



US 20070211292A1

(19) **United States**

(12) **Patent Application Publication**
Walmsley

(10) **Pub. No.: US 2007/0211292 A1**

(43) **Pub. Date: Sep. 13, 2007**

(54) **METHOD OF STORING CODE SEGEMENTS
IN PLURAL PRINTER CARTRIDGES**

Publication Classification

(75) Inventor: **Simon Robert Walmsley**, Balmain
(AU)

(51) **Int. Cl.**
B41B 1/02 (2006.01)
(52) **U.S. Cl.** **358/1.16**

Correspondence Address:
SILVERBROOK RESEARCH PTY LTD
393 DARLING STREET
BALMAIN 2041 (AU)

(57) **ABSTRACT**

A method of storing at least one functionally identical code segment in each of a plurality of printer cartridges, each having a memory, of a printing system is provided, in which a first memory location is determined which is different than a first memory location in at least a plurality of the other printer cartridges, and a first of the at least one code segments is stored in the memory at the determined first memory location. The different first memory locations are determined such that the code segment storage location of all the respective devices cannot be ascertained from the code segment storage location of any one of the printer cartridges.

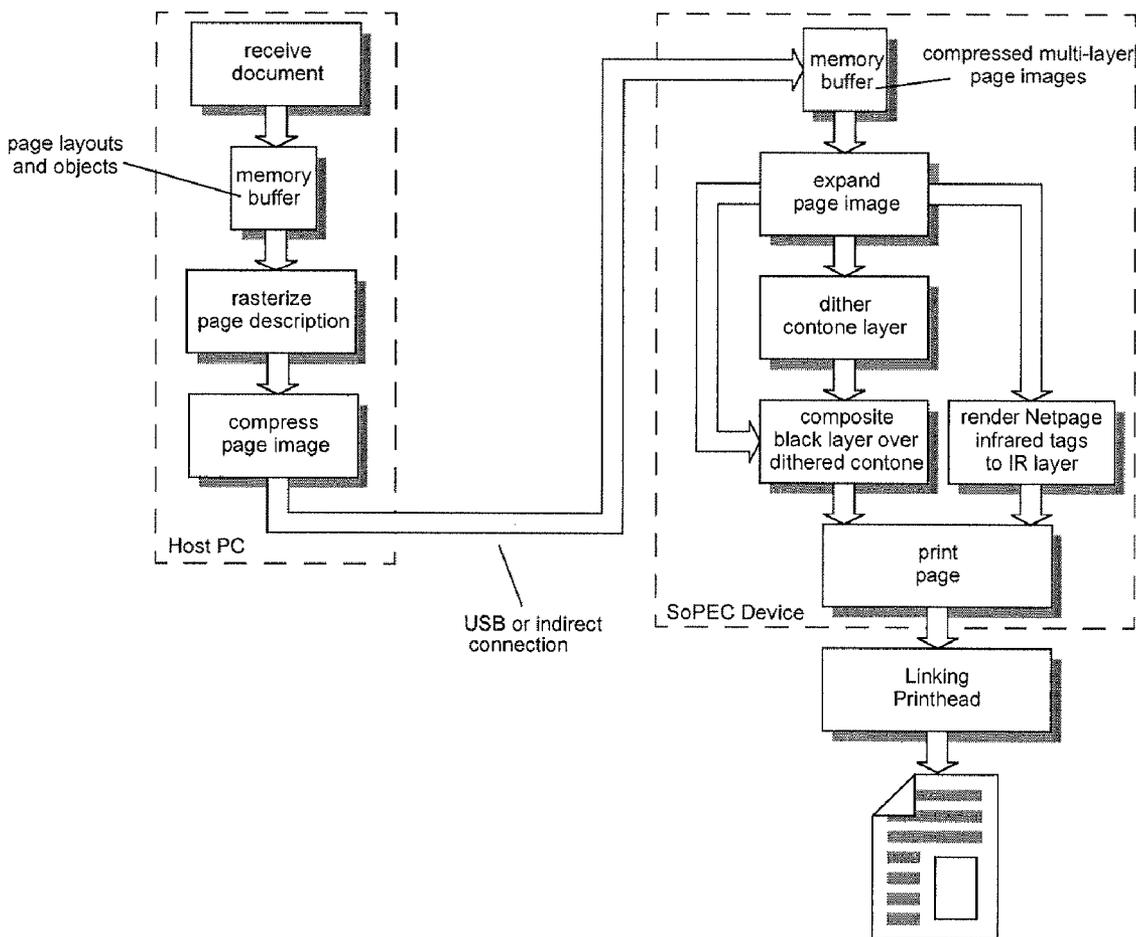
(73) Assignee: **Silverbrook Research Pty Ltd**

(21) Appl. No.: **11/749,123**

(22) Filed: **May 15, 2007**

Related U.S. Application Data

(63) Continuation of application No. 10/854,502, filed on May 27, 2004, now Pat. No. 7,243,193.



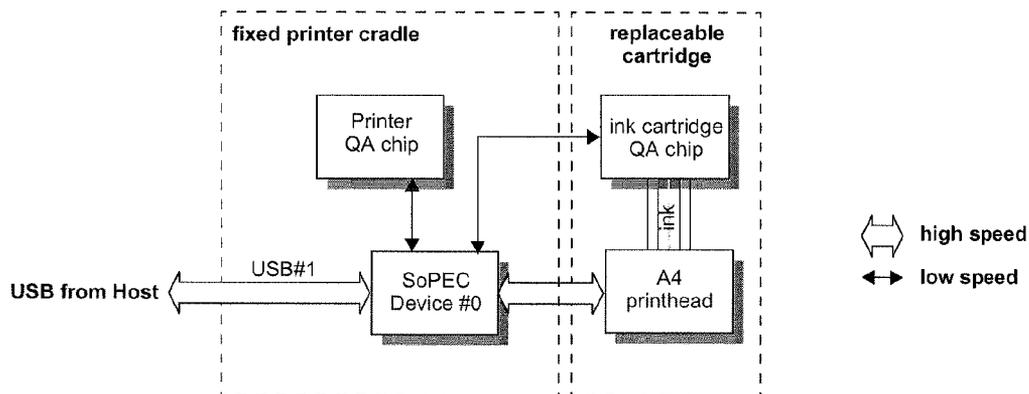


FIG. 1

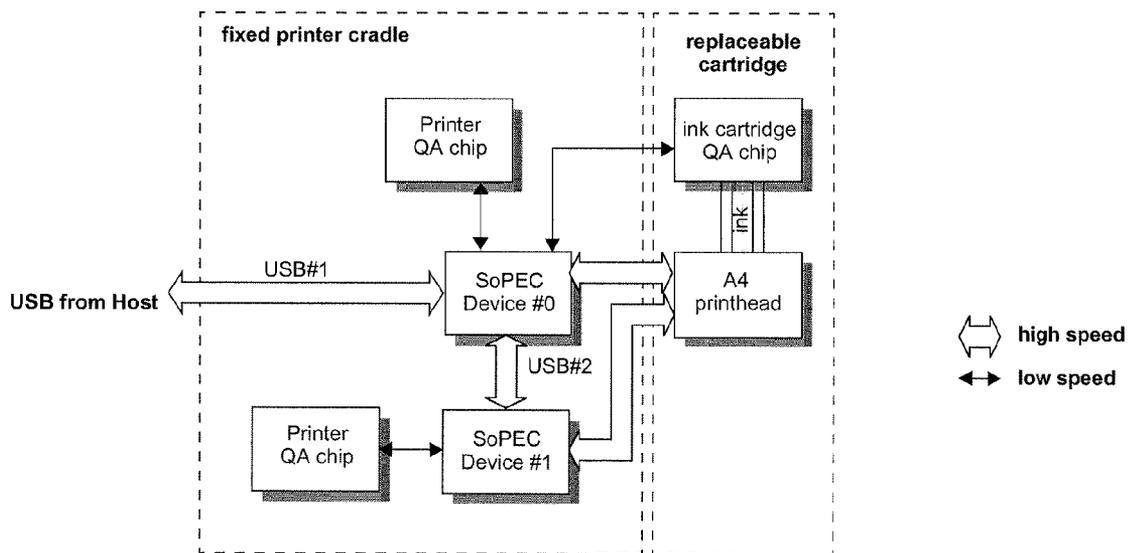


FIG. 2

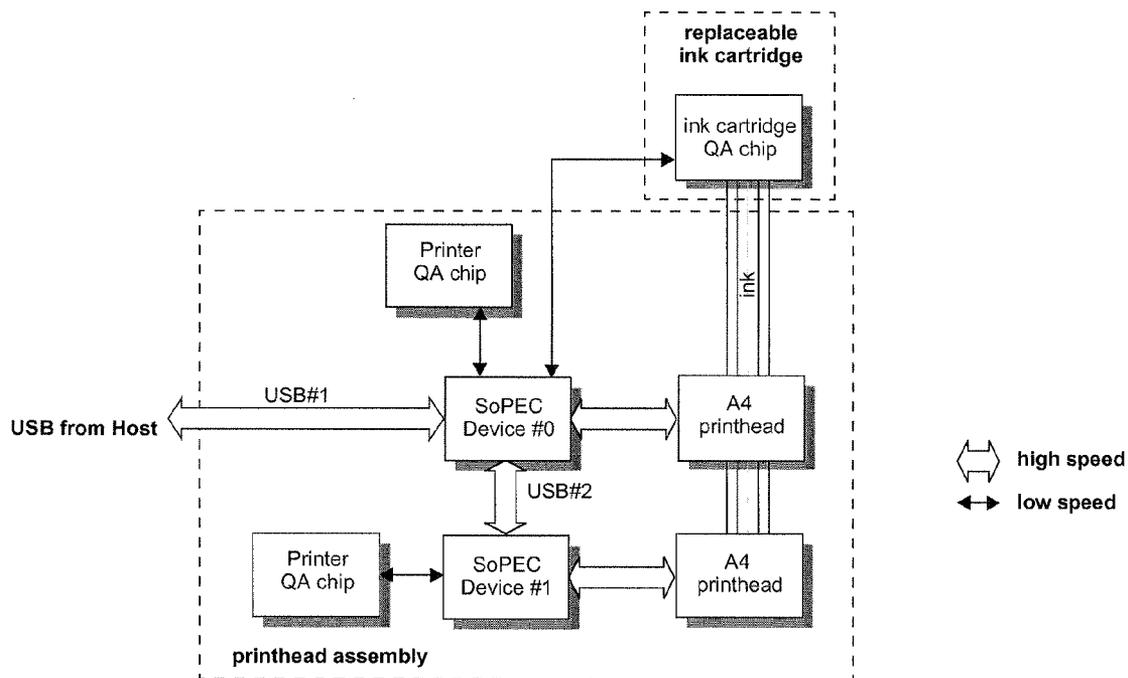


FIG. 3

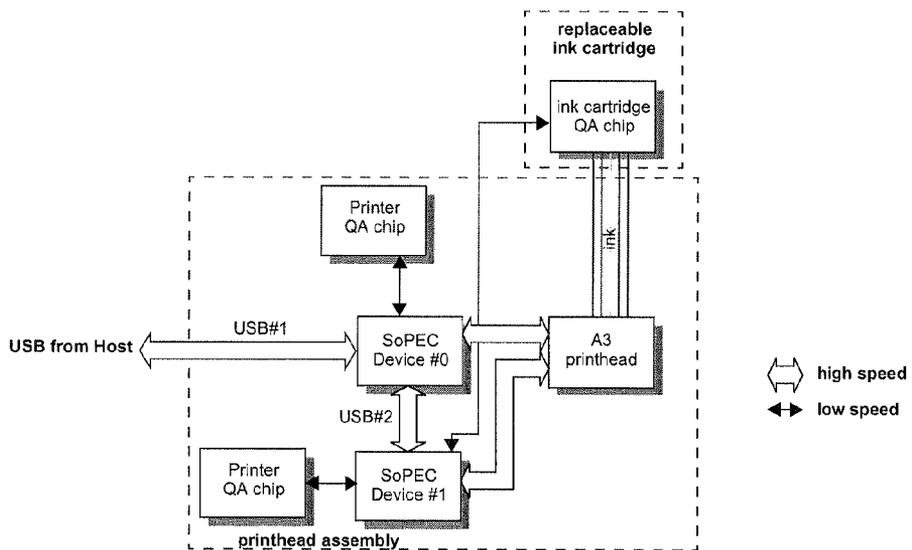


FIG. 4

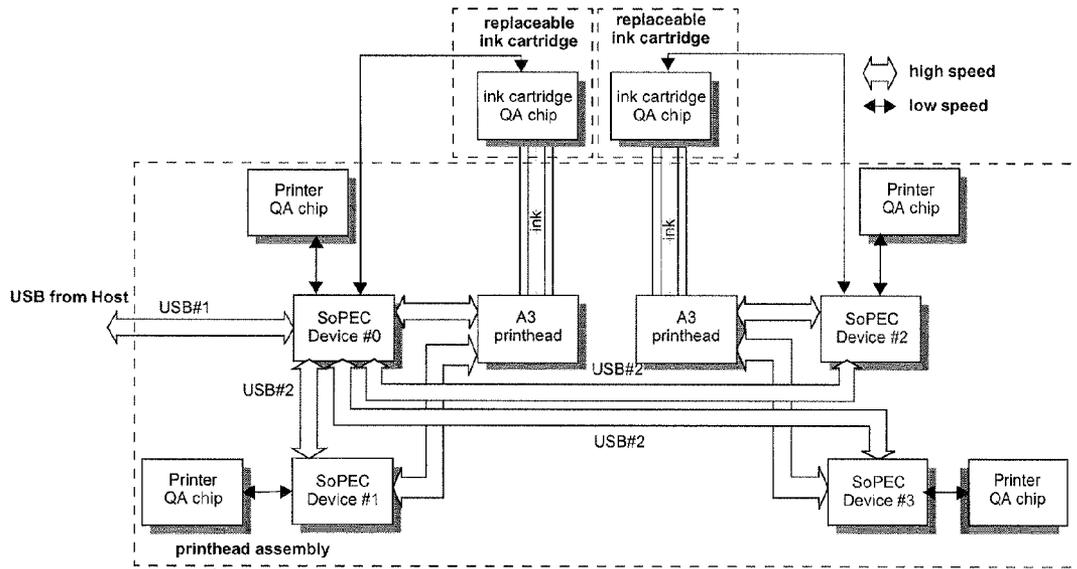


FIG. 5

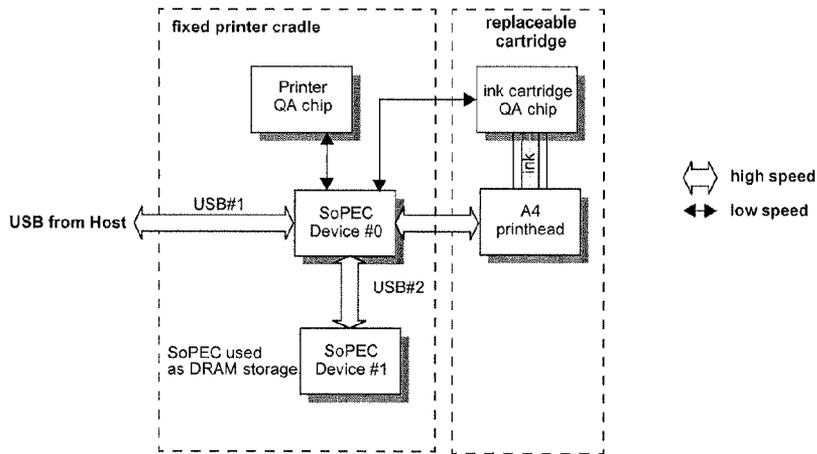


FIG. 6

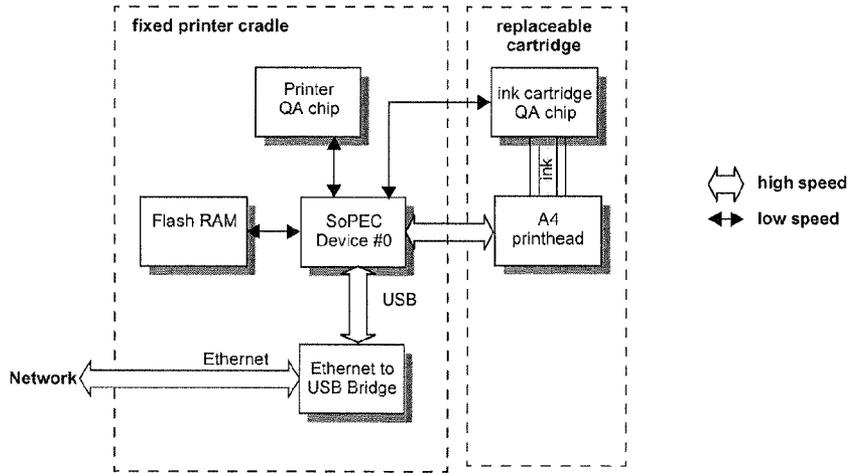


FIG. 7

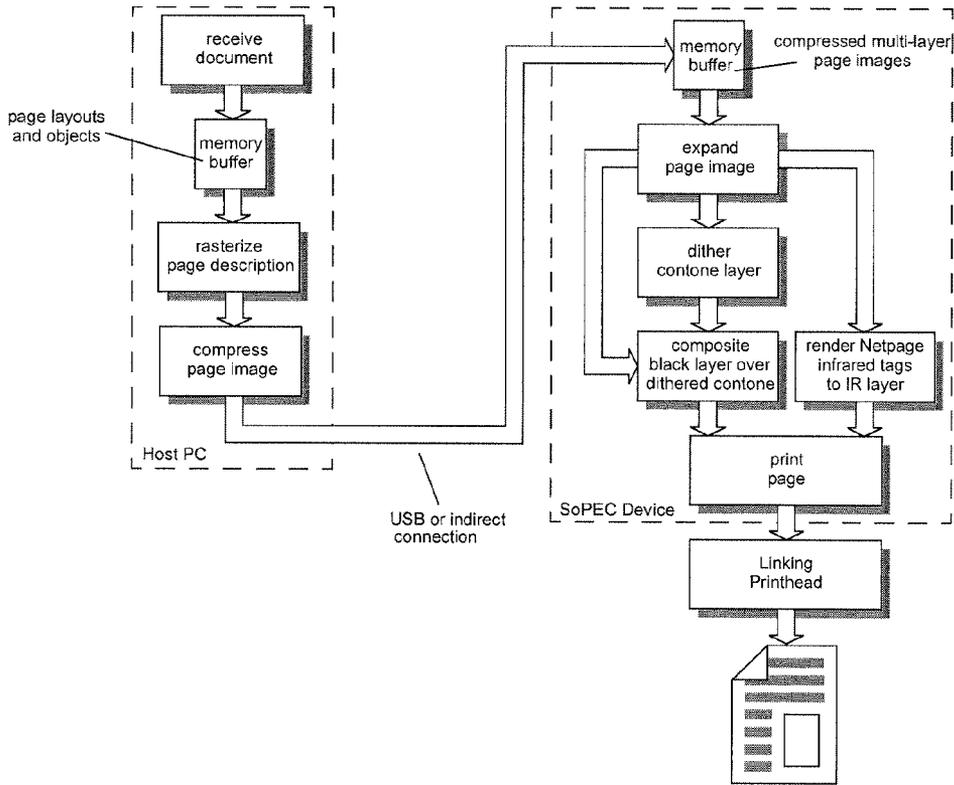


FIG. 8

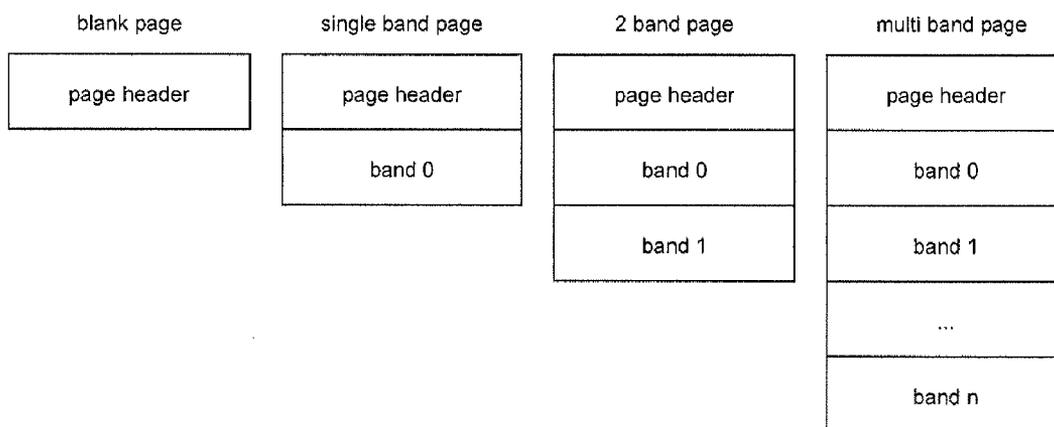


FIG. 9

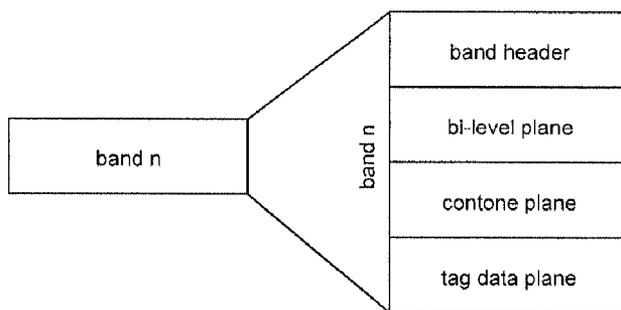


FIG. 10

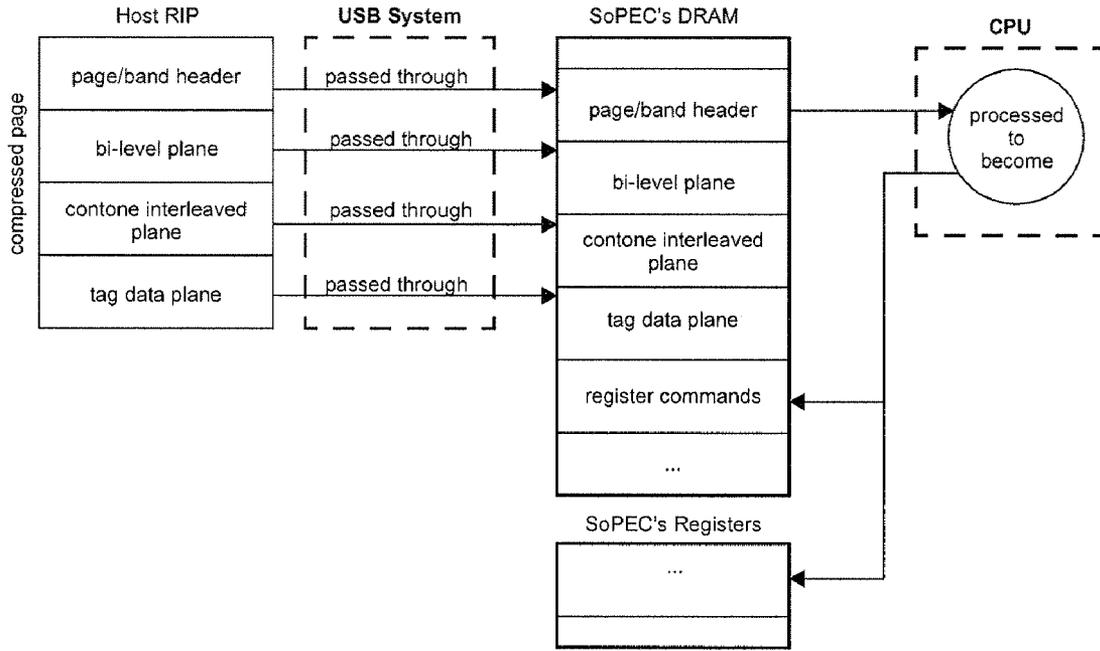


FIG. 11

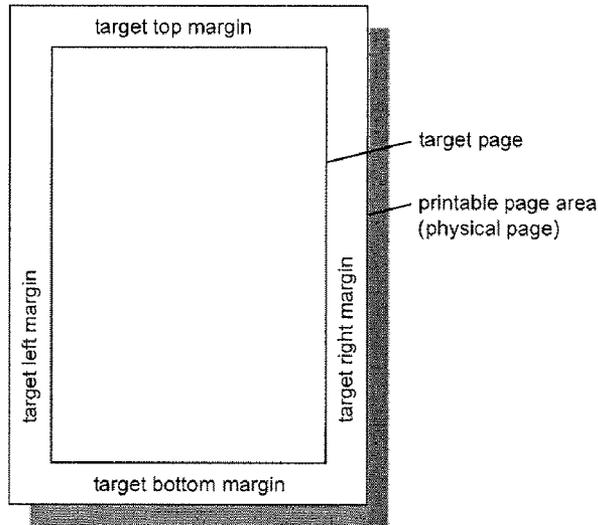


FIG. 12

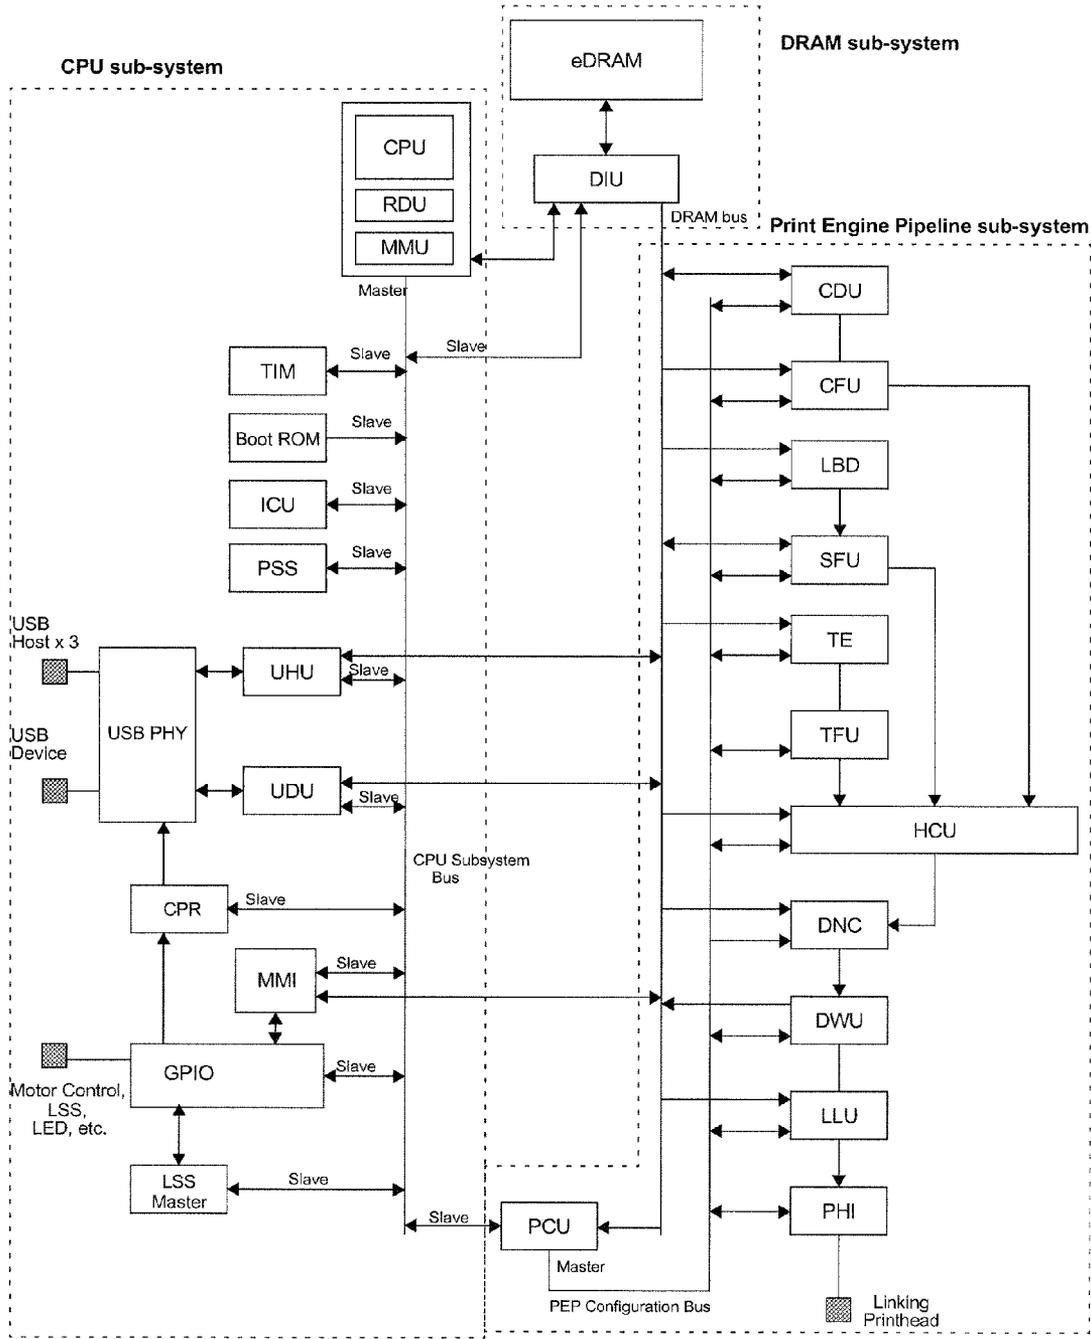


FIG. 13

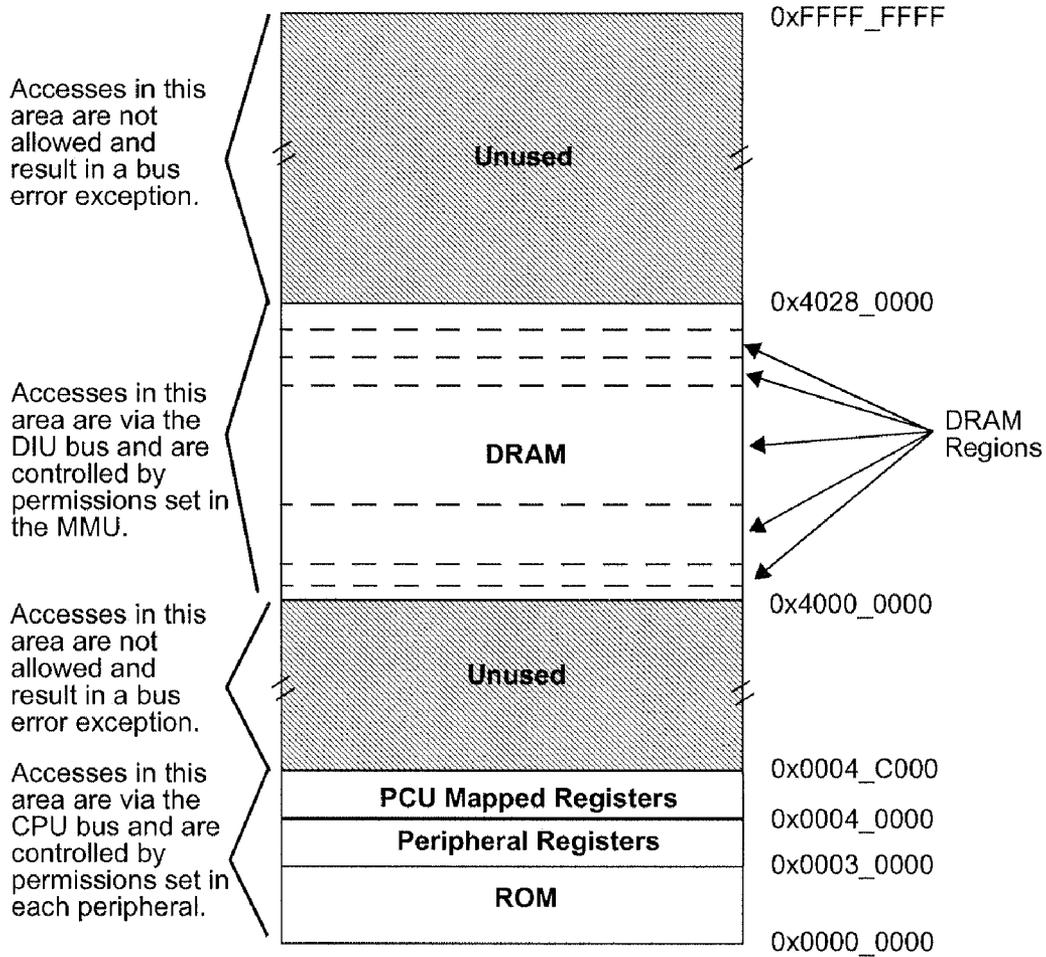
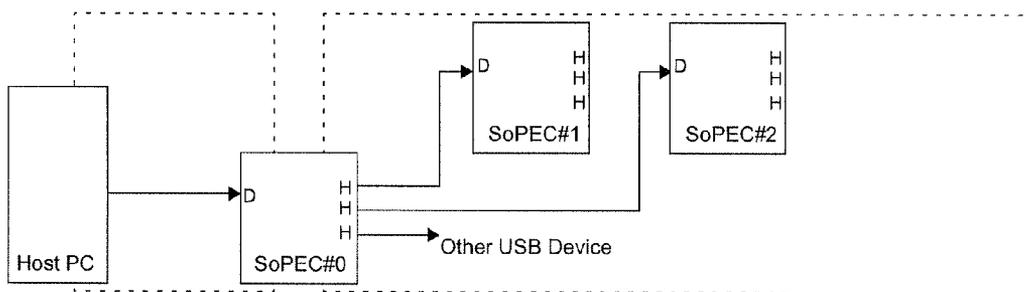
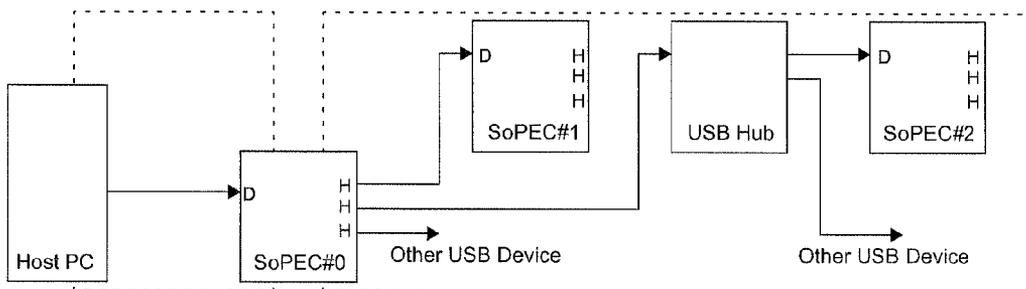


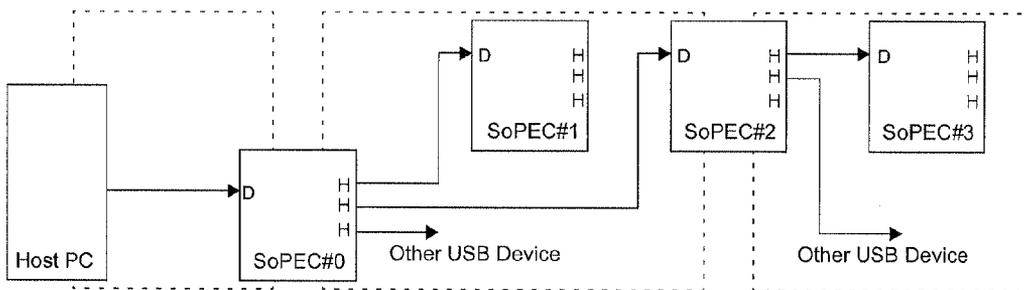
FIG. 14



Case 1: One Printer USB Bus with no Hub chips, up to 3 Devices on the bus



Case 2: One Printer USB Bus including Hub chip, more than 3 Devices on the bus



Case 3: Two Printer USB Busses, up to 3 devices on each bus

FIG. 15

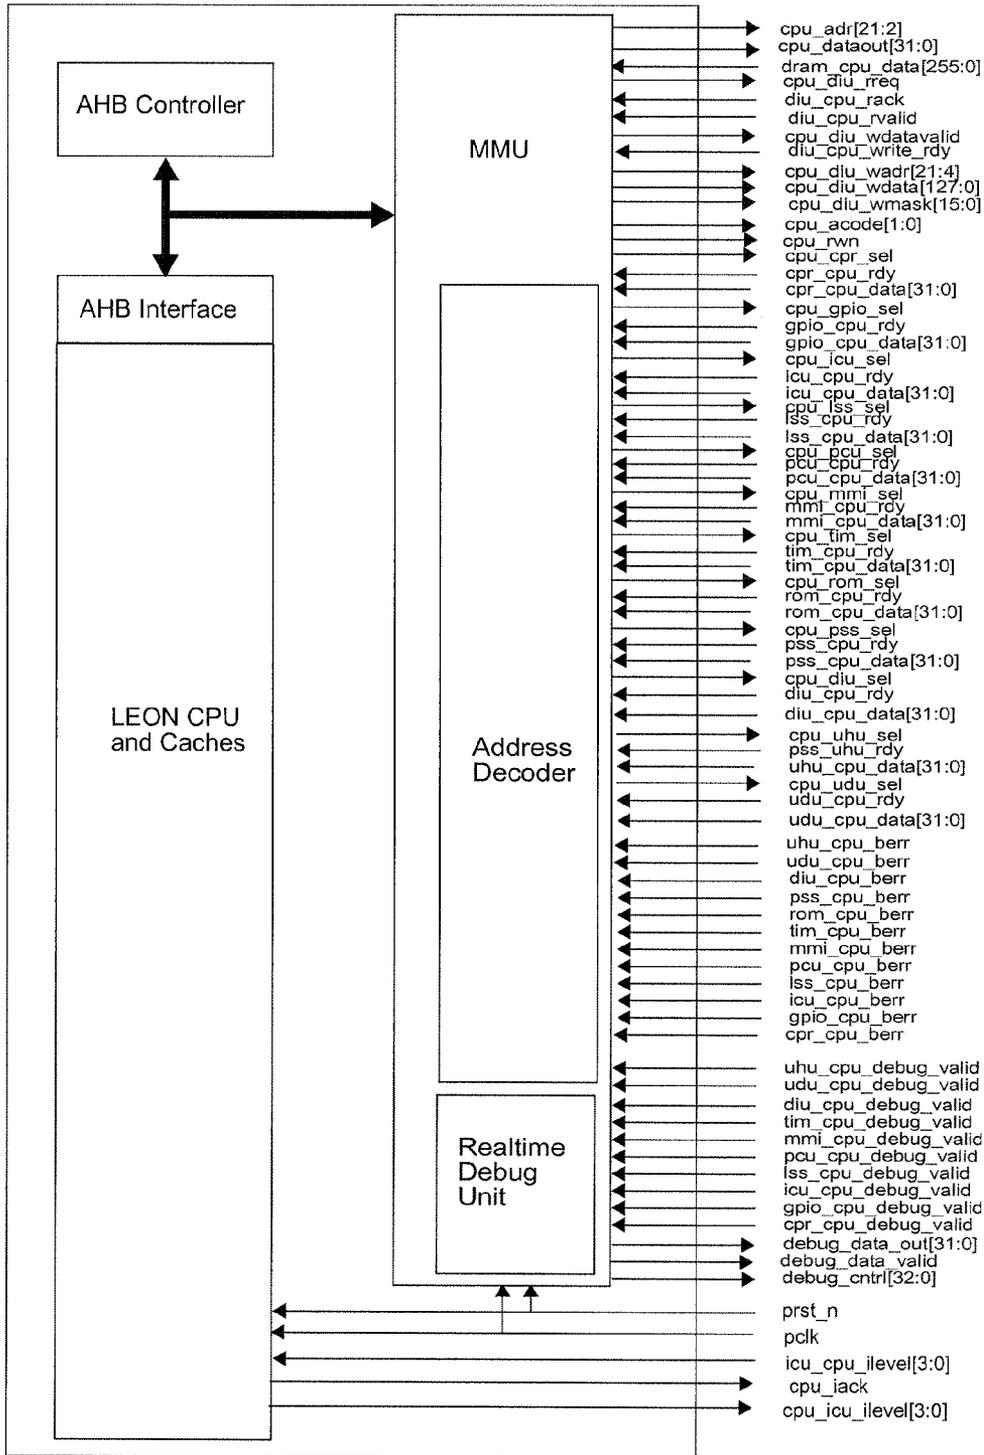


FIG. 16

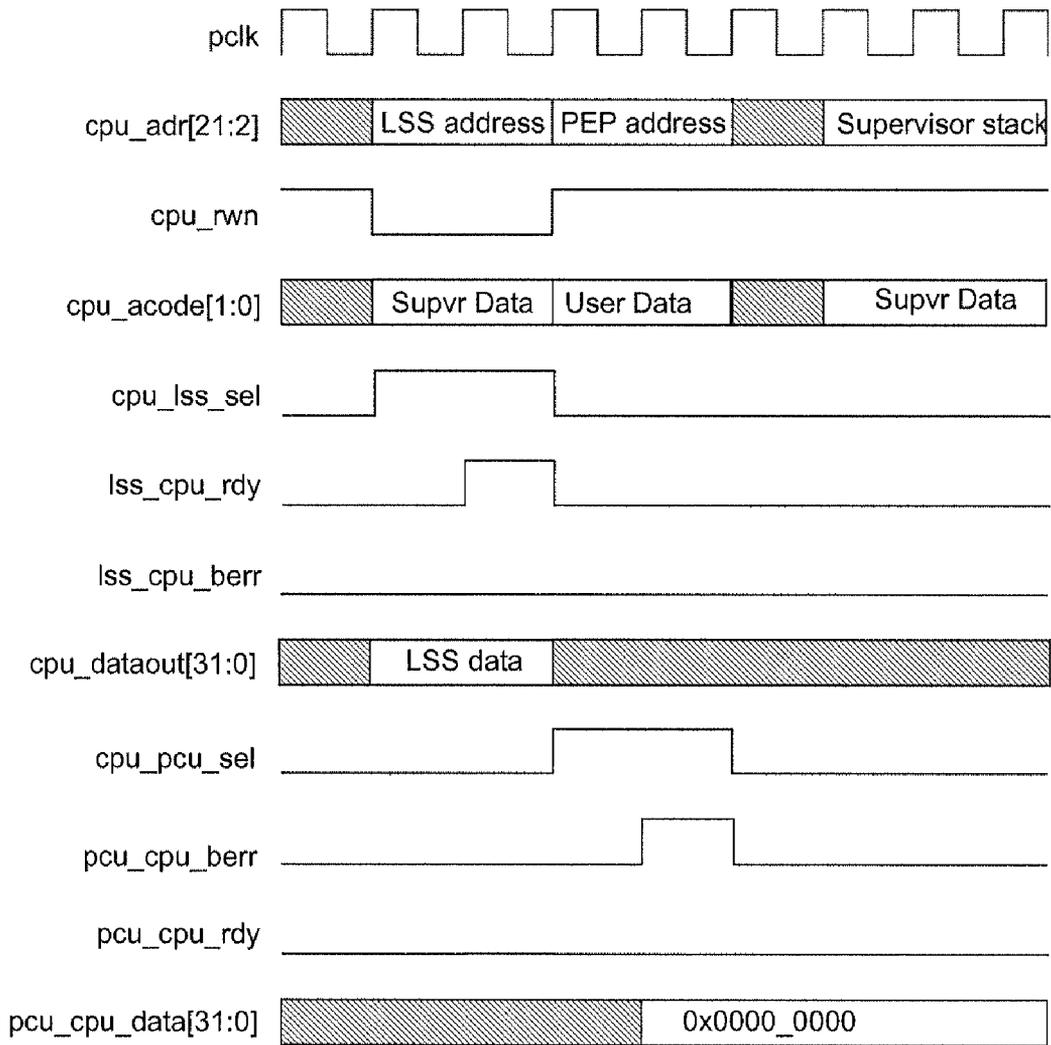


FIG. 17

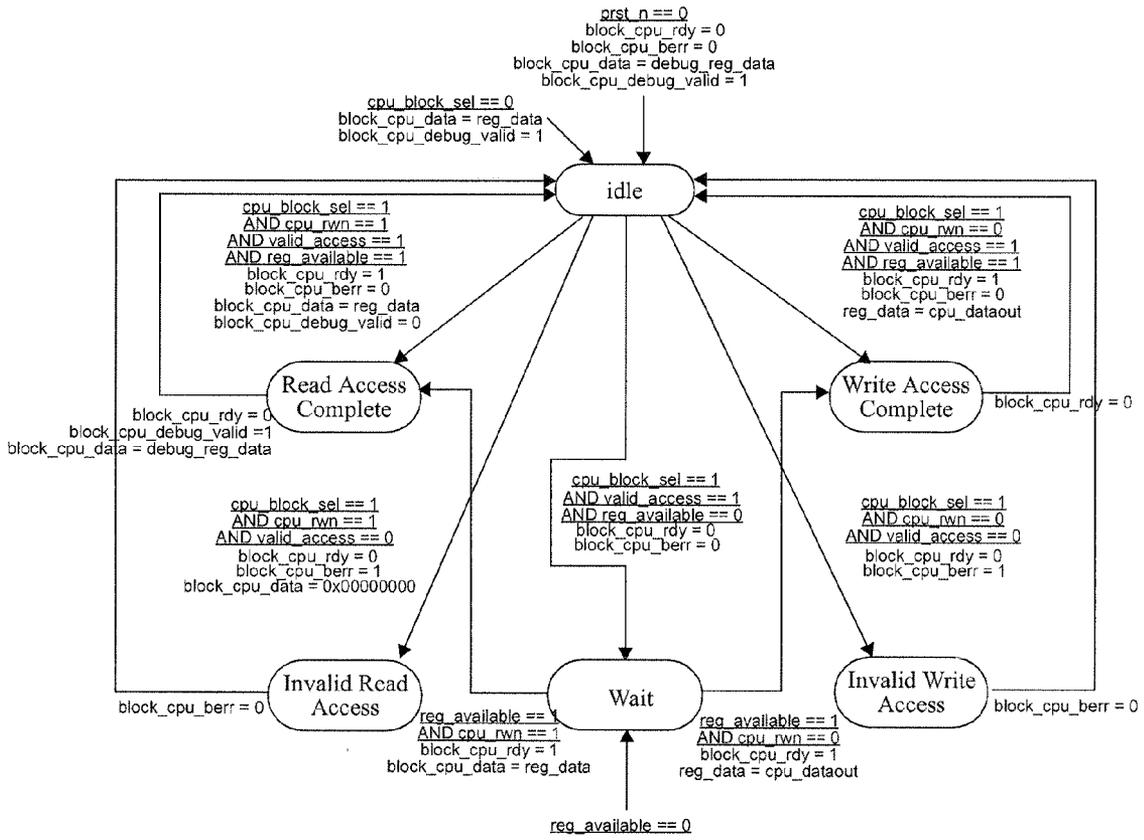


FIG. 18

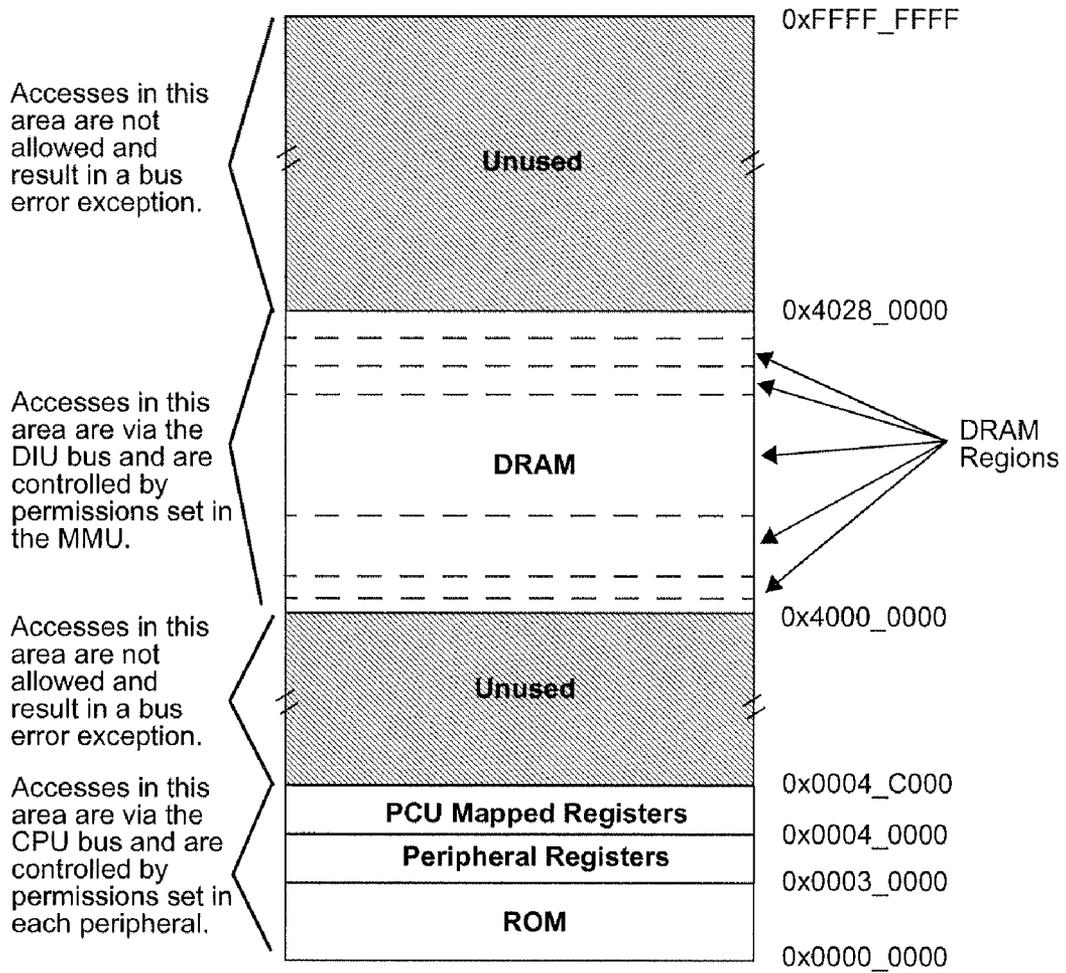


FIG. 19

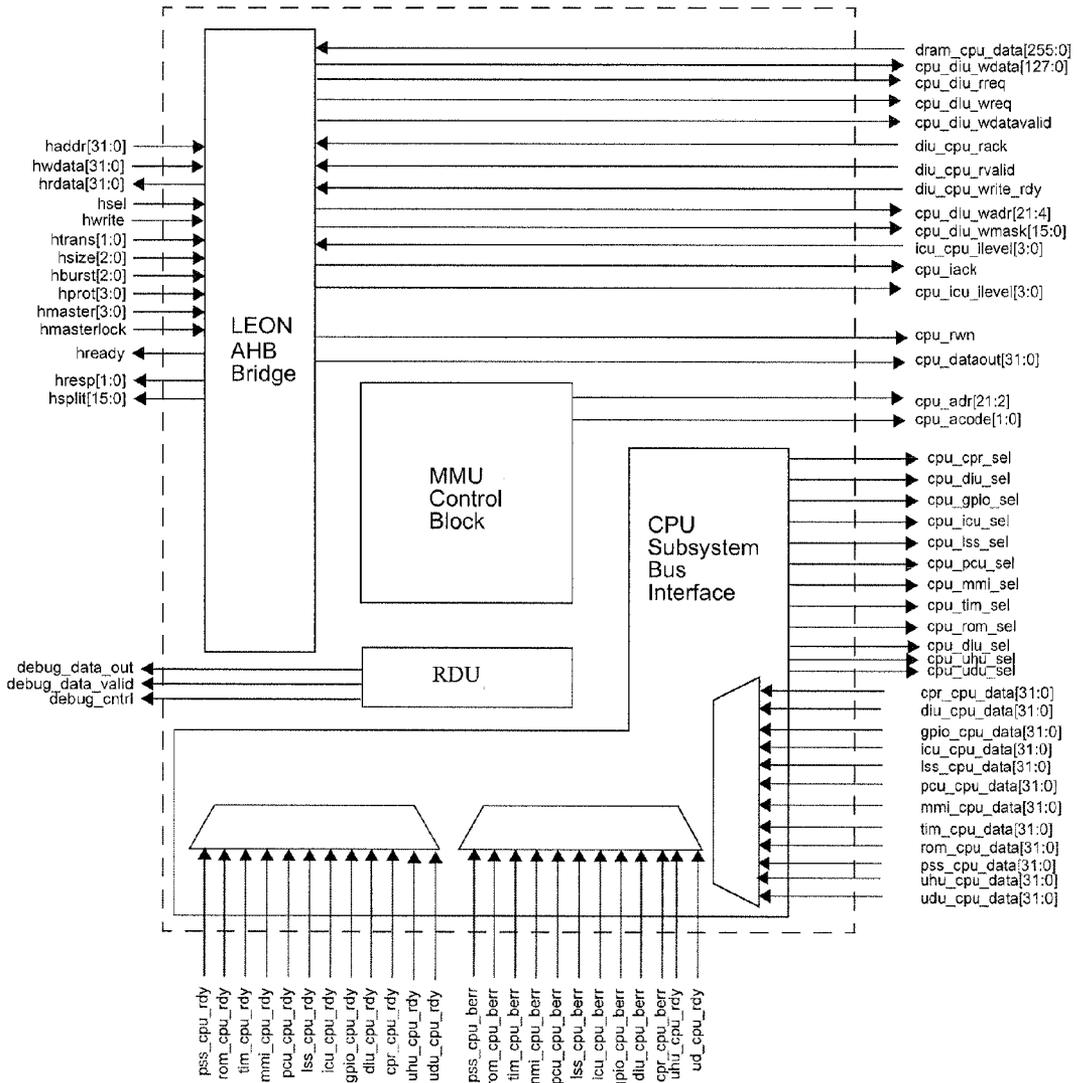


FIG. 20

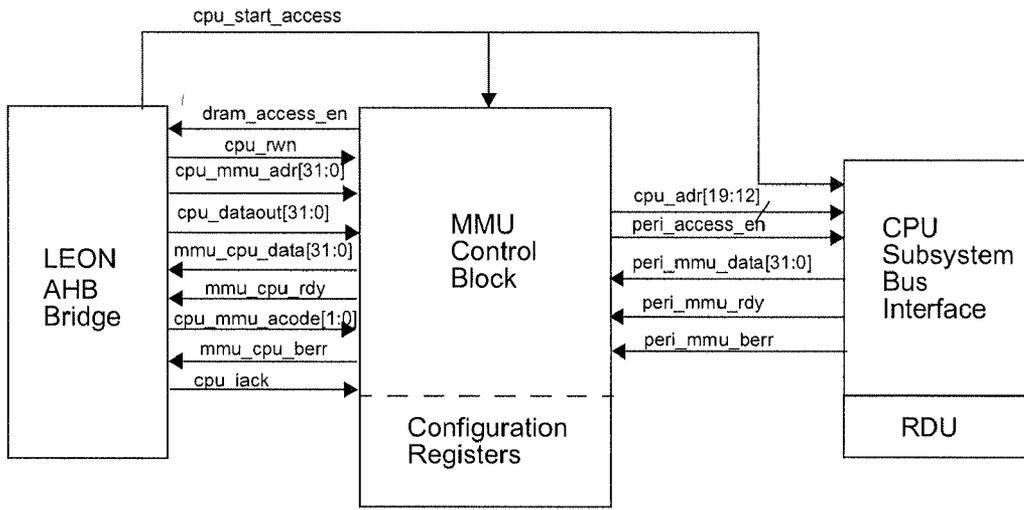


FIG. 21

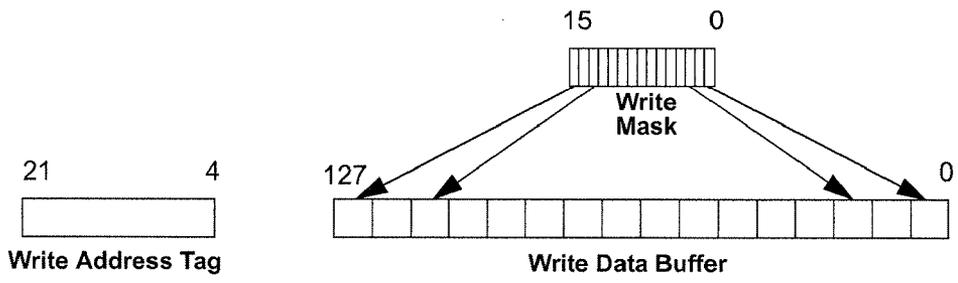


FIG. 22

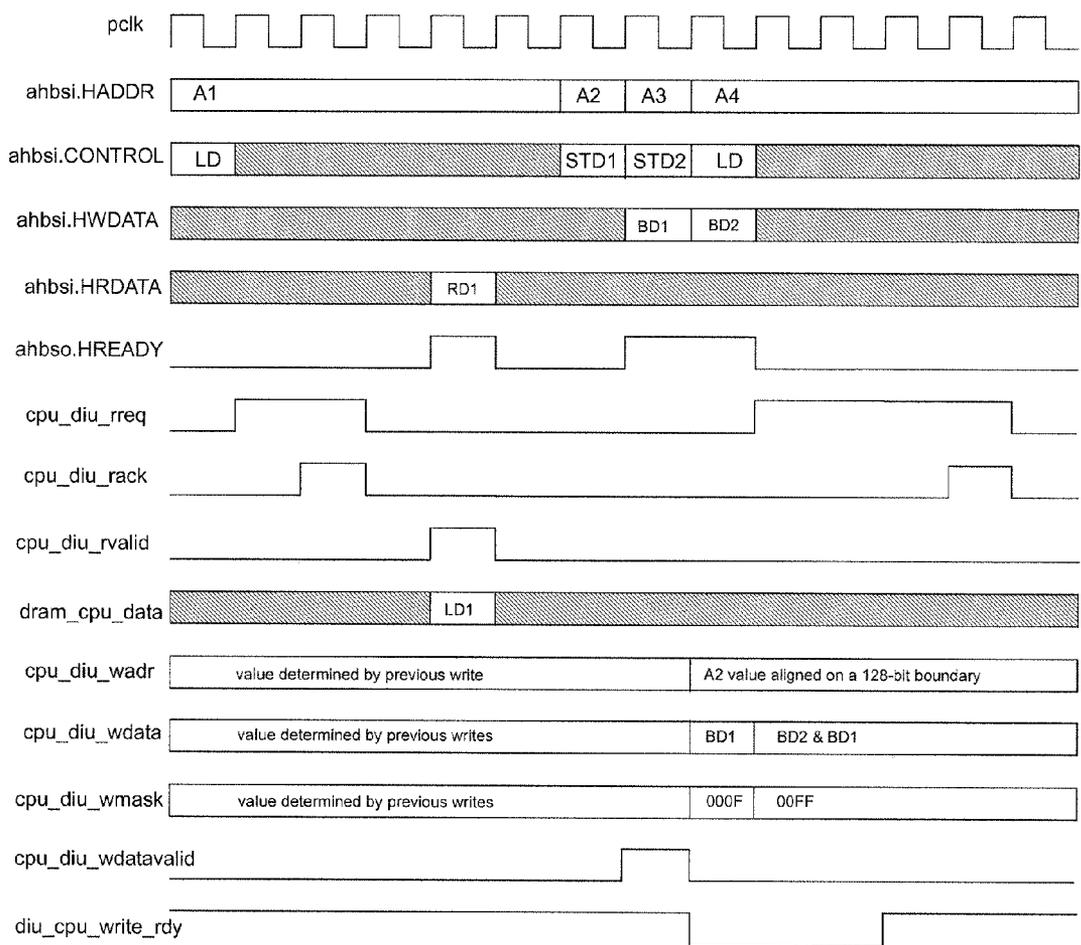


FIG. 23

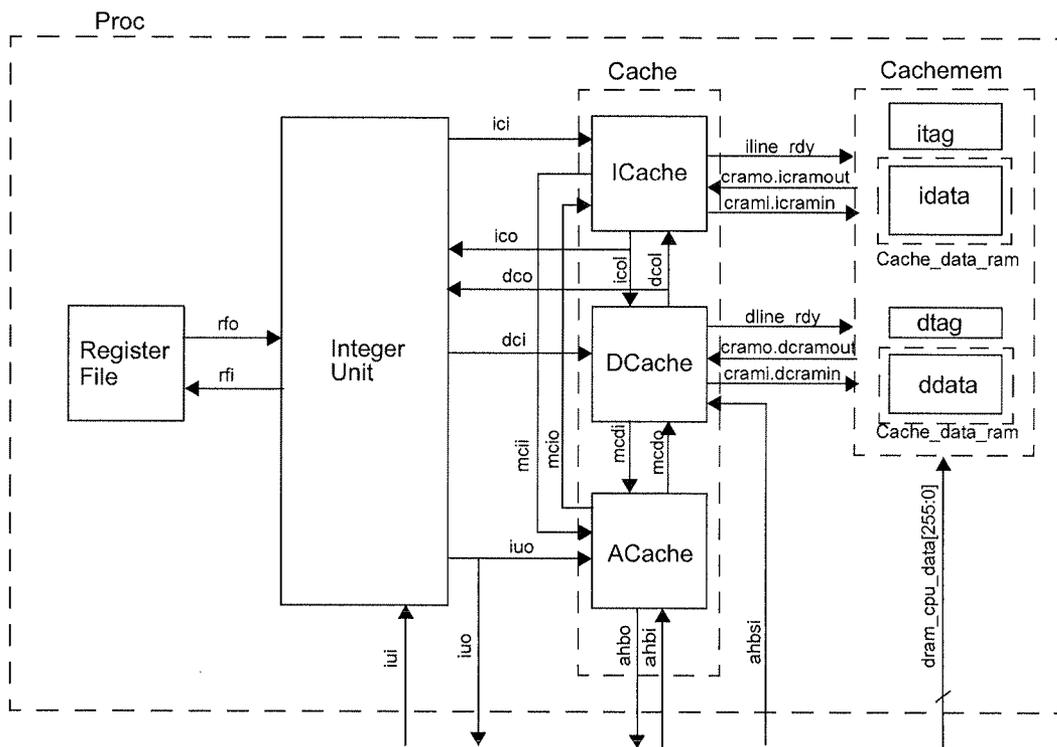


FIG. 24

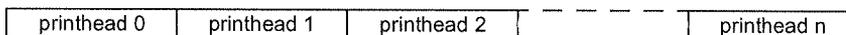


FIG. 25

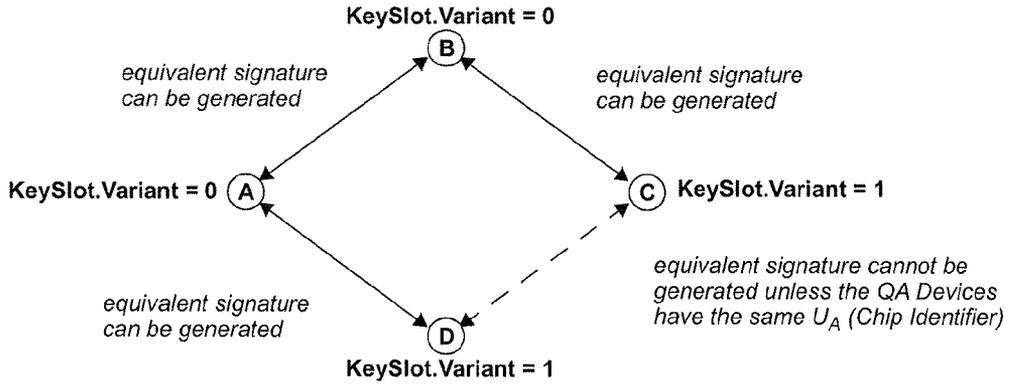


FIG. 26

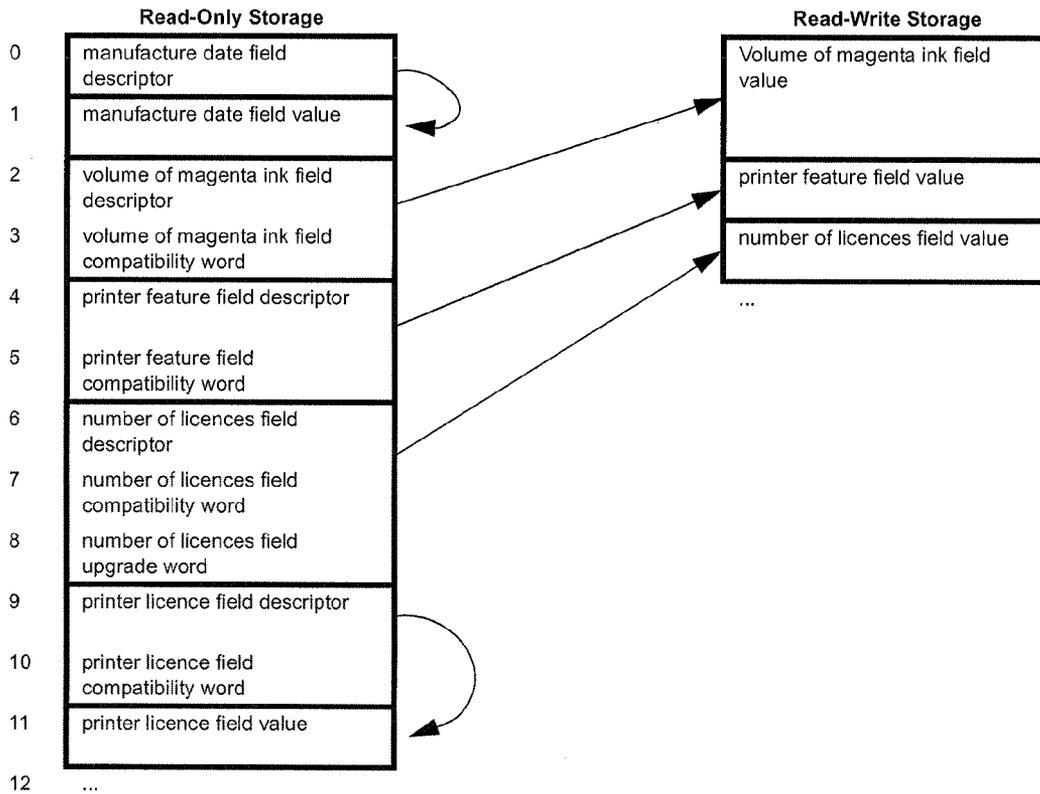


FIG. 27

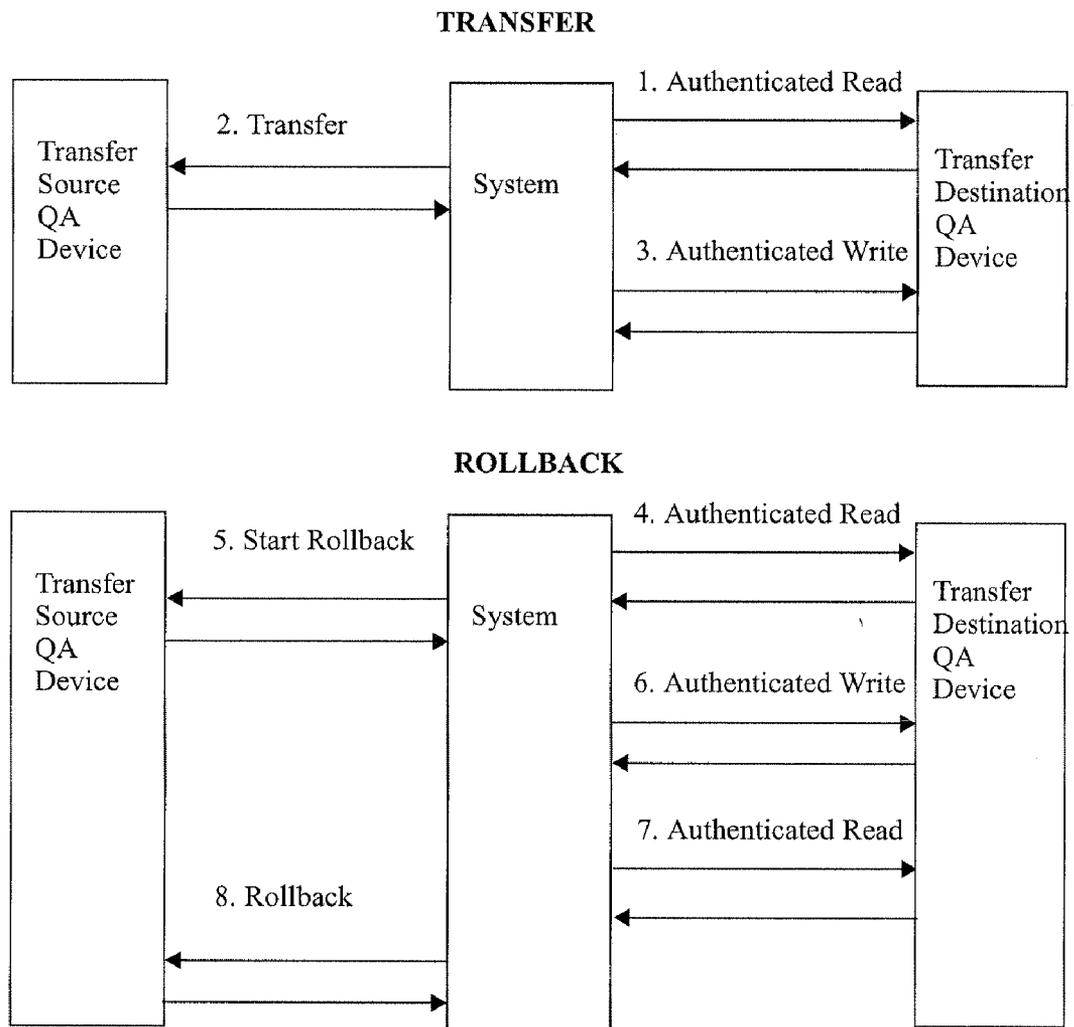


FIG. 28

METHOD OF STORING CODE SEGMENTS IN PLURAL PRINTER CARTRIDGES

CROSS REFERENCE TO RELATED APPLICATION

[0001] This application is a continuation of U.S. application Ser. No. 10/854,502 filed on May 27, 2004 all of which are herein incorporated by reference.

FIELD OF THE INVENTION

[0002] The present invention relates to the storage of bit-patterns in non-volatile memory of a printer cartridge. The invention has been developed primarily for storing one or more keys in an integrated circuit, and will be described with reference to this application. However, it will be appreciated that the invention can be applied in a number of other fields where it is desirable to store bit-patterns in non-volatile memory.

CO-PENDING APPLICATIONS

[0003] Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending applications filed by the applicant or assignee of the present invention simultaneously with the present application:

10/854521	10/854522	10/854488	10/854487	10/854503	10/854504
10/854509	7188928	7093989	10/854497	10/854495	10/854498
10/854511	10/854512	10/854525	10/854526	10/854516	10/854508
10/854507	10/854515	10/854506	10/854505	10/854493	10/854494
10/854489	10/854490	10/854492	10/854491	10/854528	10/854523
10/854527	10/854524	10/854520			

The disclosures of these co-pending applications are incorporated herein by cross-reference.

CROSS-REFERENCES

[0004] Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending applications filed by the applicant or assignee of the present invention. The disclosures of all of these co-pending applications are incorporated herein by cross-reference.

09/517539	6566858	6331946	6246970	6442525	09/517384
09/505951	6374354	09/517608	6816968	6757832	6334190
6745331	09/517541	10/636263	10/636283	10/407212	10/407207
10/683064	10/683041	10/727181	10/727162	10/727163	10/727245
7121639	7165824	7152942	10/727157	7181572	7096137
10/727257	10/727238	7188282	10/727159	10/727180	10/727179
10/727192	10/727274	10/727164	10/727161	10/727198	10/727158
10/754536	10/754938	10/727227	10/727160	6795215	6859289
6977751	6398332	6394573	6622923	6747760	6921144
10/780624	10/791792	7182267	7025279	6857571	6817539
6830198	6992791	7038809	6980323	7148992	7139091
6947173					

BACKGROUND

[0005] In embedded applications, it is often necessary to store a secret key in non-volatile memory (such as flash memory on an integrated circuit) in products that are widely

distributed. In certain applications, the same key needs to be stored in multiple integrated circuits, many of which are available to a potential attacker. For example, the integrated circuit can form part of a consumable such as an ink cartridge, which are widely distributed as replacements for empty cartridges.

[0006] One way in which an attacker can probe an integrated circuit (or chip) for a key or other secret information is to use a focussed ion beam FIB write attack. In this attack, encapsulant is carefully removed from the circuitry and a FIB used to change one or more bits in flash memory from an unknown state into a known state. Based on the effect the change has on the behaviour of the chip, an attacker may be able to deduce certain information about the state of the attacked bit or bits. For example, if the chip no longer works, it may be determined that the state of the bit or bits was changed by the FIB.

[0007] If the chip is disabled by the attack, the attacker merely obtains another chip that has an identical secret key, and attempts a similar attack on a different bit or combination of bits. By repeating the attack on different bits over a number of the chips, the attacker can either directly determine the key, or can build up a statistical model that vastly reduces the number of attempts needed to crack the security offered by the key on the chip. Of course, once the key is compromised in this way, it is compromised for all other chips having this key.

SUMMARY OF THE INVENTION

[0008] In a first aspect the present invention provides a method of storing at least one functionally identical code segment in each of a plurality of printer cartridges of a printing system, each of the printer cartridges having a memory, the method comprising, for each printer cartridge:

(a) determining a first memory location which is different than a first memory location in at least a plurality of the other printer cartridges; and

(b) storing a first of the at least one code segments in the memory at the determined first memory location;

[0009] wherein the different first memory locations are determined such that the code segment storage location of all the respective devices cannot be ascertained from the code segment storage location of any one of the printer cartridges.

[0010] Optionally, at least one of the code segments in each of the printer cartridges includes an initial instruction, the initial instruction being located at an initial instruction location, the initial instruction location being the same in all the printer cartridges. Optionally, the initial instruction in each printer cartridge is indicative of the first memory location of that printer cartridge, the initial instruction is indicative of the first memory location by including an explicit reference to the memory location, and the initial instruction is indicative of the first memory location by including an implicit reference to the memory location. Optionally, the implicit reference is a pointer to a location at which the address of the first memory location is stored, and the implicit reference is a pointer to a register that holds the address of the first memory location.

[0011] Optionally, step (a) includes randomly selecting the first memory location, includes selecting the first memory

location based on a stochastic process or mechanism, and selecting the first memory location from an existing list or sequence of memory locations.

[0012] Optionally, each memory is incorporated in a quality assurance chip of each printer cartridge, and is non-volatile memory.

[0013] Optionally, a first bit-pattern is stored in non-volatile memory of a printer cartridge, and the method comprising:

(a) applying a one way function to a second bit-pattern associated with the printer cartridge, thereby to generate a first result;

(b) applying a second function to the first result and the first bit-pattern, thereby to generate a second result; and

(c) storing the second result in the memory, thereby indirectly storing the first bit-pattern.

[0014] Optionally, a bit-pattern is stored in each of a plurality of printer cartridges, each of the printer cartridges having a memory, the method comprising, for each printer cartridge:

(a) determining a first memory location; and

(b) storing the bit-pattern at the first memory location;

[0015] wherein the first memory locations are different in at least a plurality of the respective printer cartridges.

[0016] Optionally, multiple first bit-patterns are stored in non-volatile memory of a printer cartridge, the method comprising, for each of the first bit-patterns to be stored:

(a) applying a one way function to a third bit-pattern based on a second bit-pattern associated with the printer cartridge, thereby to generate a first result;

(b) applying a second function to the first result and the first bit-pattern, thereby to generate a second result; and

(c) storing the second result in the memory, thereby indirectly storing the first bit-pattern;

[0017] wherein the third bit-patterns used for the respective first bit-patterns are relatively unique compared to each other.

[0018] In further aspects the present invention provides one or more printer cartridges having a bit-pattern and/or result stored therein in accordance with the method.

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] FIG. 1. Single SoPEC A4 Simplex system

[0020] FIG. 2. Dual SoPEC A4 Simplex system

[0021] FIG. 3. Dual SoPEC A4 Duplex system

[0022] FIG. 4. Dual SoPEC A3 simplex system

[0023] FIG. 5. Quad SoPEC A3 duplex system

[0024] FIG. 6. SoPEC A4 Simplex system with extra SoPEC used as DRAM storage

[0025] FIG. 7. SoPEC A4 Simplex system with network connection to Host PC

[0026] FIG. 8. Document data flow

[0027] FIG. 9. Pages containing different numbers of bands

[0028] FIG. 10. Contents of a page band

[0029] FIG. 11. Page data path from host to SoPEC

[0030] FIG. 12. Page structure

[0031] FIG. 13. SoPEC System Top Level partition

[0032] FIG. 14. Proposed SoPEC CPU memory map (not to scale)

[0033] FIG. 15. Possible USB Topologies for Multi-SoPEC systems

[0034] FIG. 16. CPU block diagram

[0035] FIG. 17. CPU bus transactions

[0036] FIG. 18. State machine for a CPU subsystem slave

[0037] FIG. 19. Proposed SoPEC CPU memory map (not to scale)

[0038] FIG. 20. MMU Sub-block partition, external signal view

[0039] FIG. 21. MMU Sub-block partition, internal signal view

[0040] FIG. 22. DRAM Write buffer

[0041] FIG. 23. DIU waveforms for multiple transactions

[0042] FIG. 24. SoPEC LEON CPU core

[0043] FIG. 25. Linking Printhead Concept

[0044] FIG. 26. Equivalent signature generation

[0045] FIG. 27. An allocation of words in memory vectors

[0046] FIG. 28. Transfer and rollback process

DETAILED DESCRIPTION OF PREFERRED EMBODIMENT

[0047] Also throughout this description, "printhead module" and "printhead" are used somewhat interchangeably. Technically, a "printhead" comprises one or more "printhead modules", but occasionally the former is used to refer to the latter. It should be clear from the context which meaning should be allocated to any use of the word "printhead".

[0048] A SoPEC ASIC (Small office home office Print Engine Controller) suitable for use in price sensitive SoHo printer products is described. The SoPEC ASIC is intended to be a relatively low cost solution for linking printhead control, replacing the multichip solutions in larger more professional systems with a single chip. The increased cost competitiveness is achieved by integrating several systems such as a modified PEC1 printing pipeline, CPU control system, peripherals and memory sub-system onto one SoC ASIC, reducing component count and simplifying board design. SoPEC contains features making it suitable for multifunction or "all-in-one" devices as well as dedicated printing systems.

[0049] Basic features of the preferred embodiment of SoPEC include:

[0050] Continuous 30 ppm operation for 1600 dpi output at A4/Letter.

- [0051] Linearly scalable (multiple SoPECs) for increased print speed and/or page width.
- [0052] 192 MHz internal system clock derived from low-speed crystal input
- [0053] PEP processing pipeline, supports up to 6 color channels at 1 dot per channel per clock cycle
- [0054] Hardware color plane decompression, tag rendering, halftoning and compositing
- [0055] Data formatting for Linking Printhead
- [0056] Flexible compensation for dead nozzles, print-head misalignment etc.
- [0057] Integrated 20 Mbit (2.5 MByte) DRAM for print data and CPU program store
- [0058] LEON SPARC v8 32-bit RISC CPU
- [0059] Supervisor and user modes to support multi-threaded software and security
- [0060] 1 kB each of I-cache and D-cache, both direct mapped, with optimized 256-bit fast cache update.
- [0061] 1×USB2.0 device port and 3×USB2.0 host ports (including integrated PHYs)
- [0062] Support high speed (480 Mbit/sec) and full speed (12 Mbit/sec) modes of USB2.0
- [0063] Provide interface to host PC, other SoPECs, and external devices e.g. digital camera
- [0064] Enable alternative host PC interfaces e.g. via external USB/ethernet bridge
- [0065] Glueless high-speed serial LVDS interface to multiple Linking Printhead chips
- [0066] 64 remappable GPIOs, selectable between combinations of integrated system control components:
- [0067] 2×LSS interfaces for QA chip or serial EEPROM
- [0068] LED drivers, sensor inputs, switch control outputs
- [0069] Motor controllers for stepper and brushless DC motors
- [0070] Microprogrammed multi-protocol media interface for scanner, external RAM/Flash, etc.
- [0071] 112-bit unique ID plus 112-bit random number on each device, combined for security protocol support
- [0072] IBM Cu-11 0.13 micron CMOS process, 1.5V core supply, 3.3V IO.
- [0073] 208 pin Plastic Quad Flat Pack
- [0074] The SoPEC device can be used in several printer configurations and architectures.
- [0075] In the general sense, every preferred embodiment SoPEC-based printer architecture will contain:
 - [0076] One or more SoPEC devices.
 - [0077] One or more linking printheads.
 - [0078] Two or more LSS busses.
 - [0079] Two or more QA chips.
 - [0080] Connection to host, directly via USB2.0 or indirectly.
 - [0081] Connections between SoPECs (when multiple SoPECs are used).
 - [0082] The Host PC rasterizes and compresses the incoming document on a page by page basis. The page is restructured into bands with one or more bands used to construct a page. The compressed data is then transferred to the SoPEC device directly via a USB link, or via an external bridge e.g. from ethernet to USB. A complete band is stored in SoPEC embedded memory. Once the band transfer is complete the SoPEC device reads the compressed data, expands the band, normalizes contone, bi-level and tag data to 1600 dpi and transfers the resultant calculated dots to the linking print-head.
 - [0083] The SoPEC device can print a full resolution page with 6 color planes. Each of the color planes can be generated from compressed data through any channel (either JPEG compressed, bi-level SMG4 fax compressed, tag data generated, or fixative channel created) with a maximum number of 6 data channels from page RIP to linking print-head color planes.
 - [0084] The mapping of data channels to color planes is programmable. This allows for multiple color planes in the printhead to map to the same data channel to provide for redundancy in the printhead to assist dead nozzle compensation.
 - [0085] Also a data channel could be used to gate data from another data channel. For example in stencil mode, data from the bilevel data channel at 1600 dpi can be used to filter the contone data channel at 320 dpi, giving the effect of 1600 dpi edged contone images, such as 1600 dpi color text.
 - [0086] The SoPEC device typically stores a complete page of document data on chip. The amount of storage available for compressed pages is limited to 2 Mbytes, imposing a fixed maximum on compressed page size. SoPEC would not be capable of printing worst case pages unless they are split into bands and printing commences before all the bands for the page have been downloaded. The page sizes in the table are shown for comparison purposes and would be considered reasonable for a professional level printing system. The SoPEC device is aimed at the consumer level and would not be required to print pages of that complexity. If a document with more complex pages is required, the page RIP software in the host PC can determine that there is insufficient memory storage in the SoPEC for that document. In such cases the RIP software can take two courses of action:
 - [0087] It can increase the compression ratio until the compressed page size will fit in the SoPEC device, at the expense of print quality, or
 - [0088] It can divide the page into bands and allow SoPEC to begin printing a page band before all bands for that page are downloaded.
 - [0089] Once SoPEC starts printing a page it cannot stop; if SoPEC consumes compressed data faster than the bands can be downloaded a buffer underrun error could occur causing the print to fail. A buffer underrun occurs if a line

synchronisation pulse is received before a line of data has been transferred to the printhead.

[0090] Other options which can be considered if the page does not fit completely into the compressed page store are to slow the printing or to use multiple SoPECs to print parts of the page. Alternatively, a number of methods are available to provide additional local page data storage with guaranteed bandwidth to SoPEC, for example a Storage SoPEC.

[0091] The SoPEC is a page rendering engine ASIC that takes compressed page images as input, and produces decompressed page images at up to 6 channels of bi-level dot data as output. The bi-level dot data is generated for the Memjet linking printhead. The dot generation process takes account of printhead construction, dead nozzles, and allows for fixative generation.

[0092] A single SoPEC can control up to 12 linking printheads and up to 6 color channels at >10,000 lines/sec, equating to 30 pages per minute. A single SoPEC can perform full-bleed printing of A4 and Letter pages. The 6 channels of colored ink are the expected maximum in a consumer SOHO, or office Memjet printing environment:

[0093] CMY, for regular color printing.

[0094] K, for black text, line graphics and gray-scale printing.

[0095] IR (infrared), for Netpage-enabled applications.

[0096] F (fixative), to enable printing at high speed. Because the Memjet printer is capable of printing so fast, a fixative may be required on specific media types (such as calendared paper) to enable the ink to dry before the page touches a previously printed page. Otherwise the pages may bleed on each other. In low speed printing environments, and for plain and photo paper, the fixative is not be required.

[0097] SoPEC is color space agnostic. Although it can accept contone data as CMYX or RGBX, where X is an optional 4th channel (such as black), it also can accept contone data in any print color space. Additionally, SoPEC provides a mechanism for arbitrary mapping of input channels to output channels, including combining dots for ink optimization, generation of channels based on any number of other channels etc. However, inputs are typically CMYK for contone input, K for the bi-level input, and the optional Netpage tag dots are typically rendered to an infra-red layer. A fixative channel is typically only generated for fast printing applications.

[0098] SoPEC is resolution agnostic. It merely provides a mapping between input resolutions and output resolutions by means of scale factors. The expected output resolution is 1600 dpi, but SoPEC actually has no knowledge of the physical resolution of the linking printhead.

[0099] SoPEC is page-length agnostic. Successive pages are typically split into bands and downloaded into the page store as each band of information is consumed and becomes free.

[0100] SoPEC provides mechanisms for synchronization with other SoPECs. This allows simple multi-SoPEC solutions for simultaneous A3/A4/Letter duplex printing. However, SoPEC is also capable of printing only a portion of a page image. Combining synchronization functionality with

partial page rendering allows multiple SoPECs to be readily combined for alternative printing requirements including simultaneous duplex printing and wide format printing.

[0101] The required printing rate for a single SoPEC is 30 sheets per minute with an inter-sheet spacing of 4 cm. To achieve a 30 sheets per minute print rate, this requires:

[0102] $300 \text{ mm} \times 63 \text{ (dot/mm)} / 2 \text{ sec} = 105.8 \text{ } \square \text{seconds per line, with no inter-sheet gap.}$

[0103] $340 \text{ mm} \times 63 \text{ (dot/mm)} / 2 \text{ sec} = 93.3 \text{ } \square \text{seconds per line, with a 4 cm inter-sheet gap.}$

[0104] A printline for an A4 page consists of 13824 nozzles across the page. At a system clock rate of 192 MHz, 13824 dots of data can be generated in 69.2 \square seconds. Therefore data can be generated fast enough to meet the printing speed requirement.

[0105] Once generated, the data must be transferred to the printhead. Data is transferred to the printhead ICs using a 288 MHz clock (3/2 times the system clock rate). SoPEC has 6 printhead interface ports running at this clock rate. Data is 8b/10b encoded, so the throughput per port is $0.8 \times 288 = 230.4 \text{ Mb/sec}$. For 6 color planes, the total number of dots per printhead IC is $1280 \times 6 = 7680$, which takes 33.3 \square seconds to transfer. With 6 ports and 11 printhead ICs, 5 of the ports address 2 ICs sequentially, while one port addresses one IC and is idle otherwise. This means all data is transferred on 66.7 \square seconds (plus a slight overhead). Therefore one SoPEC can transfer data to the printhead fast enough for 30 ppm printing.

[0106] From the highest point of view the SoPEC device consists of 3 distinct subsystems

[0107] CPU Subsystem

[0108] DRAM Subsystem

[0109] Print Engine Pipeline (PEP) Subsystem

[0110] See FIG. 13 for a block level diagram of SoPEC.

[0111] The CPU subsystem controls and configures all aspects of the other subsystems. It provides general support for interfacing and synchronising the external printer with the internal print engine. It also controls the low speed communication to the QA chips. The CPU subsystem contains various peripherals to aid the CPU, such as GPIO (includes motor control), interrupt controller, LSS Master, MMI and general timers. The CPR block provides a mechanism for the CPU to powerdown and reset individual sections of SoPEC. The UDU and UHU provide high-speed USB2.0 interfaces to the host, other SoPEC devices, and other external devices. For security, the CPU supports user and supervisor mode operation, while the CPU subsystem contains some dedicated security components.

[0112] The DRAM subsystem accepts requests from the CPU, UHU, UDU, MMI and blocks within the PEP subsystem. The DRAM subsystem (in particular the DIU) arbitrates the various requests and determines which request should win access to the DRAM. The DIU arbitrates based on configured parameters, to allow sufficient access to DRAM for all requesters. The DIU also hides the implementation specifics of the DRAM such as page size, number of banks, refresh rates etc.

[0113] The Print Engine Pipeline (PEP) subsystem accepts compressed pages from DRAM and renders them to bi-level dots for a given print line destined for a printhead interface that communicates directly with up to 12 linking printhead ICs.

[0114] The first stage of the page expansion pipeline is the CDU, LBD and TE. The CDU expands the JPEG-compressed contone (typically CMYK) layer, the LBD expands the compressed bi-level layer (typically K), and the TE encodes Netpage tags for later rendering (typically in IR, Y or K ink). The output from the first stage is a set of buffers: the CFU, SFU, and TFU. The CFU and SFU buffers are implemented in DRAM.

[0115] The second stage is the HCU, which dithers the contone layer, and composites position tags and the bi-level spot0 layer over the resulting bi-level dithered layer. A number of options exist for the way in which compositing occurs. Up to 6 channels of bi-level data are produced from this stage. Note that not all 6 channels may be present on the printhead. For example, the printhead may be CMY only, with K pushed into the CMY channels and IR ignored. Alternatively, the position tags may be printed in K or Y if IR ink is not available (or for testing purposes).

[0116] The third stage (DNC) compensates for dead nozzles in the printhead by color redundancy and error diffusing dead nozzle data into surrounding dots.

[0117] The resultant bi-level 6 channel dot-data (typically CMYK-IRF) is buffered and written out to a set of line buffers stored in DRAM via the DWU.

[0118] Finally, the dot-data is loaded back from DRAM, and passed to the printhead interface via a dot FIFO. The dot FIFO accepts data from the LLU up to 2 dots per system clock cycle, while the PHI removes data from the FIFO and sends it to the printhead at a maximum rate of 1.5 dots per system clock cycle.

[0119] Manufacturers of systems that require consumables (such as laser printers that require toner cartridges) have addressed the problem of authenticating consumables with varying levels of success. Most have resorted to specialized packaging that involves a patent. However this does not stop home refill operations or clone manufacture in countries with weak industrial property protection. The prevention of copying is important to prevent poorly manufactured substitute consumables from damaging the base system. For example, poorly filtered ink may clog print nozzles in an ink jet printer, causing the consumer to blame the system manufacturer and not admit the use of non-authorized consumables.

[0120] In addition, some systems have operating parameters that may be governed by a license. For example, while a specific printer hardware setup might be capable of printing continuously, the license for use may only authorise a particular print rate. The printing system would ideally be able to access and update the operating parameters in a secure, authenticated way, knowing that the user could not subvert the license agreement.

[0121] Furthermore, legislation in certain countries requires consumables to be reusable. This slightly complicates matters in that refilling must be possible, but not via unauthorized home refill or clone refill means.

[0122] To address these authentication problems, this document defines the QA Chip Logical Interface, which provides authenticated manipulation of a system's operating and consumable parameters. The interface is described in terms of data structures and the functions that manipulate them, together with examples of use. While the descriptions and examples are targeted towards the printer application, they are equally applicable in other domains. The QA Chip Logical Interface is now described.

[0123] The QA Chip Logical Interface is a logical interface, and is therefore implementation independent. Although this document does not cover implementation details on particular platforms, expected implementations include:

[0124] Software only

[0125] Off-the-shelf cryptographic hardware

[0126] ASICs, such as SBR4320 [2] and SOPEC [5] for physical insertion into printers and ink cartridges

[0127] Smart cards

[0128] An instance of a QA Chip Logical Interface (on any platform) is a QA Device.

[0129] QA Devices cannot talk directly to each other. A System is a logical entity which has one or more QA Devices connected logically (or physically) to it, and calls the functions on those QA Devices.

[0130] From the point of view of a QA Device receiving commands, System cannot inherently be trusted i.e. a given QA Device cannot tell if the System is trustworthy or not. System can, however, be constructed within a trustworthy environment (such as a SoPEC or within another physically secure computer system), and in these cases System can trust itself.

[0131] Digital signatures are used throughout the authentication protocols of the QA Chip Logical Interface. A signature is produced by passing data plus a secret key through a keyed hash function. The signature proves that the data was signed by someone who knew the secret key.

[0132] The signature function used throughout the QA Chip Logical Interface is HMAC-SHA1.

[0133] When a System is constructed within a physically/logically secure environment, then System itself is trusted, and any software/hardware running within that secure environment is trusted. A Trusted QA Device is simply a QA Device that resides within the same secure environment that System also resides in, and can therefore be trusted by System. This means that it is not possible for an attacker to subvert the communication between the System and the Trusted QA Device, or to replace the functionality of a QA Device by some other functionality. A Trusted QA Device enables a System to extend trust to external QA Devices. An example of a Trusted QA Device is a body of software inside a digitally signed program.

[0134] An External untrusted QA Device is a QA Device that resides external to the trusted environment of the system and is therefore untrusted. The purpose of the QA Chip Logical Interface is to allow the external untrusted QA Devices to become effectively trusted. This is accomplished when a Trusted QA Device shares a secret key with the external untrusted QA Device, or with a Translation QA Device (see below).

[0135] In a printing application, external untrusted QA Devices would typically be instances of SBR4320 implementations located in a consumable or the printer.

[0136] A Translation QA Device is used to translate signatures between QA Devices and extend effective trust when secret keys are not directly shared between QA Devices.

[0137] As an example, if a message is sent from QA Device A to QA Device C, but A and C don't share a secret key, then under normal circumstances C cannot trust the message because a signature generated by A cannot be verified by C. However if A and B share secret 1, and B and C share secret 2, and B is allowed to translate signatures for certain messages sent between secret 1 and secret 2, then B can be used as a Translation QA Device to allow those messages to be sent between A and C.

[0138] A Consumable QA Device is an external untrusted QA Device located in a consumable. It typically contains details about the consumable, including how much of the consumable remains.

[0139] In a printing application the consumable QA Device is typically found in an ink cartridge and is referred to as an Ink QA Device, or simply Ink QA since ink is the most common consumable for printing applications. However, other consumables in printing applications include media and impression counts, so consumable QA Device is more generic.

[0140] An Operating Parameter QA Device is an external untrusted device located within the infrastructure of a product, and contains at least some of the operating parameters of the application. Unlike the Trusted QA Device, an Operating Parameter QA Device is in a physically/logically untrusted section of the overall hardware/software.

[0141] An example of an Operating Parameter QA Device in a SoPEC-based printer system is the PrinterQA Device (or simply PrinterQA), that contains the operating parameters of the printer. The PrinterQA contains OEM and printer model information that indirectly specifies the non-upgradeable operating parameters of the printer, and also contains the upgradeable operating parameters themselves.

[0142] A Value Upgrader QA Device contains the necessary functions to allow a system to write an initial value (e.g. an ink amount) into another QA Device, typically a consumable QA Device. It also allows a system to refill/replenish a value in a consumable QA Device after use.

[0143] Whenever a value upgrader QA Device increases the amount of value in another QA Device, the value in the value upgrader QA Device is correspondingly decreased. This means the value upgrader QA Device cannot create value—it can only pass on whatever value it itself has been issued with. Thus a value upgrader QA Device can itself be replenished or topped up by another value upgrader QA Device.

[0144] An example of a value upgrader is an Ink Refill QA Device, which is used to fill/refill ink amount in an Ink QA Device.

[0145] A Parameter Upgrader QA Device contains the necessary functions to allow a system to write an initial parameter value (e.g. a print speed) into another QA Device,

e.g. an Operating Parameter QA Device. It also allows a system to change that parameter value at some later date.

[0146] A parameter upgrader QA Device is able to perform a fixed number of upgrades, and this number is effectively a consumable value. Thus the number of available upgrades decreases by 1 with each upgrade, and can be replenished by a value upgrader QA Device.

[0147] Secret transport keys are inserted into QA Devices during instantiation (e.g. manufacture). These keys must be replaced by the final secret keys when the purpose of the QA Device is known. The Key Replacement QA Device implements all necessary functions for replacing keys in other QA Devices.

[0148] An Authenticated Read is a read of data from a non-trusted QA Device that also includes a check of the signature. When the System determines that the signature is correct for the returned data (e.g. by asking a Trusted QA Device to test the signature) then the System is able to determine that the data has not been tampered en route from the read, and was actually stored on the non-trusted QA Device.

[0149] An authenticated write is a write to the data storage area in a QA Device where the write request includes both the new data and a signature. The signature is based on a key that has write access permission to the region of data in the QA Device, and proves to the receiving QA Device that the writer has the authority to perform the write. For example, a Value Upgrader Refilling Device is able to authorize a system to perform an authenticated write to upgrade a Consumable QA Device (e.g. to increase the amount of ink in an Ink QA Device).

[0150] The QA Device that receives the write request checks that the signature matches the data (so that it hasn't been tampered with en route) and also that the signature is based on the correct authorization key.

[0151] An authenticated write can be followed by an authenticated read to ensure (from the system's point of view) that the write was successful.

[0152] A non-authenticated write is a write to the data storage area in a QA Device where the write request includes only the new data (and no signature). This kind of write is used when the system wants to update areas of the QA Device that have no access-protection.

[0153] The QA Device verifies that the destination of the write request has access permissions that permit anyone to write to it. If access is permitted, the QA Device simply performs the write as requested.

[0154] A non-authenticated write can be followed by an authenticated read to ensure (from the system's point of view) that the write was successful.

[0155] Authorized modification of data refers to modification of data via authenticated writes.

[0156] The primary purpose of a QA Device is to securely hold application-specific data. For example if the QA Device is a Consumable QA Device for a printing application it may store ink characteristics and the amount of ink remaining.

[0157] For secure manipulation of data:

[0158] Data must be clearly identified (includes typing of data).

[0159] Data must have clearly defined access criteria and permissions.

[0160] Data must be able to be transferred securely from one QA Device to another, through a potentially insecure environment.

[0161] In addition, each QA Device must be capable of storing multiple data elements, where each data element is capable of being manipulated in a different way to represent the intended use of that data element. For convenience, a data element is referred to as a field.

[0162] Each QA Device requires an identifier that allows unique identification of that QA Device by external systems, ensures that messages are received by the correct QA Device, and ensures that the same device can be used across multiple transactions.

[0163] Strictly speaking, the identifier only needs to be unique within the context of a key, since QA Devices only accept messages that are appropriately signed. However it is more convenient to have the instance identifier completely unique, as is the case with this design.

[0164] In certain circumstances it is useful for a Trusted QA Device to assume the instance identifier of an external untrusted QA Device in order to build a local trusted form of the external QA Device. It is the responsibility of the System to ensure that the correct device is used for particular messages. As an example, a Trusted QA Device in a SoPEC-based printing system has the same instance identifier as the external (untrusted) Printer QA so that the System can access functionality in the Trusted QA instead of the external untrusted Printer QA. The identifier functionality is provided by ChipId.

[0165] ChipId is the unique 64-bit QA Device identifier. The ChipId is set when the QA Device is instantiated, and cannot be changed during the lifetime of the QA Device.

[0166] A 64-bit ChipId gives a maximum of 1844674 trillion unique QA Devices.

[0167] Each QA Device contains a number of secret keys that are used for signature generation and verification. These keys serve three basic functions:

[0168] For reading, where they are used to verify that the read data came from the particular QA Device and was not altered en route.

[0169] For writing, where they are used to authorise modification of data.

[0170] For transporting keys, where they are used in the process of encrypting and transporting new keys into the QA Device.

[0171] All of these functions are achieved by signature generation; a key is used to generate a signature for subsequent transmission from the device, and to generate a signature to compare against a received signature. The transportation function is additionally achieved by encryption.

[0172] The number of secret keys in a QA Device is given by NumKeys, and has a maximum value of 256, i.e. the number of keys for a particular implementation may be less than this. For convenience, we refer to a QA Device as having NumKeys keyslots, where each keyslot contains a single key. Thus the nth keyslot contains the nth key (where n has the range 0 to NumKeys-1). The keyslot concept is useful because a keyslot contains not only the bit-pattern of the secret key, but also additional information related to the secret key and its use within the QA Device. The term KeySlot[n].xxx is used to describe the element named xxx within KeySlot n.

[0173] Each key is referred to as K, and the subscripted form K_n refers to the key in the nth keyslot. Thus $K_n = \text{KeySlot}[n].K$.

[0174] The length of each key is 160 bits. 160 bits was chosen because the output signature length from the signature generation function (HMAC-SHA1) is 160 bits, and a key longer than 160-bits does not add to the security of the function.

[0175] The security of the digital signatures relies upon keys being kept secret. To safeguard the security of each key, keys should be generated in a way that is not deterministic. Ideally the bit pattern representing a particular key should be a physically generated random number, gathered from a physically random phenomenon. Each key is initially programmed during QA Device instantiation.

[0176] For the convenience of the System, each key has a corresponding 18-bit KeyId which can be read to determine the identity or label of the key without revealing the value of the key itself. Since the relationship between keys and KeyIds is 1:1 (they are both stored in the same keyslot), a system can read all the KeyIds from a QA Device and know what key is stored in each of the keyslots. A KeyId of INVALID_KEYID (=0) is the only predefined id, and indicates that the key is invalid and should not be used, although the QA Device itself will not specifically prevent its use. From a system perspective, the bit pattern of a key is undefined when KeyId=INVALID_KEYID, and so cannot be guaranteed to match another key whose KeyId is also INVALID_KEYID. The bit pattern for such a key should be set to a random bit pattern for the physical security of any other keys present in the QA Device.

[0177] To create a digital signature, the data to be signed (d) is passed together with a secret key (k) through a key dependent one-way hash function (SIG). i.e. signature=SIG_k(d). The key dependent one-way hash function used throughout the QA Chip Logical Interface is HMAC-SHA1, although from a theoretical sense any key dependent one-way hash function could be used.

[0178] Signatures are only of use if they can be validated i.e. QA Device A produces a signature for data and QA Device B can check if the signature is valid for that particular data. This implies that A and B must share some secret information so that they can generate equivalent signatures.

[0179] Common key signature generation is when QA Device A and QA Device B share the exact same key i.e. key $K_A = \text{key } K_B$. Thus the signature for a message produced by A using K_A can be equivalently produced by B using K_B . In other words $\text{SIG}_{K_A}(d) = \text{SIG}_{K_B}(d)$ because key $K_A = \text{key } K_B$.

[0180] Variant key signature generation is when QA Device B holds a base key, and QA Device A holds a variant of that key such that $K_A = \text{owf}(K_B, U_A)$ where owf is a one-way function based upon the base key (K_B) and a unique number in A (U_A). A one-way function is required to create K_A from K_B or it would be possible to derive K_B if K_A were exposed. Thus A can produce $\text{SIG}_{K_A}(\text{message})$, but for B to produce an equivalent signature B must produce K_A by being told U_A from A and using B's base key K_B . K_A is referred to as a variant key and K_B is referred to as the base key. Therefore, B can produce equivalent signatures from many QA Devices, each of which has its own unique variant of K_B . Since ChipId is unique to a given QA Device, we conveniently use that as U_A .

[0181] Common key signature generation is used when A and B are effectively equally available¹ to an attacker. Variant key signature generation is used when B is not readily available to an attacker, and A is readily available to an attacker. If an attacker is able to determine K_A , they do not know K_B for any other QA Device of class A, and they are not able to determine K_B .

¹The term "equally available" is relative. It typically means that the ease of availability of both are the effectively the same, regardless of price (e.g. both A and B are commercially available and effectively equally easy to come by).

[0182] When two or more devices share U_A (in our implementation, U_A is ChipId), then their variant keys can be effectively treated as common keys for signatures passed between them, but as variant keys when passed to other devices.

[0183] The QA Device producing or testing a signature needs to know if it must use the common or variant means of signature generation. Likewise, when a key is stored in a QA Device, the status of the key (whether it is a base or variant key) must be stored in the keyslot along with the key for future reference.

[0184] Therefore each keyslot contains a 1-bit Variant flag to hold the status of the key in that keyslot:

[0185] Variant=0 means the key in the keyslot is a base/common key

[0186] Variant=1 means the key in the keyslot is a variant key

[0187] The QA Device itself doesn't directly use the Variant setting. Instead, the System reads the value of variant from the desired keyslots in the two QA Devices (one QA Device will produce the signature, the other will check the signature) and informs the signature generation function and signature checking functions whether or not to use base or variant signature generation for a particular operation.

[0188] It is assumed in equivalent signature generation between 4 QA Devices A, B, C that each device has a single keyslot. KeySlot.KeyId of all four keys are the same i.e. $\text{KeySlot[A].KeyId} = \text{KeySlot[B].KeyId} = \text{KeySlot[C].KeyId} = \text{KeySlot[D].KeyId}$.

[0189] If $\text{KeySlot[A].Variant} = 0$ and $\text{KeySlot[B].Variant} = 0$, then a signature produced by A, can be equivalently produced by B because $K_A = K_B$.

[0190] If $\text{KeySlot[B].Variant} = 0$ and $\text{KeySlot[C].Variant} = 1$, then a signature produced by C, can be equivalently produced by B because $K_C = f(K_B, \text{ChipId}_C)$. Note that B must be told ChipId_C for this to be possible.

[0191] If $\text{KeySlot[C].Variant} = 1$ and $\text{KeySlot[D].Variant} = 1$, then a signature produced by C, cannot be equivalently produced by D unless both QA Devices have the same U_A (i.e. they must share the same Chip Identifier) While C and D will typically not share a ChipId, in certain circumstances the System can read a QA Device's Chip Identifier and install it into another QA Device. Then, using key transport mechanisms, the two QA Devices can come to share a common variant key, and can thence generate and check signatures with each other.

[0192] If $\text{KeySlot[D].Variant} = 1$ and $\text{KeySlot[A].Variant} = 0$, then a signature produced by D, can be equivalently produced by A because $K_D = f(K_A, \text{ChipId}_D)$.

[0193] While it is theoretically possible that a system could permit each key to be used to perform all of these tasks, in most cases it is a security risk to allow this.

[0194] If any key can be used to transport any other key out of a QA Device, then a compromise of a single key means a compromise of all keys. The reason is that the compromised key can be used by an attacker to transport all other keys out of a QA Device. Some QA Devices (such as Key Replacement QA Devices) are specifically required to transport keys, while others (such as those devices used in consumables) should not ever transport their keys out.

[0195] During manufacture it is not always possible to know the final intended application for a given QA Device. For example, one may end up at OEM1 while another is destined for OEM2. To decouple manufacture from installation of QA Devices, it is useful to place temporary batch keys into the QA Devices. Each of these keys should be replaceable by a different batch key or a final application key, but during their temporary existence these keys must not be capable of authenticating signatures writes of data. Thus they act as a transport key.

[0196] Likewise, in the Key Replacement QA Device, there is a need to differentiate between final use for a key in a QA Device, and storage of a key in one QA Device for subsequent injection into another. For example, a key may be a transport key when stored in QA Device A, and although we want to store that same key in a Key Replacement QA Device B for future injection into A, we do not want that key to be used to transport keys from B. Thus, if a key is not in its final intended keyslot, then it should have no abilities in that QA Device other than being transported out, and the intended use of the key (for example whether or not it will be a transport key when installed in its final destination) needs to be associated with that key.

[0197] From a security point of view there should be a time when a key in a given keyslot can be guaranteed to be in its final intended form i.e. it cannot be replaced later. If a key could be replaced at any time, attackers could potentially launch a denial of service attack by replacing keys with garbage, or could replace a key with one of their own choice. As an example, suppose keys k1 and k2 are both used to read value from a QA Device, write value to the QA Device, and to transport new keys into the QA Device. If either k1 or k2 is compromised, then the compromised key could be used to transport keys of choice to replace both keys and create value in the QA Device.

[0198] Therefore each keyslot contains 3 □1-bit flags as follows:

[0199] KeyType: whether the key is a TransportKey (0) to be used for key transport and signing reads of key meta-information, or if it is a DataKey (1) to be used for signing data as well as key meta-information

[0200] TransportOut: whether or not the key can be transported out from this QA Device

[0201] UseLocally: whether or not the key is for use locally within this QA Device or not. For transport keys this means whether or not the transport key can be used to transport another key out from this QA Device.

[0202] The following examples assume 3 bits xyz are interpreted as:

[0203] x=KeyType

[0204] y=TransportOut

[0205] z=UseLocally

[0206] A freshly manufactured QA Device A will most likely have the 3 bits for each keyslot set to 000 so that all the keys are replaceable.

[0207] To replace one of A's keys (k1) by another batch key (k2), key replacement QA Device B is required where B typically contains k1 with 3 bits set to 001, and k2 with 3 bits set to 010. After k2 has been transferred into A, the 3 bits within A will be now set to 000. Thus k2 cannot be used or replaced within B, but can be replaced within A.

[0208] To replace one of A's keys (k1) by a final use data key (k2), key replacement QA Device B is required where B typically contains k1 with 3 bits set to 001, and k2 with 3 bits set to 110. After k2 has been transferred into A, the 3 bits within A will be now set to 101. Thus k2 can be used within A but not B, and cannot be transported out of A.

[0209] Although there are KeyNum keyslots in a QA Device, not all keyslots may be required for a given application. For example, a QA Device may supply 256 keyslots, but only 2 keys may be required for a particular application. The remaining keyslots need to be invalidated so they cannot be used as a reference for signature checking or signature generation.

[0210] When QA Device A has a keyslot with KeyType, TransportOut, and UseLocally set to 000, then the key in that keyslot can be replaced.

[0211] To invalidate the keyslot in A where k1 is currently residing so that no further keys can ever be stored in that keyslot, key replacement QA Device B is required where B typically contains:

[0212] k1 with 3 bits set to 001

[0213] a base key k2 with 3 bits set to 110 and a KeyId of 0

[0214] After k2 has been transferred into A as a variant key, the 3 bits within A will be now set to 100. Thus k2 cannot be used within A, cannot be transported out of A, and cannot be replaced. Moreover, being a variant key in A, k2 will be different for each instance of A and will therefore be contribute to the entropy of A. Any system reading the KeyIds that are present in A will see that the keyslot contains

a key whose keyId is 0 (and is therefore invalid) and whose 2-bits specify that the key cannot be used.

[0215] Over the lifetime of a product, it may be desirable to retire a given key from use, either because of compromise or simply because it has been used for a specific length of time (and therefore to reduce the risk of compromise). Therefore the key in a keyslot needs to be invalidated by some means so that it cannot be used any more as a reference for signature checking or signature generation. From an audit-trail point of view, although a key has been retired from use, it is convenient to retain the key meta-information so that a System can know which keys have been retired.

[0216] In theory, a special command could be available in each QA Device to allow the caller to transform the Key-Type, TransportOut, and UseLocally settings for a keyslot from some value to 100. The key in that slot would then be non-transportable non-usable, and therefore invalid. However it would not be possible to know the previous setting for the 3 bits once the key had become invalid.

[0217] It is therefore desirable to have a boolean in each keyslot that can be set to make a particular key invalid. If a key has been marked as invalid, then TransportOut and UseLocally are ignored and treated as 0, and the key cannot be replaced.

[0218] However, a single bit representation of this boolean over-complicates 4320-based implementations of QA Devices in that it is not possible to set a single bit in shadowed mode on a 4320 device (to change a key from valid to invalid). Instead, the page containing the key would need to be erased and the key reconstructed, tasks which need to take place during initial key replacement during manufacture, but which should not need to take place after the keys are all finalised.

[0219] Therefore each keyslot contains a 4-bit boolean (which should be nybble-aligned within the keyslot data structure) referred to as Invalid, where 0000 represents a valid key in the keyslot, and non-zero represents an invalid key. A specific command (Invalidate Key) exists in the QA Logical Interface to allow a caller to invalidate a previously valid key.

[0220] If Invalid is set to a non-zero value, then the key is not used regardless of the settings for KeyType, TransportOut, and UseLocally.

[0221] In general each QA Device contains a number of data elements (each element referred to as a field), each of which can be operated upon by one or more keys. In the general case of an arbitrary device containing keys and fields, it is useful to have a set of permissions for each key on each field. For example, key 1 may have read-only permissions on field 1, but read/write permissions on field 2 and read/decrement-only permissions on field 3.

[0222] Although it can cater for all possibilities, a general scheme has size and complexity difficulties when implemented on a device with low storage capacity. In addition, the complexity of such a scheme is increased, if the device has to operate correctly with power-failures e.g. an operation must not create a logical inconsistency if power is removed partway through the operation.

[0223] Since the actual number of keys that can be stored in a low storage capacity QA Device depends on the

complexity of the program code and the size of the data structures, it is useful to minimise the functional complexity and minimise the size of the structures while not knowing the final number of keys.

[0224] In particular, the scheme must cope with multiple keys having the same permissions for a field to support the following situations:

each of the various users of the QA Device has access to a different key, such that different users can be individually included or excluded from access

only a subset of keys are in use at any one time

[0225] The concept that supports this requirement is the keygroup. A keygroup contains a number of keys, and each field has a set of permissions with respect to the keygroups. Thus keygroup 1 (containing some number of keys) may have read-only permissions on field 1, but read/write permissions on field 2 and read/decrement-only permissions on field 3.

[0226] In the limit case of 1 key per keygroup, with an arbitrary number of keygroups, the storage requirements for the permissions on each field would be the same as the general case without keygroups, but by limiting the number of keygroups, the storage requirements for the permissions on each field can be pre-known, constant, and is decoupled from the actual number of keys in the device.

[0227] The number of keygroups in a QA Device is 4. This allows for 2 different keygroups that can transfer value into the QA Device, and for 2 different keygroups that can transfer value out of a QA Device, where each of the 4 keygroups is independent of the others. Note that transport keys do not need to be allocated a keygroup since they cannot be used to authorise reads or writes of data.

[0228] Thus each keyslot contains a 2-bit KeyGroup identifier. The value of KeyGroup is relevant only when the KeyType=DataKey.

[0229] For security concerns it is important that a field not be created until all the keys for a keygroup have been created. Otherwise an attacker may be able to add a known new key to an existing keygroup and thereby subvert the value associated with the field.

[0230] However it is not possible to simply not allow the creation of fields until all of the keys have been created. It may be that two distinct phases of programming occur, with creation of keys and data based on each phase. For example a stamp franking system may contain value in the form of ink plus a dollar amount. The keys and fields relating to ink may be injected at one physical location, while the keys and fields relating to dollars may be injected at a separate location some time later.

[0231] It is therefore desirable to have a boolean indicator that indicates whether a particular keygroup is locked. Once a keygroup is locked, then no more keys can be added to that keygroup. The boolean indicator is accessible per keyslot rather than as a single indicator for each keygroup in order that someone reading the keyslot information can know:

whether they can add any more keys to a keygroup

whether they can create fields with write-permissions for the keygroup

[0232] When a key is replaced, the keygroup for that key can be locked at the same time. This will cause the QA Device to change the status of all keys with the same KeyGroup value from keygroup-unlocked to keygroup-locked, thereby preventing the addition of any more keys in the keygroup.

[0233] However, a single bit representation of this boolean over-complicates 4320-based implementations of QA Devices in that it is not possible to set a single bit in shadowed mode on a 4320 device (to change a locked status from unlocked to locked). Instead, the page containing the key would need to be erased and the key reconstructed, and this would need to take place per key (where the KeyGroup matched).

[0234] Therefore each keyslot contains a 4-bit boolean (which should be nybble-aligned within the keyslot data structure) referred to as KeyGroupLocked, where 0000 represents that the keygroup to which the key in the keyslot belongs is unlocked (i.e. more keys can be added to the keygroup), and non-zero represents that the keygroup to which the key in the keyslot belongs is locked (i.e. more keys cannot be added to the keygroup).

[0235] It is finally worth noting that a Key Replacement QA Device does not need to check whether or not there are fields on the target device with write permissions related to a particular keygroup. The reason is that the target QA Device only allows field creation related to a keygroup if the keygroup is locked. Therefore if there was such a field in the target device one of the following is true:

[0236] the target QA Device is a fake one created by an attacker. If so, and if the attacker does not know the original key, then the replaced key will be of no value. If the attacker does know the original key, then they can determine the replacement key (since the replacement key is encrypted using the original key for transport) without creating a fake QA, and can therefore generate fake value as desired.

[0237] the target QA Device has come under physical attack (it's a real QA Device). If an attacker can do this, it's easier to allow the key replacement first, and then create a fake field. This situation cannot ever be detected by the Key Replacement QA Device.

[0238] In an ideal world (for the owner of a secret key at least), a given secret key will remain secret forever. However it is prudent to minimise the loss that could occur should a key be compromised.

[0239] This is further complicated in a system where all of the components of a system are stored at the user site, potentially without direct connection to a central server that could appropriately update all components after a particular time period or if a compromise is known to have occurred.

[0240] To create rolling keys, two QA Devices A and B are required such that A and B are intended to work together via a conceptual key k. While a single key could be used for k, it is more secure to limit the lifetime of any particular key, and to have a plan in place to remove a key from use should it be compromised.

[0241] Rolling keys are where multiple keys are stored in at least one of A and B such that different keys can be used at different times during the life of A and B, different instances of A and B at differing manufacture times can be

programmed with different keys yet still work together, and keys can be retired from use in A and/or B.

[0242] In the simplest example of the problem, suppose A is embedded in a printer system that works with ink cartridges containing B. If A contains a single key k for working with B, then k is required for all Bs as long as A is deployed. A compromise of k lasts for the lifetime of A.

[0243] A rolling key example system for this example is where A contains multiple keys k_1, k_2, \dots, k_n , each with a different KeyId, where each of these keys has the same permissions on datafields within A (typically they will all belong to the same keygroup in A). At initial manufacture, B contains a single key k_1 (that is also present in A). For a given time period k_1 can be used between A and B. At some later time (or if k_1 is compromised), Bs are manufactured only containing k_2 , and new As are manufactured only containing $k_2, k_3, \dots, k_n, k_{n+1}$. At a later time, Bs are manufactured only containing k_3 and new As are manufactured only containing $k_3, k_4, \dots, k_n, k_{n+1}, k_{n+2}$ etc.

[0244] Note that if the keys shared by A and B are all common keys, then a compromise of keys from A will compromise all future value in Bs. However if A contains the variant key form and B contains a base form of each key, then compromise of keys in A does not permit an attacker to know future keys in B and the attacker can therefore not create clone Bs until a real B is released and the base key is obtained from B. This means that the more variant keys that can be injected into A the more changes in B can be coped with out any loss of security.

[0245] In the example above, note that if k_1 is compromised, an attacker can still manufacture clone Bs that will work on older As. It is therefore desirable to somehow invalidate k_1 on older As at some point to reduce the impact of clone Bs. However it is not usually the case that an immediate cut-off point can be introduced. For example, once Bs are being manufactured with k_2 , existing Bs containing k_1 may still be in use and are still valid. Just because k_2 is used with A doesn't mean that k_1 should be invalidated in A immediately. Otherwise a valid user could not then use an older valid B in A after using a newer B in A. Likewise, new As typically need to be able to work with valid old Bs. Our example assumes that newer As won't work with older Bs.

[0246] Therefore if overlapping timing is required, then several valid keys in use at a time instead of having only a single valid key in use at a time. Once valid Bs are known to be out of circulation (e.g. due to an expiry date associated with a B) then a key can be officially retired from being included in the manufacture of new As, and can be invalidated in old As. The more keys that can be used, the finer-grained the resolution of timing for invalidating a particular key, and hence the greater the reduction in exposure.

[0247] For example, B may be an ink cartridge that has a use-by date of 12 months while A is a printer that must last for 5 years:

[0248] If A contains 5 keys, B is issued with a new key each year, and a new A is released each year, then k_1 will be in B during year1, k_2 will be in B during year2 etc. As produced in year 2 will need to contain k_1 since old Bs from the previous year are still valid. Only in

year 3 can As be manufactured without k_1 , and old As can have their k_1 invalidated. Clone Bs can therefore be manufactured by an attacker causing loss during year 1 and 2. After year 2, those clone Bs won't work on new As, but will continue to work on old As until k_1 has been invalidated on the old As.

[0249] If A contains 10 keys, B is issued with a new key every 6 months, and a new A is released every 6 months, then k_1 will be in B during the first 6 months, k_2 will be in B during the second six months etc. As produced in the second and third 6-months will need to contain k_1 since old Bs from the previous year are still valid. Only in the fourth 6-month can As be manufactured without k_1 , and old As can have their k_1 invalidated. Clone Bs can therefore be manufactured by an attacker causing loss during year 1 and the first half of year 2. After this time, those clone Bs won't work on new As, but will continue to work on old As until k_1 has been invalidated on the old As. Thus the addition of keys in A and the changing of keys at a faster rate (every 6 months compared to every year) has reduced the exposure of a compromised key without increasing any risk due to exposure of keys in A.

[0250] Of course if A is used with B and a B-like entity called C, then A can have 1 set of rolling keys with B, and can have a different set of rolling keys with C. This requires 1 key in B, 1 key in C, and two sets of multiple keys in A.

[0251] The rolling key structure can be extended to work with value hierarchy. Suppose A uses value from B, and value in B is replenished by C, then A and B can have one set of rolling keys, and B and C can have a different set of rolling keys and each set of rolling keys can roll at different times and rates. In this example:

[0252] A contains multiple variants for use with B

[0253] B contains 1 base key for use with A, and multiple variants for use with C

[0254] C contains 1 base key for use with B

[0255] A compromise of key(s) in A does not allow an attacker to manufacture clone Bs

[0256] A compromise of key(s) in B does not allow an attacker to manufacture clone Cs

[0257] A compromise of the keys in A allows free B resources on that particular A only—no other As are affected

[0258] A compromise of the base key in B has a limited exposure of effect—free B resources are available to attackers for a limited time, and with each new release of A and C, the amount of exposure is reduced.

[0259] A compromise of the base key in C has a limited exposure of effect—free C resources are available to attackers for a limited time, and with each new release of B the amount of exposure is reduced.

[0260] In the general case, each of the keys in a set of rolling keys has exactly the same purpose as the others in the set, and is used in the same way in the same QA Devices, but at different times in a product's life span. Each of the keys has a different KeyId. Typically when a set of rolling keys is held in a QA Device, they all belong to the same keygroup.

[0261] When the variant/base form of rolling keys is used, at any given time, only one base key is injected during manufacture. This is the current manufactured instance of the rolling key. Several of the key instances can be used in manufacture, in their variant forms. One by one, the current manufactured instance of the rolling key is replaced by subsequent instances of the rolling key.

[0262] After a period, or after the discovery of a key compromise, a particular current manufactured instance of a key is replaced by the next instance in the rolling key set in all of the QA Devices where it is used.

[0263] A set of rolling keys has the following characteristics:

[0264] The number of instances in the set of rolling keys, N. The rolling key instances are from 0 to N-1.

[0265] The current manufactured instance of the rolling key. This is the rolling key instance which is currently being inserted into manufactured products, in base form. The current manufactured instance is rolled to the next instance when a suitable length of time has elapsed, or there is the discovery of a key compromise.

[0266] The first and last valid instances of the rolling key set. There is likely to be a number of valid key instances either side of the current manufactured instance at any given time.

[0267] Rolling key instances which are before the first valid instance are considered to be invalid, and they should be invalidated in any manufactured product in the field whenever they are found. The question is how to enforce the eradication process, especially if the QA Devices are not in direct contact with a central authority of some kind.

[0268] The QA Logical Interface allows a particular key in a keyslot to be invalidated. An external entity needs to know which keys are invalid (for example by knowing the invalid keys' KeyIds). Assuming that the entity can read the KeyIds present in a QA Device the entity can invalidate the appropriate keys in the QA Device. The entity could refuse to operate on a QA Device until the appropriate keys have been invalidated.

[0269] For example, suppose a printer system has an ink cartridge and a refill cartridge. The printer system uses rolling key set 1 to communicate with the ink cartridge, and the ink cartridge is refilled from the refill cartridge via rolling key set 2. Whenever a refill cartridge is attached to the system, the refill cartridge contains a specific field containing an invalid key list. The system software in the printer knows that this field contains an invalid key list, and refuses to transfer the ink value from the refill cartridge to the ink cartridge until it has invalidated the appropriate keys on the ink cartridge. Alternatively, every time the system software for the printer is delivered/updated to the printer (e.g. downloaded off the internet), it can contain a list of known invalid keys and can apply these to anything it is connected to, including ink cartridges and refill cartridges. Likewise, if value is injected into a QA Device over the internet, the value server can invalidate the appropriate keys on the QA Device before injection of value. Done correctly,

the invalid keys will be deleted from use in all valid systems, thereby reducing the effect of a clone product.

[0270] The methods just discussed do not apply if a user exclusively uses fake QA Devices, and never comes into contact with valid QA Devices that have lists of invalid keys. However it is possible that a system can invalidate a key by itself after a particular amount of time, but this requires the system to know the current time, and the time period between invalidating keys. While this provides the feature required, it should not be possible under normal circumstances for a user to lie about the time or to accidentally have the time set to an incorrect one. For example, suppose a user accidentally sets a clock on their computer to the wrong year in the future, the printer attached to the computer should not suddenly invalidate all of the keys for the next 12 months. Likewise, if the user changes the clock back to the previous year, previously invalid keys should not suddenly become valid. This implies the system needs to know a Most Recent Validated Date i.e. a date/time that is completely trustworthy.

[0271] If system is in a trusted environment and has an appropriate time keeping mechanism, then MostRecentValidatedDate can be obtained locally. Otherwise the MostRecentValidatedDate can only be obtained when the system comes into contact with another trusted component. The trusted component could be software that runs on system, with a particular build date (and this date is therefore trusted), or a date stored on a QA Device (providing the date is read from the QA Device via keys and can only be set by a trusted source).

[0272] It is therefore convenient that at least one of the QA Devices in systems that support rolling keys should define at least two fields for the purposes of key invalidation: a field that contains the invalid key list (a list of invalid keyIds), and a field that contains a date that can contribute to a MostRecentValidatedDate. The Logical QA Interface currently supports a field type specifically for the former (see Appendix B), while the latter depends on the specifics of a particular application.

[0273] When allocating KeyIds in a system, it may be convenient to be able to tell if two keys are in the same set of rolling keys simply from based on their KeyIds (therefore independent of instantiation in a keygroup). One way of doing this is to compose the KeyId as 2 parts:

[0274] the RollingKeySetId, which would be unique for a given purpose within a QA Device infrastructure

[0275] the RollingKeyInstance, which specifies the keys within the rolling key set

[0276] So, for example, if the 18-bit KeyId could be composed of a 10-bit RollingKeySetId, and an 8-bit RollingKeyInstance. Thus each set of rolling keys would have 256 unique key values to be used in the sequence.

[0277] Suppose we have a configuration that consists of a system A that communicates with a QA Device B. For example, a printer system that communicates with an Operating Parameter QA Device (e.g. containing the print speed). The system reads the print speed before printing a page.

[0278] The only way that A and B can securely communicate is if A and B share a key.

[0279] If B has physical security since it is a QA Device, and A does not have such high security, then it is desirable to store the variant form of the key in A and the base form of the key in B. If the key is extracted from A (having less security than B), then at least other systems cannot be subverted with clone Bs.

[0280] However there is the question of injecting the variant key into A. If A can be programmed with a variant key after B has been attached (e.g. A contains non-volatile memory), then this is desirable. If A cannot be programmed after B has been attached (such as is the case with the SoPEC ASIC) then A must be programmed with a random number and after attachment to A, the random number must be transported into B.

[0281] A can now create a Trusted QA Device and communicate with B using A's variant key.

[0282] However if A requires to communicate with additional components such as C and D which are not connected to A or B during initial manufacture, there is a requirement to allow the communication but additionally minimise loss due to key compromise, especially since A is known to be less secure than QA Devices B, C and D. Examples of C and D include a Consumable QA Device such as an ink cartridge, and a Parameter Upgrader QA Device such as a permanent speed-upgrade dongle.

[0283] If the base key that is used in B is also used in C and D, then A can communicate securely with C and D. The risk of loss from a key compromise is higher since C and D share the same key.

[0284] If A can hold many keys, i.e. can be programmed with many keys during manufacture, then A can be programmed with appropriate variant keys for C and D using the same scheme as described above for B.

[0285] However, if the cost of injecting multiple keys into A is high (for example SoPEC has very little non-volatile memory), then an alternative is required that only uses a single key stored in A. There are two approaches to secure communication in this case: communication via key transport, and communication via signature translation.

[0286] The protocol for communicating with a QA Device is now described. Although the implementation of a QA Device varies, with one implementation having different capabilities from another, the same interface applies to all.

[0287] QA Devices are passive: commands are issued to them by the System, which is an entity mediating the communications between the QA Devices.

[0288] There are up to three QA Devices that are relevant to each command:

[0289] The Commanded QA Device, i.e. the QA Device receiving the command. This QA Device checks any incoming signature (if present), performs the command, and generates the output parameters and any outgoing signature as required.

[0290] The Incoming Signature QA Device, that generated the incoming signature (if it is present). This is usually a QA Device that produces and signs the input for the command as its output, but it might be a Translation QA Device.

[0291] The Outgoing Signature QA Device, that checks the outgoing signature (if it is present).

[0292] This is usually a QA Device that accepts as input the output of the command, but it might be a Translation QA Device.

[0293] The QA Device Protocol lists a set of commands that can be sent to a QA Device, and for each command, there is a set of valid responses. The protocol defines the features that are common to the commands.

[0294] A command consists of a number of 32-bit words where the first byte of the first word contains a command byte, and subsequent words contain up to four of the following blocks of data:

[0295] An UnsignedInputParameterBlock. This is a set of input parameters with no accompanying signature.

[0296] An InputSignatureCheckingBlock. This is a block of data that tells the QA Device how to check if the SignedInputParameterBlock is correctly signed. It includes the signature, and information about how it was constructed.

[0297] A SignedInputParameterBlock. This is a set of input parameters. It is often a list of entities, or entity descriptors. The signature in the InputSignatureCheckingBlock is over this block and the generator's and checker's nonces. A SignedInputParameterBlock has a QA Device's ChipId as its first element. If the SignedInputParameterBlock is list of entities with the modify bit set, then the ChipId must be the identifier of the chip being addressed (this ensures that a signed block for one QA Device cannot be applied to another)

[0298] An OutputSignatureGenerationBlock. This is a block of data that tells the QA Device how to generate a signature on the outgoing data.

[0299] The response to a command consists of a number of 32-bit words, where the first byte of the first word contains a response byte, and subsequent words contain up to two of the following blocks of data:

[0300] An OutputParameterBlock. This is often a list of entities. It may or may not be signed.

[0301] If it is signed, it has a QA Device's ChipId as its first element. If the OutputParameterBlock is list of entities with the modify bit clear, then the ChipId must be the identifier of the chip responding to the command.

[0302] An OutputSignatureCheckingBlock. This is present if the OutputParameterBlock is signed. The signature is generated according to the OutputSignatureGenerationBlock.

[0303] The arrangement of data within each 32-bit word is arranged in big-endian format. The assumption is that the System and the QA Device are processing the commands and responses in big-endian format.

[0304] All of the blocks in both command and response are length-tagged: the first 32-bit word contains a two-byte length that indicates the block length in 32-bit words, followed by the block data itself. The length is inclusive. Thus the length for a parameter block with no data content is 1.

[0305] The QA Device identifier ChipId is present in all SignedInputParameterBlock and signed OutputParameterBlock entity lists. This ensures that a signature over the block of data uniquely identifies the QA Device that the list is for or came from. This prevents attacks where commands that are intended for one QA Device are redirected to another, or when responses from one QA Device are passed off as being from another.

[0306] If the list is an incoming modify-entity list or an outgoing read-entity list, then the list ChipId must be the ChipId of the Commanded QA Device. If it is not, then the command fails.

that QA Device's ChipId. Thus ensures that the data cannot be mistaken for data from another QA Device.

[0309] Similarly, a signed incoming list of entities being written to a QA Device has a signature over a block of data that includes that QA Device's ChipId. This ensures that the data cannot be wrongly applied to any other QA Device.

[0310] In the operation of some commands, a Commanded QA Device accepts a signed Entity List as input, where the Entity List was generated by another QA Device A, and produces a signed Entity List as output where the output is suitable to be subsequently applied to A as an incoming Entity List. These commands include: Get Key, Transfer Delta, Transfer Assign, and Start Rollback.

[0311] Commands in the QA Device command set are distinguished by CommandByte.

[0312] Table 1 describes the CommanByte values:

Values and Interpretation for CommandByte	
CommandByte	Value Description
GET INFO	1 Get summary of information from the QA Device
GET CHALLENGE	2 Get a nonce from the QA Device.
LOCK KEY GROUPS	3 Lock a specified set of keygroups. This prevents any keys in the keygroups from being subsequently replaced.
LOCK FIELD CREATION	4 Lock all field creation in the QA Device. Locking field creation prevents any fields from subsequently being created.
READ	5 Read a group of key descriptors, field descriptors and/or field values from a QA Device.
AUTHENTICATED READ	6 Read a group of key descriptors, field descriptors and/or field values from a QA Device. The results are accompanied by a signature to authenticate the results.
AUTHENTICATED READ WITH SIGNATURE ONLY	7 Specify a group of key descriptors, field descriptors and/or field values in a QA Device, and read the signature over that data.
WRITE	8 Write a group of field values to fields in the QA Device.
AUTHENTICATED WRITE	9 Write a group of field values to fields in the QA Device. The write command is authenticated by a signature over the list of field values.
CREATE FIELDS	10 Create a group of fields in a QA Device.
REPLACE KEY	11 Replace a key in a QA Device.
INVALIDATE KEY	12 Make a key in a QA Device invalid.
GETKEY	13 Get an encrypted key from a QA Device.
TEST	14 Request a QA Device to test the signature over an arbitrary block of data.
SIGN	15 Request a QA Device to create a signature over an arbitrary block of data.
TRANSFER DELTA	16 Request a QA Device to transfer some value from it to another QA Device where the value is correspondingly reduced in the Commanded QA Device).
TRANSFER ASSIGN	17 Request a QA Device to transfer an assignment of value to another QA Device . . .
START ROLLBACK	18 Request a QA Device to begin rollback proceedings to ensure that a previously transferred value has not and can never be used.
ROLLBACK	19 Request a QA Device to undo a previously requested transfer of value to another QA Device.

[0307] If the list is an incoming read-entity list or an outgoing modify-entity list, then the list ChipId is typically the ChipId of some other QA Device.

[0308] A signed outgoing list of entities being read from a QA Device has a signature over a block of data that includes

[0313] The ResultFlag is a byte that indicates the return status from a function. Callers can use the value of ResultFlag to determine whether a call to a function succeeded or failed, and if the call failed, the specific error condition.

[0314] Table 2 describes the ResultFlag values and the mnemonics used in the pseudocode

<u>ResultFlag value description</u>		
Mnemonic	Value	Description
Pass	0	Function completed successfully.
Fail	1	Function successfully completed requested task.
QA NotPresent	2	General failure. An error occurred during function processing.
Invalid Command	3	QA Device is not contactable
Bad Signature	4	The QA Device does not support the command
Invalid Key	5	Signature mismatch. The input signature didn't match the generated signature.
Invalid Key Type	6	Invalid keyslot number. The keyslot specified is greater than the number of keyslots supported in the QA Device, or the key in the specified keyslot is invalid.
Key Number Out Of Range	7	The key in the requested keyslot is the wrong type for the particular operation. For example, a TransportKey was requested for a data-based signature, or a DataKey was requested for a key-based signature.
Key Not Locked	8	A key was specified for a signature which had a key slot number out of range
Signature Generation Block Absent	9	A command was received, authenticated by an unlocked key. Unlocked keys may not be used to authenticate any operations, with the exception of the transport of keys, to authenticate and encrypt new key values.
Signature Generation Block Wrongly Present	10	A OutputSignatureGenerationBlock was not received in a command which requires an outgoing signature
Signature Block Absent	11	A OutputSignatureGenerationBlock was received in a command which does not require an outgoing signature
Signature Block Wrongly Present	12	A InputSignatureCheckingBlock was not received in a command which requires an incoming signature
Parameter Block Absent	13	A InputSignatureCheckingBlock was received in a command which does not require an incoming signature
Parameter Block Wrongly Present	14	An Input Parameter Block wasn't received in a command which requires that block, or an Output Parameter Block was not generated by a command which requires one.
Too Many Entities	15	An Input Parameter Block was received in a command which does not require that block, or an Output Parameter Block was generated in a command that does not require one.
Too Few Entities	16	The Input Parameter Block of the command has a list of more entities than the QA Device supports
Illegal Field Number	17	An Entity List or an Entity Descriptor List was received in a command or sent in a response with no entities.
Illegal Entity Descriptor	18	Field Number incorrect. The field number specified in an entity descriptor does not exist.
Modify Bit Wrong	19	An entity descriptor in an input or output parameter block list was set wrongly: it was "modify" when it needed to be "read", or "read" when it needed to be "modify".
Wrong ChipId	20	The QA Device was given a command which had a SignedInputParameterBlock with modify-entities, or generated a signed OutputParameterBlock with read-entities, and the ChipId in the signed block was incorrect, i.e. not the ChipId of the QA Device.
Illegal Entity	21	An entity in an Input Parameter Block of a command was received that is not legal for that command.
No Shared Key	22	An operation was requested in a command to a QA Device which requires a key to be shared between it and another QA Device. If there is no shared key, this error is returned.
Invalid Write Permission	23	Permission not adequate to perform operation. For example, trying to perform a Write or WriteAuth with incorrect permissions.
Field Is Read Only	24	A Write or an Authenticated Write command was applied to a read-only field that had already been written once.
Only Decrements Allowed	25	A Write or an Authenticated Write command was applied to a decrement-only field, which was not a decrement.
Key Already Locked	26	Key already locked. A key cannot be replaced if it has already been locked.
Illegal Key Entity	27	An Entity Descriptor in an Entity List wrongly specified a key value or descriptor that is not a legal entity for that command.
Illegal Field Entity	28	An Entity Descriptor in an Entity List wrongly specified a field value or descriptor that is not a legal entity for that command.

-continued

<u>ResultFlag value description</u>		
Mnemonic	Value	Description
Key Not Unlocked	28	A Replace Key command was received that was attempting to change a locked key.
Field Creation Not Allowed	29	Field creation was attempted in this QA Device, after it has been locked or there was an attempt to lock field creation after it had been already locked.
Field Storage Overflow	30	The QA Device is out of storage space for new fields.
Type Mismatch	31	Type of the data from which the amount is being transferred in the Upgrading QA Device, doesn't match the Type of data to which the amount in being transferred in the Device being upgraded.
Transfer Dest Field Invalid	32	A transfer was attempted on a field which is not capable of supporting a transfer.
Rollback Enable Field Invalid	33	The rollback enable field for the QA Device being transferred to is invalid.
No Transfer Source Field	34	There is no transfer source field available to do the transfer from.
Transfer Source Field Amount Insufficient	35	The transfer source field doesn't have the amount required for the transfer.
Invalid Operand	36	One of the command operands was invalid.
Field Over Maximum Allowed	37	A Write or an Authenticated Write command was applied to a field which would have made the field value exceed the limit implied by its "maximum allowed" bit field.
Transfer Fields Incompatible	38	The "who I am" and "who I accept" fields in the transfer source and transfer destination fields are not compatible.
Transfer Rolled Back	39	A transfer was attempted which failed. The transfer was successfully rolled back, so the source and transfer fields are unchanged.
No Matching Previous Transfer	40	A Rollback was attempted on a QA Device which had no record of having done a corresponding transfer (loss of previous record may occur depending on the depth of the rollback cache)
Key Not For Local Use	41	An operation was requested using a data key for which local use is not permitted.

[0315] Users of QA Devices must call the GetInfo function on each QA Device before calling any other functions on that device.

[0316] The GetInfo function tells the caller what kind of QA Device this is, what functions are available and what properties this QA Device has. The caller can use this information to correctly call functions with appropriately formatted parameters.

[0317] The first value returned, QA Device type, effectively identifies what kind of QA Device this is, and there-

fore what functions are available to callers. Source code control identifier tells the caller which software version the QA Device has. There must be a unique mapping of the source code control identifier to a body of source code, under source code control, in any released QA Device.

[0318] Additional information may be returned depending on the type of QA Device. The additional data fields of the output hold this additional information.

[0319] Table 3 describes each of the output parameters.

<u>Description of output parameters for GetInfo function</u>		
Parameter	#bytes	Description
ResultFlag	1	Indicates whether the function completed successfully or not. If it did not complete successfully, the reason for the failure is returned here.
QA Device type	1	This defines the function set that is available on this QA Device.
Source Code Control Identifier	4	This uniquely defines the source code for the QA Device, as controlled by a source code control system.
Key Replacement Allowed	1	Bit mask of keygroups which are not locked. Key replacement is allowed to add keys to those keygroups.
Maximum number of keys	1	The number of keyslots the QA Device can support

-continued

<u>Description of output parameters for GetInfo function</u>		
Parameter	#bytes	Description
Number of keys used	1	The number of keyslots the QA Device is currently using
Number of key groups	1	The number of keygroups that the QA Device is currently using
Field creation allowed	1	Non-zero if field creation is allowed
Number of fields	1	The number of fields which are present in the QA Device
Number of read-only words in device	2	The number of write-once then read-only (ROS) words that the QA Device supports
Number of read-only words used	2	The number of write-once then read-only (ROS) words that the QA Device is currently using
Number of writeable words in device	2	The number of writeable (RWS) words that the QA Device supports
Number of writeable words used	2	The number of writeable (RWS) words that the QA Device is currently using
ChipId	8	This QA Device's ChipId
VarDataLen	1	Length of bytes to follow.
VarData	(VarDataLen bytes)	This is additional application specific data, and is of length VarDataLen (i.e. may be 0).

[0320] Table 4 shows the mapping of QA Device Type:

<u>QA Device Types</u>	
QA Device Type\	Description
1	Base QA Device
2	Value Upgrader QA Device
3	Parameter Upgrader QA Device
4	Key Replacement QA Device
5	Trusted QA Device

[0321] Table 5 shows the mapping between the QA Device type and the available device functions on that QA Device.

<u>Mapping between QA Device Type and available device functions</u>		
QA Device Function	Supported on QA Device Types	Device description
Get Info	all	Base QA Device
Get Challenge		
Lock Key Groups		
Lock Field Creation		
Authenticated Read		

-continued

<u>Mapping between QA Device Type and available device functions</u>		
QA Device Function	Supported on QA Device Types	Device description
Authenticated Write		
Non-authenticated Write		
Create Fields		
Replace Key		
Invalidate Key		
Transfer Delta	2	Value Upgrader QA Device
Start Rollback		(e.g. Ink Refill QA Device)
Roll Back Amount		
Transfer Amount	3	Parameter Upgrader QA Device
Start Rollback		(e.g. Local Upgrader QA Device)
Rollback Field		
GetKey	4	Key Replacement QA Device
Sign	5	Trusted Device
Test		

[0322] Table 6 shows the VarData components for Value Upgrader and Parameter Upgrader QA Devices.

<u>VarData for Value and Parameter Upgrader QA Devices</u>		
VarData Components	Length in bytes	Description
DepthOfRollBackCache	1	The number of data sets that can be accommodated in the Xfer Entry cache of the device.

[0323] An Authenticated Transfer is the process where a store of value is securely transferred from one QA Device to another.

[0324] A Rollback is where a previous attempted transfer is annulled, when the transferring QA Device is given evidence that the transfer never succeeded, and can never succeed in the future.

[0325] When a transfer is taking place from one QA Device to another, the QA Device from which the value is being transferred is called the Source QA Device, and the QA Device to which the value is being transferred is called the Destination QA Device.

[0326] The stores of values can be either consumables, or properties.

[0327] In a printing application, consumables are things like picolitres of ink, millimetres of paper, page impressions etc. They are things that are consumed as the printing process is taking place.

[0328] In a printing application, properties are things like printer features, such as the right to print at a certain number of pages per second, or the right to interwork with a certain bit of equipment, such as a larger ink cartridge, (which may be cheaper to buy per litre of ink).

[0329] A property can also be a printer license, which has an implied printer feature set. That is, if a printer has a license, it has a certain feature set, and other non-selectable printer features have certain default values.

[0330] Properties are things which are not consumed as the printing takes place, but which can be assigned to a printer and which remain as attributes of that printer.

[0331] Fields in QA Devices have a transfer mode, which can be one of:

[0332] Quantity of Consumables: the field represents a volume of consumables. It can be the destination of a transfer, and if it has TxDE enabled, then it can be the source of a transfer of consumables,

[0333] Single Property: this field represents a single property of a printer, such as a printer feature or a license. This field can be assigned to, as the destination of a transfer, but cannot be the source of a transfer. Once a property has been assigned, it becomes operative, and it cannot be transferred any more.

[0334] Quantity of Properties: this field represents a quantity of properties, which are in transmit to their final destination. It can be the destination of a transfer, and also the source of a transfer. A quantity of properties does not confer any property to the QA Device which has them: they are in transmit to the place where they can be used as properties.

[0335] Other: this field cannot have value transferred from or to it.

[0336] In general, the flow of virtual consumables is from QACo, via the OEM factories, to the consumable containers, such as ink cartridges in the home or office. The virtual consumables are created ex nihil in QACo, transferred without being created or destroyed to the home or office, and then consumed. When virtual consumables are assigned to a

consumable container to be used in SOHO, it should be done in tandem with physically filling the container, so that the two are in agreement.

[0337] In general, the flow of properties is from QACo, via the OEM factories or OEM internet resellers, to printers and dongles, for use in the home and office. The properties are stored as quantities of properties until they get to their final destination, where they are assigned as single properties.

[0338] There are three general kinds of transfers, each with their corresponding rollbacks:

[0339] The transfer of a quantity of consumables. This is where a volume of consumables is transferred from source to destination. The transfer source field is decreased by the transfer delta amount, and the transfer destination field is increased by the same amount. This is a transfer delta.

[0340] The transfer of a quantity of properties. This is where a quantity of properties is transferred from source to destination. The transfer source field is decreased by the transfer delta amount, and the transfer destination field is increased by the same amount. This is also a transfer delta.

[0341] The assignment of a single property. This is where a single property is transferred from source to destination. The transfer source field is decreased by 1, and the transfer destination field is assigned with the property value. This is also a transfer assignment.

[0342] The transfer process has two basic requirements:

[0343] The transfer can only be performed if the transfer request is valid. The validity of the transfer request must be completely checked by the Source QA Device before it produces the required output for the transfer. It must not be possible to apply the transfer output to the Destination QA Device if the Source QA Device has already been rolled back for that particular transfer.

[0344] A process of rollback is available if the transfer was not received by the Destination QA Device. A rollback is performed only if the rollback request is valid. The validity of the rollback request must be completely checked by the Source QA Device, before it adjusts its value to a previous value before the transfer request was issued. It must not be possible to rollback an Source QA Device for a transfer which has already been applied to the Destination QA Device i.e the Source QA Device must only be rolled back for transfers that have actually failed. Similarly, it must not be possible to apply a transfer to the Destination QA Device after the rollback has been applied.

[0345] The transfer and rollback process is shown in FIG. 28.

[0346] The steps shown in FIG. 28 for a transfer and rollback process are:

[0347] The System performs an Authenticated Read of fields and keys in the destination QA Device. The output from the read includes field data, field descriptors, and the key descriptor of the key being used to authenticate the transfer, and a signature. It is essential that the fields are read together. This ensures that the

fields are correct, and have not been modified, or substituted from another device.

[0348] The System requests a Transfer from the Source QA Device with the amount that must be transferred, the field in the Source QA Device the amount must be transferred from, and the field in Destination QA Device the amount must be transferred to. The Transfer also includes the output from (1). The Source QA Device validates the Transfer based on the Authenticated Read output, checks that it has enough value for a successful transfer, and then produces the necessary transfer output. The transfer output typically consists of new field data for the field being refilled or upgraded, additional field data required to ensure the correctness of the transfer/rollback, along with a signature.

[0349] The System then applies the transfer output to the Destination QA Device, by calling an Authenticated Write function on it, passing in the transfer output from (2). The Write is either successful or not. If the Write is not successful, then the System may repeat calling the Write function using the same transfer output, which may be successful or not. If unsuccessful, the System initiates a Rollback of the transfer. The Rollback must be performed on the Source QA Device, so that it can adjust its value to a previous value before the current Transfer was initiated. It is not necessary to perform a rollback immediately after a failed Transfer. The Destination QA Device can still be used.

[0350] The System starts a Rollback by reading the fields and keys of the Destination QA Device.

[0351] The System makes a Start RollBack request to the Source QA Device with same input parameters as the Transfer, and the output from Read in (4). The Source QA Device validates the Start RollBack Request based on the Read output, and then produces the necessary Start Rollback output. The Start Rollback output consists only of additional field data along with a signature.

[0352] The System then applies the Start Rollback output to the Destination QA Device, by calling an Authenticated Write function on it, passing in the Start Rollback output. The Write is either successful or not. If the Write is not successful, then either (6), or (5) and (6) must be repeated.

[0353] The System then does an Authenticated Read of the fields of the Destination QA Device.

[0354] The System makes a RollBack request to the Source QA Device with same input parameters as the Transfer request, and the output from Read (7). The Source QA Device validates the RollBack request based on the Authenticated Read output, and then rolls back its field corresponding to the transfer.

[0355] There are two fields in every QA Device which can be the destination of a transfer, called the rollback enable fields.

[0356] The rollback enable fields are called RollbackEnable1 and RollbackEnable2 with field types=TYPE_ROLLBACK_ENABLE_1 and TYPE_ROLLBACK_ENABLE_2 respectively (see Table 329). They each have a transfer mode of "other", which means that they are never the destination field of a transfer, that is, they never get value

transferred to them. However, they take part in the authenticated writes which transfer value to other fields.

[0357] Both rollback enable fields are decrement-only fields, initialised to 0xFFFFFFFF when they are created, and they can only be decreased via authenticated writes.

[0358] When a transfer is requested, the authenticated read contains the field descriptors and field values for the rollback enable fields. The transfer source QA Device checks that they are present, and remembers their values.

[0359] The authenticated write for the transfer includes:

[0360] An assignment to the destination field being updated,

[0361] A decrement of -1 to RollbackEnable1, and

[0362] A decrement of -2 to RollbackEnable2.

[0363] If a rollback is requested, then the transfer source QA Device generates the arguments for an authenticated write to the transfer destination which include:

[0364] A decrement of -2 to RollbackEnable1, and

[0365] A decrement of -1 to RollbackEnable2.

[0366] This authenticated write only works if the transfer write had never been applied, (because otherwise the rollback write would be incrementing RollbackEnable2, which is not allowed; it is a decrement-only field.)

[0367] The pattern of "rollback enable value -1" and "rollback enable value -2" means that only one of the authenticated writes can be applied, not both. If the Transfer write has succeeded, then the Rollback write can never be applied, and if the Rollback write has succeeded, then the Transfer write can never be applied.

[0368] If the rollback write is successfully applied to the transfer destination, then another Authenticated Read is made to the rollback enable fields. This is presented as evidence to the transfer source QA Device, and if it can see that the rollback write has been successfully applied, it rolls back the transfer, and increments its source field.

[0369] The basic authorisation for a transfer comes from a key that has authenticated ReadWrite permission (stored in field information as KeyNum) to the destination fields in the Destination QA Device. This key is referred to as the transfer key.

[0370] After validating the input transfer request, the Source QA Device decrements the amount to be transferred from its source field, and produces the arguments for an authenticated write, and a signature using the transfer key.

[0371] The signature produced by the Source QA Device is subsequently applied to the Destination QA Device. The Destination QA Device accepts the transfer amount only if the signature is valid. Note that the signature is only valid if it was produced using the transfer key which has write permission to the destination field being written.

[0372] The Source QA Device validates the transfer request by matching the Type of the data in the destination field of Destination QA Device to the Type of data in the source field of the Source QA Device. This ensures that equivalent data Types are transferred e.g. a quantity of type

Network_OEM1_infrared ink is not transferred into a field of type Network_OEM1_cyan ink.

[0373] Each field which may be transferred from or to has a compatibility word in its field descriptor. The compatibility word consists of two 16-bit fields, called “who I am” and “who I accept”. For the transfer to take place, each side must accept the other. That is expressed in this way: if (the source “who I am” bitwise-ANDed with the destination “who I accept” is non-zero) AND (the destination “who I am” bitwise-ANDed with the source “who I accept” is non-zero) are both non-zero, then the transfer can take place, otherwise it can't.

[0374] In addition, when a quantity of properties is being transferred, the source field's “upgrade to/from” word is used as follows:

[0375] If the assignment is a “transfer delta”, then the “upgrade to/from” words in the source and destination fields must match, and

[0376] The transfer is a “transfer assignment”, then the previous value of the property must have been the “upgrade from” value, and then the assignment is of the “upgrade to” value.

[0377] This is the complete list of checks that must be made by the transfer source QA Device, before a transfer is authorised.

[0378] The signature for the authenticated read matches

[0379] The keygroup for the incoming data is locked, and the key is valid, is of type DataKey, and has a UseLocally set to 1.

[0380] All of the incoming fields can be written or at least decremented by the incoming key.

[0381] The transfer source QA Device has the appropriate key for the transfer

[0382] The rollback enable fields are present

[0383] The rollback enable field descriptors are decrement-only, type=rollback enable, transfer mode=other

[0384] The rollback enable values are ≥ 2

[0385] Source and destination field types match

[0386] Source and destination compatibility fields are compatible

[0387] If the transfer operation is “transfer delta”, then

[0388] Destination volume+delta \leq maximum allowed at destination

[0389] Source volume \geq delta

[0390] The source and destination fields either both have or both do not have an “upgrade option from/to” value

[0391] If the source field has an “upgrade option from/to” value, then it matches the destination field's value

[0392] The source and destination fields' transfer modes must be the same, and they must be either “quantity of consumables” or “quantity of properties”

[0393] If the transfer operation is “decrement and assign”, then

[0394] The source field's transfer modes must be “quantity of properties”, and the destination field's transfer mode must be “single property”

[0395] Destination value=“option from” value of the “upgrade option from/to” value

[0396] If any of these tests fail, then the transfer cannot proceed.

[0397] The Authenticated Write arguments should have these values:

[0398] The RollbackEnable1 field should have an authenticated write of its previous value—1

[0399] The RollbackEnable2 field should have an authenticated write of its previous value—2

[0400] If the transfer operation is Transfer Delta, then:

[0401] Destination volume should be set to original volume+delta.

[0402] If the transfer operation is “decrement and assign”, then

[0403] Destination value=“option to” value of the “upgrade option from/to” value

[0404] The implied delta value is 1.

[0405] The arguments of the Authenticated Write should have the “write/add” bit in the entity descriptors set to “add”, for the rollback enables, and the field value in the Transfer Delta case. It should be set to “write” for the field value in the Transfer Assign case. The use of the “add” option in the Authenticated Write eliminates a class of race conditions.

[0406] The Transfer Delta function is to transfer value, the value being a quantity of consumables or a quantity of properties. This distinction (compared to a Transfer Assign) is above.

[0407] It produces as its output the data and signature for updating given fields in a destination QA Device with an Authenticated Write. The data and signature when applied to the appropriate device through the Authenticated Write function, updates the fields of the device.

[0408] The system calls the Transfer Delta function on the upgrade device with a certain Delta. This Delta is validated by the Transfer Delta function for various rules, the function then produces the data and signature for the passing into the Authenticated Write function for the device being upgraded.

[0409] The Transfer Delta output consists of the new data for the field being upgraded, field data of the two rollback enable fields, and a signature using the transfer key.

[0410] The following data is saved in the transfer Source QA Device's Rollback Buffer:

[0411] The field number in the transfer source,

[0412] The field number in the transfer destination,

[0413] The key slot number in the transfer source,

[0414] The key slot number in the transfer destination,

[0415] The destination ChipId,

- [0416] The destination rollback enable counters, values and descriptors,
- [0417] The destination key descriptor.
- [0418] The delta.
- [0419] Non-volatile memory is memory that retains its state after power is removed. For example, flash memory is a form of non-volatile memory. The terms flash memory and non-volatile memory are used interchangeably in the detailed description.
- [0420] In a flash memory, a bit can either be in its erased state or in its programmed state. These states are referred to as E and P. For a particular flash memory technology, E may be 0 or 1, and P is the inverse of E.
- [0421] Depending on the flash technology, a FIB (Focused Ion Beam) can be used to change chosen bits of flash memory from E to P, or from P to E. Thus a FIB may be used to set a bit from an unknown state to a known state, where the known state depends on the flash memory technology.
- [0422] An integrated circuit (IC or chip) may be manufactured with flash memory, and may contain an embedded processor for running application program code.
- [0423] XOR is the bitwise exclusive-or function. The symbol \square is used for XOR in equations.
- [0424] A Key, referred to as K, is an integer (typically large) that is used to digitally sign messages or encrypt secrets. K is N bits long, and the bits of K are referred to as K_0 to K_{N-1} , or K_i , where i may run from 0 to N-1.
- [0425] The Binary Inverse of a Key is referred to as $\sim K$. The bits of $\sim K$ are referred to as $\sim K_i$, where i may run from 0 to N-1.
- [0426] A Random Number used for the purposes of hiding the value of a key when stored in non-volatile memory is referred to as R. The bits of R are referred to as R_i , where i may run from 0 to N-1.
- [0427] If a function of a key K is stored in non-volatile memory, it is referred to as X. The bits of X are referred to as X_i , where i may run from 0 to N-1.
- [0428] In embedded applications, it is often necessary to store a secret key in non-volatile memory such as flash on an integrated circuit (IC), in products that are widely distributed.
- [0429] In certain applications, the same key is stored in multiple ICs, all available to an attacker. For example, the IC may be manufactured into a consumable and the consumable is sold to the mass market.
- [0430] The problem is to ensure that the secret key remains secret, against a variety of attacks.
- [0431] This document is concerned with FIB (Focussed Ion Beam) attacks on flash-based memory products. Typically a FIB attack involves changing a number of bits of flash memory from an unknown state (either E or P) into a known state (E or P). Based on the effect of the change, the attacker can deduce information about the state of the bits of the key.
- [0432] After an attack, if the chip no longer works, it is disposed of. It is assumed that this is no impediment to the

attacker, because the chips are widely distributed, and the attackers can use as many of them as they like.

[0433] Note that the FIB attack is a write-only attack—the attacker modifies flash memory and tests for changes of the chip behaviour.

[0434] Attacks that involve reading the contents of flash memory are much more difficult, given the current state of flash memory technology. However, if an attacker were able to read from the flash memory, then it would be straightforward to read the entire contents, then to disassemble the program and calculate what operations are being performed to obtain the key value. In short, all keys would be compromised if an attacker is capable of arbitrary reads of flash memory.

[0435] Note that this document is addressing direct attacks on the keys stored in flash memory. Indirect attacks are also possible. For example, an attacker may modify an instruction code in flash memory so that the contents of the accumulator are sent out an output port. Indirect attacks are not addressed in this document.

[0436] If a key K consisting of N bits is stored directly in non-volatile memory, and an attacker knows both N and the location of where K is stored within the non-volatile memory, then the attacker can use a simple FIB attack to obtain K.

[0437] For each bit i in K:

[0438] The attacker uses the FIB to set K_i to P,

[0439] If the chip still works the attacker can deduce that the bit was originally P.

[0440] If the chip no longer works, then the attacker can deduce that the bit was originally E.

[0441] A series of FIB attacks allows the attacker to obtain the entire key. At most, an attacker requires N chips to obtain all N bits, but on average only N/2 chips are required.

[0442] If the attacker cannot set a bit to P, but can set it to E, then an equivalent attack is possible. i.e. For each bit i in K:

[0443] The attacker uses the FIB to set K_i to E.

[0444] If the chip still works the attacker can deduce that the bit was originally E.

[0445] If the chip no longer works, then the attacker can deduce that the bit was originally P.

[0446] Thus storing a key directly in non-volatile memory is not secure, because it is easy for an attacker to use a FIB to retrieve the key.

[0447] Instead of storing K directly in flash, it is possible to store R and X, where R is a random number essentially different on each chip, and X is calculated as $X = K \square R$. Thus K can be reconstructed by the inverse operation i.e. $K = X \square R$.

[0448] In this case, a simple FIB attack as described in Section 2.1 will not work, even if the attacker knows where X and R are stored. This is because the bits of X are essentially random, and will differ from one chip to the next. If the attacker can deduce that a bit of X in one chip is a

certain state, then this will not have any relation to what the corresponding bit of X is in any other chip.

[0449] Even so, an attacker can still extract the key. For each bit i in the key:

[0450] The attacker uses the FIB to set both X_i and R_i to P,

[0451] If the chip still works, the attacker knows that X_i and R_i were originally either both P or both E. Both of these cases imply that the key bit K_i is 0.

[0452] If the chip no longer works, the attacker knows that exactly one of X_i and R_i was originally P and one was E. This implies that the key bit K_i is 1.

[0453] If the chip no longer works, replace it with a new chip.

[0454] If the attacker cannot set a bit to P, but can set it to E, then an equivalent attack is possible.

[0455] A series of FIB attacks allows an attacker to obtain the entire key. For each bit, there is a 50% chance that the chip cannot be reused because it is damaged by the attacks (this is the case where $X_i < R_i$). This means that on average it will take it will take an attacker $50\% \times N$ chips to obtain all N bits.

[0456] Therefore this method of storing a key is not considered secure, because it is easy for an attacker to use a FIB to retrieve the key.

[0457] Instead of storing K directly in flash, it is possible to store K and its binary inverse $\sim K$ in flash such that for each chip, K is stored randomly in either of 2 locations and $\sim K$ is stored in the other of the 2 locations (the program that accesses the key also needs to know the placement). As a result, given a randomly selected chip, an attacker does not know whether the bit stored at a particular location belongs to K or $\sim K$.

[0458] If the program in flash memory checks that the value read from the first location is the binary inverse of the value stored in the second location, before K is used, and the program fails if it is not, then an attacker cannot use the behaviour of the chip to determine whether a single bit attack hit a bit of K or $\sim K$.

[0459] However the chip is subject to an attacker performing multiple-bit FIB attacks, assuming that the attacker knows the two locations where K and $\sim K$ are stored, but does not know which location contains K; and that the program in the chip checks that the values stored at the two locations are inverses of each other, and fails if they are not.

[0460] For each bit $i > 0$ in the key:

[0461] The attacker chooses a positive integer T.

[0462] The attacker repeats the following experiment up to T times, on a series of chips:

[0463] The attacker uses the FIB to set bits 0 and i of the value stored at one of the 2 locations (the attacker doesn't know if the value is K or $\sim K$) to P,

[0464] If the chip still works, then the attacker can deduce that K_0 and K_i have the same value: they are either both 1 or both 0. This is because the bits that were attacked must have both been originally P, and the

FIB left them that way, and so the chip still worked. It is not clear whether the attacked bits were in K or $\sim K$, and so the attacker can't deduce whether the key bits were 0 or 1, but the attacker has discovered that K_0 and K_i are the same. If this result occurs, stop repeating the experiment.

[0465] If the chip no longer works, then the attacker can only deduce that either the bits in the key are different, (with a probability $2/3$), or the bits in the key are the same but the attack hit the bits in the key or the inverse that were both E, (with a probability of $1/3$). That is, the attacker can get no certain information from this result, but can get a probable result.

[0466] 3. After T attempts, if there have been any results that indicate that K_0 and K_i have the same value, then the attacker knows that the bits are the same. Otherwise, the attacker knows that there is a $(1/3)^T$ probability that the bits are the same. The probability that K_0 and K_i are the same can be made arbitrarily close to 0 by increasing T until the attacker has an appropriate level of comfort that the bits are different.

[0467] If the attacker cannot set a bit to P, but can set it to E, then an equivalent attack is possible.

[0468] At the end of the experiments, the relation of K_0 to all of the other key bits K_i ($i=1$ to $N-1$), is either known or almost certainly known. This means that the key value is almost certainly known to within two guesses: one where $K_0=0$, and the other where $K_0=1$. For each guess, the other key bits K_i are implied by the known relations. The attacker can try both combinations, and at worst may need to try other combinations of keys based on the probabilities returned for each bit position during the experiment.

[0469] An attacker can use a series of FIB attacks to obtain the entire key. For each K_i , there is a 75% chance that the chip cannot be reused because it is damaged by the attacks: this is the case where the tested bits K_0 and K_i were not both P. On average, it will take 1.5 attempts to determine that K_0 and K_i are identical, and T attempts to find that K_0 and K_i are different. This means that on average it will take it will take an attacker $75\% \times (T+1.5) / 2 \times (N-1)$ chips to obtain the relations between K_0 and the other $N-1$ bits.

[0470] Therefore this method of storing a key is not considered secure, because it is easy for an attacker to use a FIB to retrieve the key.

[0471] It is possible to store X, $\sim X$ and R in flash memory where R is a random number, K is the key, $X=K \square R$, and $\sim X = \sim K \square R$.

[0472] X, $\sim X$ and R are stored in memory randomly with respect to each other, and the program that accesses the key also needs to know the placement. Thus, for a randomly selected chip it is not clear to an attacker whether a bit at a particular location belongs to X, $\sim X$ or R.

[0473] It is assumed that the attacker knows where X, $\sim X$ and R are stored, but does not know which one is stored in each of the 3 locations; and that the program in the chip checks that the stored value for X is indeed the binary inverse of the stored value for $\sim X$, and fails if it is not.

[0474] An attacker cannot extract the key using the method described in Section 2.3 because that method will

reveal whether X_0 is the same as X_i , (where X is one of X , $\sim X$ and R), for an individual chip, but this can give no information about the relationship of K_0 and K_i , because they are XORed with the random R that differs from chip to chip.

[0475] So a “pairs of bits” FIB attack cannot get the attacker any information about K .

[0476] However, K still susceptible to attack, by an attacker performing FIB attacks on pairs of bit pairs.

[0477] It is assumed that the chip is programmed with X , $\sim X$ and R , and they are in known locations, but it is not known by the attacker what order they are in; and that the program in the chip checks that stored value for X is indeed the binary inverse of the stored value for $\sim X$, and fails if it is not.

[0478] For each bit $i>0$ in the key:

[0479] Choose a positive integer T .

[0480] Repeat this experiment up to T times, on a series of chips:

[0481] The attacker uses the FIB to set bits 0 and i of two of the entities (X , $\sim X$ or R), to P .

[0482] The attacker does not know which of the entities were hit.

[0483] If the attacker hits bits in X and R , and all 4 of them were P , or if the attacker hits bits in $\sim X$ and R , and all 4 of them were P , then the program will always pass. In these events, the attacker can deduce that K_0 and K_i are the same. The probability of this outcome is $1/6$. If this result occurs, stop repeating the experiment.

[0484] If the attacker hits bits in X and R , and not all 4 of them were P , or if the attacker hits bits in $\sim X$ and R , and not all 4 of them were P , then the program will always fail. In this case the attacker can only deduce that either the bits in the key are different, or the bits in the key are the same but the attack hit the bits in the key or the inverse that were both E . That is, the attacker can get no certain information from this result, but can get a probable result. The probability of this outcome is $1/2$. The probability of this outcome when $K_0=K_i$ is $1/6$. The probability of this outcome when $K_0 < > K_i$ is $1/3$.

[0485] If the attacker hits bits in X and $\sim X$, then the program will always fail, because the corresponding bits in X and $\sim X$ must be different (by definition). One bit from each bit pair must have been changed from P to E by the attack, and the program checks will fail. In this event, the attacker cannot find out any information about the bits of the key K . The probability of this outcome is $1/3$. The probability of this outcome when $K_0=K_i$ is $1/6$. The probability of this outcome when $K_0 < > K_i$ is $1/6$.

[0486] After T attempts, if there have been any results that indicate that K_0 and K_i have the same value, then the attacker knows that the bits are the same. Otherwise, the attacker knows that there is a $(2/3)^T$ probability that the bits are the same. The probability that K_0 and

K_i are the same can be made arbitrarily close to 0 by increasing T . That is, the attacker can be almost certain that the bits are different.

[0487] If the attacker cannot set a bit to P , but can set it to E , then an equivalent attack is possible.

[0488] At the end of the experiments, the relation of K_0 to all of the other key bits K_i ($i=1$ to $N-1$), is either known or almost certainly known. This means that the key value is almost certainly known to within two guesses: one where $K_0=0$, and the other where $K_0=1$. For each guess, the other key bits K_i will be implied by the known relations. The attacker can try both combinations, and at worst may need to try other combinations of keys based on the probabilities returned for each key position during the experiment.

[0489] Thus an attacker can use a series of FIB attacks to obtain the entire key.

[0490] Therefore this method of storing a key is not considered secure because it is not difficult for an attacker to use a FIB to retrieve the key.

[0491] The above-described attacks rely on the attacker having knowledge of where the key K and related key information are placed within flash memory.

[0492] If the program insertion re-links the program every time a chip is programmed, then the key and key-related information can be placed in an arbitrary random places in memory, on a per-chip basis. For any given chip, the attacker will not know where the key could be.

[0493] This will slow but not stop the attacker. It is still possible to launch statistical attacks to discover the key.

[0494] This section shows how any attack that can succeed against keys in known locations can be modified to succeed against keys that are placed in non-overlapping random locations, different for every programmed chip. The following assumptions are made:

[0495] That the places where the key information may be stored do not overlap with each other. That is, if a FIB attack hits a bit of key information, the attacker knows which bit of the key was hit, and

[0496] That the attacker knows the possible locations of the key information, and their alignment, and

[0497] That if a FIB attack leaves a chip reporting that the key was wrong, then it is more likely that this was because the key was corrupted, than because some part of the program code that manipulates the key was hit.

[0498] When an attacker attacks a bit in flash memory with a FIB attack to set its state to P there are a number of possibilities:

[0499] A bit can be hit that is already in the state P , and is therefore not changed. There is no change of behaviour of the chip. In some circumstances this can provide the attacker with some information.

[0500] A bit that is part of some key-related information can be hit, and the bit changes from state E to P . This will cause the program to fail, reporting an incorrect key value.

- [0501] A bit that is not part of some key-related information can be hit, and the bit changes from state E to P. This may or may not cause the chip to fail for some other reason.
- [0502] There are an equivalent set of possibilities if the attacker uses a FIB attack to set the state of a bit to E.
- [0503] It is important to distinguish between the two kinds of failures: (a) failures where the program either reports an incorrect key value, or it is clear that the key value is incorrect, because it is unable to encrypt, and (b) other kinds of failures. If the program becomes unable to do key-related functions (encrypt, decrypt, digitally sign or check digital signatures, etc), but is otherwise functioning well, then the attacker can deduce that the most recent attack probably hit some key-related information.
- [0504] If a program stops working, or comes up with some other unrelated error condition, then the most recent attack hit some part of the flash memory that was not key information, but was necessary for something else.
- [0505] In the situation where K is placed into a random location in flash memory for each chip, and that the possible locations for the key cannot overlap with each other, then an attacker can extract the key.
- [0506] For each bit i in $N-1$:
- [0507] Choose a positive integer T .
- [0508] Repeat the following experiment T times, on a series of chips:
- [0509] The attacker chooses the address A of a potential key.
- [0510] The attacker uses the FIB to set the A_i to P.
- [0511] If the chip gets an error that implies that it has an incorrect key value, then probably K was actually at address A . In this case, the attacker records a hit, and records that K_i is probably E.
- [0512] Otherwise the attacker records a miss.
- [0513] The attacker would do well to discard the chip, whether or not the chip failed. This is because there might be some silent damage to the chip, that could interact in unexpected ways with subsequent FIB attacks. It is safer to start each new experiment with a new chip.
- [0514] After T attempts, the attacker has a record of how many hits H_i were recorded for bit i in the key.
- [0515] Since there are N key bits in flash memory, out of a total of M total bits of flash memory, the attacker can expect that a key bit was hit N out of M times. Sometimes this hit would have changed a bit from E to P, and other times it would leave the bit unchanged at P.
- [0516] The attacker is now able to observe that for each bit i , the H_i/T converge to two values: N/M and 0. If $H_i/T=N/M$, then K_i is probably E, and if $H_i/T=0$, then K_i is probably P.
- [0517] To launch this attack, an attacker requires $T \times N$ chips. Note that for the experiments to be useful, T needs to be large enough to launch an attack on M .
- [0518] If the attacker cannot set a bit to P, but can set it to E, then an equivalent attack is possible.
- [0519] This method of storing a key is not considered secure, because it is difficult, though not impossible, for an attacker to use an FIB to retrieve the key.
- [0520] In the situation where for each chip, K and $\sim K$ are each placed into a random location in flash memory such that the possible locations for storage do not overlap with each other, and that the program in the chip checks that the stored values at the two locations are inverses of one another and fails if it is not, then an attacker can extract the key.
- [0521] For each bit i in $N-1$:
- [0522] Choose a positive integer T .
- [0523] Repeat this experiment T times, on a series of chips:
- [0524] The attacker chooses an address A (hoping it will be the address of K or $\sim K$).
- [0525] The attacker uses the FIB to set bits A_0 and A_i to P.
- [0526] If the chip gets an error that implies that it has an incorrect key value, then probably either K or $\sim K$ was actually at address A . In this case, the attacker records a hit. The attacker can also deduce that bits A_0 and A_i were not both P. This can mean one of 2 things:
- [0527] A_0 and A_i were different, and they were part of K or $\sim K$. This implies that $K_0 < > K_i$. This happens $2/3$ of the time.
- [0528] A_0 and A_i were both E, and they were part of K or $\sim K$. This implies that $K_0 = K_i$. This happens $1/3$ of the time.
- [0529] Otherwise the attacker records a miss.
- [0530] The attacker would do well to discard the chip, whether or not the chip failed. This is because there might be some silent damage to the chip, that could interact in unexpected ways with subsequent FIB attacks. It is safer to start each new experiment with a new chip.
- [0531] After T attempts, there will be a record of how many hits H_i were recorded for bit i in the key.
- [0532] Since there are $2N$ bits in flash memory containing K and $\sim K$, out of a total of M total bits of flash memory, the attacker can expect that key-related bits were hit $2N$ out of M times.
- [0533] The attacker should observe that for each bit i , the H_i/T converge to two values: N/M and $N/2M$. If $H_i/T=N/M$, then K_i is probably $\sim K_0$, and if $H_i/T=N/2M$, then K_i is probably K_0 .
- [0534] At the end of the experiments, the relation of K_0 to all of the other key bits K_i ($i=1$ to $N-1$), is probably known. This means that the key value is probably known to within two guesses: one where $K_0=0$, and the other where $K_0=1$. For each guess, the other key bits K_i will be implied by the known relations. The attacker should try both combinations.
- [0535] To launch this attack, an attacker requires $T \times N$ chips. Note that for the experiments to be useful, T needs to be large enough to launch an attack on M .
- [0536] If the attacker cannot set a bit to P, but can set it to E, then an equivalent attack is possible.

[0537] Therefore this method of storing a key is not considered secure, because although it is difficult, it is not impossible for an attacker to use a FIB to retrieve the key.

[0538] Storing a key in arbitrary non-overlapping places in flash memory will slow but not stop a determined attacker.

[0539] The same methods of attack that work for keys in known locations, work for keys in unknown locations. They are slower because they rely on statistics that are confounded with the failures that occur because of reasons other than corruption of keys.

[0540] A sufficient number of experiments allows the attacker to isolate the failures caused by differences in the value of the bits of keys from other failures.

[0541] The above-described attacks rely on the attacker having knowledge of where the key K and related key information are placed within flash memory, or knowledge that the locations where the key information may be placed do not overlap each other.

[0542] It is possible to place the key and key-related information in random locations in memory on a per-chip (assuming the program that references the information knows where the information is stored). For a randomly selected chip, the attacker will not know exactly where the key is stored. This will slow but not stop the attacker. It is still possible to launch statistical attacks that discover the key.

[0543] Any attack that can succeed against keys in known locations can be modified to succeed against keys that are placed in random locations, different for every programmed chip. The following assumptions are made:

[0544] If a FIB attack leaves a chip reporting that the key was wrong, then it is more likely that this was because the key was corrupted, than because some part of the program code that manipulates the key was hit.

[0545] Some inside information is helpful for the attack.

[0546] For a given computer architecture and software design, the keys will be held in memory in units of a particular word size, and those words will be held in an array of words, aligned with the word size. So, for example, a particular key might be 512 bits long, and held in an array of 32-bit words, and the words are aligned in flash memory at 32-bit boundaries. Similarly, another system might have a key that is 160 bits long, held in an array of bytes, aligned on byte boundaries.

[0547] Additional useful information for the attacker is the minimum alignment in flash memory for the key, denoted by W .

[0548] If a key is N bits long, aligned with a word-size of W , and placed in flash memory starting at an arbitrary word address, then there will be N/W bits that are aliased together from the point of view of the attacker. This is called the aliased bit group. This is because an attack on bit x in flash could be a hit to K_x, K_{x+W}, K_{x+2W} , etc, depending on which word in memory the key started.

[0549] For example, if a particular key is 512 bits long, and is held in an array of 32-bit words, then there are 16 elements (512/32) in each aliased bit group. Similarly, if

another system's key is 160 bits long, held in an array of bytes, then there are 20 elements (160/8) in each aliased bit group.

[0550] When an attacker discovers something about a particular chip's key by attacking a bit of flash memory, the attacker can generally only deduce some bulk characteristics of the aliased bit group, rather than individual bits of the key. For small enough aliased bit groups, however, this can dramatically reduce the search size necessary to compromise the key.

[0551] The boundary conditions of aliased bit groups allows an attacker to gather particular types of statistics:

[0552] If a flash memory stores key related information on arbitrary bit boundaries, then the word size is 1, and the aliased bit group size is the key size. In this situation, the attacker can only gather statistics about the key bits as a whole.

[0553] If a flash memory stores key related information in words with an alignment greater than or equal to the key size, then the aliased bit group size is 1. In this situation, each bit of flash memory can only be a unique bit of the key, and any key-related information the attacker finds about that bit of flash memory can be applied to exactly that key bit.

[0554] It is in the attacker's interest for the word size to be as large as possible, so that there is a minimum of aliasing of bits.

[0555] When an attacker attacks a bit in flash memory with a FIB attack, there are a number of possible outcomes:

[0556] A bit can be hit that is already in the state P , and is therefore not changed. There is no change of behaviour of the chip. In some circumstances this can provide the attacker with some information.

[0557] A bit that is part of some key-related information can be hit, and the bit changes from state E to P . This will cause the chip to become unable to use its key correctly, and the program will fail.

[0558] A bit that is not part of some key-related information can be hit, and the bit changes from state E to P . This may or may not cause the chip to fail for some other reason.

[0559] There are an equivalent set of possible outcomes if the attacker uses a FIB attack to set the state of a bit to E .

[0560] It is important to distinguish between the two kinds of failures: (a) failures where the program becomes unable to use its key, and (b) other kinds of failures. If the program becomes unable to do key-related functions (encrypt, decrypt, digitally sign or check digital signatures, etc), but is otherwise functioning well, then the attacker can deduce that the most recent attack hit some key-related information.

[0561] If a program stops working, or comes up with some other unrelated error condition, then the most recent attack hit some part of the flash memory that was not key information, but was necessary for something else.

[0562] If the key K is placed into a random location in flash memory for each chip, then an attacker can extract the key.

[0563] For each bit i in $0-W-1$, where W =the word size:

[0564] Choose a positive integer T .

[0565] The attacker repeat the following experiment T times, on a series of chips:

[0566] The attacker chooses the address A of a word in flash memory.

[0567] The attacker uses the FIB to set the A_i to P .

[0568] If the chip becomes unable to use the key K , then clearly the word at address A was in K . That is, $A_i=K_{i+jW}$, where $(i+jW)<N$. In this case, the attacker records a hit.

[0569] Otherwise the attacker records a miss.

[0570] The attacker would do well to discard the chip, whether or not the chip failed. This is because there might be some silent damage to the chip, that could interact in unexpected ways with subsequent FIB attacks. It is safer to start each new experiment with a new chip.

[0571] After T attempts, there will be a record of how many hits H_i were recorded for bit i in the word size.

[0572] At the end of the experiment, the attacker has W fractions H_i/T , one for every bit in the flash memory's words.

[0573] Since there are N key bits in flash memory, out of a total of M total bits of flash memory, the attacker can expect that a key bit was hit N out of M times. Sometimes this hit would have changed a bit from E to P , and other times it would leave the bit unchanged at P .

[0574] If all of the bits in the key's aliased bit group were E , then the attacker should expect that $H_i/T=N/M$. That is, all of the bits of a particular word bit i that hit a key bit changed it from E to P .

[0575] If all of the bits in the key's aliased bit group were P , then the attacker should expect that $H_i/T=0$. That is, all of the bits of a particular word bit i that hit a key bit left it unchanged at P .

[0576] If there are k bits in the aliased bit group, then the attacker should be able to observe that $B_i=k(H_i/T)/(N/M)$ takes on $k+1$ values, from 0 to k , for each bit i in the flash memory words.

[0577] B_i is the number of bits in the aliased bit group that are E in the key. $k-B_i$ is the number of bits in the aliased bit group that are P in the key. So the attacker knows to within a permutation what the key bit values are.

[0578] To launch this attack, an attacker requires $T \times W$ chips. Note that for the experiments to be useful, T needs to be large enough to launch an attack on M .

[0579] If the attacker cannot set a bit to P , but can set it to E , then an equivalent attack is possible.

[0580] Therefore this method of storing a key is not considered secure, because it is difficult, though not impossible, for an attacker to use a FIB to retrieve the key.

[0581] If K and $\sim K$ are each placed into one of two random locations in flash memory for each chip, and the program checks that the stored values in both locations are

binary inverses of each other and fails if they are not, then an attacker can extract the key.

[0582] For each bit i in $1-W-1$, where W =the word size:

[0583] Choose a positive integer T .

[0584] The attacker repeat the following experiment T times, on a series of chips:

[0585] The attacker chooses the address A of a word in flash memory.

[0586] The attacker uses the FIB to set bits A_0 and A_i to P .

[0587] If the chip becomes unable to use the key K , then clearly the word at address A was either in K or $\sim K$. That is, $A_i=K_{i+jW}$, or $A_i=\sim K_{i+jW}$, where $(i+jW)<N$. In this case, the attacker records a hit. The attacker can also deduce that bits A_0 and A_i were not both P .

[0588] This can mean one of 2 things:

[0589] A_0 and A_i were different, and they were part of K or $\sim K$. This implies that $K_{i+jW} > K_{jW}$, for some j . This happens $2/3$ of the time.

[0590] A_0 and A_i were both E , and they were part of K or $\sim K$. This implies that $K_{i+jW}=K_{jW}$, for some j . This happens $1/3$ of the time.

[0591] Otherwise the attacker records a miss.

[0592] The attacker would do well to discard the chip, whether or not the chip failed. This is because there might be some silent damage to the chip, that could interact in unexpected ways with subsequent FIB attacks. It is safer to start each new experiment with a new chip.

[0593] After T attempts, there will be a record of how many hits H_i were recorded for bit i in the word size.

[0594] At the end of the experiment, the attacker has $W-1$ fractions H_i/T , one for each bit $1-W-1$ in the flash memory's words.

[0595] If an attack hits bits K_{i+jW} and K_{jW} , for some j , and those key bits are different, this will always cause a failure. If those key bits are the same, this will cause a failure half the time, on average.

[0596] So the attacker should expect that

[0597] $H_i/T=(N/M) \times \text{Sum}(j=0 \text{ to } k-1, (\text{if } (K_{i+jW}=K_{jW}) \text{ then } 1/2 \text{ else } 1))$ where k is the number of elements in the aliased key group.

[0598] If we define $B_i=(H_i/T)/(N/M)$, for $i=1$ to $W-1$, then the attacker finds $B_i=(k-1)$ for the case where key bit $K_{i+jW} > K_{jW}$, for j in 0 to $k-1$. The attacker finds $B_i=(k-1)/2$ for the case where key bit $K_{i+jW}=K_{jW}$, for j in 0 to $k-1$.

[0599] The attacker should try various combinations of K_i that make these equalities true. This dramatically decreases the search space necessary to compromise the key.

[0600] If the attacker cannot set a bit to P , but can set it to E , then an equivalent attack is possible.

[0601] Storing a key in arbitrary places in flash memory will slow but not stop a determined attacker.

[0602] The same methods of attack that work for keys in known locations work for keys in unknown locations. They are slower, because they rely on statistics that are confounded with the failures that occur because of reasons other than corruption of keys.

[0603] A sufficient number of experiments will allow the attacker to isolate the failures caused by differences in the value of the bits of keys, from other failures.

[0604] When keys are stored in flash, the key bits can be guarded by an increasingly elaborate set of operations to confound attackers. Examples of such operations include the XORing of key bits with random numbers, the storing of inverted keys, the random positioning of keys in flash memory, and so on.

[0605] Based on previous discussion, it seems likely that this increasingly elaborate series of guards can be attacked by an increasingly elaborate series of FIB attacks. Note however, that the number of chip samples required by an attacker to make a success likely may be prohibitively large, and thus a previously discussed storage method may be appropriately secure.

[0606] The basic problem of the storing and checking of keys is that the bits of the key-related entities ($\sim K$, R , etc) can be directly correlated to the bits of the key.

[0607] Assuming a single key, a method of solving the problem is to guard the key bits using a value that has no correlation with the key bits as follows:

[0608] R and X are stored in the flash memory where R is a random number different for each chip, and $X=K\oplus\text{owf}(R)$, where $\text{owf}()$ is a one-way function such as SHA1 (see [1]).

[0609] R and X may be stored at known addresses

[0610] For the program to use the key, it must calculate $K=X\oplus\text{owf}(R)$

[0611] The one-way function should have the property that if there is any bit difference in the function input, there are on average differences in about half of the function output bits. SHA1 has this property.

[0612] If an attacker modifies even a single bit of R , it will affect multiple bits of the $\text{owf}()$ output and thus multiple bits of the calculated K .

[0613] This property makes it impossible to make use of multiple bit attacks, because if bit 0 and bit i of R are modified, this will affect on average $N/2$ bits of K , that may or may not include bits 0 and i . The attacker cannot deduce any information about bits of K .

[0614] Similarly, if bit 0 and bit i of X are modified, the attacker is able to tell if X_0 and X_i were both P in this particular chip, but this will give the attacker no information about key bits K_i , because the attacker will not know the whole of R , and hence the attacker doesn't know any bits of $\text{owf}(R)$.

[0615] If the attacker is restricted to FIB attacks, it doesn't matter if R and X are stored in fixed known locations, because these FIB attacks cannot extract any information about K .

[0616] A chip may need to hold multiple keys in flash memory. For this discussion it is assumed that a chip holds NumKeys keys, named $K[0]-K[\text{NumKeys}-1]$

[0617] These keys can be held in a number of ways.

[0618] They can be stored as NumKey instances of any of the insecure key storage algorithms discussed above. These key storage methods are insecure for the storage of multiple keys for the same reasons that they are insecure for the storage of single keys.

[0619] If the keys are stored as processed keys using the method introduced in Section 5 then there is an issue of how many random numbers are required for same storage. The two basic cases are:

[0620] Processed keys are stored along with a single random number R as $X[0]-X[\text{NumKeys}-1]$, where $X[i]=K[i]\oplus\text{owf}(R)$

[0621] Processed keys are stored along with a set of random numbers $R[0]-R[\text{NumKeys}-1]$, in the form $X[0]-X[\text{NumKeys}-1]$, where $X[i]=K[i]\oplus\text{owf}(R[i])$.

[0622] Both storage techniques are immune to FIB attacks, as long as no keys have been compromised.

[0623] If storage technique (1) is used, and an attacker knows one of the keys, then that knowledge can be used with a FIB attack to obtain the value of another keys and hence all keys. The attack assumes that the attacker knows:

[0624] the location of R and $X[0]-X[\text{NumKeys}-1]$, where $X[i]=K[i]\oplus\text{owf}(R)$.

[0625] the value of $K[a]$, and wishes to discover the value of $K[b]$.

[0626] For each bit i in the key $K[b]$:

[0627] The attacker uses the FIB to set R_i and $X[a]_i$ to P ,

[0628] If the chip still works when it uses $K[a]$,

[0629] The attacker knows that R_i and $X[a]_i$ in this particular chip were originally P ,

[0630] The attacker uses the FIB to set $X[b]_i$ to P ,

[0631] If the chip still works when it uses $K[b]$, then the attacker can deduce that $X[b]_i$ was originally P , in which case $K[b]_i$ is 0 .

[0632] If the chip no longer works when it uses $K[b]$, then the attacker can deduce that $X[b]_i$ was originally E , in which case $K[b]_i$ is 1 .

[0633] If the chip no longer works, then

[0634] repeat this procedure for $K[b]_i$ with a new chip.

[0635] If the attacker cannot set a bit to P , but can set it to E , then an equivalent attack is possible.

[0636] The attack relies on the fact that even if the attacker does not know the value of R , the same value $\text{owf}(R)$ is used to guard all of the keys and there is known correlation between corresponding bits of each X .

[0637] Note that if the locations of R and $X[0]-X[\text{NumKeys}-1]$, are randomised during program insertion, it will slow but not stop this kind of attack, for the reasons described in Section 4.

[0638] Therefore storage technique (2) is more secure, as it uses a set of different owf(R[i]) values to guard the keys. However storage technique (2) requires additional storage over storage technique (1).

[0639] The problem with storage technique (1) is that there is a single value (owf(R)) used to guard the keys, and there is known correlation between corresponding bits of each stored form of key. i.e. XOR is a poor encryption function.

[0640] Storage technique (2) relies on storing a different R for each key so that the values used to protect each key are uncorrelated on a single chip, and are uncorrelated between chips. The problem with storage technique (2) is that additional storage is required—one R per key.

[0641] However, it is possible to use a single base-value such that the bit-pattern used to protect each K is different. i.e.: storage technique (3) is as follows:

[0642] Processed keys are stored with a single random number R in the form X[0]-X[NumKeys-1], where X[i]=K[i]⊕owf(Ri), where owf() is a one-way function such as SHA1.

[0643] For the program to use a key, it must calculate K[i]=X[i]⊕owf(Ri).

[0644] The keys may be stored at known addresses.

[0645] In general, technique (3) stores X[i] where X[i]=Encrypt(K[i]) using key Q. The Encrypt function is XOR, and Q is obtained by owf(Ri) where R is an effectively random number per chip. Normally XOR is not a strong encryption technique (as can be seen by the attack in Section 2.2), but it is strong when applied to an uncorrelated data, as is the case with this method. The technique used to generate Q is such that uncorrelated Qs are obtained to protect the keys, each Q is uncorrelated from the stored R, and both Rs and Qs are uncorrelated between chips. It isn't quite a pure one-time-pad, since the same stored R is used each time the key is decrypted, but it is a one-time-pad with respect to the fact that each Q is different on a single chip, and each R (and hence the Qs) is different between chips.

[0646] The following terminology is now used:

[0647] A nonce is a parameter that varies with time. A nonce can be a generated random number, a time stamp, and so on. Because a nonce changes with time, an entity can use it to manage its interactions with other entities.

[0648] A session is an interaction between two entities. A nonce can be used to identify components of the interaction with a particular session. A new nonce must be issued for each session.

[0649] A replay attack is an attack on a system which relies on replaying components of previous legitimate interactions.

[0650] In the generation of non-deterministic sequences, nonces are useful in challenge-response systems to protect against replay attacks.

[0651] A entity, referred to as a challenger, can issue a nonce for each new session, and then require that the nonce be incorporated into the encrypted response or be included with the message in the signature generated from the other party in the interaction. The incorporation of a challenger's

nonce ensures that the other party in the interaction is not replaying components of a previous legitimate session, and authenticates that the message is indeed part of the session they claim to be part of.

[0652] However, if an attacker can predict future nonces, then they can potentially launch attacks on the security of the system. For example, an attacker may be able to determine the distance in nonce-sequence-space from the current nonce to a nonce that has particular properties or can be used in a man-in-the-middle attack.

[0653] Therefore security is enhanced by an attacker not being able to predict future nonces.

[0654] To prevent these kinds of attacks, it is useful for the sequence of nonces to be hard to predict. However, it is often difficult to generate a sequence of unpredictable random numbers.

[0655] Generation of sequences is typically done in one of two ways:

[0656] An entity can use a source of genuinely random numbers, such as a physical process which is non-deterministic;

[0657] An entity can use a means of generating pseudo-random numbers which is computationally difficult to predict, such as the Blum Blum Shub pseudo-random sequence algorithm [1].

[0658] For certain entities, neither of these sources of random numbers may be feasible. For example, the entity may not have access to a non-deterministic physical phenomenon. Alternatively, the entity may not have the computational power required for complex calculations.

[0659] What is needed for small entities is a method of generating a sequence of random numbers which has the property that the next number in the sequence is computationally difficult to predict.

[0660] At a starting time, for example when the entity is programmed or manufactured, a random number called x₀ is injected into the entity. The random number acts as the initial seed for a sequence, and should be generated from a strong source of random numbers (e.g. a non-deterministic physically generated source).

[0661] When the entity publishes a nonce R, the value it publishes is a strong one-way function (owf) of the current value for x: i.e:

$$R=owf(x)$$

[0662] The strong one-way function owf() can be a strong one-way hash function, such as SHA-1, or a strong non-compressing one-way function.

[0663] Characteristics of a good one-way function for this purpose are that it:

[0664] is easy to compute

[0665] produces a sufficiently large dynamic range as output for the application

[0666] is computationally infeasible to find an input which produces a pre-specified output (i.e.

[0667] it is preimage resistant). This means an attacker can't determine x_n from R_n.

- [0668] is computationally infeasible to find a second input which has the same output as any pre-specified input (i.e. it is 2nd-preimage resistant).
- [0669] produces a large variance in the output for minimally different inputs is collision resistant over the output bit range i.e. is computationally infeasible to find any two distinct inputs x_1 and x_2 which produce the same output
- [0670] The number of bits n in x needs to be sufficiently large with respect to the chosen one-way function. For example, n should be at least 160 when owf is SHA-1.
- [0671] To advance to the next nonce, the seed is advanced by a simple means. For example, it may be incremented as an n -bit integer, or passed through an n -bit linear feedback shift register.
- [0672] The entity publishes a sequence of nonces $R_0, R_1, R_2, R_3, \dots$ based on a sequence of seeds $x_0, x_1, x_2, x_3, \dots$
- [0673] Because the nonce is generated by a one-way function, the exported sequence, $R_0, R_1, R_2, R_3, \dots$ etc., is not predictable (or deterministic) from an attacker's point of view. It is computationally difficult to predict the next number in the sequence.
- [0674] The advantages of this approach are:
- [0675] The calculation of the next seed, and the generation of a nonce from the seed are not computationally difficult.
- [0676] A true non-deterministic number is only required once, during entity instantiation. This moves the cost and complexity of the difficult generation process out of the entity.
- [0677] There is no need for a source of random numbers from a non-deterministic physical process in the running system.
- [0678] Note that the security of this sequence generation system relies on keeping the current value for x secret. If any of the x values is known, then all future values for x can be predicted and hence all future R values can be known.
- [0679] Note that the random sequence produced from this is not a strong random sequence e.g. from the view of guaranteeing particular distribution probabilities. The behaviour is more akin to random permutations. Nonetheless, it is still useful for the purpose of generating a sequence for use as a nonce in such applications as a SoC-based implementation of the QA Logical Interface.
- [0680] In one embodiment, functionally identical code segments are stored in each of multiple devices. The device can be, for example, a series of printer cartridges, and more specifically the QA printer chip attached to such cartridges.
- [0681] The programs stored in the devices are functionally identical to each other, which is to say that they implement the same instructions in the same way, although the individual instances of the programs may operate on different data and using different keys.
- [0682] Whilst the program instances are functionally identical, they are broken up into code segments that are each stored at different locations in the flash memory. For convenience, each code segment can be a function or other relatively self-contained subset of instructions, although this is not required.
- [0683] After the chip has been manufactured, the program code is injected such that the position of particular code segments varies across the devices. The memory location at which each code segment starts can be selected in any convenient manner. It is not strictly necessary that every segment be placed in a truly random or unique location in the memory from device to device. Rather, it is enough that a potential attacker cannot rely on the same code being in the same place in a series of different integrated circuits.
- [0684] It is still, however, desirable that the location of particular code segments be selected at least pseudo-randomly, and preferably randomly.
- [0685] In the preferred embodiment, an initial instruction is located at an initial memory location that is the same across all of the devices. This means that a common boot program can be used at startup, since it always looks to the initial location to commence the program. Somewhere in the code segment following the initial location, the program jumps to one of the random or pseudo-random memory locations. From this point in the program, the instructions are effectively unknown to an attacker. Of course, it is possible that only a relatively small (but preferably important) code section is located at this random or pseudo-random location. The rest of the code can be at common locations across the devices.
- [0686] The reference to the random or pseudo-random location in the program code can be explicit (as above) or implicit. For example, the program code can refer to a pointer or register that contains the location of interest. The location is stored in the pointer or register during program instantiation. The location of interest can also be stored in a jump table.
- [0687] Multiple random or pseudo random locations can be used. The program can jump to multiple locations during its execution, each of the locations being different across several devices. The code segments themselves can be different to each other, such that even the segments themselves (in number or size) vary from device to device.
- [0688] Terms: A number of terms are used in the specification and claims. The following list includes some definitions that are to be used when these terms appear, unless a contrary meaning is clearly intended in context:
- [0689] "Relatively unique"—Depending upon the context, this phrase generally means that a value or bit-pattern is rarely repeated across multiple devices. It is usually preferable that the value or bit-pattern is selected in a random or at least pseudo-random way. However, in some applications it is sufficient to ensure that the value or bit-pattern is merely not frequently repeated from device to device. Sometimes, a relatively small number of potential values or bit-patterns will be sufficient to make attacking a chip or other device sufficiently hard that it will not be worth attempting
- [0690] "Associated with a base key"—A variant key is associated with a base key when it is the result of applying a one way function to the base key and a bit-pattern.
- [0691] "Cryptographically strong"—Whilst this is a relative term, it has some use when comparing the ease with

which functions can be broken when used in cryptography. For example, an XOR function, whilst useful in some circumstances in cryptography, is considerably easier to “crack” than, say, a hash function or sufficient length. Also, a hash function combined with a key into a MAC (i.e. “message authentication code”) such as HMAC-SHA1 used with a certain length of key will be cryptographically stronger if the key length is increased, up to a certain length of key.

[0692] “Bit-pattern”—A generic term that can refer to keys, nonces, random numbers, pseudo-random numbers, serial numbers, and any other strings of interest.

[0693] “Functionally identical”—Code segments that are functionally identical operate in the same way, using the same functions and subroutines as each other where each of the functions and subroutines are also functionally identical. However they may use different keys, constants or variables, and/or operate on different stored data or data and program segment code stored at different locations in memory. For example, two functionally identical code segments may each load a particular constant into a register for use in evaluating an expression, and although the order of steps taken to load the constant may differ between segments, the value of the constant may differ between segments, and the address of the constant in memory may differ between segments, the functional intent of the code segment is the same for both.

[0694] It will be appreciated by those skilled in the art that the foregoing represents only a preferred embodiment of the present invention. Those skilled in the relevant field will immediately appreciate that the invention can be embodied in many other forms.

1. A method of storing at least one functionally identical code segment in each of a plurality of printer cartridges of a printing system, each of the printer cartridges having a memory, the method comprising, for each printer cartridge:

- (a) determining a first memory location which is different than a first memory location in at least a plurality of the other printer cartridges; and
- (b) storing a first of the at least one code segments in the memory at the determined first memory location;

wherein the different first memory locations are determined such that the code segment storage location of all the respective devices cannot be ascertained from the code segment storage location of any one of the printer cartridges.

2. A method according to claim 1, wherein at least one of the code segments in each of the printer cartridges includes an initial instruction, the initial instruction being located at an initial instruction location, the initial instruction location being the same in all the printer cartridges.

3. A method according to claim 2, wherein the initial instruction in each printer cartridge is indicative of the first memory location of that printer cartridge.

4. A method according to claim 3, wherein the initial instruction is indicative of the first memory location by including an explicit reference to the memory location.

5. A method according to claim 4, wherein the initial instruction is indicative of the first memory location by including an implicit reference to the memory location.

6. A method according to claim 5, wherein the implicit reference is a pointer to a location at which the address of the first memory location is stored.

7. A method according to claim 6, wherein the implicit reference is a pointer to a register that holds the address of the first memory location.

8. A method according to claim 1, wherein step (a) includes randomly selecting the first memory location.

9. A method according to claim 8, wherein step (a) includes selecting the first memory location based on a stochastic process or mechanism.

10. A method according to claim 1, wherein step (a) includes selecting the first memory location from an existing list or sequence of memory locations.

11. A method according to claim 1, wherein each memory is incorporated in a quality assurance integrated circuit of each printer cartridge.

12. A method according to claim 1, wherein the memory is non-volatile memory.

13. A method according to claim 1, for storing a first bit-pattern in non-volatile memory of a printer cartridge, the method comprising:

- (a) applying a one way function to a second bit-pattern associated with the printer cartridge, thereby to generate a first result;
- (b) applying a second function to the first result and the first bit-pattern, thereby to generate a second result; and
- (c) storing the second result in the memory, thereby indirectly storing the first bit-pattern.

14. A method according to claim 1, for storing a bit-pattern in each of a plurality of printer cartridges, each of the printer cartridges having a memory, the method comprising, for each printer cartridge:

- (a) determining a first memory location; and
- (b) storing the bit-pattern at the first memory location;

wherein the first memory locations are different in at least a plurality of the respective printer cartridges.

15. A method according to claim 1, for storing multiple first bit-patterns in non-volatile memory of a printer cartridge, the method comprising, for each of the first bit-patterns to be stored:

- (a) applying a one way function to a third bit-pattern based on a second bit-pattern associated with the printer cartridge, thereby to generate a first result;
- (b) applying a second function to the first result and the first bit-pattern, thereby to generate a second result; and
- (c) storing the second result in the memory, thereby indirectly storing the first bit-pattern;

wherein the third bit-patterns used for the respective first bit-patterns are relatively unique compared to each other.

16. A printer cartridge having a code segment stored in it in accordance with the method of claim 1.

17. A plurality of printer cartridges having respective code segments stored in them in accordance with the method of claim 1.