open the file. The class CCompositeMoniker (a generic composite

moniker) is a moniker class that identifies a composition of two monikers (a left and a right moniker). When a generic composite moniker is bound to, it invokes the binding method of the right moniker indicating that the left moniker is composed with the right moniker. The right moniker performs its binding behavior, which may include invoking the binding method of the left moniker. The class CItemMoniker (an item moniker) is a moniker class that implements behavior common to the identification of containee objects. An item moniker can be used to identify, for example, a chart contained within a chart file or a range within a spreadsheet. An item moniker uses the IOleItemContainer interface (described in detail in Appendix A) to interact with the container. Code Table 4 contains the class definitions for a file moniker, a generic composite moniker, an item moniker, an anti moniker, and a pointer moniker. An anti moniker and a pointer moniker are described below in detail. A file moniker contains a string (m_szPath) indicating a path name and a count of anti monikers (m_cAnti). A generic composite moniker contains a pointer to the left moniker (m_pmkLeft) and a pointer to the right moniker (m_pmkRight) of the generic composite and a flag (m_fReduced) indicating whether the composite is reduced. An item moniker contains a pointer to a string (m_lpszItem) that defines the item.

CODE TABLE 4

class CFileMoniker: IMoniker	
{ char FAR *	m_szPath;
UINT m_cAnti;	
}	
class CCompositeMoniker: IMoniker	
{ LPMONIKER	m_pmkLeft;
LPMONIKER	m_pmkRight;
BOOL	m_fReduced;
}	
class CItemMoniker: IMoniker	
{ char FAR *	m_lpszItem;
}	
class CAntiMoniker: IMoniker	
{ ULONG	m_count;
}	
class CPointerMoniker: IMoniker	
{ LPUNKNOWN	m_punk;
}	
IMoniker::BindToObject	
HRESULT IMoniker::BindToObject(pbc, pmkToLeft, iidResult,	
ppvResult)	

[0124] The method BindToObject locates and loads the object semantically referred to by this moniker according to the interface specified by iidResult and returns a pointer to the object through ppvResult. In the following, the term "this moniker" refers to the moniker for which a method is invoked. In general, each class of moniker is designed to be used as one component in a generic composite moniker which gives the complete path to the referenced object. In a generic composite, any component moniker has a certain prefix of the generic composite to its left, and a certain suffix to its right. If the method BindToObject is invoked on a component moniker, then the implementation of BindToObject typically requires certain services of the object indicated by the prefix to its left. Item monikers, for example, require the IOleItemContainer interface of the object to their left. The Item Moniker implementation of the method BindToObject (as described below) recursively calls pmkToLeft->BindToObject in order to obtain this interface. If the

moniker does not need services of the object to its left, yet one is provided by the caller nevertheless, no error occurs. Rather, the moniker ignores the object to its left. If the object indicated by the moniker does not exist, then the error MK_E_NOOBJECT is returned.

[0125] In general, binding a moniker can be a complicated process, since it may need to launch servers, open files, etc. This may involve binding to other objects, and the binding components of a generic composite to the right of certain components will require the same other objects. In order to avoid loading the object, releasing it, then having it loaded again later, the method BindToObject can use the bind context passed through the pbc parameter to defer releasing objects until the binding process overall is complete. The bind context is described in detail in Appendix A.

[0126] Binding to a moniker a second time typically returns the same running object as binding the first time, rather than reloading it again from storage. This functionality is supported with a running object table. The running object table is a lookup table keyed by a moniker whose values are pointers to the corresponding now-running object. As objects become running, they register themselves in this table. Implementations of the method BindToObject uses this table to determine if the object to which they point is already running. More precisely, if the passed pmkToLeft parameter is NULL (and this is not an error; that is, the moniker does not require something to its left), then the moniker fully reduces itself, then looks itself up in the running object table, and returns the pointer to the object found there. The running object table is described in detail in Appendix A.

[0127] The following table describes the parameters of the method BindToObject:

Argument	Type	Description
pbc	IBindCtx*	the bind context to be used for this binding operation.
pmkToLeft	IMoniker*	the moniker of the object to the left of this moniker.
iidResult	IID	the requested interface by which the caller wishes to connect to the object.
ppvResult	void**	on successful return, a pointer to the instantiated object is placed here, unless BINDFLAGS_JUSTTESTEXISTENCE was specified in the binding options, in which case NULL may be returned instead.
return value	HRESULT	S_OK, MK_E_NOOBJECT, STG_E_ACCESSDENIED, MK_E_EXCEEDEDDEADLINE, MK_E_CONNECTMANUALLY, MK_E_INTERMEDIATEINTERFACENOTSUPPORTED, E_OUTOFMEMORY, E_NOINTERFACE

[0128] FIG. 11 is a flow diagram of the method BindToObject of the class CFileMoniker. This method determines the class identifier of the file, determines the server for that file, launches the server (if necessary), and requests the server to open and bind to the file. In step 1101, if a moniker to the left of this moniker is specified, then the method continues at step 1105, else the method continues at step 1102. In steps 1102 through 1104, the method determines whether the object identified by the moniker to the left is in the running object table. If the object to the left is in the running object table, then the requested interface (iidResult) is retrieved from that object and returned. In step 1102, if the

object to the left is in the running object table, then the method continues at step 1103, else the method continues at step 1105. In step 1103, the method retrieves a pointer to the object to the left from the running object table. In step 1104, the method retrieves the requested interface from the object to the left by invoking the method QueryInterface of the object to the left and then returns. In step 1105, the method retrieves the class identifier corresponding to the path (m_szPath) of this moniker. The class identifier is preferably retrieved from a persistent global registry that maps file name suffixes to class identifiers. In step 1106, the method invokes the function FileBindToObject to bind to the file and returns the requested interface. The method then returns.

[0129] FIG. 12 is a flow diagram of the function FileBindToObject. This function is passed a class identifier and a requested interface. This function instantiates an object of the passed class identifier and returns a pointer to the requested interface. In step 1201, the function creates an instance of an object of the passed class identifier and retrieves the IUnknown interface. In step 1202, the function retrieves the IPersistFile interface from the instantiated object. The IPersistFile interface provides methods to load and save files and is described in detail in Appendix A. In step 1203, the function initializes binding options, which are described in detail in Appendix A. In step 1204, the function invokes the method Load of the IPersistFile interface to load the data for the created object. In step 1205, the function retrieves the requested interface from the created object and returns.

[0130] FIG. 13 is a flow diagram of the method BindToObject of the class CItemMoniker. In step 1301, if no moniker to the left of this moniker is specified, then the method returns an error, else the method continues at 1302. An item moniker identifies a containee object and requires

a moniker to a container object to its left. In step 1302, the method invokes the method BindToObject of the moniker of the object to the left. The method requests the IOleItemContainer interface from the object to the left. In step 1303, the method invokes the method GetObject of the IOleItemContainer interface and passes the item name (m_lpszItem) and the requested interface. The method then returns with the interface retrieved by the method GetObject.

[0131] FIG. 14 is a flow diagram of the method BindToObject of the class CCompositeMoniker. The method binds to a generic composite moniker in a right-to-left manner. Conceptually; a generic composite moniker for-

wards the bind request to its last component moniker informing the last component moniker of the moniker to its left in the composite. The last component moniker, if it needs to, recursively binds to the object to its left. In step 1401, if no moniker to the left of this moniker is specified, the method continues at step 1402, else the method continues at step 1405. In step 1402, if this moniker is in the running object table, then the method continues at step 1403, else the method continues at step 1405. In step 1403, the method retrieves a pointer to the object from the running object table. In step 1404, the method retrieves the requested interface of the object by invoking the method QueryInterface of the object and returns. In steps 1405 through 1408, the method invokes the method BindToObject of the last component moniker of this moniker passing a moniker comprising the prefix component monikers as the moniker to the left. In step 1405, if no monikers to the left of this moniker is specified, then the method continues at step 1406, else the method continues at step 1407. In step 1406, the method creates a new left moniker that contains all but the last component moniker of this moniker. The method then invokes the method BindToObject of the last component moniker of this moniker passing it the newly-created left moniker and the requested interface in step 1408 and returns. In step 1407, the method composes the moniker to the left with all but the last component moniker of this moniker by invoking the method ComposeWith of the moniker to the left. The method then invokes the method BindToObject of the last component moniker of this moniker passing it the newly-created composed moniker and the identifier of the requested interface in step 1408 and returns.

[0132] FIGS. 15A through 15G are block diagrams illustrating the binding to an object identified by a generic composite moniker. FIG. 15A illustrates the generic composite moniker that is to be bound. The generic composite moniker comprises components 1501, 1502, and 1503. The component monikers 1501, 1502, and 1503 represent a reference to an object identified by "C:\Q3RPT.DOC\SALESTBL\R2C2:R7C4". The component moniker 1501 is the first component moniker of the generic composite moniker, and the component moniker 1503 is the last component moniker of the generic composite moniker. The component monikers 1502 and 1503 are all but the first component monikers of the generic composite moniker, and the component monikers 1501 and 1502 are all but the last component monikers of the generic composite moniker. These component monikers are composed using generic composite monikers 1504 and 1505. To bind to the object identified by the generic composite moniker 1504, the method BindToObject is invoked indicating that there is no moniker to the left and indicating the identifier of a requested interface. Since moniker 1504 is a generic composite moniker, the method represented by the flow diagram of FIG. 14 is executed. Since there is no moniker to the left and since for this example the generic composite moniker 1504 is not in the running object table, the method continues at step 1406. In step 1406, the method creates the generic composite moniker 1506 that contains all but the last component moniker of the generic composite moniker 1504. In step 1408, the method invokes the method BindToObject of the last component moniker 1503 passing the generic composite moniker 1506 as the moniker to the left and the identifier of the requested interface. Since the component moniker 1503 is an item moniker, the method represented by

the flow diagram of FIG. 13 is executed. Since a moniker to the left is specified, step 1302 is executed. In step 1302, the method invokes the method BindToObject of the moniker to the left (generic composite moniker 1506) passing no moniker to the left and requesting the IOleItemContainer interface. Since the generic composite moniker 1506 is a generic composite moniker, the method represented by the flow diagram of FIG. 14 is executed. Since no moniker to the left is specified, the method continues at step 1406. In step 1406, the method sets the new left moniker to the moniker 1501. In step 1408, the method invokes the method BindToObject of the item moniker 1502 passing the new left moniker 1501 and requesting the passed interface, which is the interface to the IOleItemContainer. Since the item moniker 1502 is an item moniker, the method represented by the flow diagram of FIG. 13 is executed. Since there is a moniker to the left is specified, step 1302 is executed. In step 1302, the method invokes the method BindToObject of the moniker to the left (file moniker 1501). Since the file moniker 1501 is a file moniker, the method represented by the flow diagram of FIG. 11 is executed. Since no moniker to the left is specified, the method continues at step 1105. In steps 1105 and 1106, the method binds to the file object and returns the requested interface as shown in FIG. 15G. The invocation of the method BindToObject of the item moniker 1503 eventually in step 1303 invokes the method GetObject of the IOleItemContainer interface returned in step 1302 to retrieve the requested interface.

IMoniker::ComposeWith

RESULT IMoniker::ComposeWith(pmkRight, fOnlyIfNotGeneric, ppmkComposite)

[0133] This method ComposeWith returns a new moniker which is a composite formed with this moniker on the left and the passed moniker (pmkRight) on the right. There are two kinds of composite monikers: those composite monikers that know nothing about their component monikers other than that they are monikers (a generic composite moniker), and those composite monikers that know more (a special composite moniker). For example, a file moniker containing a relative path may be composed on to the end of another file moniker. The resulting composite moniker could be a new file moniker containing the complete path. The new file moniker is a special composition. A special composition is useful for monikers that are capable of collapsing a path within a storage domain to a more efficient representation in a subsequent reduction.

[0134] Each moniker class may have a set of other kinds of special monikers that can be composed onto the end of it in a non-generic way. Each implementation of the method ComposeWith examines the passed moniker on the right (pmkRight) to see if it is such a special moniker for the implementation. If the specified moniker on the right is special, then the implementation does whatever is appropriate for that special case. If it is not, then the passed flag fOnlyIfNotGeneric controls what occurs. If flag fOnlyIfNotGeneric is true, then NULL is passed back through parameter ppmkComposite and the status MK_E_NEEDGENERIC returned; if fOnlyIfNotGeneric is false, then a generic composite moniker is created using the function CreateGenericComposite and returned.

[0135] If the specified moniker on the right (pmkright) completely annihilates this moniker, the resulting composite

is empty and the parameter ppmkComposite is set to NULL and the status S_OK returned.

[0136] Composition of monikers is an associative operation. That is, if A, B, and C are monikers, then

$$(A \circ B) \circ C$$

is always equal to

$$A \circ (B \circ C)$$

where \circ represents the composition operation.

[0137] The following table describes the parameters of the method ComposeWith:

Argument	Type	Description
pmkRight	IMoniker*	the moniker to compose onto the end of the receiver.
fOnlyIfNotGeneric	BOOL	controls whether a composite moniker should be returned when the right moniker is not a special moniker for this moniker.
ppmkComposite	IMoniker*	on exit, the resulting composite moniker. Possibly NULL.
return value	HRESULT	S_OK, MK_E_NEEDGENERIC

[0138] FIG. 16 is a flow diagram illustrating the overall behavior of implementations of the method ComposeWith. In step 1601, if the moniker to the right (pmkRight) is special, then the method continues at step 1602, else the method continues at step 1603. In step 1602, the method performs a composition associated with the special moniker and returns. In step 1603, if the caller wants a generic composition when no special composition occurs (fOnlyIfNotGeneric==FALSE), then the method continues at step 1604, else the method continues at step 1605. In step 1604, the method creates a generic composite moniker by invoking the function CreateGenericMoniker and returns. In step 1605, the method returns an indication that composite moniker is NULL and sets the return flag to indicate that no special or generic composition occurred.

[0139] FIG. 17 is a flow diagram of the method ComposeWith of the class CCompositeMoniker. The method invokes the function CreateGenericComposite (described below) and returns.

CODE TABLE 4

```

CreateGenericComposite (pmkFirst, pmkRest, ppmkComposite)
{
Case1:
    pmkFirst->ComposeWith (pmkRest, TRUE, ppmkComposite)
    if (no composition occurred)
        CCompositeMoniker::Create (pmkFirst, pmkRest,
            ppmkComposite)
Case2:
    pmkFirst->ComposeWith (pmkFirstOfRest, TRUE, &pmk)
    if (no composition occurred)
        CCompositeMoniker::Create (pmkFirst, pmkRest,
            ppmkComposite)
    else
        if (pmk! = NULL)
            CreateGenericComposite (pmk,
                pmkAllButFirstOfRest, ppmkComposite)
        else
            *ppmkComposite = pmkAllButFirstOfRest
}
    
```

CODE TABLE 4-continued

```

Case3:
    pmkLastOfFirst->ComposeWith (pmkRest, TRUE, &pmk);
    if (no composition occurred)
        CCompositeMoniker::Create (pmkFirst, pmkRest,
            ppmkComposite)
    else
        if (pmk! = NULL)
            CreateGenericComposite (pmkAllButLastOfFirst,
                pmk, ppmkComposite)
        else
            *ppmkComposite = pmkAllButLastOfFirst
Case4:
    pmkLastOfFirst->ComposeWith (pmkFirstOfRest, TRUE,
        &pmk)
    if (no composition occurred)
        CCompositeMoniker::Create (pmkFirst, pmkRest,
            ppmkComposite)
    else
        if (pmk! = NULL)
            CreateGenericComposite
                (pmkAllButLastOfFirst, pmk, &pmk2)
            CreateGenericComposite (pmk2,
                pmkAllButFirstOfRest, ppmkComposite)
        else
            CreateGenericComposite
                (pmkAllButLastOfFirst,
                pmkAllButFirstOfRest, ppmkComposite)
}
    
```

CreateGenericComposite

HRESULT CreateGenericComposite (pmkFirst, pmkRest, ppmkComposite)

[0140] The function CreateGenericComposite allocates and returns a new generic composite moniker. The parameters pmkFirst and pmkRest point to the first and trailing monikers that are to comprise the generic composite monikers, respectively. Either pmkFirst or pmkRest may be a generic composite moniker, or another kind of moniker. The following table describes the parameters of the function CreateGenericComposite:

Argument	Type	Description
pmkFirst	IMoniker*	the first moniker in the new composite.
pmkRest	IMoniker*	the trailing (rest) moniker in the new composite.
ppmkComposite	IMoniker*	a pointer to the new composite.
return value	HRESULT	S_OK, E_OUTOFMEMORY

[0141] Code Table 4 contains C++ pseudocode for the function CreateGenericComposite. The function handles four specific cases. The first case occurs when neither the first moniker (pmkFirst) nor the rest moniker (pmkRest) are generic composite monikers. The second case occurs when the first moniker is not a generic composite moniker, but the rest moniker is a generic composite moniker. The third case occurs when the first moniker is a generic composite moniker, but the rest moniker is not a generic composite moniker. The fourth case occurs when both the first moniker and the rest moniker are generic composite monikers.

[0142] In the first case, the function CreateGenericComposite invokes the method ComposeWith of the first moniker passing the rest moniker and specifying that a composition

should occur only if not generic. If the rest moniker is not a special moniker for the first moniker, then no composition occurs and the function creates a composite moniker by invoking the method Create of the class CCompositeMoniker passing the first moniker and the rest moniker. The method Create of the class CCompositeMoniker creates a generic composite moniker that points to the specified monikers and returns a pointer to the created moniker. FIG. 18 illustrates the resulting generic composite moniker 1803 of the first case when the rest moniker 1802 is not a special moniker of the first moniker 1801. In the second case, the function CreateGenericComposite invokes the method ComposeWith of the first moniker passing the first component moniker of the rest moniker. If the first component moniker of the rest moniker is not a special moniker for the first moniker, then no composition occurs and the method creates a composite moniker by invoking the method Create of the class CCompositeMoniker passing the first moniker and the rest moniker. FIG. 19A illustrates a representative resulting generic composite moniker 1905 when the first component moniker 1903 of the rest moniker 1902 is not a special moniker of the first moniker 1901. If, however, a composition of the first moniker and the first component of the rest moniker occurs and a moniker is returned, then the function recursively calls the function CreateGenericComposite to compose the moniker returned with all but the first component moniker of the rest moniker. FIG. 19B illustrates the resulting generic composite moniker 1907 that contains the composite moniker 1906, which is the composition of the first moniker and the first component moniker of the rest moniker, and contains the moniker 1904, which includes all but the first component moniker of the rest moniker. If the composition of the first moniker and the first component moniker of the rest moniker resulted in an annihilation of the monikers (pmk==NULL), then the function returns a pointer to a moniker formed by all but the first component moniker of the rest moniker as illustrated in FIG. 19C. In the third case, the function CreateGenericComposite invokes the method ComposeWith of the last component moniker of the first moniker passing the rest moniker. If the rest moniker is not a special moniker for the last component moniker of the first moniker, then no composition occurs and the function creates a composite moniker by invoking the method Create of the class CCompositeMoniker passing the first moniker and the rest moniker as illustrated by the representative generic composite moniker in FIG. 20A. If, however, a composition occurs, and a moniker is returned, then the method recursively invokes the function CreateGenericComposite passing all but the last component moniker of the first component moniker and the returned moniker as indicated in FIG. 20B. If, however, the composition of the first moniker with the first component moniker of the rest moniker resulted in an annihilation of the monikers, then the function returns a moniker comprising all but the last component moniker of the first moniker as the composite moniker as illustrated by FIG. 20C. In the fourth case, the function CreateGenericComposite invokes the method ComposeWith of the last component moniker of the first moniker passing the first component moniker of the rest moniker. If the first component moniker of the rest moniker is not a special moniker for the last component moniker of the first moniker, then no composition occurs and the function invokes the method Create of the class CCompositeMoniker passing the first moniker and the rest moniker

resulting in the sample generic composite moniker of FIG. 21A. If, however, a composition does occur without annihilation, then the function recursively invokes the function CreateGenericComposite passing the composed moniker and all but the last component moniker of the first moniker. The function then recursively invokes the function CreateGenericComposite passing the resulting composite moniker and all but the first component moniker of the rest moniker resulting in the representative composite moniker of FIG. 21B. If the composition results in annihilation of the last component moniker of the first moniker and the first component moniker of the rest moniker, then the function recursively invokes the function CreateGenericComposite passing all but the last component moniker of the first moniker and all but the first component moniker of the rest moniker resulting in the representative generic composite moniker of FIG. 21C.

IMoniker::Reduce

```
HRESULT IMoniker::Reduce(pbc, dwReduceHowFar,
ppmkToLeft, ppmkReduced)
```

[0143] The method Reduce requests a moniker to re-write itself into another equivalent moniker. This method returns a new moniker that will bind to the same object, but does so in a more efficient way. This capability has several uses:

[0144] It enables the construction of user-defined macros or aliases as new kinds of moniker classes. When reduced, the moniker to which the macro evaluates is returned.

[0145] It enables the construction of a kind of moniker which tracks data as it moves about. When reduced, the moniker of the data in its current location is returned.

[0146] On certain file systems which support an ID-based method of accessing files that is independent of file names, a file moniker could be reduced to a moniker which contains one of these IDs.

[0147] FIG. 22 shows an example of moniker reduction. This example illustrates the reduction of a moniker which names the net income entry for this year's report in the "Projects" directory of the current user's home directory. (Note that the particular classes of monikers used here are for illustrative purposes only.) Several monikers in this example are reduced to something completely different, and some bind to something during their reduction, but some do not. For example, to reduce the alias "Home", the reduction must access the information that "Home" was an alias for "\\server\share\fred". Monikers may reduce to themselves, when they cannot be rewritten any further. A moniker which reduces to itself indicates this by returning itself through parameter ppmkReduced and the returning status code MK_S_REduced_TO_SELF. A moniker which reduces to nothing returns NULL in parameter ppmkReduced and the status code S_OK. If a moniker does not reduce to itself, then this method does not reduce this moniker in-place; instead, it returns a new moniker.

[0148] The reduction of a moniker which is a composite of other monikers repeatedly reduces the component monikers of which it is composed until they all reduce to themselves, and then returns the composite of the reduced components. The parameter dwReduceHowFar controls the stopping point of the reduction process. It controls to what extent the reduction should be carried out. It has the following values.

```
[0149] typedef enum tagMKRREDUCE {
    [0150] MKRREDUCE_ONE=3<<16,
    [0151] MKRREDUCE_TOUSER=2<<16,
    [0152] MKRREDUCE_THROUGUSER=1<<16,
    [0153] MKRREDUCE_ALL=0
    [0154] } MKRREDUCE;
```

[0155] These values have the following semantics.

Value	Description
MKRREDUCE_ONE	Perform only one step of reduction on this moniker. In general, the caller will have to have specific knowledge as to the particular kind of moniker in question in order to be able to usefully take advantage of this option.
MKRREDUCE_TOUSER	Reduce this moniker to the first point where it first is of the form where it represents something that the user conceptualizes as being the identity of a persistent object. For example, a file name would qualify, but a macro or an alias would not. If no such point exists, then this option should be treated as MKRREDUCE_ALL.
MKRREDUCE_THROUGUSER	Reduce this moniker to the point where any further reduction would reduce it to a form which the user does not conceptualize as being the identity of a persistent object. Often, this is the same stage as MKRREDUCE_TOUSER.
MKRREDUCE_ALL	Reduce the entire moniker, then, if needed reduce it again and again to the point where it reduces to simply itself.

[0156] The following table describes the parameters of the method Reduce:

Argument	Type	Description
pbcb	IBindCtx*	The bind context to use in this operation.
dwReduceHowFar	DWORD	Indicates to what degree this moniker should be reduced.
ppmkToLeft	IMoniker**	On entry, the moniker which is the prefix of this one in the composite in which it is found. On exit, the pointer is either NULL or non-NULL. Non-NULL indicates that what was previously thought of as the prefix should be disregarded and the moniker returned through ppmkToLeft considered the prefix in its place. NULL indicates that the prefix should not be so replaced. Thus, most monikers will NULL out this parameter before returning.
ppmkReduced	IMoniker**	On exit, the reduced form of this moniker. Possibly NULL.

-continued

Argument	Type	Description
return value	HRESULT	S_OK, MK_S_REDUCE_TO_SELF, MK_E_EXCEEDEDDEADLINE.

[0157] FIG. 23 is a flow diagram of the method Reduce of the class CCompositeMoniker. The method reduces each of the component monikers in a left-to-right manner and creates a composite of the result. If any of the component monikers do not reduce to themselves (and thus, the generic composite moniker overall does not reduce to itself), then the process of reduction is repeated. In an alternate embodiment, the method tracks component monikers that reduce to themselves and suppresses their re-reduction. In step 2301, if this moniker is already reduced as indicated by the data member m_fReduced, then the method continues at step 2302, else the method continues at step 2303. In step 2302, the method sets the pointer to the reduced moniker to point to this moniker and returns indicating that the moniker reduced to itself. In step 2303, if the left moniker is NULL, then the method continues at step 2306, else the method continues at step 2304. In step 2304, the method invokes the method Reduce of the left moniker passing the moniker to the left of this moniker and returning a left reduced moniker (pmkLeftReduced). In step 2305, if no error occurred or the left moniker reduced to itself, then the method continues at step 2306, else the method returns an error. If this moniker has no right moniker (m_pmkRight), then the method continues at step 2309, else the method continues at step 2307. In step 2307, the method invokes the method Reduce of the right moniker passing an indicator of a NULL moniker to the left and returning a right reduced moniker (pmkRightReduced). In step 2308, if no error occurred or the right moniker reduced to itself, then the method continues at step 2309, else the method returns an error. In step 2309, if both the left and right monikers reduced to themselves, then the method continues at step 2310, else the method continues at step 2312. In step 2310, the method sets the state of this moniker to reduced (m_fReduced). In step 2311, the method sets the pointer to the reduced moniker to point to this moniker and the method returns with an indication that this moniker reduced to itself. In step 2312, the method invokes the method Create of the class CCompositeMoniker passing the left reduced moniker and the right reduced moniker and returning the result as a composite reduced moniker (pmkCompositeReduced). In step 2313, the method sets the state of the composite reduced moniker to reduced (m_fReduced). In step 2314, the method sets the pointer to the reduced moniker to point to the composite reduced moniker and returns.

[0158] FIG. 24 is a flow diagram of the method Reduce of the sample class CAliasMoniker. The reduction of an alias moniker is illustrated in FIG. 22. In step 2401, if the alias moniker indicates the home directory, then the method continues at step 2402, else the method tests for other types of alias indicated by the ellipsis. In step 2402, the method retrieves the network volume associated with the user. In step 2403, the method retrieves the home directory for the user. In step 2404, the method creates a net volume moniker passing it the name of the network volume. In step 2405, the

method creates a file moniker passing it the name of the user's home directory. In step 2406, the method creates a generic composite moniker passing it the net volume moniker and the file moniker. In step 2407, the method sets the generic composite moniker to indicate that it is reduced. In step 2408, the method sets the pointer to the reduced moniker to the composite moniker (pmkCompositeReduced) and returns.

CODE TABLE 5

```

{
    A = CreateMoniker (cFileMoniker,
        "c:\reports\expenses\weekly")
    B = CreateMoniker (cFileMoniker,
        GetCurrentUserName)
    C = CreateMoniker (cFileMoniker, dayofweek
        (GetCurrentDate-oneDay))
    Result = AoBoC
}
    
```

[0159]

CODE TABLE 6

```

{
    A = CreateMoniker (cFileMoniker, "c:\taxes")
    Prompt "Enter year:", year
    B = CreateMoniker (cFileMoniker, year)
    C = CreateMoniker (cFileMoniker, "\1040.XLS")
    D = CreateMoniker (cItemMoniker, "R1C1:R10C10")
    Result AoBoCoD
}
    
```

[0160] In a preferred embodiment, a macro moniker allows for arbitrary moniker creation. A macro moniker contains a macro script that controls the reduction of a macro moniker to another moniker. During reduction, the script is parsed and processed by the method Reduce. One skilled in the art would appreciate that parsing and processing macro scripts are well known. The result of the processing is another moniker that is returned as the reduced moniker. For example, Code Table 5 contains a script that directs the macro moniker to reduce to a moniker referencing the directory "c:\reports\expenses\weekly\user\dayofweek", where user is the current user name (e.g., "Smith") and dayofweek is the day of week of yesterday (e.g., "Thursday"). The macro moniker with the script of Code Table 5 may reduce to a file moniker with a path name of "c:\reports\expenses\weekly\smith\thursday". The macro moniker may contain a pointer to the reduced moniker. The method BindToObject of a macro moniker would typically invoke the method Reduce and then invoke the method BindToObject of the reduced moniker. Code Table 6 contains a macro script that directs the macro moniker to reduce to a moniker and in the process prompts the user for a portion of the path.

CODE TABLE 7

```

{ SELECT FIRST printer.name
  FROM CampusPrinter
  WHERE (printerType == PostScript OR
        printerType == PCL)
  AND
    
```

CODE TABLE 7-continued

```

        printerLocation
        INCLUDES "Building 1"
    ORDER BY printQueueLength
}
    
```

[0161] In a preferred embodiment, a query moniker allows for arbitrary reduction to a moniker identified by a query. A query moniker contains a query that controls the reduction. The query is evaluated to produce a file moniker that satisfies the query. For example, Code Table 7 contains a query (in a structured query language) that may reduce to the file moniker with path name "\\printserver10\printer2". The query evaluates to a list of printers that can accommodate either PostScript or PCL documents and that is in a certain building. The list is sorted by the length of the print queue, and the printer with the shortest print queue is selected.

IMoniker::IsEqual

HRESULT IMoniker::IsEqual(pmkOtherMoniker)

[0162] This method determines whether this moniker and the specified other moniker (pmkOtherMoniker) reference the same object. This method is used in a preferred implementation of a running object table. The following table describes the parameters of the method IsEqual:

Argument	Type	Description
pmkOtherMoniker	IMoniker*	the other moniker with whom this moniker is compared.
return value	HRESULT	S_OK, S_FALSE

[0163] FIG. 25 is a flow diagram of the method IsEqual of the class CFileMoniker. In step 2501, if the other moniker (pmkOtherMoniker) is a file moniker, then the method continues at step 2502, else the monikers are not equal and the method returns a false. In step 2502, if the count of anti-monikers for this moniker (m_cAnti) is equal to the count of anti-monikers for the other moniker, then the method continues at step 2503, else the monikers are not equal and the method returns a false. In step 2503, if the path for this moniker (m_lpszPath) is equal to the path for the other moniker, then the method returns indicating that the monikers are equal, else the method returns indicating that the monikers are not equal.

[0164] FIG. 26 is a flow diagram of the method IsEqual of the class CCompositeMoniker. In step 2601, if the other moniker is a composite moniker, then the method continues at step 2602, the method returns indicating that the monikers are not equal. In step 2602, if the left moniker of this moniker (m_pmkLeft) is equal to the left moniker of the other moniker, then the method continues at step 2603, else the method returns an indication that the monikers are not equal. In step 2603, if the right moniker of this moniker (m_pmkRight) is equal the right moniker of the other moniker, then the method returns an indication that the monikers are equal, else the method returns an indication that the monikers are not equal. In an alternate embodiment of the method IsEqual, the method checks each component moniker of this moniker and the other moniker to determine if the monikers are equal.

IMoniker::Hash

HRESULT IMoniker::Hash(pdwHash)

[0165] This method returns a 32-bit integer associated with this moniker. This integer is used for maintaining tables of monikers: the moniker can be hashed to determine a hash bucket in the table, then compared with the method IsEqual against all the monikers presently in that hash bucket. Two monikers that compare as equal have the same hash value. The following table describes the parameters of the method Hash:

Argument	Type	Description
pdwHash	DWORD *	the place in which to put the returned hash value.
return value	HRESULT	S_OK

[0166] FIG. 27 is a flow diagram of the method Hash of the class CCompositeMoniker. In step 2701, the method invokes the method Hash of the left moniker. In step 2702, the method invokes the method hash of the right component moniker. In step 2703, the method generates the exclusive-or of the left hash value and the right hash value and returns that as the hash value of the method. The method Hash of the class CItemMoniker performs a hash function on the item name and returns the value.

IMoniker::Inverse

HRESULT IMoniker::Inverse(ppmk)

[0167] The method Inverse returns a moniker that when composed onto the end of this moniker or one of similar structure annihilates it; that is, composes to NULL. The method Inverse is an abstract generalization of the “..” operation in traditional file systems. For example, a file moniker that represents the path “a\b\c\d” has as its inverse a moniker containing the path “..\..\..\.”, since “a\b\c\d” composed with “..\..\..\.” yields nothing. The inverse of a moniker does not annihilate just that particular moniker, but all monikers with a similar structure. Thus, the inverse of a generic composite moniker is the reverse composite of the inverse of its component monikers. Certain classes of monikers may have trivial inverses. If a moniker adds one more component moniker to an existing structure; its inverse is a moniker that removes the last component of the existing structure. A moniker that when composed onto the end of a moniker removes the last component is referred to as anti moniker. One skilled in the art would appreciate that not all monikers have inverses. The inverse of an anti moniker, for example, does not exist. The following table describes the parameters of the method Inverse.

Argument	Type	Description
ppmk	IMoniker**	the place to return the inverse moniker.
return value	HRESULT	S_OK, MK_E_NOINVERSE.

[0168] An anti moniker is a moniker that when composed onto the end of a generic composite moniker removes the last component moniker. Composing an anti moniker onto

the end of another kind of moniker preferably annihilates the other moniker. The class CAntiMoniker contains a data member that is a count of the number of anti monikers (m_cAnti). Whenever an anti moniker is composed with another anti moniker, the resulting composition is an anti moniker with its count equal to the sum of the counts of the composed anti monikers.

[0169] FIGS. 28A, 28B, and 28C are block diagrams illustrating composition with inverse monikers. In FIG. 28A, an item moniker 2801 is composed with anti moniker 2802 which results in annihilation of the monikers. In FIG. 28B, the generic composite moniker 2803 which comprises component moniker 2804 and 2801 is composed with anti moniker 2802. Since anti moniker 2802 is the inverse of item moniker 2801, the item moniker 2801 and the anti moniker 2802 is annihilated. The result of the composition is file moniker 2805. In FIG. 28C, generic composite moniker 2806 is composed with generic composite moniker 2807 resulting in annihilation of generic composition monikers 2806 and 2807. Generic composite moniker 2806 comprises file moniker 2804 and item moniker 2801. Generic composite moniker 2807 comprises anti moniker 2802 and file moniker 2808. The anti moniker 2802 is the inverse of item moniker 2801, and file moniker 2808 is the inverse of file moniker 2804.

[0170] FIG. 29 is a flow diagram of the method Inverse of the class CCompositeMoniker. In step 2901, the method invokes the method Inverse of the left moniker of this moniker to retrieve its inverse moniker. In step 2902, the method invokes the method Inverse of the right moniker of this moniker to retrieve its inverse moniker. In step 2903, the method invokes the function CreateGenericComposite passing it the right inverse and the left inverse and returns. The left inverse moniker is composed onto the end of the right inverse moniker so that during composition the right moniker will compose with the right inverse and the left moniker will compose with the left inverse moniker resulting in annihilation.

[0171] FIG. 30 is a flow diagram of the method Inverse of the class CItemMoniker. Since the inverse of an item moniker is an anti moniker, the method creates an instance of an anti moniker and returns it.

[0172] FIG. 31 is a flow diagram of the method Compose-With of the class CItemMoniker. This method illustrates the annihilation of monikers. In step 3101, if the moniker on the right to be composed with is an anti moniker, then the method continues at step 3102, else the method continues at step 3103. In step 3102, the method invokes the method AnnihilateOne of the anti moniker and returns the result as the composite moniker. Consequently, an item moniker that is composed with an anti moniker with a count greater than one results in an anti moniker. In step 3103, if the caller wants a generic composition returned (fOnlyIfNotGeneric==FALSE), then the method continues at step 3104, else the method returns an indicator that a generic composite moniker is needed. In step 3104, the method invokes the function CreateGenericComposite passing this moniker and the right moniker and returns with the composite moniker.

[0173] FIG. 32 is a flow diagram of the method AnnihilateOne of the class CAntiMoniker. This method annihilates one of the counts of this moniker. In step 3201, if the count of the anti moniker is equal to one, then the method

continues at step 3202, else the method continues at step 3203. In step 3202, the method sets the moniker to return to NULL and returns. In step 3203, the method creates a new anti moniker and sets its count to the count of this moniker minus one and returns that newly-created moniker.

[0174] FIG. 33 is flow diagram of the method ComposeWith of the class CAntiMoniker. In step 3301, if the right moniker (pmkRight) is an anti moniker, then the method continues at step 3302, else the method continues at step 3304. In step 3302, the method calculates the sum of the count of this moniker and the count of the right moniker. In step 3303, the method creates an anti moniker and sets its count to the sum and returns it as the composite moniker. In step 3304, if the caller wants a generic composition, then the method continues at step 3305, else the method returns an indication that a generic composition is needed. In step 3305, the method invokes a function CreateGenericComposite passing it this moniker and the right moniker and returns the composite.

IMoniker::CommonPrefixWith

HRESULT IMoniker::CommonPrefixWith(pmkOther, ppmkPrefix)

[0175] This method determines the longest common prefix that this moniker shares with the other moniker (pmkOther). The following table describes the parameters of the method CommonPrefixWith.

Argument	Type	Description
pmkOther	IMoniker*	the moniker with whom the common prefix is to be determined.
ppmkPrefix	IMoniker*	the place to return the common prefix moniker. NULL is returned only in the case that the common prefix does not exist.
return value	HRESULT	MK_S_ME, indicating that this moniker is the common prefix. MK_S_HIM, indicating that the other moniker (pmkOther) is the common prefix. MK_S_US, indicating that the two monikers are equal. S_OK, indicating that the common prefix exists but is neither this moniker nor the other moniker. MK_S_NOPREFIX indicating that no common prefix exists.

[0176] FIG. 34 is a block diagram illustrating a common prefix of generic composite monikers. The generic composite moniker 3401 and 3402 represent generic composite monikers for whom a common prefix is to be determined. Generic composite moniker 3401 includes component item monikers A, B, C, and D. Generic composite moniker 3402 comprises component item monikers A, B, and D. The resulting moniker is a generic composite moniker 3403 with a component moniker for Item A and a component moniker for Item B. The common prefix of (A○B○C○D) and (A○B○D) is (A○B).

[0177] FIG. 35 is a flow diagram of the method CommonPrefixWith of the class CCompositeMoniker. In step 3501, if the other moniker (pmkOther) is a generic composite moniker, then the method continues at step 3505, else the

method continues at step 3502. In step 3502, the method invokes the method CommonPrefixWith of the first component moniker of this moniker passing the other moniker and returning the common prefix. In step 3503, if the first component moniker of this component moniker is a common prefix with the other moniker, then the method returns an indication that a common prefix exists but neither this moniker nor the other moniker is the common prefix, else the method continues at step 3504. In step 3504, if the method returns an indication that the first component moniker of this moniker and the other moniker are common prefixes, then the method returns an indication that the other moniker is a common prefix of this moniker, else the method returns with the indication returned in step 3502. In step 3505, the method invokes the method CommonPrefixWith for the first component moniker of this moniker passing the first component moniker of the other moniker and returning the result. If the first component moniker of this moniker and the first component moniker of the other moniker are common prefixes of each other, then the method continues at step 3507, else the method continues at step 3509. In step 3507, the method invokes the method CommonPrefixWith for a moniker comprising all but the first component moniker of this moniker and passing it all but the first component moniker of the other moniker. This recursively invokes the method CommonPrefixWith to determine the extent of the common prefix. In step 3508, the method composes the result of step 3505 and step 3507 and returns. In step 3509, if the first component moniker of the other moniker is a prefix of the first component moniker of this moniker, then the method continues at step 3510, else the method continues at step 3511. In step 3510, the method indicates that the prefix is the first component moniker of the other moniker and returns. In step 3511, if the first component moniker of this moniker is the prefix of the first component moniker of the other moniker, then the method continues at step 3512, else the method returns the result that was returned in step 3505. In step 3511, the method sets the prefix to return to this moniker and returns.

[0178] FIG. 36 is a flow diagram of the method CommonPrefixWith of the class CFileMoniker. The method scans the path names of this moniker and the other moniker to determine the common prefix. The count of anti monikers is assumed to precede the path names. In step 3601, if the other moniker (pmkOther) is a file moniker, then the method continues at step 3603, else the method continues at step 3602. In step 3602, the method invokes the function MonikerCommonPrefix (described below) and returns. In step 3603, if the count of anti monikers for this moniker (m_cAnti) is not equal to the count of anti monikers for the other moniker, then the method continues at step 3606, else the method continues at step 3604. In step 3604, the method scans the file path of the other moniker (m_lpszPath) and the file path of this moniker to determine the common prefix. In step 3605, the method sets the result and returns. In step 3606, the method determines the minimum of the anti moniker count of this moniker and the anti moniker count of the other moniker. In step 3607, the method creates a file moniker as the prefix moniker to return. In step 3608, the method sets the count of the anti monikers of the prefix to the minimum count of anti monikers. In step 3609, if the minimum count of anti monikers is in this moniker, then the method continues at step 3611, else the method continues at step 3610. In step 3610, if the path of the other moniker is

NULL, then the method returns an indication that the other moniker is a common prefix, else the method returns an indication that neither moniker is the common prefix. In step 3611, if the path of this moniker is NULL, then the method returns indicating that this moniker is a common prefix, else the method returns indicating that neither moniker is the common prefix.

[0179] FIG. 37 is a flow diagram of the method CommonPrefixWith of the class CItemMoniker. In step 3701, if this moniker is equal to the other moniker, then the method continues at step 3702, else the method returns an indication that there is no prefix in common. In step 3702, the method sets the prefix moniker equal to this moniker and returns an indication that both monikers are common prefixes.

[0180] FIG. 38 is a flow diagram of the method CommonPrefixWith of the class CAntiMoniker. In step 3801, if the other moniker is an anti moniker, then the method continues at step 3803, else the method continues at step 3802. In step 3802, the method invokes the function MonikerCommonPrefix and returns. If the count of this moniker is less than or equal to the count of the other moniker, then the method continues at step 3805, else the method continues at step 3804. In step 3804, the method sets the prefix moniker to point to the other moniker and returns an indication that the other moniker is the common prefix. In step 3805, the method sets the prefix moniker equal to this moniker. In step 3806, if the count of this moniker is equal to the count of the other moniker, then the method returns an indication that both monikers are common prefixes, else the method returns an indication that this moniker is a common prefix.

MonikerCommonPrefixWith

HRESULT MonikerCommonPrefixWith(pmkThis, pmkOther, ppmkPrefix)

[0181] This function is invoked by implementations of the method CommonPrefixWith. This function handles the situation when the implementation does not recognize the type of the other moniker. The following table describes the parameters of the function MonikerCommonPrefixWith:

Argument	Type	Description
pmkThis	IMoniker *	one moniker for the computation of the common prefix.
pmkOther	IMoniker *	the other moniker for the computation of the common prefix.
ppmkPrefix	IMoniker **	pointer to the common prefix.
return value	HRESULT	S_OK, MK_S_HIM, MK_S_ME, MK_S_US, MK_S_NOPREFIX

[0182] FIG. 39 is a flow diagram of the function MonikerCommonPrefixWith. In step 3901, if this moniker (pmkThis) is a generic composite moniker, then the method continues at step 3902, else the method continues at step 3903. In step 3902, the method invokes the method CommonPrefixWith of this moniker passing it the other moniker (pmkOther) and returns. In step 3903, if the other moniker is a generic composite, then the method continues at step 3905, else the method continues at step 3904. In step 3904,

if neither moniker is a generic composite moniker, then the method sets the prefix to NULL and returns an indication that there is no prefix in common. In step 3905, the method invokes the method CommonPrefixWith of the other moniker passing this moniker. In step 3906, if this moniker is the common prefix, then the method returns with an indication, else the method continues at step 3907. In step 3907, if the other Moniker is the common prefix, then the method returns with an indication, else the method returns with the indication returned in 3905.

IMoniker::RelativePathTo

HRESULT IMoniker::RelativePathTo(pmkOther, ppmkRelPath)

[0183] This method returns as a moniker that when composed onto the end of this moniker or one with a similar structure yields the other moniker (pmkOther). Implementations of this method preferably determine the longest prefix that this moniker and the other moniker have in common. This separates this moniker and the other into two parts, say (P, T_{this}) and (P, T_{other}) respectively, where P is the common prefix. T_{this} and T_{other} represent the trailing components. The relative path result is then T⁻¹_{this} ◦ T_{other}, where T⁻¹ indicates the inverse. Thus, (P, T_{this}) ◦ (T⁻¹_{this} ◦ T_{other}) = (P, T_{other}).

Argument	Type	Description
pmkOther	IMoniker*	the other moniker to which a relative path should be taken.
ppmkRelPath	IMoniker*	May not be NULL. The place at which the relative path is returned.
return value	HRESULT	MK_S_HIM, indicating that the only form of relative path is in fact just the other moniker (pmkOther). S_OK, indicating that a non-trivial relative path exists.

[0184] FIG. 40 is a block diagram illustrating a relative path to moniker. When generic composite moniker 4001 is composed with relative path to moniker 4002, the result is generic composite moniker 4003. During composition, the component moniker Item C of generic composite moniker 4001 and the anti moniker of generic composite moniker 4002 annihilate. The generic composite moniker 4002 is the relative path to moniker to go from the composite generic moniker 4001 to 4003. More precisely, the relative path to moniker is the inverse of the portion that is not the common prefix of generic composite monikers 4001 and 4003 composed with the portion of generic composite moniker 4003 that is not the common prefix.

[0185] FIGS. 41A and 41B are block diagrams illustrating a usage of the method RelativePathTo. The document 4101 has a path name of "c:\reports\monthly\june.doc". The document 4101 contains a link to a range within the spreadsheet document 4102. The spreadsheet document has the path "c:\reports\monthly\data\june.xls". The range is identified by "R4C7:R8C12". The document 4101 preferably stores a moniker indicating the absolute path name of the range and the relative path name of the range. In this way, when the document 4101 is moved to a new directory with a similar structure, the correct range can be retrieved from that directory structure using the relative path. FIG. 41B illus-

trates a generic composite moniker **4103** describing the absolute path to the spreadsheet range. The generic composite moniker **4103** includes a file moniker **4104** with the absolute path name of the spreadsheet file and an item moniker **4105** which describes the spreadsheet range. The relative path to moniker is generic composite moniker **4106**. Generic composite moniker **4106** includes file monikers **4107** and **4108**, and item moniker **4109**. The generic composite moniker **4106** is the relative path to moniker for the absolute path name of the document **4101**. When the moniker for the absolute path name of the document **4101** is composed with the generic composite moniker **4106**, the result is the proper composite moniker needed to access the spreadsheet file.

[0186] FIGS. 42A, 42B, and 42C comprise a flow diagram of the method RelativePathTo of the class CCompositeMoniker. In step **4201**, if the other moniker (pmkOther) is a generic composite moniker, then the method continues at step **4202**, else the method continues at step **4227**. In step **4202** through **4207**, the method loops determining the common prefix of this moniker and the other moniker. In step **4202**, the method sets a variable nCount equal to 0. The variable nCount keeps track of the number of component monikers in the common prefix. In step **4203**, the method selects the next component moniker of the other moniker, starting with the first. The enumeration of generic composite monikers is described below. In step **4204**, the method selects the next component moniker of this moniker, starting with the first. In step **4205**, if a component moniker of the other moniker and a component moniker of this moniker are selected, then the method continues at step **4206**, else the method continues at step **4208**. In step **4206**, if the selected component monikers are equal, then the method continues at step **4207**, else the method continues at step **4208**. In step **4207**, the method increments the variable nCount and loops to step **4203** to select the next component monikers. In step **4208**, if the variable nCount equals 0, then this moniker and the other moniker have no common prefix and the method continues at step **4209**, else the method continues at step **4214**. In step **4209**, the method invokes the method RelativePathTo of the first component moniker of this moniker passing it the first component moniker of the other moniker and indicating to return the result as the inverse of the first component of this moniker. In step **4210**, if the invocation of the method RelativePathTo in step **4209** indicates that a non-trivial relative path exists, then the method continues at step **4211**, else the method continues at step **4213**. In step **4211**, the method indicates that the first component moniker of this moniker is equal to the result of step **4209**. In step **4212**, the method selects the next component moniker of the other moniker and continues at step **4216**. In step **4213**, the method sets the relative path to point to the other moniker and returns an indication that the relative path is the other moniker. In step **4214**, if not all the component monikers of the this moniker have been selected, then the method continues at step **4215**, else the method continues at step **4216**. In step **4215**, the method determines the inverse of the selected component moniker of this moniker and continues at step **4216**. In step **4216**, if the selected component moniker of this moniker is not equal to NULL, then the method continues at step **4217**, else the method continues at step **4221**. In step **4217**, the method determines the inverse of the component monikers of this moniker that have not been selected. In step **4218**, if the inverse is not NULL, then

the method continues at step **4219**, else the method continues at step **4220**. In step **4219**, the method invokes the method ComposeWith of the inverse passing it the inverse of the last selected component moniker of this moniker to create an inverse for the tale portion of this moniker and the method continues at step **4221**. In step **4220**, the method sets the inverse of the tale portion of this moniker equal to the inverse of the last selected component moniker of this moniker and continues at step **4221**. In step **4221**, if the selected component moniker of the other moniker is not equal to NULL, then the method continues at step **4223**, else the method continues at **4222**. In step **4222**, the method sets the tale of the other moniker equal to NULL and continues at step **4224**. In step **4223**, the method composes the selected component moniker of the other moniker with tale of the other moniker and continues at step **4224**. In step **4224**, if the inverse of the tale of this moniker is not equal to NULL, then the method continues at step **4226**, else the method continues at step **4225**. In step **4225**, the method sets the relative path to the tale of the other moniker and returns. In step **4226**, the method composes with the inverse of the tale of this moniker with the tale of the other moniker and returns that as a relative path. In step **4227**, the method selects the next component moniker of this moniker. In step **4228**, if the selected component moniker of this moniker is equal to the other moniker, then the method continues at step **4229**, else the method continues at step **4230**. In step **4229**, the method gets the inverse of this moniker and returns it as the relative moniker. In step **4230**, the method invokes a method RelativePathTo of the selected component moniker of this moniker passing it the other moniker. In step **4232**, if the method invoked in step **4230** indicates that the other moniker is the relative path to moniker, then the method continues at step **4232**, else the method continues at step **4233**. In step **4232**, the method sets the relative path to moniker to point to the other moniker and returns. In step **4233**, the method gets the inverse of the tale of this moniker. In step **4234**, the method composes with the inverse of the tale of this moniker with the relative path to moniker returned in step **4230** and returns.

MonikerRelativePathTo

HRESULT MonikerRelativePathTo(pmSrc, pmkDest, ppmkRelPath, reserved)

[0187] This function is invoked by implementations of the method RelativePathTo. This method handles the situation when the implementation does not recognize the type of the other moniker.

Argument	Type	Description
pmkSrc	IMoniker *	the starting moniker for the computation of the relative path.
pmkDest	IMoniker *	the moniker to which a relative path should be taken.
ppmkRelPath	IMoniker **	May not be NULL. The place at which the moniker of pmkDest relative to pmkSrc is to be returned.
reserved	BOOL	must be non-zero
return value	HRESULT	S_OK, MK_S_HIM

[0188] FIG. 43 is a flow diagram of the function MonikerRelativePathTo. In step **4301** if the source moniker or the

destination moniker is a generic composite moniker, then the method continues at step 4304, else the method continues at step 4302. In step 4302, if the parameter fCalledFromMethod is true, then the method continues at step 4303, else the method continues at step 4305. In step 4303, the method sets the relative path to the destination moniker and returns an indication moniker that the destination moniker is the relative path to. In step 4304, if the source moniker is a generic composite moniker, then the method continues at step 4305, else the method continues at step 4306. In step 4305, the method invokes the method RelativePathTo of the source moniker passing the destination moniker and returns the relative path from that invocation. In step 4306, if the source moniker is equal to the first component moniker of the destination moniker, then the method continues at step 4307, else the method continues at step 4308. In step 4307, the method sets the relative path to moniker equal to a moniker comprising all but the first component moniker of the destination moniker and returns. In step 4308, the method invokes the method RelativePathTo of the source moniker passing the first component moniker of the destination moniker. In step 4309, if no error is returned, then the method continues at step 4310, else the method continues at step 4311. In step 4310, the method invokes the function CreateGenericComposite passing the relative path to moniker returned in step 4308 and a moniker comprising all but the first component moniker of the destination moniker and returns. In step 4311, the method sets the relative path to moniker equal to the destination moniker and returns an indication that the destination moniker a relative path to moniker.

IMoniker::Enum

HRESULT IMoniker::Enum(fForward, ppenmMoniker)

[0189] This method returns an interface that permits the enumeration of the component monikers of which this moniker is logically a composite. For a generic composite moniker, this enumerates the components of which the composite is composed. For other monikers, the semantics of the components of which it is a composite are implementation-defined. For example, enumerating the components of a file moniker returns each portion of the internally stored path name, even though they are not stored internally as actual separate monikers. Monikers that have no discernible internal structure return NULL instead of an enumerator.

[0190] The IEnumMoniker interface is an enumerator that supports the enumeration of items which are monikers.

```
interface IEnumMoniker : IUnknown {
    virtual HRESULT Next(ULONG celt, IMoniker* rgelt[ ],
        ULONG* pceltFetched) = 0;
    virtual HRESULT Skip(ULONG celt) = 0;
    virtual HRESULT Reset() = 0;
    virtual HRESULT Clone(IEnumMoniker** ppenm) = 0;
};
```

[0191] The following table describes the parameters of the method Enum:

Argument	Type	Description
fForward	BOOL	If true, then the enumeration should be done in the normal order. If false, then the order should be the reverse of the order enumerated by the normal order.
ppenmMoniker	IEnumMoniker**	On exit, the returned enumerator. May be NULL, signifying that there is nothing to enumerate.
return value	HRESULT	S_OK.

[0192]

CODE TABLE 8

```
class CCompositeMonikerEnum: IEnum
{
    CCompositeMoniker FAR * m_pcm;
    BOOL m_fForward;
    se m_pBase;
    se m_pTop;
    LPMONIKER m_pNext;
}
```

[0193] FIG. 44 is a flow diagram of the method Enum of the class CCompositeMoniker. In step 4401, the method instantiates an enumerator object of type CCompositeMonikerEnum for this composite object. Code Table 44 illustrates the data members of a class definition for the enumerator. The object contains a pointer to the moniker being enumerated (m_pcm), a flag indicating whether the enumeration is in the forward direction (m_fForward), pointers indicating the base and the top of a stack (m_pBase, m_pTop), and a pointer to the next component moniker to be returned (m_pNext). In steps 4402 through 4406, the method initializes the data members of the instantiated object. In step 4406, the method invokes the method GetNext passing this moniker and sets data member m_pNext to the result.

[0194] FIG. 45 is a flow diagram of the private method GetNext of the class CCompositeMonikerEnum. The enumeration of a composite moniker is a left-to-right, depth-first traversal of the tree formed by the composite moniker. The enumerator maintains a stack to track the traversal of the tree. In step 4501, the method sets a pointer to a rover (pmkRover) equal to the passed moniker. The rover pointer is used to traverse the tree. In step 4502, if the passed moniker is NULL, then the method returns NULL, else the method continues at step 4503. In step 4503, if the moniker pointed to by the rover is a generic composite moniker, then the method continues at step 4504, else a component moniker is located and the method returns the moniker pointed to by the rover. In step 4504, the method invokes the method Push to push the generic composite moniker indicated by the rover onto the stack. In step 4505, if the enumeration is being performed in the forward direction, the method continues at step 4506, else the method continues at step 4507. In step 4506, the method sets the rover to point to the left moniker of the moniker pointed to by the rover and loops to step 4503. In step 4507, the method sets the rover to point to the right moniker of the moniker pointed to by the rover and loops to step 4503.

[0195] FIG. 46 is a flow diagram of the method Next of the class CCompositeMonikerEnum. Each time this method is invoked it returns the next component moniker of the composite moniker. The method Next returns an array of component monikers up to a maximum number specified as a parameter. In step 4601, the method initializes a count variable to 0. In step 4602, if the count is less than the requested count of component monikers and not all the component monikers have been returned, then the method continues at step 4603, else the method continues at 4607. In step 4603, the method stores the next component moniker in the return array. In steps 4604 and 4605, the method increments the count and index into the return array. In step 4606, the method sets the next component moniker to point to the moniker returned by the method Pop and loops to step 4602. In step 4607, if a pointer to the count of fetched elements is not equal to NULL, then the method sets the pointer equal to the count in step 4608. The method then returns.

[0196] The methods Push and Pop of the class CCompositeMonikerEnum implement a stack. The method Push pushes the passed generic composite moniker onto a stack. The method Pop removes a top generic composite moniker from the stack and invokes the method GetNext passing it the right moniker if a search is being performed in the forward direction and a left moniker if the search is being performed in the reverse direction. The method Pop returns the component moniker returned by the method GetNext.

Pointer Moniker Class

[0197] In a preferred embodiment, a pointer moniker is a class of moniker that wraps an existing pointer to an object in a moniker so that it may participate as a component moniker in the moniker binding process. A pointer is a reference into "active space," that is, memory of a process. A moniker typically is a reference into "passive space," that is, the representation of an object on disk. Pointer monikers provide a means by which a given use of a moniker can transparently reference either active or passive space.

[0198] In a preferred embodiment, the method BindToObject of a pointer moniker invokes the method QueryInterface of the pointed to object. The method BindToStorage returns MK_E_NOSTORAGE. The method Reduce reduces the moniker to itself. The method ComposeWith does a generic composition. The method Enum returns NULL. The method IsSystemMoniker returns MKSYS_NONE. The method IsEqual uses the identity test paradigm on pointers after first checking that the other moniker for the right class. The method Hash returns a constant. The method GetTimeOfLastChange returns MK_E_UNAVAILABLE. The method Inverse returns an anti moniker. The method RelativePathTo returns the other moniker. The method GetDisplayName returns NULL. The method ParseDisplayName binds to the punk pointer using IParseDisplayName interface. Pointer monikers do not serialize; that is, the Save of the IPersistStream interface returns an error.

CreatePointerMoniker

HRESULT CreatePointerMoniker(punk, ppmk)

[0199] This function wraps a pointer in a pointer moniker so that it can be presented to interfaces that require monikers for generality, but specific uses of which can usefully deal with a moniker which cannot be saved to backing store. The

following table describes the parameters of the method CreatePointerMoniker.

Argument	Type	Description
punk	IUnknown*	the pointer that we are wrapping up in a moniker.
ppmk	IMoniker**	the returned Pointer Moniker.
return value	HRESULT	S_OK, E_OUTOFMEMORY

[0200] FIG. 47 is a flow diagram of the method Create of the class CPointerMoniker. In step 4701, the method instantiates a new pointer moniker. In step 4702, the method sets a data member (m_Punk) the new pointer moniker to point to the passed object and returns.

[0201] FIG. 48 is a flow diagram of the method BindToObject of the class CPointerMoniker. In step 4801, the method invokes the method QueryInterface of the pointed to object passing the identifier of the requested interface. The method then returns with the requested interface.

[0202] FIG. 49 is a diagram illustrating a pointer moniker. Object 4901 is wrapped in pointer moniker 4902. Pointer moniker 4902 contains a pointer to the object 4901.

[0203] Although the present invention has been described in terms of preferred embodiments, it is not intended that the invention be limited to these embodiments. Modifications within the spirit of the invention will be apparent to those skilled in the art. The scope of the present invention is defined by the claims which follow.

1.-13. (canceled)

14. A computer-readable medium storing computer-executable instructions for causing a computer programmed thereby to perform a method of evaluating a moniker object, the method comprising:

receiving an invocation of a hashing function member of a first moniker object, the first moniker object including naming information for source data; and

returning a hash value for the first moniker object.

15. The computer-readable medium of claim 14, wherein the method further comprises

for each of one or more other moniker objects having the hash value,

receiving an invocation a comparison function member of the first moniker object, wherein the invocation includes passing a reference to the other moniker object; and

returning an indicator whether the first moniker object is equal to the other moniker object.

16. The computer-readable medium of claim 15 wherein a running object table includes an entry for each of the one or more other moniker objects.

17. The computer-readable medium of claim 15, wherein the method further comprises:

if none of one or more other moniker objects have the hash value, returning an indicator that the first moniker object is equal to none of the one or more other moniker objects.

18. A computer-readable medium storing computer-executable instructions for causing a computer programmed thereby to perform a method of providing a client with access to a resource identified by a server moniker object, the server moniker object having state information identifying the resource, the method comprising:

under control of a client,

receiving from a server an indication of unmarshalling code and state information for use by the unmarshalling code, the state information identifying a resource; and

executing the unmarshalling code to instantiate a client moniker object as a proxy for the server moniker object, the client moniker object including the state information.

19. The computer-readable medium of claim 18 wherein the method further comprises:

binding to the resource by invoking binding code of the client moniker object, wherein the binding code uses the state information to bind to the resource without accessing a server moniker object.

20. The computer-readable medium of claim 18 wherein the state information is immutable.

21. A computer-readable medium storing computer-executable instructions for causing a computer programmed thereby to perform a method of providing a client with access to a resource identified by a server moniker object,

the server moniker object having state information identifying the resource, the method comprising:

receiving an indication of unmarshalling code and state information for use by the unmarshalling code, the state information identifying a resource; and

passing the indication of the unmarshalling code and the state information to the client for use in instantiating a client moniker object as a proxy for the server moniker object, the client moniker object including the state information.

22. The computer-readable medium of claim 21 wherein the receiving the indication of unmarshalling code comprises getting class identifying information for the proxy.

23. The computer-readable medium of claim 21 wherein the passing comprises passing class identifying information for the proxy and a marshalling packet to the client.

24. The computer-readable medium of claim 23 wherein the method further comprises:

creating the marshalling packet; and

buffering in the marshalling packet information used in instantiating the client moniker object as the proxy for the server moniker object.

25. The computer-readable medium of claim 21 wherein the state information is immutable.

* * * * *