

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
3 October 2002 (03.10.2002)

PCT

(10) International Publication Number  
WO 02/077878 A1

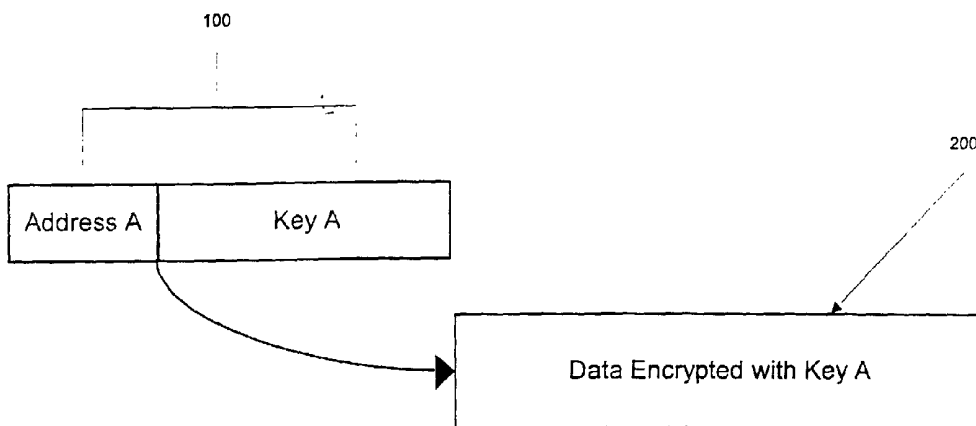
- (51) International Patent Classification<sup>7</sup>: **G06F 17/30**
- (21) International Application Number: PCT/US02/09426
- (22) International Filing Date: 26 March 2002 (26.03.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
60/278,823                      26 March 2001 (26.03.2001)    US
- (71) Applicant (for all designated States except US): **GALOIS CONNECTIONS INC** [US/US]; 3875 SW Hall Boulevard, Beaverton, OR 97005 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **LAUNCHBURY, John** [GB/US]; 13075 SW Davies Road, Beaverton, OR 97008 (US). **NORDIN, Thomas** [SE/US]; 420 NW 11th #1020, Portland, OR 97209 (US).
- (74) Agent: **MUNSON, Steven, J.**; Stoel Rives LLP, Suite 2600, 900 SW Fifth Avenue, Portland, OR 97204-1268 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: CRYPTO-POINTERS FOR SECURE DATA STORAGE



Crypto-pointer always points to cell encrypted with associated key

(57) Abstract: The present invention relates to pairing of a different cryptographic key with each pointer in a data structure to form a crypto-pointer (100). The cryptographic key is used to encrypt the contents of all data stored at the physical location on the storage device indicated by the pointer (100). Preferably the only data accessible in an unencrypted form is contained in cells that are reachable from root-set crypto-pointers (100). Once the crypto-pointer (100) associated with a particular memory cell (200) is deleted, normally by overwriting or explicitly zeroing the crypto-pointer (100), the contents of the memory cell (200) become inaccessible because the data stored at that cell is in encrypted form (cipher text) and the crypto-pointer (100) that included the cryptographic key for decrypting the cipher text has been deleted from the system.



WO 02/077878 A1



---

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## CRYPTO-POINTERS FOR SECURE DATA STORAGE

### Related Applications

[0001] This application is a continuation of U.S. Provisional Patent Application 60/278,823 filed March 26, 2001 and incorporated by reference herein.

### Technical Field

[0002] The present invention relates to secure data storage and to preventing recovery of deleted data.

### Background of the Invention

[0003] One way of storing data in a computer system is to store files in a directory structure. The directory structure normally includes a variety of information about any files and sub-directories within the directory structure. One piece of information included in the directory for each file or sub-directory is a directory listing that includes a field containing the physical location, or a pointer to a physical location, on a storage device, such as a disk drive or other non-volatile storage medium, where the file or sub-directory is located. Deletion of a file or sub-directory from the directory typically consists of removing only the directory listing. Once the directory listing has been deleted the system can no longer access the file or sub-directory because it does not know the physical location on the storage device for the file or sub-directory. However, the file or sub-directory is usually still stored on the storage device and can be recovered through the use of well known data recovery techniques.

[0004] One way to prevent recovery of deleted data is to explicitly zero, or overwrite, all data at the physical location pointed to by the directory listing. However, this can be a cumbersome process. Particularly, if there are large amounts of data stored on the storage device. The problem can be even more difficult in a heap data structure because it is generally not clear when data within the structure is discarded. In a conventional heap, data is only discarded when the data becomes inaccessible via the heap structure, *i.e.* an

access path within the heap structure no longer includes any pointers to the data. Discarding data in this environment typically consists of deleting any pointer that includes the physical location of the data on a storage device. The data itself is generally not deleted. Due to the uncertainty as to when the pointers, which include the physical location of the data, is discarded it can be difficult and cumbersome to overwrite or zero the data at its physical location on the storage device.

#### Summary of the Invention

**[0005]** The present invention relates to hierarchical data structures employing crypto-pointers in place of normal pointers. Crypto-pointers enable fine-grain crypto-protection by the use of vast numbers of distinct keys within a single data structure. Crypto-pointers solve key-management problems in terms of the data structure itself. The use of crypto-pointers helps to ensure that data becomes unrecoverable whenever it becomes inaccessible, *i.e.* data becomes unrecoverable whenever the crypto-pointer associated with the data is deleted.

**[0006]** The present invention relates to the pairing of a unique cryptographic key with each pointer in a data structure to form a crypto-pointer. The cryptographic key is used to encrypt all data stored at the physical location on the storage device indicated by the pointer. In accordance with the present invention, preferably the only data accessible in an unencrypted form is a root set of crypto pointers that provide entry points into the data structure. Once the crypto-pointer associated with a particular data cell is deleted, normally by overwriting or explicitly zeroing the crypto-pointer, the particular data cell or cells that were pointed to by the crypto pointer become inaccessible because the data at those cells is in encrypted form (cipher text) and the crypto-pointer that included the cryptographic key that could have been used to decrypt the data has been deleted from the system.

**[0007]** Additional aspects and advantages of this invention will be apparent from the following detailed description of preferred embodiments thereof, which proceeds with reference to the accompanying drawings.

#### Brief Description of the Drawings

**[0008]** Fig. 1 is a block diagram depicting the logical structure of a crypto-pointer.

**[0009]** Fig. 2 is a block diagram depicting the association of the crypto-pointer of Fig. 1 with encrypted data stored in memory at a location indicated by an address field of the crypto-pointer.

[0010] Fig. 3 is a block diagram depicting how a crypto-pointer A enables access to data associated with crypto-pointers B and C.

[0011] Fig. 4 is a block diagram of a simple data structure with three access points, each of which is associated with a plain text version of a different crypto-pointer.

[0012] Fig. 5 shows the data structure of Fig. 4 after the plain text version of crypto-pointer C has been deleted.

[0013] Fig. 6 shows the data structure of Fig. 5 after the plain text version of crypto-pointer A has been deleted.

[0014] Fig. 7 is a flow chart depicting a method of using crypto-pointers.

[0015] Fig. 8 is a block diagram depicting a data structure using crypto-pointers to limit access to portions of the data structure based on the point of entry into the data structure.

#### Detailed Description of Preferred Embodiments

[0016] A heap is a collection of cells within a portion of memory, where the cells may contain data, pointers (which are memory addresses to other cells), or a combination of both. In a normal heap structure, blocks of memory in one or more storage devices are allocated and freed in an arbitrary manner. The heap structure typically includes a set of root pointers. Each of the root pointers includes information about a memory block within the heap structure including the physical location of the memory block on a storage device. The root pointer provide entry points into the heap structure. The memory blocks pointed to by the root pointers can include data, additional pointers, or both. In that way the heap structure can take on a tree structure, or a graph structure. A memory block is freed when no other memory blocks accessible from one of the entry points includes a pointer to that memory block. At that point, a garbage collector can identify that the memory block is no longer needed and the memory block can be allocated for storage of new data. However, until new data is written to a particular memory block the old data stored at that physical location is preserved and can be accessed via known data recovery techniques.

[0017] In accordance with the present invention a secure data structure can be formed by using "crypto-pointers" instead of normal pointers. A crypto-pointer includes an address, much like a normal pointer, which indicates the physical location on a non-volatile storage device where data has been, or will be, stored. In a crypto-pointer the address is paired with a cryptographic key. Fig. 1 is a block diagram depicting a crypto-pointer 100 in accordance with the present invention. With reference to Fig. 1, an address A is paired

with an encryption Key A to form crypto-pointer 100. Crypto-pointer 100 is a single logical entity and should preferably always be manipulated or handled as a single entity.

**[0018]** Fig. 2 is a block diagram depicting the association between crypto-pointer 100 and data stored in a memory cell 200 on a storage device (not shown). The physical location of memory cell 200 on the storage device corresponds to the location indicated by address A. With reference to Fig. 2, data to be stored at memory cell 200 is always encrypted using encryption key A and an appropriate encryption algorithm to form a cipher-text version of the data. The cipher-text version of the data is then stored at memory cell 200. In order to access a plain text version of the data stored at memory cell 200 the cipher-text version of the data must be decrypted using encryption Key A and an appropriate decryption algorithm.

**[0019]** Ensuring that only cipher-text versions of data are stored in memory cells on a storage device has implications for allocating, updating, and reading the data stored in a memory cell. First, whenever a new memory cell is allocated, the address is recorded, a new random encryption key is generated, and the random encryption key is associated with the address to form a new crypto-pointer. The new random encryption key is used to encrypt all initial contents of the new memory cell. Second, whenever the memory cell content needs to be updated, the random encryption key of the crypto-pointer is used to encrypt the new content (whether data, other crypto-pointers, or a combination of both), and the memory cell is overwritten with the updated encrypted information. Conversely, whenever the memory cell pointed to by the crypto-pointer is to be read, the random cryptographic key is used to decrypt the data stored at the address, also called the cipher-text, thus revealing the plain text data stored at the address. Whatever information this cell contains (again, whether data or other crypto-pointers), becomes known only after decryption.

**[0020]** In a preferred embodiment each memory cell has its own unique encryption key as shown in Fig. 3. With reference to Fig. 3, the plain text version of the data stored at memory cell 200 can only be accessed by decrypting the stored cipher-text using encryption key A and an appropriate decryption algorithm. The plain text contents include data 300, crypto-pointer 302 and crypto pointer 304. Crypto pointer 302 comprises an Address B and an encryption Key B and is associated with a memory cell 306. Crypto pointer 304 comprises an Address C and an encryption Key C and is associated with a memory cell

308. By associating each memory cell with its own crypto-pointer a user can only access the plain text version of data stored at a memory cell if the user has the appropriate crypto-pointer for that memory cell.

**[0021]** If the memory cell is small enough, the key itself may be used directly for encryption and decryption by combining it with the cell contents using an exclusive-OR operation. Preferably, however, a stream cipher should be used to generate a stream of bits long enough for encrypting the data to be stored at the memory cell. Depending on the performance requirements of a particular application, the cryptographic operations may be implemented in either software, or hardware, or a combination of both. The ideal size of any encryption keys will depend on the desired security constraints for the application. In the most sensitive applications, a block-cipher method such as Rijndael or RC6 may be desirable.

**[0022]** If an additional reference to a cell is required, it can be produced only by copying an existing crypto-pointer (both the address and the cryptographic key fields are copied). Simply generating a regular pointer (computing the address of a memory cell by pointer arithmetic for example), will allow only the encrypted contents of that cell to be accessed. Because it lacks the associated cryptographic key, the generated regular pointer does not allow a user to obtain a plain text version of the contents of the memory cell.

**[0023]** Whenever a crypto-pointer is discarded, both the address and the cryptographic key are physically overwritten. Overwriting in this sense is typically accomplished by either zeroing both fields, the address and the cryptographic key, if the crypto-pointer is being set to nil, or by overwriting both fields with another crypto-pointer.

**[0024]** Heap structures, as mentioned above, typically use a root-set to define any entry points into the structure. Similarly, a secure heap utilizing the present invention has a root set consisting of one or more crypto-pointers, which are maintained in an unencrypted form (in contrast to any crypto-pointers that may be embedded within the heap). The heap can be traversed starting from any one of these root crypto-pointers by successively retrieving keys and decrypting cells at each step within the heap as shown in Fig. 4. In this way, transitive access to the contents of all descendant cells is possible. With reference to Fig. 4, the root set of crypto-pointers includes crypto-pointer 100, crypto-pointer 302, and crypto-pointer 304. A user that begins traversing the data structure with crypto-pointer 100, of the root set, can gain access to the data associated with crypto-pointer 302 and crypto-pointer 304

because the plain text version of the data associated with crypto-pointer 100 includes, as discussed above, a plain text copy of both crypto-pointer 302 and crypto pointer 304. A user that begins traversing the data structure with crypto-pointer 302 of the root set of crypto-pointers can only gain access to a plain text version of the encrypted data associated with crypto-pointer 302, which is stored at memory cell 306. Likewise, a user that begins traversing the data structure with crypto-pointer 304, of the root set of crypto-pointers, can only gain access to the plain text version of the encrypted data associated with crypto-pointer 304, which is stored at memory cell 308.

**[0025]** It is often the case that a cell in the heap is connected to the root-set by many access paths. As crypto-pointers in these access paths are duplicated, the number of alternative access paths to any given cell may increase. Conversely, as crypto-pointers are deleted the number of possible access paths decreases, though the cell may still remain accessible

**[0026]** Fig. 5 shows the data structure of Fig. 4 after crypto-pointer 304, of the root set of crypto-pointers, has been deleted. With reference to Fig. 5, the plain text data associated with crypto-pointer 304 is now only accessible to a user who begins traversing the data structure with crypto-pointer 100, because the root set version of crypto-pointer 304 is no longer available. The access paths to the data associated with crypto-pointer 302 remain unchanged from Fig. 4.

**[0027]** Ultimately, deletion of crypto-pointers may eliminate all access paths, causing a cell to become disconnected from the root-set. Once this happens, no plain text version of its associated cryptographic key exists, nor can a plain text version of the associated cryptographic key be obtained, short of breaking the cryptography. Therefore, the contents of the cell can no longer be decrypted, short of breaking the encryption algorithm or performing a brute force search of the key space. Fig. 6 show the data structure of Fig. 5 after crypto-pointer 100 has been deleted from the root set of crypto-pointers. With reference to Fig. 6, the data associated with crypto-pointer 304 is no longer available because the data structure no longer includes a path through which a user can access the data that was associated with crypto-pointer 304. Once crypto-pointer 100 was deleted the encrypted data associated with crypto-pointer 100 became inaccessible because encryption key A was no longer available. As a result a plain text version of crypto-pointer 304 can no longer be obtained from the encrypted data associated with crypto-pointer 100 and a user



is left with no way to access encryption key C and no way to decrypt the encrypted data stored at memory cell 308. Likewise, a user can no longer access data 300 that was stored at memory cell 200 because the user no longer has access to encryption key A. However, the data stored at memory cell 306 is still accessible because the root set still includes a plain text version of crypto-pointer 302 giving the user access to encryption key B which can be used with an appropriate encryption algorithm to generate a plain text version of the data stored at memory cell 306.

**[0028]** Fig. 7 is a flow chart showing a method of using crypto-pointers to create a secure data structure. With reference to Fig. 7, a computer system, or user, obtains a first address of a first memory cell in the data structure. The first address is then associated with a first cryptographic key to form a first crypto-pointer. The cryptographic key can be of any length deemed appropriate for the sensitivity of the data to be stored in the data structure. When the system, or user, receives a first block of data to be stored at the first memory cell it encrypts the first block of data using the first encryption key, and an appropriate encryption algorithm, prior to storing the first block of data at the first memory cell. In order to access a plain text version of the first block of data, the system must use the first crypto pointer to access the first memory cell. The system retrieves the encrypted first block of data using the first address and decrypts the encrypted first block of data using the first encryption key and an appropriate encryption algorithm. When the first block of data is no longer needed, then the system deletes the first crypto pointer, preferably by either setting both the address and key fields to zero or by overwriting the first crypto-pointer with a new crypto pointer having a new address and a new cryptographic key. Once the first crypto-pointer has been deleted in this manner the encrypted first block of data can not be accessed by the system. In addition, the first block of data cannot be recovered using data recovery techniques because those techniques would only recover the encrypted version of the first block of data. In order to have access to a plain text version of the first block of data one would need to break the encryption scheme.

**[0029]** Taking this to a more complex example, if a memory cell X (not shown) is disconnected from a root set of pointers (not shown), then any other memory cell (not shown) which, when decrypted has a plain text crypto-pointer pointing to memory cell X must also be disconnected from the data structure (not shown), such that there are only encrypted versions of the crypto-pointer pointing to memory cell X stored on the system

and there are no paths in the data structure through which a user may obtain a plain-text version of the crypto-pointer pointing to memory cell X. Even though encrypted versions of crypto-pointers pointing to memory cell X may remain in memory, those encrypted versions of crypto-pointers pointing to memory cell X can not be decrypted, so the encryption key to decrypt the data stored at memory cell X can not be recovered unless the encryption scheme is broken. In effect, any memory cell whose corresponding crypto-pointer has been deleted from the heap appears to be random noise to known data recovery techniques.

**[0030]** The contents of a normal computer storage disk (hard disk drive, floppy disk, CD, DVD, etc.) are normally arranged as a directory tree. A directory of this type can be viewed as a memory cell that contains a directory listing for each object (file or sub-directory) "within" the directory. The directory listing for an object includes information about the object such as the object's physical location on the storage disk plus other information particular to the object itself. The physical location of the object corresponds to a normal pointer, as discussed above, and plays a similar role. Deletion of an object normally consists only of removal of the object's directory listing from the directory. With deletion in this manner, the object itself is still stored at the physical location on the storage disk that was pointed to by the directory listing and can be recovered by well known data recovery techniques.

**[0031]** The contents of a computer storage disk can be made secure by modifying the organization of information on the storage disk such that all directory pointers (*i.e.* information about each object's physical location on a storage device) are replaced with crypto-pointers, as described above. In accordance with a presently preferred embodiment of the invention all data stored at the physical location on the computer storage disk pointed to by the address field of a particular crypto-pointer is encrypted with the encryption key field of that particular crypto-pointer. In this way, when the directory listing for an object, a program or subdirectory, is deleted and overwritten the data stored at the physical location on the computer storage disk pointed to by the crypto-pointer associated with the directory listing for the object is rendered unrecoverable because the key to decrypt the data is not available to the system. Unlike in other approaches, there is preferably no master key for the disk that would reveal all the (current and old) data it may contain. Unless the

encryption is broken, only data accessible by an existing path in the directory structure can be decrypted.

**[0032]** To gain access to the contents of the file-store, a user must have access to a plain text version of a root crypto-pointer. Preferably the root crypto-pointer should be treated with a level of care similar to that used for normal cryptographic keys. Preferably the root crypto-pointer will only be distributed according to a known certificated key-management process.

**[0033]** Because file deletion within a crypto-pointer computer storage system is permanent, this form of organizing disk storage should be used with caution and preferably only in situations where accidental deletion of files is unlikely to occur, or where backup copies of deleted directory listings are maintained for some period of time. A preferred way to minimize unintentional deletion is to copy all deleted directories into a trash directory, while deleting the actual directory entry. Deleted listings would then remain in the trash directory for a period of time, thus allowing users an opportunity to recover unintentionally deleted files. Preferably users would need to manually empty their trash directory, at which point the user should be advised that non-reversible deletion is about to occur. Alternatively, the trash directory could be set to automatically empty itself of deleted directory listings on a periodic basis.

**[0034]** In addition, in order to maintain the unrecoverability of files whose directory listing has been deleted, symbolic links, such as shortcuts or aliases, must preferably not be implemented as crypto-pointers. If symbolic links were implemented as crypto-pointers then the existence of a symbolic link to a file whose directory listing has been deleted would maintain a decryption path for the otherwise deleted files. To address this concern, symbolic links should be represented as a path description that a file manager application uses to traverse the directory hierarchy, starting from the root set of pointers for the disk. Of course, if hard links across the directory hierarchy are to be permitted by design, then the directory structure becomes a graph rather than a tree, and deletion occurs only when a file is no longer accessible from a root entry point into the disk.

**[0035]** The organizational structure described above provides an additional benefit when distributing a large body of information (whether program source-code, object-code, business documents, data, or families of task-specific cryptographic keys etc.) to a variety of users who may need access to different portions of the information. The body of

information can be graph-structured according to an appropriate information-release policy, laid out on a DVD (or CD), and distributed widely. Access to portions of the information graph is enabled by providing each individual user with the relevant crypto-pointer that unlocks the information they are permitted to see, and therefore, to all the descendent information also. The desired access policy is embodied in the graph structure. With suitable care in the design of that structure, each user would need only one single entry point into the graph. Any portion of the DVD not accessible via a users entry point into the graph remains cryptographically secure from that user. If a user's permissions change (through purchasing an upgrade, for example) then the vendor need only supply another crypto-pointer to provide access to a previously unreadable portion of the DVD.

**[0036]** Fig. 8 is a block diagram illustrating a graph data structure limiting access based on a user's entry point into the data structure, as described above. With reference to Fig. 8, the arrows between boxes represent an access path between the data corresponding to the boxes that only works in the direction of the arrow, *i.e.* the box attached to the tail of the arrow represents data that when decrypted includes a crypto-pointer pointing to the data that corresponds to the box attached to the head of the arrow. An access path in this context means that the unencrypted data corresponding to a box (representing a directory listing or crypto-pointer) attached to the tail of the arrow includes a plain text version of a crypto-pointer pointing to the encrypted data corresponding to a box in the graph attached to the head of the arrow. For example, data stored in a memory cell corresponding to a box labeled first feature includes a copy of a crypto-pointer pointing to a second feature, a sixth feature and a special feature A. Thus when a user has access to a plain text version of the data corresponding to first feature that user also has plain text crypto-pointers corresponding to second feature, sixth feature, and special feature A.

**[0037]** With reference to Fig. 8, a first tier user has a root crypto-pointer that allows the first tier user to access plain text crypto-pointers corresponding to a fourth feature and a fifth feature. Once data stored at a location pointed to by the crypto-pointer corresponding to the fifth feature has been decrypted based on its encryption key the first tier user will be have a plain text version of a crypto-pointer pointing to a ninth feature and a crypto-pointer pointing to a tenth feature, as shown in Fig. 8. First tier user gains access to ninth and tenth features because fifth feature includes, among other things, a plain text version of crypto-pointers that correspond to ninth and tenth features. Finally, first tier user can

access a special feature D because the decrypted version of the data corresponding to the tenth feature's crypto-pointer includes a plain text version of a crypto-pointer corresponding to special feature D. As shown in Fig. 8, first tier user has no access paths to first, second, third, sixth, seventh, and eighth features. Likewise, first tier user does not have access to special features A-C because the entry point into the data structure employed by first tier user does not allow the first tier user to gain access to a plain text version of the crypto-pointers associated with those features. Even if first tier user knows the physical location on the storage disk of any inaccessible features, first tier user can not access those features without defeating the encryption.

**[0038]** A second tier user on the other hand has access to third, seventh, and eighth features, in addition to the numbered features that the first tier user can access. The second tier user also has access to special features B and C as well as special feature D. The third tier user in this example has complete access to all the features within the data structure.

**[0039]** Other examples of applications that can be secured with a data structure as described above are as follows: Hard-disk drives with built-in clean-delete features; Hard-disk drives, for which multiple users have root access, permitting individuals private access to their own areas, group access to appropriate group spaces, and universal access to common spaces for all users; Software disk-drivers (for LINUX, Windows etc.) that implement clean-delete on standard disk drives; Software disk-drivers to enable multiple users to maintain private space on a shared disk; DVDs and CDs with enforced information-release policies.

**[0040]** Applications of crypto-pointers in heap data structures, as described above, include programs that manipulate tree-structured or graph-structured data. This includes heap-based implementations of object-oriented languages such as Java, or C++, and heap-based implementations of functional languages, such as Haskell.

**[0041]** In the course of executing a Java program, regardless of whether the platform is a smart card or a supercomputer, new objects are created all the time, and connections between objects are maintained using traditional pointers. When an object is no longer accessible, it remains physically in memory (along with all its internal storage) until a garbage collector reclaims the space. Even then, depending on the design of the particular garbage collector, the cell contents may not have been erased. Even when the physical

space that corresponded to a “deleted” object is reused some part of the “deleted object” might remain stored in memory if the new object is smaller than the “deleted” object.

**[0042]** Within this context, replacing the normal pointers associated with each object with crypto-pointers alleviates this problem. The object, including its internal storage, is stored in an encrypted format at its physical location in the heap. Only when the object is accessed using a copy of its valid crypto-pointer can the contents of the object be decrypted into a plain text version that can be used and understood. While this might impose a dramatic efficiency overhead if the cryptography is done in software, performing it in hardware can make the cost reasonable.

**[0043]** Special care is needed to deal with an execution stack in this environment. If a traditional stack is used, then a set of unencrypted crypto-pointers will exist in memory. Depending on the security level of the application, this may be fine, so long as any crypto-pointer on the top of the stack is overwritten each time the stack is popped. In more stringent situations, the stack should be broken up into a linked list of encrypted stack frames, each link in the list being a crypto-pointer. Example applications of crypto-pointers in this area include the following: Secure Java, by using crypto-pointers instead of ordinary pointers de-allocated objects will not remain accessible; Secure Java smart cards; Secure Haskell (and other functional languages), using crypto-pointers instead of normal pointers results in a heap in which discarded data can not be accessed without breaking the encryption; Secure C, where pointer arithmetic works within the context of a defined array structure only (the key is a fixed function of the base key and the increment), but pointer arithmetic fails to produce the correct key if the structure is overrun; Graph-algorithms, where the data is sensitive, and the graph evolves over time, deleting the pointer to portions of information as those portions become unneeded results in the discarded data being inaccessible unless the encryption is broken.

**[0044]** In addition, crypto-pointers may be used in an Internet, intranet, or any network context where resources are shared. In a network environment a crypto-pointer will have the same fields as described above, an address field and a key field. In this context the address field corresponds to a Uniform Resource Locator (URL) or other scheme for locating or identifying network resources. The key field, as above, is a cryptographic key used to encrypt and decrypt the information pointed to by the address field. For example, the key is used to decrypt the web page or ftp file after download, and the key would also

be used to encrypt information prior to uploading that information to the address. This provides users with a way to store sensitive data in a distributed environment when the network is untrusted, and also when the storage nodes are untrusted. The storage nodes only store encrypted information, and in a preferred embodiment the storage nodes do not receive a decryption key. Moreover, once the crypto-pointer is deleted, such as removing and physically overwriting the URL crypto-pointer from all network machines that had a copy, an untrusted storage node will have nothing to gain by maintaining a copy of the encrypted page in the hope that the key will be compromised because the encryption key for that page is unrecoverable.

**[0045]** As far as the network is concerned, the only plain text visible to potential network snoopers are hyperlinks; all other traffic is encrypted. Possession of a hyperlink is equivalent to possession of a pointer address without the corresponding encryption key. Thus, the use of crypto-URLs (URLs paired with cryptographic keys, used in the manner described above) provides a viable technique for securing sensitive data stored in a networked environment, including data such as health and financial records. The sensitive data is maintained in a heap-structured manner, and as a result any individual piece of sensitive data can only be accessed by following the official structure of the database. A user attempting to access the sensitive information will normally also have to satisfy whatever additional authorization mechanisms that an administrator chooses to impose on the database structure. Preferably, the database structure is implemented with no “back doors” through which an unauthorized user can access the sensitive information.

**[0046]** A similar technique can be implemented on a smaller scale, such as a single computer. With proper management of a root set of crypto-pointers, a set of heap traversal procedures can be implemented so that unencrypted crypto-pointers do not leave the on-chip registers of the central processing unit. In this way, a secure heap can be suitable for situations in which the internal security of a computer is questionable.

**[0047]** In an alternative implementation of the present invention crypto-pointers are combined with regular pointers to build a mixed data structure. In this embodiment, portions of the data structure are stored in plain text format and the directory listing for these unsecured portions of the data structure includes only a pointer to the physical location of the data. While, other portions of the data structure are implemented using crypto-pointers such that the corresponding data is stored in encrypted format and the

directory listings corresponding to the data stored in encrypted format include crypto-pointers corresponding to the encrypted data. In this context, regular URLs can be included in secure web pages, and may be followed in the usual way. Preferably, the unsecured portions of the mixed data structure should not contain plain text crypto-pointers because the data to which they point would then also be public, and the cryptography would provide no additional security.

**[0048]** Other examples of applications that can be secured with a data structure as described above are as follows: Distributed storage of sensitive information on untrusted hosts; Individualized record access for financial or health records, wherein by providing an additional authorization key in the encrypted page, a user who has a key to access a record (and hence can decrypt it) can also gain access to the authorization key to enable information update; Manipulation of secure data in machines vulnerable to physical probing.

**[0049]** It will be obvious to those having skill in the art that many changes may be made to the details of the above-described embodiments of this invention without departing from the underlying principles thereof. The scope of the present invention should, therefore, be determined only by the following claims.



### Claims

1. In a hierarchical digital data structure comprising a plurality of memory cells, a method of preventing recovery of deleted data comprising the steps of:

- obtaining a first address of a first memory cell in the data structure;
- providing a first cryptographic key;
- associating the first address and the first cryptographic key with one another so as to form a first logical entity defining a first crypto-pointer to the first memory cell;
- encrypting first data using the first cryptographic key to form first cipher-data;
- storing the first cipher-data data in the first memory cell; and
- when the first data is no longer needed, deleting the first crypto-pointer, thereby rendering the first data unrecoverable at the moment it becomes inaccessible.

2. The method of preventing recovery of deleted data according to claim 1 wherein said deleting the first crypto-pointer comprises physically overwriting both the first address and the first cryptographic key.

3. The method of preventing recovery of deleted data according to claim 1 wherein said deleting the first crypto-pointer comprises physically overwriting both the first address and the first cryptographic key by zeroing both fields.

4. The method of preventing recovery of deleted data of claim 1 further comprising:

- obtaining a second address of a second memory cell in the data structure;
- providing a second cryptographic key different from the first cryptographic key;
- associating the second address and the second cryptographic key with one another so as to form a second logical entity defining a second crypto-pointer to the second memory cell;
- encrypting second data using the second cryptographic key to form second cipher-data;
- storing the second cipher-data in the second memory cell; and
- when the second data is no longer needed, deleting the second crypto-pointer, thereby rendering the second data unrecoverable at the moment it becomes inaccessible.

5. The method of preventing recovery of deleted data of claim 1 further comprising creating an additional reference to the first memory cell by copying the first crypto-pointer.

6. The method of preventing recovery of deleted data of claim 1 wherein the data structure comprises a heap structure in a non-volatile storage device.

7. The method of preventing recovery of deleted data of claim 1 wherein the data structure comprises a file directory.

8. A crypto-pointer logical entity for use in secure digital data storage consisting of:

an address field that points to a corresponding cell in a digital data structure; and  
a cryptographic key associated with the address field for encrypting data to be stored in the corresponding cell.

9. The crypto-pointer entity of claim 8 wherein the data structure is a file directory stored on a non-volatile storage medium.

10. The crypto-pointer entity of claim 8 wherein the data structure is a heap structure stored on a computer readable storage medium.

11. A crypto-pointer logical entity for use in secure digital data storage comprising:  
an address field that points to a corresponding physical location available for storing digital data;

a cryptographic key associated with the address field for encrypting data to be stored at the corresponding physical location.

12. The crypto-pointer entity of claim 11 wherein the corresponding physical location is a memory cell in a computer readable storage medium and the address field is a pointer to the memory cell.

13. The crypto-pointer entity of claim 11 wherein the corresponding physical location is a resource attached to a network and the address field is a uniform resource locator.

14. A secure method of digital data storage comprising the steps of:  
allocating a first cell in a digital memory system;  
recording a first address corresponding to the first cell;  
generating a first random key associated with the first address;  
encrypting first data for storage in the first cell using the first random key;  
allocating a second cell in the digital memory system;  
recording a second address corresponding to the second cell;  
generating a second random key associated with the second address; and

encrypting second data for storage in the second cell using the second random key.

15. A secure method of digital data storage according to claim 14 and further comprising updating the first memory cell contents by:

- obtaining new data for storage in the first cell;
- encrypting the new data using the first random key; and
- overwriting the first cell contents with the encrypted new data.

16. A secure method of digital data storage according to claim 14 further comprising:

reading the first cell contents by retrieving the encrypted first data and decrypting the retrieved encrypted first data first using the first random key; and

reading the second cell contents by retrieving the encrypted second data and decrypting retrieved encrypted second data using the second random key.

17. A secure method of digital data storage according to claim 14 wherein said encrypting the first data for storage in the first cell using the first random key comprises a Boolean exclusive-OR (XOR) operation.

18. A secure method of digital data storage according to claim 14 wherein said encrypting the first data for storage in the first cell using the first random key comprises using the first random key in conjunction with a predetermined encryption algorithm to encrypt the first data, thereby producing a cipher-text version of the first data.

19. A secure method of digital data storage according to claim 14 wherein the encryption operation is implemented in software.

20. A secure method of digital data storage according to claim 14 wherein the encryption operation is implemented in hardware.

21. A secure memory structure comprising:

a root set of pointers that provide at least one entrance point into the memory structure, wherein the root set of pointers includes a first crypto-pointer

a first crypto-pointer in the root set of pointers having a first encryption key and a first address, the first address corresponding to a first memory cell in the memory structure, wherein all data stored at the first address is encrypted with the first encryption key.

22. A secure memory structure according to claim 21 wherein the root set of pointers includes a second crypto-pointer, having a second crypto-pointer and a second address corresponding to a second memory cell in the memory structure, wherein all data stored at the second address is encrypted with the second encryption key.

23. A secure memory structure according to claim 22 wherein the root set of pointers are distributed according to a key authentication scheme.

24. A secure memory structure according to claim 23 wherein the memory structure comprises a memory heap.

25. A secure memory structure according to claim 23 wherein the memory structure comprises a graph data structure.

26. A secure memory structure according to claim 23 wherein the memory structure comprises a directory tree structure.

27. A secure heap logical entity for use in secure digital data storage comprising:

a plurality of memory cells within a memory storage device;

a first crypto-pointer having a first address and a first cryptographic key and associated with one of the plurality of memory cells, such that information stored at the associated memory cell is encrypted with the first cryptographic key before it is stored at the memory cell.

28. A method of accessing encrypted data in a secure heap data structure comprising:

maintaining a root set of pointers in an unencrypted form outside of the secure heap;

within the root set of pointers, providing a first crypto-pointer associated with a first memory cell in the secure heap data structure, the first crypto pointer having a first address and a first key; and

accessing the encrypted data stored at the first memory cell by decrypting the data using the first key.

29. A method of accessing encrypted data in a secure heap data structure according to claim 28 wherein, the decrypted data includes a second crypto-pointer associated with a second memory cell in the secure heap data structure, the second crypto-pointer having a second address and a second key, further comprising accessing the encrypted data stored at the second memory cell by decrypting the data using the second key.

1/5




Figure 1. Structure of crypto-pointer

FIG. 1

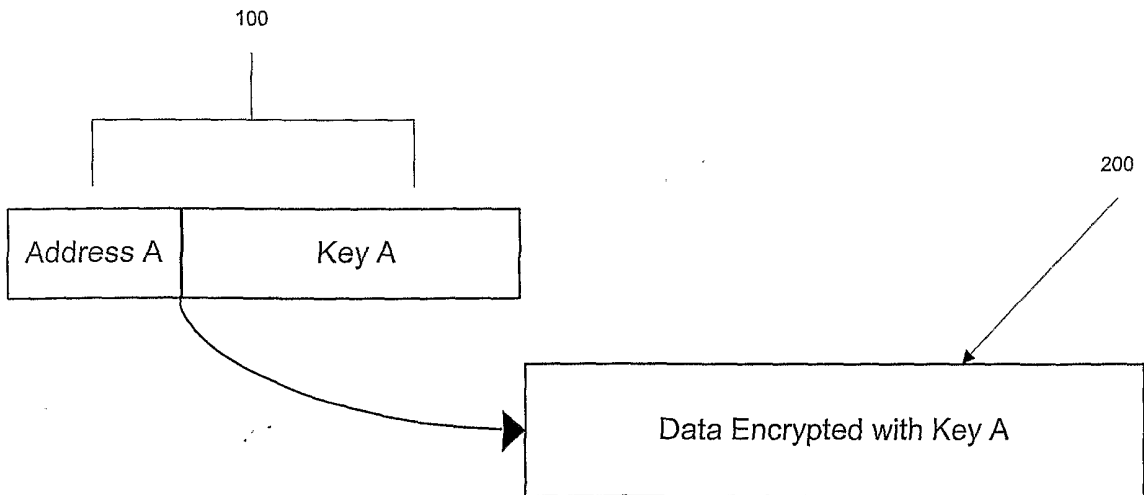


Figure 2. Crypto-pointer always points to cell encrypted with associated key

FIG. 2

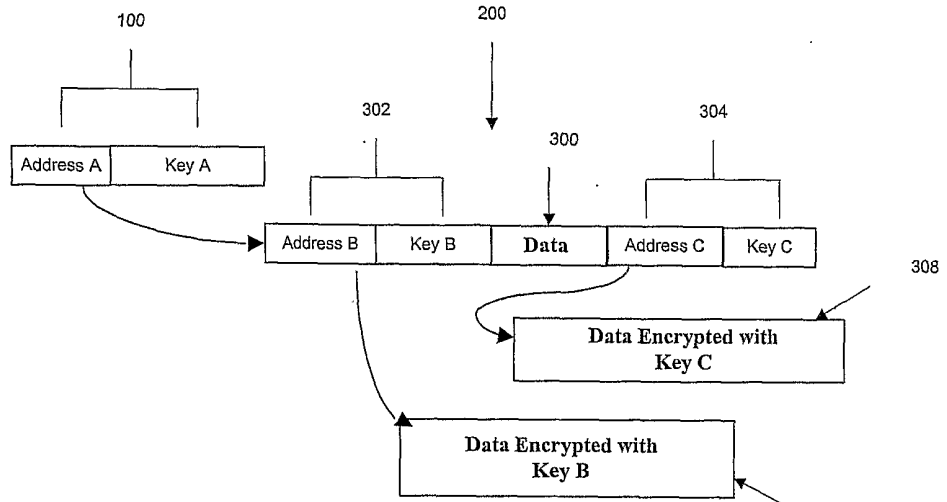


Figure 3. Each step in a chain of crypto-pointers uses a separate key.

FIG. 3

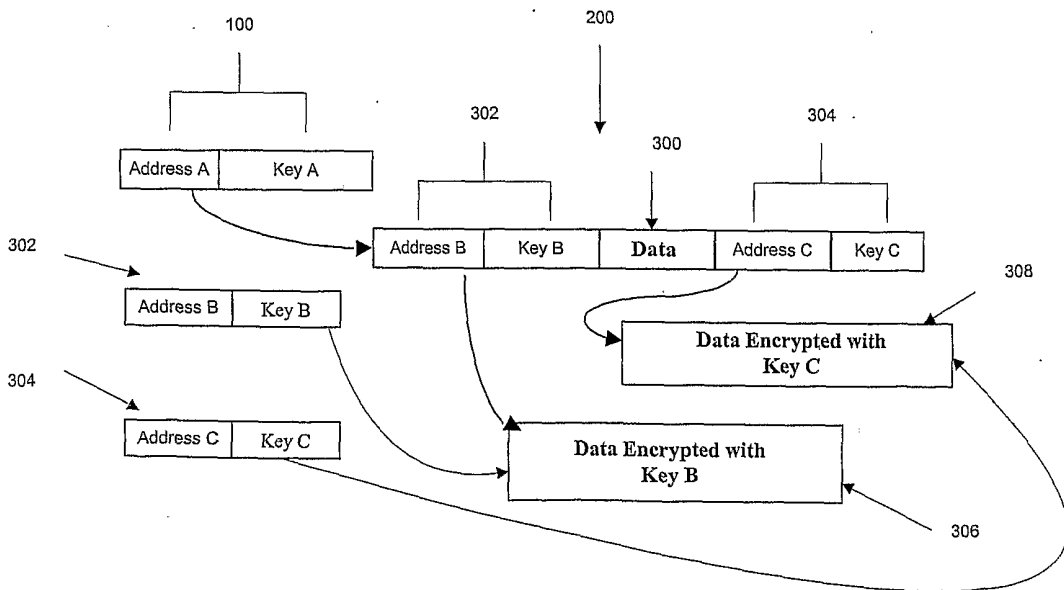


Figure 4. Decrypting cell contents with Key A reveals subsequent crypto-pointers

FIG. 4

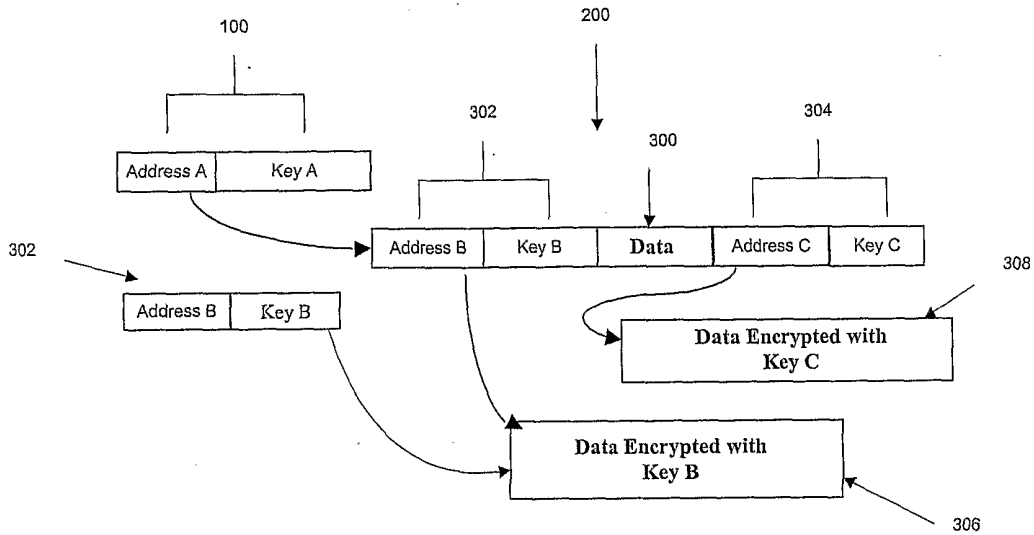


Figure 5. Discarding crypto-pointer C still permits indirect access to the C structure, via A again

FIG. 5

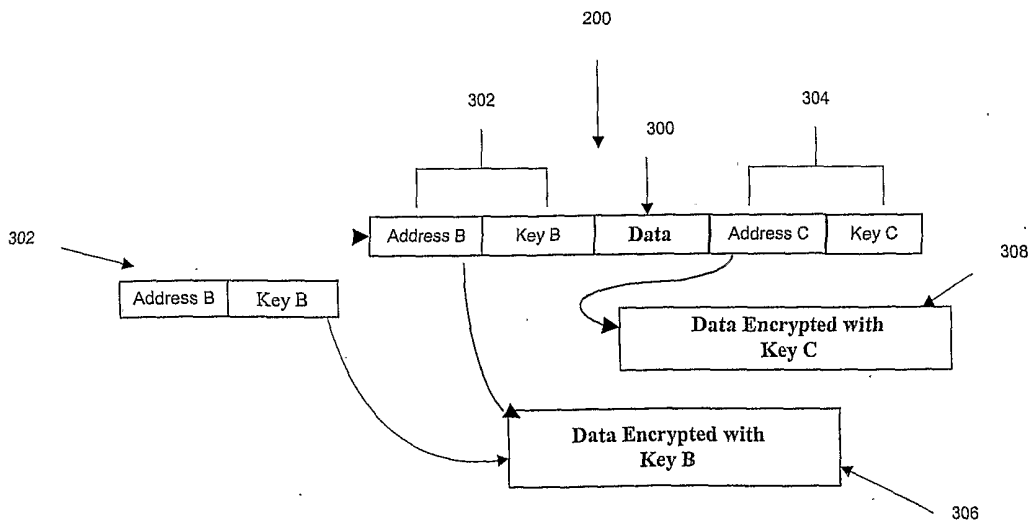


Figure 6. Discarding crypto-pointer A makes Data and the C structure forever inaccessible

FIG. 6

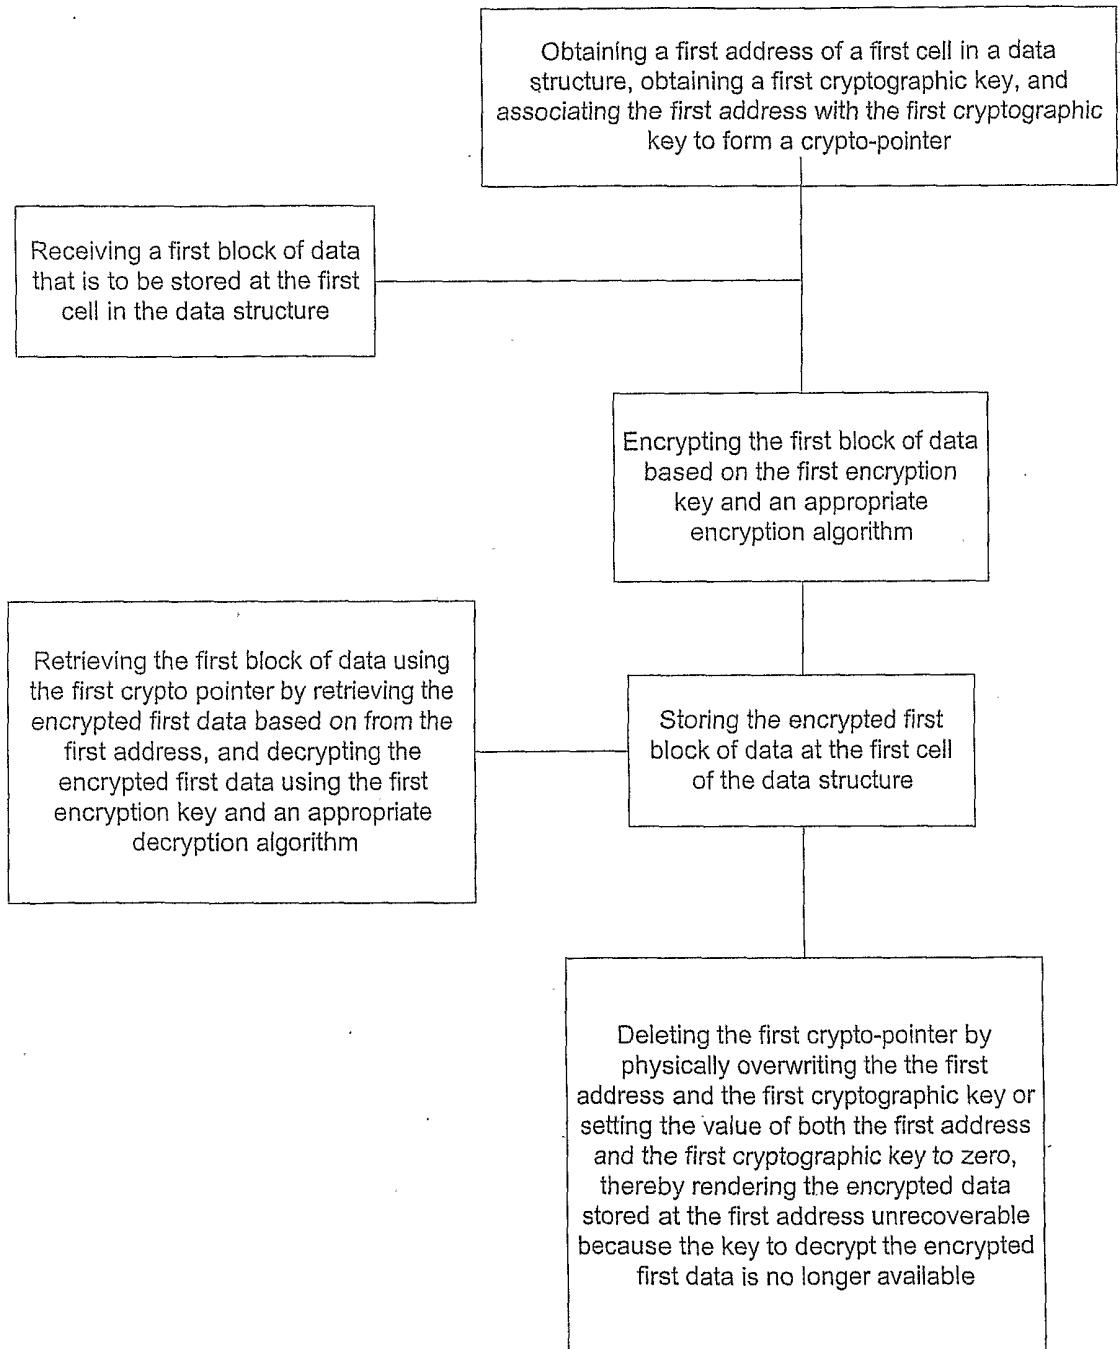


Fig. 7



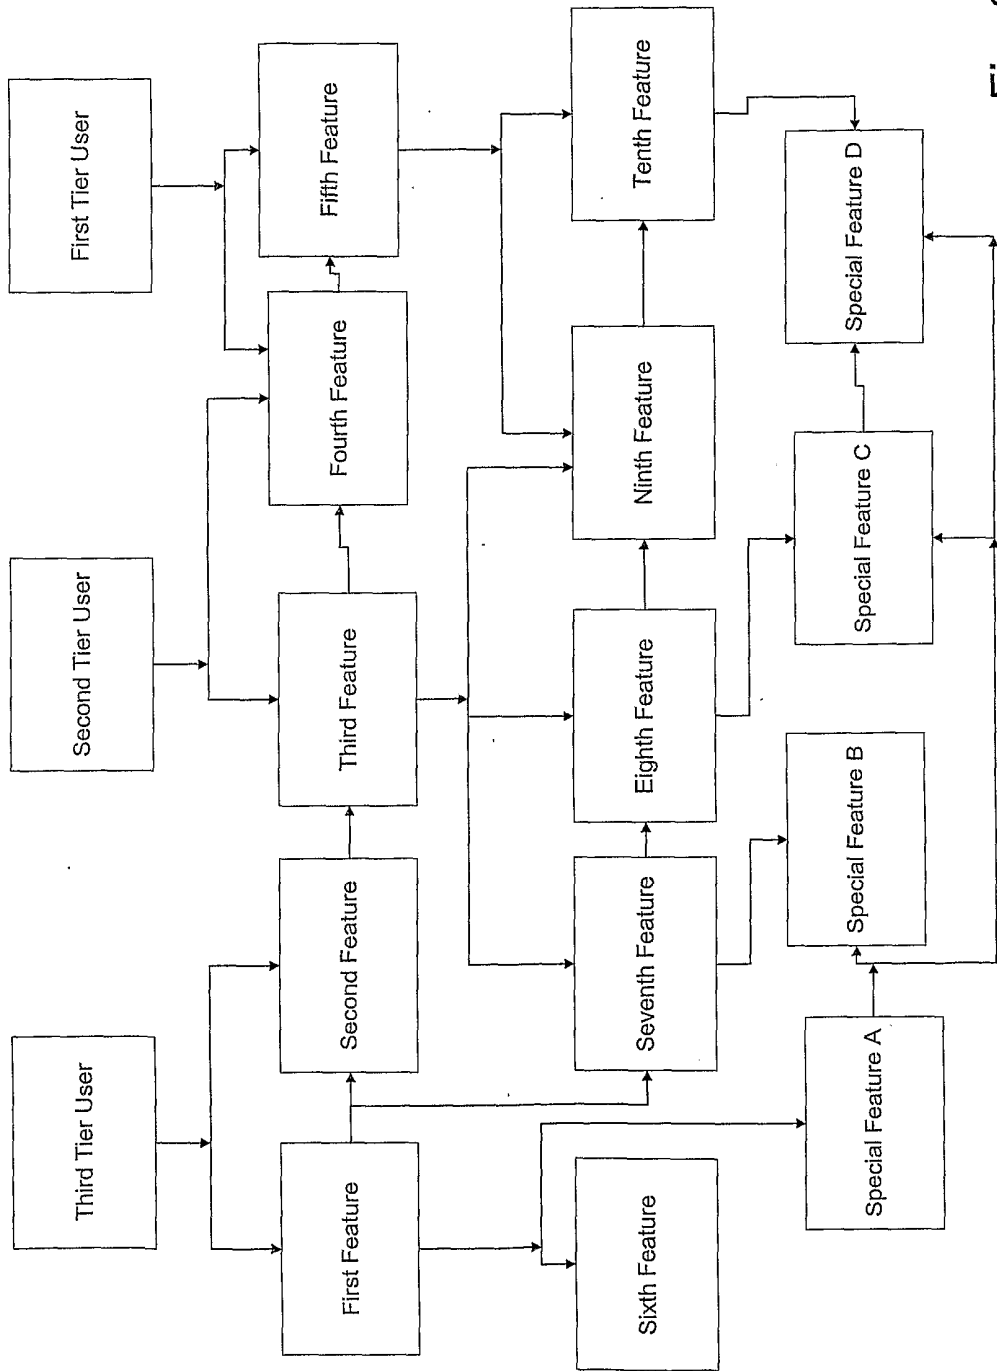


Fig. 8

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/09426

**A. CLASSIFICATION OF SUBJECT MATTER**  
 IPC(7) : G06F 17/30  
 US CL : 713/190  
 According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**  
 Minimum documentation searched (classification system followed by classification symbols)  
 U.S. : 713/190;707/9,101;380/30

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
 Please See Continuation Sheet

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,E --- Y	US 2002/0073326 A1 (FONTIJN) 13 June 2002 (13.06.2002), abstract, pg. 1, paragraph 10, pg. 2, paragraph 22,23, pg.3, paragraph 29, pg.3,claims 1-7, 11.	1,3-6,8-12,14-29 ----- 2,7,13
Y	US 4,698,617 A (BAUER) 06 October 1987 (06.10.1987), abstarct, column 1, lines 61-67, column 2, lines 60-67column 3, lines 11-54, column 6, lines 45-62 (Tables 1 and 2).	2,7,13
Y	US 5,058,164 A (ELMER et al) 15 October 1991 (15.10.1991), the entire document.	1-29
Y	US 5,563,945 A (GERCEKCI) 08 October 1996 (08.10.1996), the entire document.	1-29

Further documents are listed in the continuation of Box C.       See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier application or patent published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search: 27 June 2002 (27.06.2002)  
 Date of mailing of the international search report: 24 JUL 2002

Name and mailing address of the ISA/US: Commissioner of Patents and Trademarks, Box PCT, Washington, D.C. 20231, Facsimile No. (703)305-3230  
 Authorized officer: Gail O Hayes (with initials FM), Telephone No. (703)305-4274  
 Signature: James R. Matthews

**INTERNATIONAL SEARCH REPORT**

International application No.

PCT/US02/09426

**Continuation of B. FIELDS SEARCHED Item 3:**

WEST, DIALOG, ProQuest, Dogpile. Search terms: storage device, pointer or link or address, encryption and addressed based key, file system and address location, encryption, encipher and file system