



(19) **United States**

(12) **Patent Application Publication**

Benson

(10) **Pub. No.: US 2004/0103197 A1**

(43) **Pub. Date: May 27, 2004**

(54) **USE OF BROWSER COOKIES TO STORE STRUCTURED DATA**

(60) Provisional application No. 60/118,266, filed on Feb. 2, 1999.

(76) Inventor: **Eric A. Benson**, Seattle, WA (US)

Publication Classification

Correspondence Address:
KNOBBE MARTENS OLSON & BEAR LLP
2040 MAIN STREET
FOURTEENTH FLOOR
IRVINE, CA 92614 (US)

(51) **Int. Cl.⁷** **G06F 15/16**
(52) **U.S. Cl.** **709/227; 719/311**

(21) Appl. No.: **10/694,509**

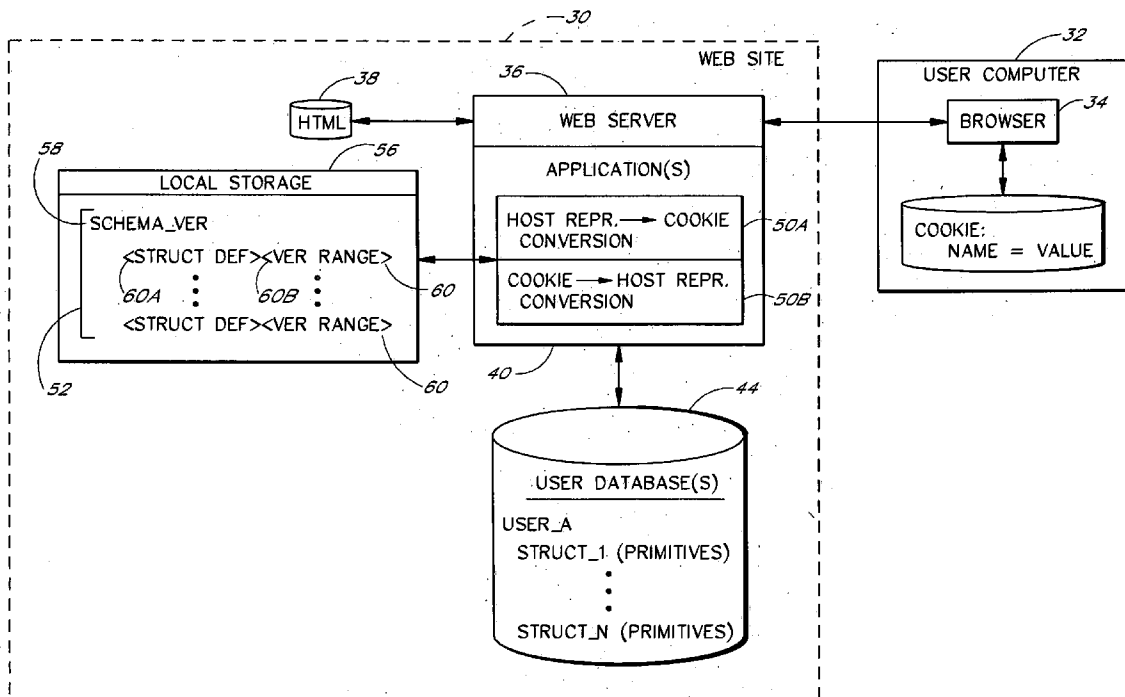
(22) Filed: **Oct. 27, 2003**


Related U.S. Application Data

(63) Continuation of application No. 09/494,712, filed on Jan. 31, 2000, now Pat. No. 6,714,926.

(57) **ABSTRACT**

A web site system implements a process for storing selected data structures within browser cookies. The data structures may contain a variety of different types of data elements, including N-bit integers and other non-character elements. A version tracking scheme provides forward and backward compatibility between client and server software. The process is implemented without the need for any browser extensions, and without the need for users to download any special code to their computers.





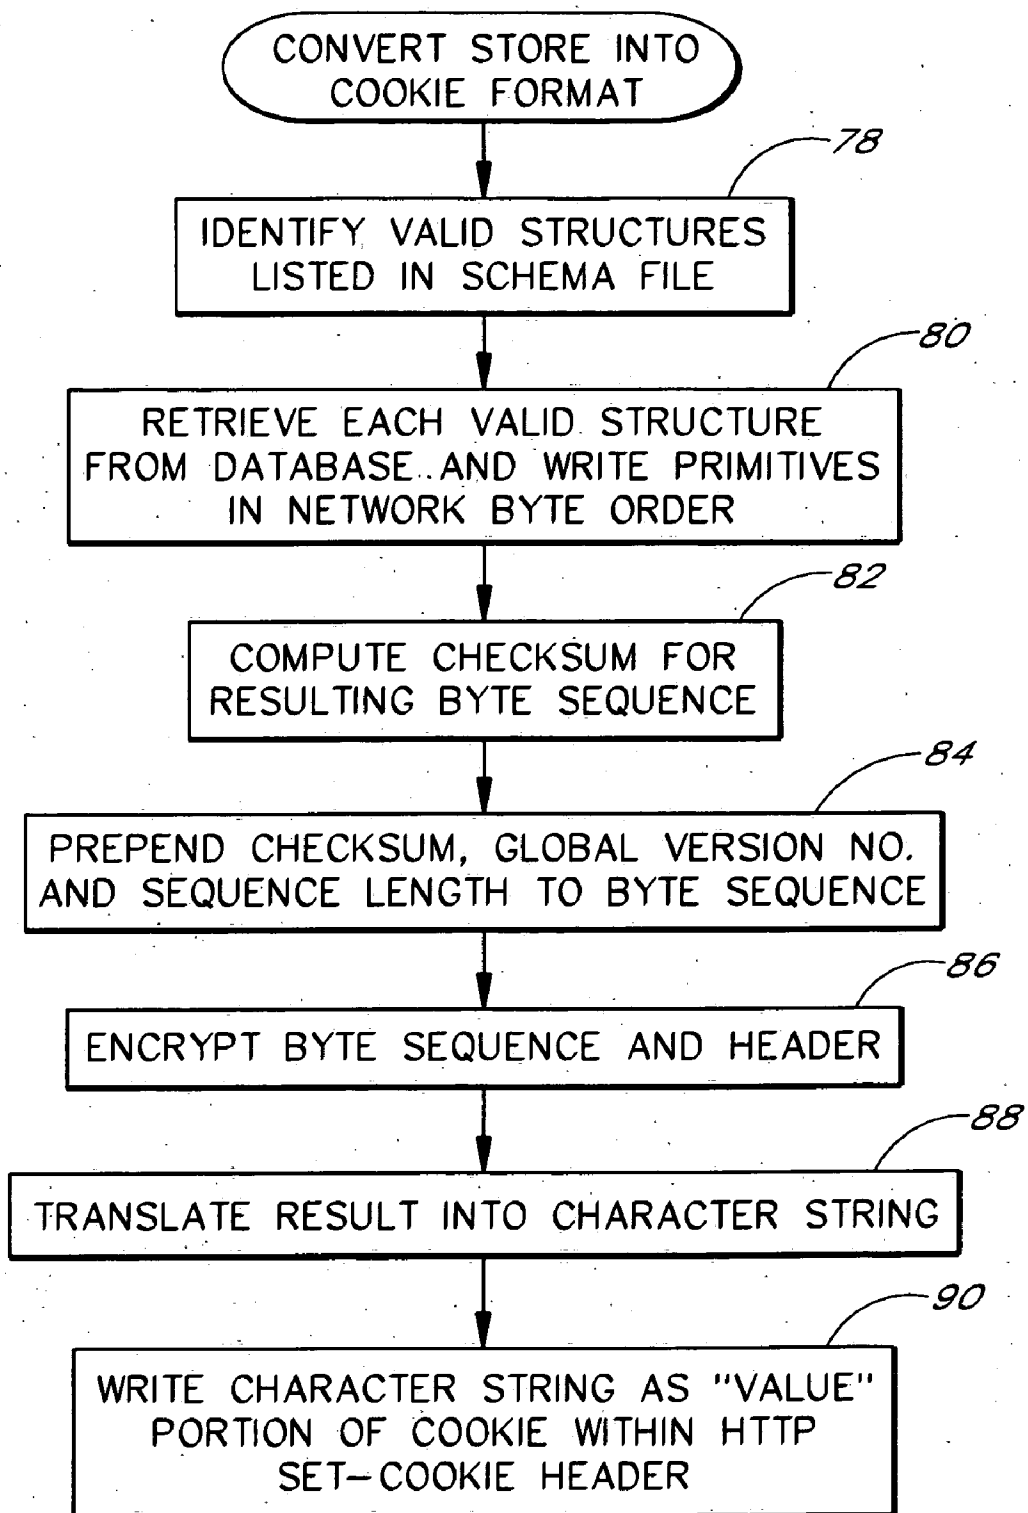


FIG. 2

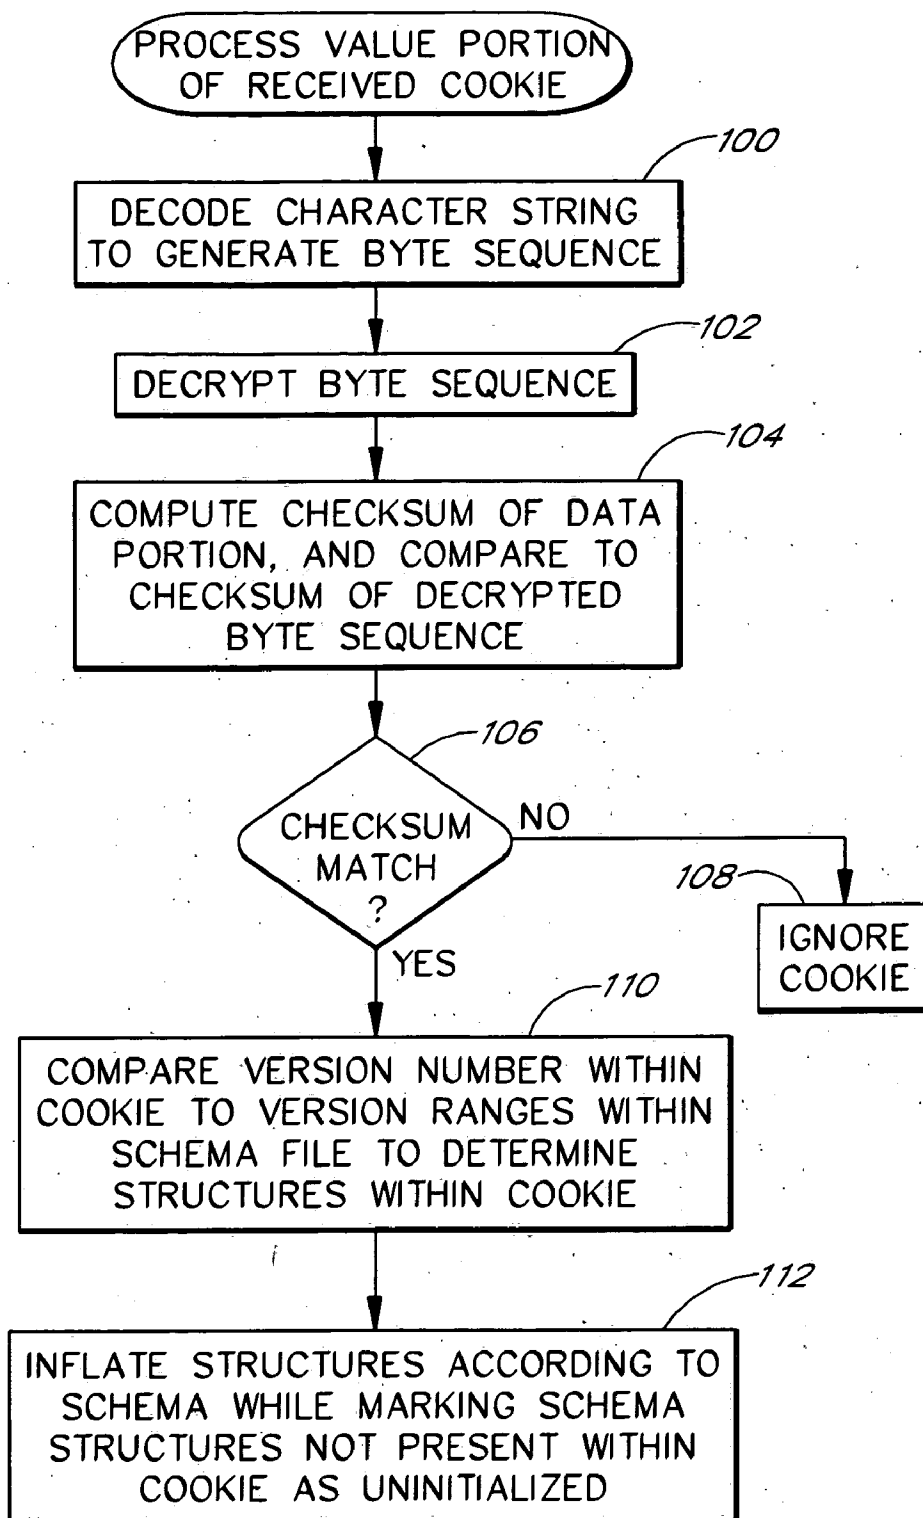


FIG. 3

USE OF BROWSER COOKIES TO STORE STRUCTURED DATA

RELATED APPLICATIONS

[0001] This application is a continuation of U.S. patent application Ser. No. 09/494,712, filed Jan. 31, 2000, which claims the benefit of U.S. provisional application No. 60/118,266, filed Feb. 2, 1999.

FIELD OF THE INVENTION

[0002] The present invention relates to Web site customization using cookies, and more particularly, to a method of extending the functionality of cookies to increase Web site performance.

BACKGROUND

[0003] Commercially available Web browsers such as Internet Explorer and Netscape Navigator support a mechanism by which a Web server can store persistent information on a user's machine for subsequent retrieval. This information is commonly referred to as a "cookie," and typically includes information about the user's identity. The cookie is stored on the user's computer as a name-value pair (NAME=VALUE), together with a URL (or range of URLs) to which the cookie corresponds and an optional expiration date. When the user initiates a request for a URL that corresponds to the cookie, the browser automatically sends the cookie (and any other cookies that correspond to the requested URL) to the host Web server with the URL request.

[0004] One common application for cookies involves dynamically customizing Web pages and functions for known users of a Web site. Typically, this involves storing known information about the user, such as a user ID, on the user's computer in response to a page request. When the user subsequently accesses the Web site, the information contained within the cookie is used to access a back-end database to retrieve additional information about the user, such as the user's preferences or account information. This database information may then be used to customize the requested Web page.

[0005] One problem with the above approach is that it requires frequent accesses to the database. For Web sites that experience many thousands of hits per day, the need to access the database can produce a significant performance degradation. The performance degradation may be the result of a limited load capacity of the database system, increased network traffic between physical Web servers and the database system, or both.

SUMMARY

[0006] One potential solution to the above problem would be to store within cookies all of the user information needed to customize the Web pages, or at least the most frequently accessed Web pages. This would allow accesses to back-end databases during page requests to be reduced or avoided. Unfortunately, the existing cookies specification does not provide a mechanism for storing non-character data. Thus, for example, where the database information used to customize Web pages includes a list or table of 16-bit integers, storing such data within browser cookies is not a viable option. Further, the existing cookies specification does not

provide a mechanism for allowing server software to keep track of versions of data elements contained within cookies.

[0007] The present invention overcomes this problem by providing a server system and process for storing selected data structures within browser cookies, preferably using a version tracking scheme to provide forward and backward compatibility between client and server software. The data structures may contain a variety of different types of data elements, including N-bit integers and other non-character elements. The process is implemented without the need for any browser extensions, and without the need for users to download any special code to their computers.

[0008] In a preferred embodiment, the system uses a schema file or other data structure that specifies past and present schemas for encoding data structures within cookies. A first server component uses the schema file to encode data structures into character strings that are stored within cookies on user computers. A checksum value, a length field, and a schema version number are also preferably encoded within the cookies. A second server component receives cookies from user computers, uses the schema file to identify the types of data structures stored in such cookies, and reproduces the data structures in local memory for temporary use. Application software uses these temporary data structures to generate personalized web pages and/or perform other custom operations.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 illustrates a Web site system according to a preferred embodiment of the invention.

[0010] FIG. 2 illustrates a preferred sequence of steps that are implemented by the host-to-cookie module of FIG. 1 to translate a sequence of structures to a cookie format.

[0011] FIG. 3 illustrates a preferred sequence of steps that are implemented by the cookie-to-host conversion module of FIG. 1 to convert a character string back to the original host representation.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

[0012] A preferred embodiment of the process and associated components will now be described with reference to FIGS. 1-3. Throughout these figures and the associated description, reference will be made to various implementation-specific details such as specific data types, algorithms, rules and conventions. These and other details are provided in order to illustrate an example implementation, and not to limit the scope of the invention.

[0013] FIG. 1 illustrates a Web site system ("Web site") 30 according to a preferred embodiment of the invention, and illustrates a representative user computer 32 that communicates with the Web site via the Internet. For purposes of illustration, it may be assumed that the Web site 30 is the Amazon.com Web site. The user computer includes a Web browser 34, such as Microsoft Internet Explorer 4.0, that supports cookies.

[0014] The Web site 30 includes a Web server 36 which accesses a store 38 of HTML (Hypertext Markup Language) documents. The Web site also includes one or more applications 40 that are used to dynamically generate customized

Web pages for users. The Web server **36** and the applications **40** may, for example, run on one or more physical servers of the Web site.

[0015] As depicted by **FIG. 1**, the applications **40** access one or more databases or other repositories **44** (collectively “user database”) that contain information about known users of the site. In the context of the Amazon.com site, this information may include, for example, user names and identifiers, payment and shipping information, shopping preference information (e.g., 1-Click settings), shopping cart contents, group discount information, product recommendations, user purchase histories, and service data (e.g., BookMatcher ratings). Different items within the user database **44** may be used by the applications **40** to customize different Web pages or functions.

[0016] The information stored in the user database **44** for each user is in the form of multiple structures, **STRUCT_1** through **STRUCT_N**. The structures may include, for example, lists, tables, arrays, variables, records, and other types of elements. Each structure includes one or more primitives, where a primitive is a data element that does not include a reference to another data element. For purposes of illustration, it may be assumed that some of these primitives are in the form of non-character elements (such as 16-bit and 32-bit integers) that are used to customize frequently-accessed Web pages.

[0017] Previously, the process of dynamically customizing Web pages for users has required the applications **40** to frequently access the user database **44**. For example, when a known user would request the site’s home page, information contained within the cookie (or cookies) transmitted with the URL request would be used to look up in the database **44** the user’s name and certain preference information; this information would then be used to generate the home page. As described above, such frequent accesses to back-end databases can significantly degrade performance as seen by end users.

[0018] In accordance with the invention, the following additional components are provided to reduce the need for accesses to the user database **44**. These components include a host representation to cookie conversion module **50A** (“host-to-cookie module”), a cookie to host representation conversion module **50B** (“cookie-to-host module”), and a schema file **52**. The conversion modules **50A**, **50B** may be in the form of executable functions that can be called by the applications.

[0019] The schema file **52**, which may alternatively be in the form of a table or other data structure, specifies the user database structures that have been selected (typically by application developers) to be included within cookies. (The term “schema,” as used herein, refers generally to the pattern or arrangement of a collection of data elements.) Typically, the structures that are included are those which contain relatively static user data that is used to customize frequently-accessed Web pages. As described below, information about structures that are no longer in use (due to application design changes) is retained within the schema file to support backward and forward compatibility. The schema file **52** thus acts as a specification of past and present schemas used to generate cookies.

[0020] The schema file **52** is preferably stored in local storage **56** of the physical server(s) that run the applications

40. For example, where the applications are replicated across multiple physical servers, the schema file **52** may be stored on the hard drive of each such server, and may be cached in the server’s RAM to increase performance.

[0021] In operation, the conversion modules **50A**, **50B** perform the task of converting selected structures of the user database **44** (preferably those used for the customization of popular Web pages) between the host representation and a cookie representation. During the cookie generation process, the host-to-cookie conversion module **50A** is invoked to convert selected database structures for a particular user into a cookie to be stored on the user’s computer; during this process, the schema file **52** is accessed to determine which of the specific structures (or primitives thereof) are to be included within the cookie.

[0022] When the cookie is subsequently returned with a URL request, the cookie-to-host module **50B** is invoked to convert the cookie back to the host representation used by the applications **40**. During the decoding process, the schema file **52** is used to determine the content and format of the embedded structure data. In addition, the schema file **52** is used to determine whether any of the structures contained within the cookie is no longer valid; if any such structure exists, it need not be fully decoded. The structures that are extracted from the cookie are stored in local storage **56** while they are used by the applications **44**, and are then deleted. Typically, the extracted information is used to customize one or more Web pages, but the information could also or alternatively be used to perform other types of user-specific functions (e.g., authentication).

[0023] In one implementation of the disclosed system, structures that use primitive datatypes of 8, 16, and 32 bit integers and null-terminated ASCII strings are supported. In other implementations of the disclosed design, more complex structures such as nested or cross-referenced structures could be stored.

[0024] The content of the schema file **52** will now be described in further detail with reference to **FIG. 1**. As illustrated, the schema file **52** includes a global schema version number **58** (**SCHEMA_VER**) which, as described below, is incremented when certain types of changes are made to the schema. The schema file **52** also includes a sequence of structure entries **60**, each of which corresponds to a particular database structure that has been selected for inclusion. Each entry **60** includes a structure definition portion **60A** and a valid version range portion **60B**. The structure definition portion **60A** contains the metadata for the corresponding structure, including a list of the primitive fields of the structure and their datatypes and address offsets.

[0025] The valid version range **60B** specifies the range of schema version numbers for which the corresponding structure is valid, and consists of an upper bound and a lower bound. Since the deletion of a structure is rarely planned in advance, the upper bound of the valid version range **60B** for a structure is preferably set to be unlimited. As discussed below, a structure can be retired by incrementing the global schema version number **58** and setting the upper bound of the valid version range **60B** of the structure to be less than the new global version number.

[0026] To support the ability to read both old and new cookies, the schema file **52** is modified according to the following set of rules.

- [0027] 1. The set of data elements specified within a structure definition **60A** cannot be modified.
- [0028] 2. No insertions into, or deletions from, the sequence of structure entries **60** are permitted.
- [0029] 3. A new structure can be added to the schema by adding a new entry **60** to the sequence and incrementing the global schema version number.
- [0030] 4. To retire a structure, the upper bound of the structure's version range **60B** is set to the current global version number **58**, and all deployed application software is updated to reflect this change. The global version number **58** is then incremented.
- [0031] These rules may be followed "by hand" by developer(s) during manual editing of the schema file **52**. Alternatively, a special management layer could be provided which enforces these rules.
- [0032] **FIG. 2** illustrates a preferred sequence of steps that are implemented by the host-to-cookie module **50A** to translate the sequence of structures to a cookie format. All but the final step **90** of this process can optionally be performed in an off-line mode, rather than in response to a URL request, to reduce the delay experienced by the user.
- [0033] The first step **78** of the process involves accessing the schema file **52** to identify the structures to be included within the cookie. This involves identifying the structures listed in the schema file **52** that are currently valid (i.e., have a valid version number range **60B** which includes the global schema version number **58**).
- [0034] In step **80**, each valid structure is read from the user database **44**, and the primitive elements of the structure are sequentially written in their network byte order. Conventional NTOHL and HTONL families of functions may be used for this purpose. Putting the primitives in network byte order allows machines having different endian architectures to be used for generating and decoding cookies. The result of step **80** is a variable-length sequence of bytes which represents the primitive elements of the valid structures.
- [0035] In step **82**, a checksum is computed for the resulting sequence of bytes. A 32-bit checksum may be used for this purpose. The checksum provides a mechanism for later determining whether the cookie has been modified. The checksum, global version number **58**, and the length of the sequence are then prepended or otherwise attached to the byte sequence (step **84**), with these values being again written in network byte order. In step **86**, the resulting sequence is encrypted using a standard encryption library (not shown).
- [0036] As indicated by step **88**, the resulting, encrypted byte sequence is then translated into a character string. This is preferably accomplished by dividing the sequence of bytes into a sequence of six-bit values (padding the sequence as needed to provide even divisibility), and converting each six-bit value into a base-64 digit. The character set used for this purpose consists of the numbers 0-9, all upper-case and lower-case alphabetical characters, and the characters "@" and "?." Any of a variety of other character sets and conversion schemes could be used for this purpose.
- [0037] Finally, in step **90**, the resulting character string is written to the corresponding user's computer **32** as the value

portion of a cookie. As is conventional, the cookie is written by including a Set-Cookie header within an HTTP response to the browser **34**. The user's structures that are written to the cookie are preferably retained within the user database **44**, but could alternatively be deleted to reduce server storage burden.

[0038] When the cookie is subsequently transmitted to the Web server **36** as part of a URL request, the cookie-to-host conversion module **50B** applies the above steps in reverse order to convert the character string back to the original host representation. As depicted by **FIG. 3**, the first step **100** of this process involves decoding the character string to reproduce the encrypted byte sequence. The encrypted byte sequence is then decrypted (step **102**) using the same encryption library and key used for encryption. At this point, the length, version number, checksum and data are available in network byte order.

[0039] As represented by step **104**, a checksum is then computed for the extracted data portion, and this checksum is compared to the extracted checksum to determine whether the cookie value was modified. If a mismatch occurs, the cookie is ignored (steps **106** and **108**); in this event, the user database **44** may be accessed to retrieve the desired user data.

[0040] Otherwise, the structures are "inflated" to their host representation according to the information contained within the schema file, and are written to local storage **56** (steps **110** and **112**) for use by the applications. During this process, the version number contained within the cookie is compared to the valid version ranges **60B** of the schema file **52** (step **110**). If the version number is within or exceeds the valid range of a schema entry **60**, that structure is inflated from the byte structure; otherwise, there is no data stored within the cookie for that structure, and the structure is skipped and marked as uninitialized (step **112**).

[0041] The above scheme allows values stored in old versions of cookies to be read by new versions of application software. If the cookie predates the addition of a structure to the end of the schema sequence, that structure is simply marked as uninitialized so that it can be treated as such by the new application software. If a structure has been deactivated (retired), the data from an old cookie will still be read but may then be ignored by the new application software. Further, cookies associated with new versions of application software can be read by old versions of the software.

[0042] As will be appreciated by the foregoing, the schema used to generate each cookie could alternatively be incorporated into that cookie. Doing so, however, increases the cookie size, and can decrease performance. The above-described approach of keeping track of the current and past schemas in a schema file or other server data structure is therefor preferred. It will also be recognized that the information stored within the schema file **52** could be distributed across multiple different files or data structures and/or represented in a different form.

[0043] Although this invention has been described in terms of certain preferred embodiments, other embodiments that are apparent to those of ordinary skill in the art are also within the scope of this invention.

What is claimed is:

1. A system for servicing web page requests, the system comprising a server that responds to user requests for web pages, said server comprising a memory;

schema data stored in the memory of the server, said schema data specifying past and present schemas used to encode data structures into cookies stored on user computers;

a conversion component executed by the server, said conversion component configured to use the schema data to identify and decode the data structures encoded within cookies received from user computers to generate temporary data structures within the memory of the server; and

application code executed by the server, said application code configured to use the temporary data structures to customize web pages requested by the user computers.

2. The system of claim 1; wherein the conversion component is not specific to a particular type of data structure.

3. The system of claim 1, wherein the conversion component supports a plurality of different types of data structures.

4. The system of claim 1, wherein the schema data specifies the content and format of each of a plurality of data structures that are encoded within the cookies.

5. The system of claim 1, wherein the schema data includes, for a given data structure that is encoded within cookies, an identification of primitive fields of the given data structure and datatypes of said primitive fields.

6. The system of claim 5, wherein the schema data further includes address offsets of the primitive fields.

7. The system of claim 1, wherein the schema data includes, for a given data structure that is encoded within cookies, an indication of a range of schema versions for which the data structure is valid, wherein the conversion component uses said range to determine whether a particular data structure encoded within a cookie is valid.

8. The system of claim 1, wherein the schema data includes information about at least one data structure that is no longer in use.

9. The system of claim 1, wherein the conversion component uses the schema data to determine which of a set of the data structures encoded within a received cookie are to be decoded for use.

10. The system of claim 1, wherein the conversion component uses the schema data to determine whether a particular data structure that is encoded within a received cookie is to be decoded for use.

11. The system of claim 1, wherein the conversion component is an executable function that is called by the application code.

12. The system of claim 1, wherein the schema data is cached with random access memory of the server.

13. The system of claim 1, wherein the schema data is stored within a file in the memory of the server.

14. The system of claim 1, wherein the schema data is stored within a table in the memory of the server.

15. The system of claim 1, wherein the conversion component uses a checksum included within a received cookie to evaluate whether the cookie has been modified.

16. The system of claim 1, wherein the conversion component applies a decryption algorithm to encrypted information contained in the cookies received from the user computers.

17. The system of claim 1, further comprising an encoding component that encodes data structures into cookies according to a current schema specified by the schema data.

18. A method of extracting information from cookies, the method comprising

(a) receiving, at a server, cookie data that has a data structure encoded therein, said cookie data received from a user computer;

(b) determining whether the data structure encoded within the cookie data is valid; and

(c) if and only if the data structure encoded within the cookie data is determined to be valid in step (b), decoding the encoded data structure to reproduce the data structure within a memory of the server.

19. The method of claim 18, wherein the validity of the data structure as determined in step (b) reflects whether the data structure is currently used by any applications running on the server.

20. The method of claim 18, wherein step (b) comprises identifying a version of a schema used to encode the data structure within the cookie data.

21. The method of claim 18, wherein step (b) comprises using schema data stored within the memory of the server to determine whether the data structure is valid.

22. The method of claim 18, wherein a plurality of data structures are encoded within the cookie data, and the method comprises determining which of the plurality of data structures should be fully decoded.

23. The method of claim 18, further comprising using the data structure as reproduced in step (c) to customize a web page requested by the user computer.

24. The method of claim 18, wherein step (c) is performed by executable code that is not specific to a particular type of data structure.

25. The method of claim 18, wherein the data structure includes primitives of at least one of the following data types: 16-bit integers, 32-bit integers.

26. A server system that performs the method of claim 18.

27. A computer program that embodies the method of claim 18 stored within a computer readable medium.

28. A method of generating cookie data for storage on a computer of a user, the method comprising:

identifying a set of data structures to be encoded within the cookie data;

encoding the set of data structures within the cookie data according to schema data stored within a computer memory, said schema data specifying how the set of data structures is to be encoded within the cookie data; and

incorporating into the cookie data at least one of the following to facilitate extraction of the set of data structures from the cookie data: (a) a schema identifier, (b) the schema data.

29. The method of claim 28, wherein the method comprises incorporating the schema data into the cookie data.

30. The method of claim 28, wherein the method comprises incorporating a schema identifier into the cookie data,

said schema identifier identifying one of a plurality of schemas used over a period of time to encode data structures within cookie data.

31. The method of claim 28, wherein the set of data structures is encoded within the cookie data using executable code that is not specific to a particular type of data structure.

32. The method of claim 28, wherein the set of data structures includes noncharacter. primitives.

33. The method of claim 28, wherein the set of data structures includes primitives of at least one of the following data types 16-bit integers, 32-bit integers.

34. The method of claim 28, further comprising incorporating a checksum into the cookie data to permit subsequent detection of whether the cookie data has been modified.

35. The method of claim 28, wherein the method is performed in an off-line mode to reduce a delay experienced by a user.

36. A computer-readable medium comprising cookie data generated according to the method of claim 28.

37. A computer system configured to perform the method of claim 28.

38. A computer-readable medium that stores a computer program embodying the method of claim 28.

39. A method for using browser cookies to store structured data, the method comprising:

storing schema data on at least one server computer of a web site system, said schema data specifying schemas used by executable software to (a) encode data structures within cookies for storage on user computers, and (b) decode said cookies to extract the data structures when the cookies are returned by the user computers; and

modifying the schema data over time to add data structures to, and remove data structures from, a set of data structures encoded within cookies by the executable software.

40. The method of claim 39, wherein the schema data is modified according to a set of rules to enable the executable software to decode cookies encoded using both past and present schemas.

41. The method of claim 40, wherein the rules are enforced by a management layer.

42. The method of claim 39, wherein the executable software is not specific to a particular type of data structure.

* * * * *