

(19) AUSTRALIAN PATENT OFFICE

(54) Title
Physical access control

(51)⁶ International Patent Classification(s)
E05B 49/00 (2006.01) 19/10
G06K 17/00 (2006.01) 20060101ALI2005122
G06K 19/10 (2006.01) ^{OBMJP} **G06Q**
G06Q 20/00 20/00
(2006.01) 20060101ALI2005100
G07C 9/00 (2006.01) ^{8BMEP} **G07C**
H04L 9/32 (2006.01) 9/00
E05B 49/00 20060101ALI2005100
20060101AFI2005122 ^{8BMEP} **H04L**
^{OBMJP} **G06K** 9/32
17/00 20060101ALI2005122
20060101ALI2005122 ^{OBMJP}
^{OBMJP} **G06K** PCT/US03/10748

(21) Application No: 2003228468

(22) Application Date: 2003 .04 .08

(87) WIPO No: W003/088166

(30) Priority Data

(31) Number	(32) Date	(33) Country
60/446,149	2003 .02 .10	US
60/370,867	2002 .04 .08	US
60/421,756	2002 .10 .28	US
60/443,407	2003 .01 .29	US
60/372,951	2002 .04 .16	US
60/421,197	2002 .10 .25	US
60/373,218	2002 .04 .17	US
60/427,504	2002 .11 .19	US
60/420,795	2002 .10 .23	US
10/395,017	2003 .03 .21	US
60/422,416	2002 .10 .30	US
60/374,861	2002 .04 .23	US

(43) Publication Date : 2003 .10 .27

(43) Publication Journal Date : 2003 .12 .04

(71) Applicant(s)
Corestreet, Ltd.

(72) Inventor(s)
Reyzin, Leo, Engberg, David, Sinelnikov, Alex, Libin, Phil, Micali, Silvio

(74) Agent/Attorney
Davies Collison Cave, 1 Nicholson Street, Melbourne, VIC, 3000

(56) Related Art
US 5742035
EP0618550
FR2774833
EP0798671
US 5768379
EP1024239
EP0723251

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

CORRECTED VERSION

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
23 October 2003 (23.10.2003)


PCT

(10) International Publication Number
WO 2003/088166 A3

- (51) International Patent Classification⁷: **G07C 9/00**, 60/443,407 29 January 2003 (29.01.2003) US
G07F 7/08 60/446,149 10 February 2003 (10.02.2003) US
10/395,017 21 March 2003 (21.03.2003) US
- (21) International Application Number: PCT/US2003/010748 (71) Applicant: **CORESTREET, LTD.** [US/US]; One Alewife Center, Suite 200, Cambridge, MA 02140 (US).
- (22) International Filing Date: 8 April 2003 (08.04.2003) (72) Inventors: **MICALI, Silvio**; 459 Chestnut Hill Avenue, Brookline, MA 02146 (US). **ENGBERG, David**; 38 Dover Street, Cambridge, MA 02140 (US). **LIBIN, Phil**; 57 Clay Street, Cambridge, MA 02140 (US). **REYZIN, Leo**; 2456 Massachusetts Avenue, Apt.7, Cambridge, Massachusetts 02116 (US). **SINELNIKOV, Alex**; 278 Clarendon Street, Apt.7, Boston, Massachusetts 02116 (US).
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/370,867 8 April 2002 (08.04.2002) US
60/372,951 16 April 2002 (16.04.2002) US
60/373,218 17 April 2002 (17.04.2002) US
60/374,861 23 April 2002 (23.04.2002) US
60/420,795 23 October 2002 (23.10.2002) US
60/421,197 25 October 2002 (25.10.2002) US
60/421,756 28 October 2002 (28.10.2002) US
60/422,416 30 October 2002 (30.10.2002) US
60/427,504 19 November 2002 (19.11.2002) US
- (74) Agents: **MUIRHEAD, Donald, W.** et al.; Choate, Hall & Stewart, Exchange Place, 53 State Street, Boston, MA 02109-2804 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,

[Continued on next page]

(54) Title: PHYSICAL ACCESS CONTROL



(57) Abstract: A system and method are disclosed for controlling physical access through a digital certificate validation process that works with standard certificate formats and that enables a certifying authority (CA) to prove the validity status of each certificate C at any time interval (e.g., every day, hour, or minute) starting with C's issue date, D_i. C's time granularity may be specified within the certificate itself, unless it is the same for all certificates. For example, all certificates may have a one-day granularity with each certificate expires 365 days after issuance. Given certain initial inputs provided by CA, a one-way hash function is utilized to compute values of a specified byte size that are included on the digital certificate and to compute other values that are kept secret and used in the validation process.

WO 2003/088166 A3



CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PI, PL, PT, RO, RU, SC, SD, SE,
SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC,
VN, YU, ZA, ZM, ZW.

Published:

— with international search report

(88) Date of publication of the international search report:

1 April 2004

(48) Date of publication of this corrected version:

5 August 2004

(84) Designated States (regional): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,
ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO,
SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM,
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(15) Information about Correction:

see PCT Gazette No. 32/2004 of 5 August 2004, Section II

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

03 Aug 2009

20032228468

- 1 -

PHYSICAL ACCESS CONTROL**Cross-reference to Related Applications**

The present application is based on: U.S. Provisional Application No. 60/370,867,
5 filed April 8, 2002, entitled SCALABLE CERTIFICATE VALIDATION AND
SIMPLIFIED PKI MANAGEMENT; U.S. Provisional Application No. 60/372,951, filed
April 16, 2002, entitled CLOCK-LESS DEVICE VALIDATION; U.S. Provisional
Application No. 60/373,218, filed April 17, 2002, entitled TECHNIQUES FOR
TRAVERSING HASH SEQUENCES; U.S. Provisional Application No. 60/374,861,
10 filed April 23, 2002, entitled PHYSICAL ACCESS CONTROL; U.S. Provisional
Application No. 60/420,795, filed October 23, 2002, entitled SECURE PHYSICAL
ACCESS; U.S. Provisional Application No. 60/421,197, filed October 25, 2002, entitled
REAL TIME CREDENTIALS OVER OCSP; U.S. Provisional Application No.
60/421,756, filed October 28, 2002, entitled REAL TIME CREDENTIALS; U.S.
15 Provisional Application No. 60/422,416, filed October 30, 2002, entitled PROTECTING
MOBILE COMPUTING RESOURCES; U.S. Provisional Application No. 60/427,504,
filed November 19, 2002, entitled PRIVATE KEY SECURE PHYSICAL ACCESS OR
REAL TIME CREDENTIALS (RTCs) IN KERBEROS-LIKE SETTINGS; U.S.
Provisional Application No. 60/443,407, filed January 29, 2003, entitled THREE-
20 FACTOR AUTHENTICATION WITH REAL-TIME VALIDATION; and U.S.
Provisional Application No. 60/446,149, filed February 10, 2003, entitled RTC
PHYSICAL ACCESS WITH LOWER-END CARDS; the teachings of all of which are
incorporated herein by reference.

25 The present application is related to U.S. Patent Application No. 10/103,541, filed
March 20, 2002, entitled SCALABLE CERTIFICATE VALIDATION AND
SIMPLIFIED MANAGEMENT, the teachings of which are incorporated herein by
reference, which itself is a continuation in part of U.S. Patent Application No. 09/915,180,
filed July 25, 2001, entitled CERTIFICATE REVOCATION SYSTEM, and which is a
30 continuation of U.S. Patent Application No. 09/483,125, filed January 14, 2000, which is
a continuation of U.S. Patent Application No. 09/356,745, filed July 19, 1999, which is a

20032228468 03 Aug 2009

- 2 -

continuation of U.S. Patent Application No. 08/823,354, filed March 24, 1997, (now U.S. Patent No. 5,960,083), which is a continuation of U.S. Patent Application No. 08/559,533, filed November 16, 1995, (now U.S. Patent No. 5,666,416), which is based on U.S. Provisional Patent Application No. 60/006,038, filed October 24, 1995. U.S. Patent

5 Application No. 10/103,541 is also a continuation in part of U.S. Patent Application No. 08/992,897, filed December 18, 1997, which is based on U.S. Provisional Application No. 60/033,415, filed December 18, 1996, and which is a continuation in part of U.S. Patent Application No. 08/715,712, filed September 19, 1996, entitled CERTIFICATE REVOCATION SYSTEM, which is based on U.S. Provisional Application No.

10 60/004,796, filed October 2, 1995, entitled CERTIFICATE REVOCATION SYSTEM. U.S. Patent Application No. 08/992,897 is also a continuation in part of U.S. Patent Application No. 08/729,619, filed October 11, 1996, entitled TREE- BASED CERTIFICATE REVOCATION SYSTEM, (now U.S. Patent No. 6,097,811), which is based on U.S. Provisional Application No. 60/006,143, filed November 2, 1995, entitled

15 TREE BASED CERTIFICATE REVOCATION SYSTEM. U.S. Patent Application No. 08/992,897 is also a continuation in part of U.S. Patent Application No. 08/804,868, filed February 24, 1997, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, which is a continuation of U.S. Patent Application No. 08/741,601, filed November 1, 1996, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, which is based

20 on U.S. Provisional Application No. 60/006,143, filed November 2, 1995, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM. U.S. Patent Application No. 08/992,897, is also a continuation in part of U.S. Patent Application No. 08/872,900, filed June 11, 1997, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, which is a continuation of U.S. Patent Application No. 08/746,007, filed November 5,

25 1996, entitled CERTIFICATE REVOCATION SYSTEM, (now U.S. Patent No. 5,793,868), which is based on U.S. Provisional Application No. 60/025,128, filed August 29, 1996, entitled CERTIFICATE REVOCATION SYSTEM. U.S. Patent Application No. 08/992,897 is also based on U.S. Provisional Application No. 60/035,119, filed

February 3, 1997, entitled CERTIFICATE REVOCATION SYSTEM, and is also a continuation in part of U.S. Patent Application No. 08/906,464, filed August 5, 1997, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, (abandoned), which is a continuation in part of U.S. Patent Application No. 5 08/763,536, filed December 9, 1996, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, (now U.S. Patent No. 5,717,758), which is based on U.S. Provisional Application No. 60/024,786, filed September 10, 1996, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, and is based on U.S. Patent Application No. 08/636,854, filed April 23, 1996, (now U.S. Patent No. 10 5,604,804), and is also based on U.S. Provisional Application No. 60/025,128, filed, August 29, 1996, entitled CERTIFICATE REVOCATION SYSTEM. U.S. Patent Application No. 08/992,897 is also a continuation in part of U.S. Patent Application No. 08/756,720, filed November 26, 1996, entitled SEGMENTED CERTIFICATE REVOCATION LISTS, (abandoned), which is based on U.S. Provisional 15 Application No. 60/025,128, filed August 29, 1996, entitled CERTIFICATE REVOCATION SYSTEM, and is also based on U.S. Patent No. 08/715,712, filed September 19, 1996, entitled CERTIFICATE REVOCATION SYSTEM, (abandoned), and is also based on U.S. Patent Application No. 08/559,533, filed November 16, 1995, (now U.S. Patent No. 5,666,416). U.S. Patent Application No. 20 08/992,897 is also a continuation in part of U.S. Patent Application No. 08/752,223, filed November 19, 1996, entitled CERTIFICATE ISSUE LISTS, (now U.S. Patent No. 5,717,757), which is based on U.S. Provisional Application No. 60/025,128, filed August 29, 1996, entitled CERTIFICATE REVOCATION SYSTEM, and is also a continuation in part of U.S. Patent Application No. 08/804,869, filed February 25 24, 1997, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, (abandoned), which is a continuation of U.S. Patent Application No. 08/741,601, filed November 1, 1996, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, (abandoned), which is based on U.S. Provisional Application No. 60/006,143, filed November 2, 1995, entitled TREE-BASED CERTIFICATE 30 REVOCATION SYSTEM. U.S. Patent Application No. 08/992,897, is also a continuation in part of U.S. Patent Application No. 08/823,354, filed March 24, 1997, entitled CERTIFICATE REVOCATION SYSTEM, (now U.S. Patent No.

2003228468 03 Aug 2009

- 4 -

5,960,083), which is a continuation of U.S. Patent Application No. 08/559,533, filed November 16, 1995, entitled CERTIFICATE REVOCATION SYSTEM, (now U.S. Patent No. 5,666,416), which is based on U.S. Provisional Application No. 60/006,038, filed October 24, 1995, entitled ENHANCED CERTIFICATE REVOCATION SYSTEM.

- 5 U.S. Patent Application No. 10/103,541 is also based on U.S. Provisional Application No. 60/277,244, filed March 20, 2001, and U.S. Provisional Application No. 60/300,621, filed June 25, 2001, and U.S. Provisional Application No. 60/344,245, filed December 27, 2001. The disclosure of all of the above documents is incorporated herein by reference.

- 10 The present application is also related to U.S. Patent Application No. 09/915,180, filed June 25, 2001, entitled CERTIFICATE REVOCATION SYSTEM, the teachings of which are incorporated herein by reference, which itself is a continuation of U.S. Patent Application No. 09/483,125, filed January 14, 2000, which is a continuation of U.S. Patent Application No. 09/356,745, filed July 19, 1999, which is a continuation of U.S.
- 15 Patent Application No. 08/823,354, filed March 24, 1997, (now U.S. Patent No. 5,960,083), which is a continuation of U.S. Patent Application No. 08/559,533, filed November 16, 1995, (now U.S. Patent No. 5,666,416), which is based on U.S. Provisional Application No. 60/006,038, filed October 24, 1995. The teachings of all of the above are incorporated herein by reference.

20

The present application is also related to U.S. Patent Application 10/395,017, filed March 21, 2003, entitled EFFICIENT CERTIFICATE REVOCATION, the teachings of which are incorporated herein by reference, which itself is a continuation of U.S. Patent Application No. 10/244,695 filed September 16, 2002, which is a continuation of U.S.

- 25 Patent Application No. 08/992,897 filed December 18, 1997, (now U.S. Patent No. 6,487,658), which is based on U.S. provisional patent application No. 60/033,415, filed December 18, 1996, and which is a continuation in part of U.S. Patent Application No. 08/715,712, filed September 19, 1996, entitled CERTIFICATION REVOCATION SYSTEM, which is based on U. S. Patent Application

No. 60/004,796, October 2, 1995, entitled CERTIFICATE REVOCATION SYSTEM, and which is also a continuation in part of U.S. Patent Application No. 08/729,619, filed October 10, 1996, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, (now U.S. Patent No. 6,097,811), which is based on U.S. Patent Application No. 60/006,143, filed November 2, 1995, entitled Tree Based CERTIFICATE REVOCATION SYSTEM, and which is also a continuation in part of U.S. Patent Application No. 08/804,868, filed February 24, 1997, entitled Tree-Based CERTIFICATE REVOCATION SYSTEM, Abandoned, which is a continuation of U.S. Patent Application No. 08/741,601, filed November 1, 1996, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, Abandoned, which is based on U.S. Patent Application No. 60/006,143, filed November 2, 1995, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, and which is also a continuation in part of U.S. Patent Application No. 08/872,900, filed June 11, 1997, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, Abandoned, which is a continuation of U.S. Patent Application No. 08/746,007 filed November 5, 1996, entitled CERTIFICATE REVOCATION SYSTEM, (Now U.S. Patent No. 5,793,868), which is based on U.S. Patent Application No. 60/025,128, filed August 29, 1996, entitled CERTIFICATE REVOCATION SYSTEM, and which is also based on U.S. Patent Application No. 60/035,119, filed February 3, 1997, entitled CERTIFICATE REVOCATION SYSTEM, and which is also a continuation in part of U.S. Patent Application No. 08/906,464, filed August 5, 1997, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, Abandoned, which is a continuation of U.S. Patent Application No. 08/763,536 filed December 9, 1996, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, (now U.S. patent No. 5,717,758), which is based on U.S. Patent Application No. 60/024,786, filed September 10, 1996, entitled WITNESS BASED CERTIFICATE REVOCATION SYSTEM, and is also based on U.S. Patent Application No. 08/636,854, filed April 23, 1997, (now U.S. Patent No. 5,604,804), and U.S. Patent Application No. 60/025,128, filed August 29, 1996, entitled CERTIFICATE REVOCATION SYSTEM, and which is also a continuation in part of U.S. Patent Application No. 08/756,720, filed November 26, 1996, entitled SEGMENTED CERTIFICATE REVOCATION LISTS, Abandoned, which is based

on U.S. Patent Application No. 60/025,128, filed August 29, 1996, entitled CERTIFICATE REVOCATION SYSTEM, and also based on U.S. Patent Application No. 08/715,712, filed September 19, 1996, entitled CERTIFICATE REVOCATION SYSTEM, Abandoned, and is also based on U.S. Patent Application No. 08/559,533, filed November 16, 1995, (now U.S. Patent No. 5,666,416), and which is also a continuation in part of U.S. Patent Application No. 08/752,223, filed November 19, 1996, entitled CERTIFICATE ISSUE LISTS, (now U.S. Patent No. 5,717,757), which is based on U.S. Patent Application No. 60/025,128, filed August 29, 1996, entitled, CERTIFICATE REVOCATION SYSTEM, and is also a continuation in part of U.S. Patent Application No. 08/804,869, filed February 24, 1997, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, Abandoned, which is a continuation of U.S. Patent Application No. 08/741,601, filed November 1, 1996, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, Abandoned, which is based on U.S. Patent Application No. 60/006,143, filed November 2, 1995, entitled TREE-BASED CERTIFICATE REVOCATION SYSTEM, and which is also a continuation in part of U.S. Patent Application No. 08/823,354 filed March 24, 1997, entitled CERTIFICATE REVOCATION SYSTEM, (now U.S. Patent No. 5,960,083) which is a continuation of U.S. Patent Application No. 08/559,533, filed November 16, 1995, entitled CERTIFICATE REVOCATION SYSTEM, (Now U.S. Patent No. 5,666,416), which is based on U.S. Patent Application No. 60/006,038, filed October 24, 1995, entitled CERTIFICATE REVOCATION SYSTEM. The teachings of all of the above are incorporated herein by reference.

Field of the Invention

The present invention relates to the field of digital certificates and more particularly to the field of digital certificate validation for controlling physical access.

Background of the Invention

In essence, a digital certificate (C) consists of a certifying authority's (CA's) digital signature securely binding together several quantities: SN, a serial number

unique to the certificate, PK, the public key of the user, U, the user's identifier, D₁, the issue date, D₂, the expiration date, and additional fields. In symbols, $C = \text{SIG}_{CA}(\text{SN}, \text{PK}, \text{U}, \text{D}_1, \text{D}_2, \dots)$.

5 It is widely recognized that digital certificates provide the best form of Internet and other access authentication. However, they are also difficult to manage. Certificates may expire after one year (i.e., $D_2 - D_1 = 1$ year), but they may be revoked prior to their expiration; for instance, because their holders leave their companies or assume different duties within them. Thus, each transaction enabled
10 by a given digital certificate needs a suitable proof of the current validity of that certificate, and that proof often needs to be archived as protection against future claims.

 Unfortunately, traditional technologies for proving the validity of issued
15 certificates do not scale well. At tomorrow's volume of digital certificates, today's validity proofs will be either too hard to obtain in a secure way, or too long and thus too costly to transmit (especially in a wireless setting). Certificate validation is universally recognized as a crucial problem. Unless efficiently solved, it will severely limit the growth and the usefulness of PKIs.

20

 Today, there are two main approaches to proving certificates' validity: Certificate Revocation Lists (CRLs) and the Online Certificate Status Protocol (OCSP).

25 CRLs

 CRLs are issued periodically. A CRL essentially consists of a CA-signed list containing all the serial numbers of the revoked certificates. The digital certificate presented with an electronic transaction is then compared to the most recent CRL. If the given certificate is not expired but is on the list, then everyone knows from the
30 CRL that the certificate is not valid and the certificate holder is no longer authorized to conduct the transaction. Else, if the certificate does not appear in the CRL, then the certificate is deduced to be valid (a double negative).

CRLs have not found much favor; for fear that they may become unmanageably long. (A fear that has been only marginally lessened by more recent CRL-partition techniques.) A few years ago, the National Institute of Standards and Technology tasked the MITRE Corporation to study the organization and cost of a Public Key Infrastructure (PKI) for the federal government. (See Public Key Infrastructure, Final Report; MITRE Corporation; National Institute of Standard and Technology, 1994). This study concluded that CRLs constitute by far the largest entry in the Federal PKI's cost list.

OCSP

In the OCSP, a CA answers a query about a certificate C by returning its own digital signature of C's validity status at the current time. The OCSP is problematic in the following areas.

Bandwidth. Each validity proof generated by the OCSP has a non-trivial length. If RSA or other factoring based signature schemes are used, such a proof in fact requires at a minimum 2,048 bits for the CA's signature.

Computation. A digital signature is a computationally complex operation. In certain large applications, at peak traffic, the OCSP may require computing millions of signatures in a short time, which is computationally very expensive to do.

Communication (if centralized). Assume a single validation server implements the OCSP in a centralized manner. Then, all certificate-validity queries would have, eventually, to be routed to it, and the server will be a major "network bottleneck" causing considerable congestion and delays. If huge numbers of honest users suddenly query the server, a disrupting "denial of service" will probably ensue.

Security (if distributed). In general, distributing the load of a single server across several (e.g., 100) servers, strategically located around the world, alleviates network congestion. In the OCSP case, however, load distribution introduces worse

2003228468 03 Aug 2009

- 9 -

problems than those it solves. In order to sign its responses to the certificate queries it receives, each of the 100 servers should have its own secret signing key. Thus, compromising any of the 100 servers is compromising the entire system. Secure vaults could protect such distributed servers, but at *great* cost.

5

Summary

Some embodiments relate to a method for an entity A to control the access of at least one user U to at least one disconnected door D having means to determine the current time, comprising the steps of:

- 10 for each time interval d of a sequence of dates, having A produce a digital signature SIGUDd, indicating that user U can access door D during time interval d;
 - causing a card of the user U to receive SIGUDd during time interval d for presentation to door D in order to pass through D;
 - having U present the card with SIGUDd to door D; and
- 15 having D open after verifying that (i) SIGUDd indeed is a digital signature of A indicating that U can access door D at time interval d, and (ii) that the current time is indeed within time interval d, wherein the door D stores, on the card of the user, information corresponding to an access attempt of a different user.

- 20 Some embodiments relate to a method for controlling the access to at least one disconnected door, comprising:

for each time interval of a sequence of dates, causing an entity to produce a digital signature, wherein the digital signature indicates that at least one user can access the disconnected door during the time interval;

- 25 causing a card of a first user to receive the digital signature during the time interval for presentation to the disconnected door in order to pass through the disconnected door; and

after the first user presents the card with the digital signature to the disconnected door, causing the disconnected door to open after verifying that: (i) the digital signature is a digital signature of the entity indicating that the first user can access the disconnected door at the time interval, and (ii) that a current time is within the time interval, wherein

30

2003228468 03 Aug 2009

- 10 -

the disconnected door stores, on the card of the first user, information that corresponds to a second user different from the first user.

5 A system and method are disclosed for controlling physical access through a digital certificate validation process that works with standard certificate formats and that enables a certifying authority (CA) to prove the validity status of each certificate C at any time interval (e.g., every day, hour, or minute) starting with C's issue date, D₁. C's time granularity may be specified within the certificate itself, unless it is the same for all certificates. For example, all certificates may have a one-day granularity with each
10 certificate expiring 365 days after issuance. Given certain initial inputs provided by the CA, a one-way hash function is utilized to compute values of a specified byte size that are included on the digital certificate and to compute other values that are kept secret and used in the validation process.

15 Controlling physical access includes reviewing real time credentials, where the real time credentials include a first part that is fixed and a second part that is modified on a periodic basis, where the second part provides a proof that the real time credentials are current, verifying validity of the real time credentials by performing an operation on the second part and comparing the result to the first part, and allowing physical access only if
20 the real time credentials are verified as valid. The first part may be digitally signed by an authority. The authority may provide the second part or the second part may be provided by an entity other than the authority. The real time credentials may be provided on a smart card. A user may obtain the second part of the real time credentials at a first location. The user may be allowed access to a second location different and separate from
25 the first location. At least a portion of the first part of the real time credentials may represent a one-way hash applied plurality of times to a portion of the second portion of the real time credentials. The plurality of times may correspond to an amount of time elapsed since the first part of the real time credentials were issued. Controlling physical access may include controlling access through a door.

2003228468 03 Aug 2009

- 10A -

Brief Description of the Drawings

Embodiments are described with reference to the several figures of the drawing, in which:

5 **FIG. 1** is a schematic illustration of how the CA sends to a Directory individual certificate revocation status information CRS_i about each of its issued, but not-yet expired certificates C₁ ... C_k, according to some embodiments;

FIG. 2 is a schematic illustration of the sequence of transactions in a trivial OCSF environment;

10 **FIG. 3** is a schematic illustration of a major "network bottleneck" in a server causing considerable congestion and delays;

FIG. 4 is a schematic illustration showing how OCSF has difficulties in servicing certificate validity requests originating from different security domains;

FIG. 5 is a schematic illustration showing the servicing of certificate validity requests originating from different security domains according to some embodiments;

15 **FIG. 6** is a schematic illustration of the RTC System according to some embodiments;

FIG. 7 is a schematic illustration showing how RTC-over-OCSF would be deployed in a cross-CA environment according to some embodiments;

20 **FIG. 8** is a schematic illustration of the system operation according to some embodiments;

FIG. 9 is a schematic illustration of a stolen computer timeline.

Detailed Description of the Preferred Embodiments**SECURE PHYSICAL ACCESS**

25 Ensuring that only authorized individuals access protected areas is crucially important (e.g., at an airport, a military installation, office building etc.). Protected areas may be defined by physical doors (in particular doors through which a human

may enter, or doors of a container, or safe, or vehicle, etc.) and walls, or may be virtually defined in other ways. For instance, a protected area may consist of an area entering which causes a detector to signal intrusion (and possibly send a signal or sound an alarm if authorization is not provided). In an airport, often entering the
5 gate area through an exit lane will trigger such a signal, even though no doors or walls have been violated. Notice also that throughout this application, doors should be construed to include all other types of access access-control devices implementable with a traditional or more modern type of a key. In particular, key mechanisms used to start engines (so that our invention becomes a novel way to
10 ensure that only currently authorized users may start a plane, a truck, or otherwise access other valuables).

Having established the generality of our context, in the sequel for concreteness, but without loss of generality intended, we shall refer to a "door" as
15 the means of controlling access or establishing the perimeter and to "entering" as the means of accessing an area which one wishes to protect.

Smart doors provide such access control. At the simplest level, a smart door may be equipped with a key pad, through which a user enters his/her PIN or
20 password. The key pad has an attached memory or elementary processor in which a list of valid PINs/passwords are stored, so that it can be checked whether the currently entered one belongs to the list. If so, the door opens, else it remains lock. Such elementary access control mechanism offers minimum security. In particular a terminated employee may no longer be authorized to go through that door; yet, if he
25 still remembers his own PIN, he would have no trouble to open such an elementary smart door. Therefore, it would be necessary to "deprogram" the PIN of terminated employees. Such a procedure, however, may be very cumbersome and costly: an airport facility may have hundreds of doors, and dispatching a special team of workers to go out and deprogram all of such doors whenever an employee leaves or
30 is terminated may be too impractical. More security is certainly needed, without incurring excessive costs and sacrificing convenience.

Of course, rather than (solely) relying on traditional keys or simple key pads, a more modern smart door may work (in alternative or in conjunction) with cards – such as smart cards and mag-strip cards-- or contactless devices. But this enhanced set of tools does not *per se* guarantee the security, convenience and low-cost of the access-control system. These crucially depend on how such tools are used in the overall security architecture.

Ideally, a smart door should *identify* the person entering and *verify* that he is currently authorized to do so. Of the two tasks, the first is perhaps easier.

10 *Identification* may be performed in a variety of ways: in particular:

1. using PINs and passwords, that can be entered at a key pad associated to the door;
2. using biometrics, that can be entered by users via special readers associated with the door;
- 15 3. using traditional signatures, provided by the user via a special pad associated to the door;
4. using a smart cards or contactless cards (e.g., sending a PIN to the door via a special reader/receiver)
5. using a digital certificate -- e.g., one stored in a smart card, contactless card or a wireless device, that can “communicate to the door” via a card reader or other receiver.
- 20

We believe that digital certificates are particularly attractive for use within the inventive system, and thus we wish to elaborate a little further on some ways to use them with smart doors which we envision incorporating within the inventive system. For concreteness, but without loss of generality intended, we will refer to the device in possession of a person wishing access as a “card.” The card may store a digital certificate and the corresponding secret key(s). Upon proper command from the cardholder (performed, for example, by punching a secret code on a key pad on the card), the card would transmit the digital certificate to the door mechanism and perform an identification protocol (e.g., decrypt a random challenge) by using the corresponding secret key. Preferably, the digital certificate, and

particularly its corresponding secret key(s), should be protected within a secure-hardware portion of the card/device.

In some cases, one wishes to have anonymous yet secure access control. In
5 this case, identification needs not be performed, but authorization still needs to be
performed. In most cases, however, identification in some form is mandated: thus
we can assume that identification can or has already been performed (e.g, by any
one of the 5 methods described above). Either way: *how can authorization be*
performed? Even if the door knows for certain that it is dealing with John Doe, how
10 can the door make sure that John Doe is currently authorized to enter now?
Traditionally, a smart door consults a database of currently (e.g., on a given
day/date) authorized users to verify that so indeed is the individual requesting
access. But *this requires that the smart door to be connected to the distant database.*
Moreover, this is not ordinary network connection: *it must be a secure network*
15 *connection.* In fact, not only one must use cryptographically protected
communication to prevent an impostor from impersonating the database to the door,
but must also prevent an enemy to cut the wire connecting the door to the database,
else once disconnected a door must choose from equally bad options: (a) always
open or (b) always remain closed. But a secure network connection easily dwarfs the
20 cost of the electromechanical component of the door lock: a top of the line
component may cost \$1,000 while the secure network connection may cost \$4,000
(more if a wire must securely connect large distance, such at an airport. Moreover,
even after spending such \$4,000, is there such a thing as a secure network
connection in a public place such as an airport? Notice that providing a smart door
25 with a wireless connection to a distant database is not a viable alternative either.
First of all, long range wireless transmitters and receivers are expensive. Second, in
certain facilities, wireless bandwidth can be severely restricted (to avoid possible
interference with other instrumentation) or banned altogether for such uses. Third,
wireless communication can be easily jammed, so as to effectively disconnect the
30 door from the database (thus forcing it to opt for two equally bad decisions). Fourth,
if the door belongs to a container in the middle of the Atlantic, most probably it
cannot wireless talk to any database on the shore.

2003228468 03 Aug 2009

- 14 -

Some embodiments thus provide low-cost, convenient and secure *disconnected* smart doors, that is low-cost, convenient and *secure smart doors having no connection (whether wired or wireless) to any database or authority.*

5 DIGITAL SIGNATURES AND CERTIFICATES

Some embodiments rely on digital signatures, and preferably on 20-byte technology. Digital signatures (such as RSA) are used to prove that a given message M originates from a given user U. To this end U produces a pair of matching keys: a verification key PK and a signature key SK. Digital signatures are produced via SK, and
 10 verified via the matching key PK. A user U should keep his own SK *secret* (so that only U can sign on U's behalf). Digital signatures work because PK does not "betray" the matching key SK, that is, knowledge of PK does not give an enemy any practical advantage in computing SK. Therefore, a user U should make his own PK as *public* as possible (so that every one can verify U's signatures). For this reason PK is preferably
 15 called the public key. We shall denote by $SIG_U(M)$ U's digital signature of the message M. Digital signature is intended to include private-key signatures, in which case signed and verifier may share a common secret key.

Alphanumeric strings called certificates enable digital signatures by guaranteeing
 20 that a given key PK is indeed the public key of a user U. A Certifying Authority (CA) generates and issues a certificate to a user, once assured of the user's identity. Thus the certificate proves to everyone that the CA has verified the holder's identity, and possibly other attributes. (E.g., if a company acts as its own CA and issues certificates for its own employees, a certificate may prove the extent to which its holder is authorized to bind
 25 his/her employer.) Certificates expire after a specified amount of time, typically one year in the case of public CAs. In essence, a digital certificate C consists of a CA's digital signature securely binding together several quantities: SN, a serial number unique to the certificate, PK, the public key of the user, U, the user's name, D₁, the issue date, D₂, the expiration date, and additional data. In symbols, $C = SIG_{CA}(SN, PK, U, D_1, D_2, \dots)$.

2003228468 03 Aug 2009

- 15 -

A certificate may also encompass the case where PK is an encryption key. In this case U may prove his identity to a verifier V by sending V the certificate C, by having V encrypt a random challenge (String) R with key PK, and then ask U to send back the decryption. If the user responds with R, then V is ensured that he is dealing with U, because only U should know the decryption key matching PK.

Described embodiments provide a much better solution for access control. Specifically, if the card contains a digital certificate according to described embodiments, then authorization can be performed much cheaper. Instead of consulting the central database about the validity of every digital certificate, the door would simply need to obtain the 20-byte validity proof according to the present invention that verifies the current validity of the card.

Example 1:

Let now A be an authority (i.e., entity) controlling a set of smart doors and U a user to whom access to a given door should be granted for a given period of time.

Each user possesses a card (in the general sense discussed before).

Each smart door has an associated card reader (in the general sense capable of communicating or at least receiving information from a user card), coupled with an electromechanical lock in the case of a really physical (rather than virtual) door. Preferably each door also has a unique identifier (and knows its own identifier). The door has a card reader and a non-easily tamperable clock and a computing device possessing A's public key PKA and capable of verifying A's signatures.

The authority decides which users can go through which doors in a given time interval. (For instance, without loss of generality intended, we may assume that each interval of time of interest consists of a day.) To this end, A may use her own

private database DB1, storing all permissions, that is who is authorized to go through which door at a given (or any foreseeable future day). Presumably, A protects this database, else an enemy could alter the permissions stored there to his advantage. However, A computes from DB a public database PDB as follows. For each user U having permission to go through door D at day d, A computes a digital signature $SUDd$ indicating that indeed this is the case. For instance A computes $SUDd = \text{SIG}_A(U, D, d)$. Notice that only A can compute these digital signatures, while all having A's public key PKA can verify them. These signatures are unforgeable by someone not knowing A's secret key SKA, nor can they be modified in any manner (e.g., by transforming U's permission into permission for an unauthorized user U') without making them invalid. Thus A can timely compute and send (eg, at the beginning of a day) these signatures to a repository PR without much worry. A repository is a place that can be accessed by users. For instance a server located at the employee entrance of a large facility (such as an employee entrance at an airport). Because A's signatures are unforgeable, the connection between A and PR needs not be secure. It suffices that A succeeds to transfer its signatures to PR within a reasonable time.

When employee U arrives at work on day d at the facility (eg, through a point of entrance in which PR is located) he can connect his card with PR (eg, he inserts his card in a card reader/writer connected with or remotely communicating with PR). By doing this he picks up on his card $SIGUDd$, the digital signature indicating that that day he is authorized to go through door D. This requires that the point of entrance, rather than hundreds of doors, be connected with A, and this connection needs not be secure either. In reality, D needs not to indicate a single door. For instance, it can indicate a set of doors (eg, baggage handling doors) and the signature of A indicates that U can go through each door indicated by D. Alternatively, a plurality of doors, $D_1, \dots, D_n$, can be indicated one by one, and the fact that U can go that day through each one of them can be indicated by more than one signature of A. For example $SIGUD_1d \dots SIGUD_nd$. In which case, all such signatures are transferred to U's card.

Assume now that during day d U walks around the facility and reaches a door D for which he has granted permission. Therefore, his card now stores $SIGUD_d$. Then U may insert his card C into a card reader at door D . The processor associated with the door then verifies that the $SIGUD_d$ indeed is valid using A 's public key. Then verifies that the current day is indeed d using its own clock. If both items are true, then door D opens. Notice that the door can check that the cardholder is indeed by performing identification in a variety of ways. In particular, U may also be required to enter his PIN on a key pad associated with the door. (Notice that, differently than before, a dismissed employee cannot enter door D even if he remembered his own PIN. In fact the door in this example would need both the PIN and the correct signature for the current day. However, after U has been fired, A no longer produces signatures $SIGUD_d$ for any subsequent day d , therefore U cannot provide the door with such a signature. Nor can he forge such a signature of A . Therefore he cannot "convince" door D to open on any day after being fired.)

Alternatively, the card can transfer $SIGUD_d$ to D 's card reader only if U inputs the right PIN on a key pad on the back of C , and the repository PR may download $SIGUD_d$ onto card C , only after the card proves that indeed it is U 's card. Alternatively, U may represent an identifier for card C , belonging to U , and when inserted in the card reader, the card indeed proves –eg, by means of a cryptographic protocol, that indeed it is card C . Alternatively, and preferably, U 's card carries a certificate for U , and after the proper PIN is entered, the card proves the identity of U by decrypting a random challenge of the door. In this case, it is preferable that $SIGUD_d$ indicates that U has permission to go through door D by indicating that U 's certificate carries that permission for his owner. For instance, $SIGUD_d = SIGu_d$, where u is an identifier for U 's certificate, such as the serial number (and issuer) of U 's certificate.

In all these ways, it should be appreciated that the door is "disconnected" from A . The door only (possibly identifies U and) checks that the U has permission of entering via an internal computation and utilizing A 's public key and its own internal clock. The system therefore, not only is very secure, but also very economical.

This validity or authorization proof can be provided in a number of different ways. The following are just examples of how this can be done.

5 Example 2:

 The card owner may "pick up" the validity proof at the appropriate time. For example, in a work environment, each person may pick up the current validity proof when reporting to work. In many work places (particularly those sensitive to security, such as airports), employees sign in when reporting to work. This "sign in" may include obtaining the 20-byte validity, SIGUDd, and storing it on the card. The card may obtain the value via a wired or a wireless connection.

 Example 3:

15 The card may obtain the validity proof via a wireless network, such the pager network, for example. At the appropriate time, if the card is authorized for access, a 20-byte value is sent to it. Note that the bandwidth requirements are minimal: the authorization value is shorter than a typical message transmitted by the pager network. At the appropriate time, if the card is authorized for access, SIGUDd is sent to it.

 Example 4:

 The door may obtain the validity proof similarly via a wired or a wireless network, for every card that it expects to encounter, in advance.

25 Example 5:

 The door may obtain the validity proof for a card on demand, when the card starts interacting with it.

30 Note that none of the above methods require any sort of secure connection between the door and a central server. This is so because the validity proof is self-authenticating, so that even if the door receives it from an untrusted source and/or

via an insecure connection, it can still ascertain its correctness. The fact that these methods require no connection at all for the door provides a much better means for access control in large and/or remote areas, areas with multiple doors and mobile areas, such as airplanes' or trucks' doors.

5

Note also that throughout this application, door and protected areas should be construed to include all other types of access points that could be protected with a traditional or more modern type of key. In particular, key mechanism that used to start engines (so that only currently authorized employees may start a plane, a truck, or other engine).

10

Those skilled in the art can realize that the 20-byte validity proof is a special, restricted type of a digital signature scheme, and while it offers unique benefits, such as compactness and efficiency, many other benefits can be derived by practicing the invention with more general digital signature schemes, possibly without validation technology. The components of the preferred embodiment of the present invention are: (1) A door mechanism capable of verifying digital signatures, coupled with means of opening the door upon successful verification; (2) An authority component, providing a digital signature signifying that authorization for entering through the door has been granted for a given time period; (3) A card or other wired/wireless device component capable of receiving a digital signature and presenting it.

15

20

The authorization of access may be accomplished by any of the following sequences of steps:

25

SEQUENCE 1:

- (1) The authority component causes the card to receive the authorizing signature;
- (2) The card receives and stores the authorizing signature;
- (3) The card presents the authorizing signature to the door, which verifies it and opens if and only if the authorizing signature is valid

30

SEQUENCE 2:

- (1) The card presents itself to the door requesting authorization for access;
- (2) The door requests the authorizing signature;
- 5 (3) The authority component causes the door to receive the authorizing signature;
- (4) The door verifies the authorizing signature and opens if and only if it is valid.

10 SEQUENCE 3:

- (1) The card requests the authorizing signature from the authority component
- (2) The authority component transmits the authorizing signature to the card;
- 15 (3) The card receives and stores the authorizing signature;
- (4) The card presents the authorizing signature to the door, which verifies it and opens if and only if the authorizing signature is valid.

SEQUENCE 4:

- 20 (1) The door receives in advance (either at its own request or not) authorizing signatures for a plurality of cards it is expected to encounter from the authority component;
- (2) The card presents itself to the door requesting authorization for access;
- 25 (3) The door verifies the card's authorizing signature and opens if and only if it is valid.

These sequences are only some of the multitude of examples. In addition, these sequences may be combined. For example, the door may receive part of the
30 information/authorization (e.g., the 20-byte value), while the card may receive another part (e.g., the digital certificate). They may also be separated in time: the

card may receive part of the information/authorization (e.g., the digital certificate) at first, and then receive other parts (e.g., the 20-byte value for each hour) later.

Moreover, the authorizing digital signatures may be tied to the long-term certificate of the cardholder. For example, the card may contain a long-term certificate valid for each year, and the authority component may issue daily signatures verifying that the certificate is still valid on the current day.

The authority component may generate authorizations automatically, without any requests. For example, every night the authority component may generate authorizing signatures for the employees that will be authorized for the next day. This approach enables the authorization component to be non-interactive and thus easier to build securely.

In addition, the authority component may use separate, possibly insecure devices, for dissemination of authorizing signatures to cards and/or doors. This will enable the authorization component to focus on only one task: generation of authorizations. It will remove the need for the cumbersome direct connections between the secure authorization component and the (possibly less secure) doors and cards. Specifically, dissemination of authorizations may occur as follows: (1) The authority component generates authorizations; (2) The authority component transmits authorization, over possibly insecure connections, to dissemination databases. These databases may be at multiple locations, and need not be secured. For example, in a company with five employee entrances, there may be one dissemination database at each entrance; (3) The dissemination databases transmit authorizations to cards and/or doors, either upon request ("pull") or automatically ("push").

The property enabling the above methods is that the authorization itself is unforgeable – it can be produced *only* by the authority component. Therefore, once produced, it can be disseminated over possibly untrusted lines and devices without any risks to security. This removes the need for any other party or device to interact

with the authority component, thus resulting in a much cheaper solution than any requiring secure connections.

5 In fact, no connections among any of the components in this system need to be secured. (Only the authority component itself has to be secured, so that inappropriate authorizations are not produced.) Thus, a fault-tolerant, distributed access authorization infrastructure can be much more easily built. Moreover, as stated above, it is possible to build such an infrastructure without any connections needed for the doors.

10

It should be appreciated that the inventive access control system can be combined with the tenant CAs of Section 3. For instance, several authorities (e.g., in an office building, the parking authority, the cleaning authority, or the multiple companies sharing occupancy in the building) may utilize the same certificate while retaining individual control over the ability of its holder to access the various protected areas.

Example 6:

The system could operate as follows. A user U (or his card) has a certificate
20 CERT that contains a validation field ---say D365--- for each door D of interest. Permission that U can go through door D at day j can be proved by releasing the unforgettable 20-byte value X365-j. Door D can check that permission by hashing it j times and checking whether the result coincides with the validity field D365 of CERT. In case A must deal with a plurality of doors (eg, 1000 doors, then CERT
25 may contain 1000 different validity fields, each corresponding to different doors, and each door Dj checks its computations relative to the jth validity field. In this case, even if permission for a user to go through each door is proved separately, each user has at most 1000 proofs on a given day. Thus at most 20K bytes need to be loaded on his card on a given day.

30

Notice that because cards are general cards here, the card can be a contactless card, the card reader may be a receiver, and the card need not be

inserted into the reader but transmit to it. Notice that such a “wireless” card-reader interaction is still quite local, and very different from a card-authority/database interaction when A or the database is far away.

5 Moreover, the authorizing digital signatures may be tied to the long-term certificate of the cardholder. For example, the card may contain a long-term certificate valid for each year, and the authority component may issue daily signatures verifying that the certificate is still valid on the current day.

10 The authority component may generate authorizations automatically, without any requests. For example, every night the authority component may generate authorizing signatures for the employees that will be authorized for the next day. This approach enables the authorization component to be non-interactive and thus easier to build securely.

15 In fact, no connections among any of the components in this system need to be secured. (Only the authority component itself has to be secured, so that inappropriate authorizations are not produced.) Thus, a fault-tolerant, distributed access authorization infrastructure can be much more easily built. Moreover, as
20 stated above, it is possible to build such an infrastructure without any connections needed for the doors.

 It should be appreciated that the inventive access control system can be combined with the tenant CAs as described. For instance, several authorities (e.g., in
25 an office building, the parking authority, the cleaning authority, or the multiple companies sharing occupancy in the building) may utilize the same certificate while retaining individual control over the ability of its holder to access the various protected areas.

30 Logging Proofs of Access With Disconnected Doors

 While disconnected (from authorities and databases) and yet very secure, low-cost and convenient smart doors are preferable to connected smart doors, the

latter provide for the ability of logging access through a given door. For instance, it can be important to know who went through a given door on a given day. Connected doors may easily do this by sending proper "access information to a distant database or authority. But disconnected doors cannot quite do that. Access information can be
 5 gather by sending proper personal to collect such information from door to door. This may not be always convenient to do. However, the following system provides a very viable alternative.

When a user U passes (or attempts to pass) through a door D at a time t, the
 10 door may produce a proper string LOGUDt, and locally store it (at least temporarily). To ensure that this info reaches a proper database, the door may use the cards used to enter through it. For instance, D may write LOGUDt (or cause LOGUDt to be written) on the card(s) of other user(s) U' (possibly including U himself). Whenever U' makes a connection with PR (eg the next day of work) or
 15 with any other wired or well connected device, then PR or said device transmits LOGUDt to the proper database. This way a proper database will eventually receive and then store more permanently and in an easily auditable way LOGUDt. Possibly the database will receive redundant copies of LOGUDt, but it is easy for it to clear any unwanted redundancy and keep clean copies only.

20

A preferable LOGUDt, is one that consists or includes a digital signature of U himself. This way, U cannot easily deny that he went through a given door at a given time and claim that the access information of the door is a fabrication. Indeed, only he has the secret signing key for producing LOGUDt. For instance LOGUDt e
 25 consist of $SIG_U(D,t)$, indicating that U went through door D at time t. This is very easy to accomplish if user U's card carries the secret signing key SKU matching a public key PKU. Preferably the card also carries a digital certificate for PKU, and thus LOGUD may include not only $SIG_U(D,t)$, but also U's certificate. Preferably too, the user card may produce $SIG_U(D,t)$ according to the time t shown on its own
 30 clock, and the door may let U in only after he provides such a good access proof $SIG_U(D,t)$ (possibly in addition to other authorization proofs such as those discussed above), and provided that the time certified by U is sufficiently close to the current

time t' as measured by the door clock. Still the user may claim that he entered at time t door D , but that this door was somewhere else altogether, and thus that $SIG_U(D,t)$ does not at all prove that he went through ---say---the second door of the third floor of a given building: someone went through the trouble to transfer to said
5 location the door reader etc. To prevent this claim too, or to protect the user against such fraud, the user card (device) may incorporate a GPS mechanism, and $SIG_U(D,t)$ may actually include the local position lp as measured by the card. In which case, the user may tend to the door the proof of access $SIG_U(D,t, ps)$, and the door may accept it and let the user in only if not only the time looks correct but also the local
10 position. Rather than computing ps inside the card/device, the user may use some one or more components, which he trusts, and which can compute his position from information they receive from him (and possibly their own positions).

IMPLEMENTATION

THE BASIC SYSTEM

As seen in the FIG. 1, the CA sends to a Directory individual certificate revocation status information CRS_i about each of its issued, but not-yet expired
 5 certificates $C_1 \dots C_k$. The Directory sends CRS_i to a requesting user who has inquired about certificate serial number "i" of that certifying authority.

A system and method are disclosed for controlling physical access through a digital certificate validation process that works with standard certificate formats
 10 (e.g., X.509v3) and that enables a certifying authority (CA) to prove the validity status of each certificate C at any time interval (e.g., every day, hour, or minute) starting with C's issue date, D_1 . C's time granularity may be specified within the certificate itself, unless it is the same for all certificates. To be concrete, but without limitation intended, below we assume a one-day granularity for all certificates, and
 15 that each certificate expires 365 days after issuance.

Making a certificate C. In addition to traditional quantities such as a serial number SN, a public key PK, a user name U, an issue date D_1 , an expiration date D_2 ($=D_1+365$), a certificate C also includes two 20-byte values unique to it.
 20 Specifically, before issuing a certificate C, a CA randomly selects two different 20-byte values, Y_0 and X_0 , and from them computes two corresponding 20-byte values, Y_1 and X_{365} , using a one-way hash function H enjoying the following properties: H is at least 10,000 times faster to compute than a digital signature; H produces 20-byte outputs, no matter how long its inputs; and H is hard to invert: given Y, finding
 25 X such that $H(X)=Y$ is practically impossible. (See, for example, Secure Hash Standard; FIPS PUB 180, Revised July 11, 94 (Federal Register, Vol. 59, No. 131, pp. 35211-34460); revised August 5, 1994 (Federal Register Vol. 59, No. 150, pp. 39937-40204). Value Y_1 is computed by hashing Y_0 once: $Y_1=H(Y_0)$; and X_{365} by hashing X_0 365 times: $X_1=H(X_0)$, $X_2=H(X_1)$, ..., $X_{365}=H(X_{364})$. Because H always
 30 produces 20-byte outputs, Y_1 , X_{365} , and all intermediate values X_j are 20-byte long. The values $Y_0, X_0, X_1, \dots, X_{364}$ are kept secret, while Y_1 and X_{365} are included in the

certificate: $C = \text{SIG}_{\text{CA}}(\text{SN}, \text{PK}, \text{U}, \text{D}_1, \text{D}_2, \dots, \text{Y}_1, \text{X}_{365})$. We shall call Y_1 the *revocation target* and X_{365} the *validity target*.

- Revoking and validating a not-yet-expired certificate C.* On the i -th day after
- 5 C's issuance (i.e., on day $\text{D}_1 + i$), the CA computes and releases a 20-byte proof of status for C as follows. If C is revoked, then, as a proof of C's revocation, the CA releases Y_0 , that is, the H-inverse of the revocation target Y_1 . Else, as a proof of C's validity on that day, the CA releases X_{365-i} , that is, the i -th H-inverse of the validity target X_{365} . (E.g., the proof that C is valid 100 days after issuance consists of X_{265} .)
- 10 The CA may release Y_0 or X_{365-i} by providing the value in response to a query or by posting it on the World Wide Web.

- Verifying the status of a not-yet-expired certificate C.* On any day, C's revocation proof, Y_0 , is verified by hashing Y_0 once and checking that the result
- 15 equals C's revocation target, Y_1 . (I.e., the verifier tests for himself that Y_0 really is the H-inverse of Y_1 .) Note that Y_1 is guaranteed to be C's revocation target, because Y_1 is certified within C. On the i -th day after C's issuance, C's validity proof on that day, X_{365-i} , is verified by hashing i times the value X_{365-i} and checking that the result equals C's validity target, X_{365} . (I.e., the verifier tests for himself that X_{365-i} really is
- 20 the i -th H-inverse of X_{365} .) Note that a verifier knows the current day D as well as C's issuance date D_1 (because D_1 is certified within C), and thus immediately computes $i = \text{D} - \text{D}_1$.

Security

- 25 *A proof of revocation cannot be forged.* The proof of revocation of a certificate C consists of the H-inverse of C's revocation target Y_1 . Because H is essentially impossible to invert, once a verifier checks that a given 20-byte value Y_0 is indeed C's proof of revocation, it knows that Y_0 must have been released by the CA. In fact, only the CA can compute the H-inverse of Y_1 : not because the CA can
- 30 invert H better than anyone else, but because it computed Y_1 by starting with Y_0 and hashing it! Because the CA never releases C's revocation proof as long as C remains valid, an enemy cannot fake a revocation proof.

A proof of validity cannot be forged. On day i , the proof of validity of a certificate C consists of the i -th H-inverse of C 's validity target X_{365} . Because H is essentially impossible to invert, once a verifier checks that a given 20-byte value X_{365-i} is indeed C 's proof of validity on day i , it knows that the CA must have released X_{365-i} . In fact, only the CA can compute the i -th H-inverse of X_{365} : not because the CA can invert H better than anyone else, but because it computed X_{365} by starting with X_0 and hashing it 365 times, thus computing along the way all the first 365 inverses of X_{365} ! If certificate C become revoked on day $i+1$, the CA has already released the values $X_{365-1}, \dots, X_{365-i}$ in the preceding i days (when C was still valid) but has not released and will never release the value $X_{365-i+1}$ (or any other value X_j for $j < 365-i$) in the future. Consequently, to forge C 's validity proof on day $i+1$, an enemy should compute on his own the $i+1$ st H-inverse of X_{365} (i.e., the H-inverse of X_{365-i}), which is very hard to do! Similarly, an enemy cannot compute a validity proof for C on any day after $i+1$. To do so, it should again be able to invert H on input X_{365-i} . For instance, if it could compute C 's validity proof on day $i+2$, $X_{365-i-2}$, then by hashing it once it would easily obtain $X_{365-i-1}$, the H-inverse of X_{365-i} .

Efficiency

A certificate C includes only two additional 20-byte values, Y_1 and X_{365} . This is a negligible cost. Recall that C already consists of a CA signature (at least 2048-bit long) of data that includes a public key PK (at least 1024-bit long), and that C may include comments and plenty of other data in addition to SN , PK , U , $D1$ and $D2$.

Generating Y_1 and X_{365} requires only 366 hashings total. This too is a negligible cost. Recall that issuing a certificate already requires computing a signature.

Proofs of revocation and proofs of validity are only 20-bytes long. Our 20-byte proofs are trivial to transmit and trivial to store, making the 20-byte technology

ideal for wireless applications (because here bandwidth is still limited, and so is the storage capacity of many cellular phones and other wireless devices).

Proofs according to embodiments of the present invention can be so short
 5 because they derive their security from elementary cryptographic components, such as one-way functions, which should exhibit an exponential amount of security. (Quite differently, digital signature schemes have complex security requirements. Their typical number-theoretic implementations offer at best a sub-exponential amount of security, and thus necessitate much longer keys.)

10

The proofs remain 20-bytes long whether the total number of certificates is a few hundred or a few billion. In fact there are 2^{160} possible 20-byte strings, and the probability that two certificates may happen to have a common proof of revocation or validity is negligible.

15

Note too that the length of our 20-byte proofs does not increase due to encryption or authentication. Our 20-byte proofs are intended to be public and thus need not be encrypted. Similarly, our 20-byte proofs are self-authenticating: by hashing them the proper number of times they yield either the validity target or the
 20 revocation target specified within the certificate. They will not work if faked or altered, and thus need not be signed or authenticated in any manner.

Finally, a 20-byte proof of validity on day i , X_{365-i} , need not additionally include the value i : in a sense, it already includes its own time stamp! Indeed, as
 25 discussed before, i is the difference between the current day and the certificate's issue day, and if hashing X_{365-i} i times yields the validity target of certificate C , then this proves that X_{365-i} is C 's proof of validity on day i .

The 20-byte proofs are computed instantly. A proof of revocation Y_0 or a
 30 proof of validity X_{365-i} is just retrieved from memory. (Alternatively, each X_{365-i} could be recomputed on the fly on day i ; for instance by at most 364 hashings, if just

2003228468 03 Aug 2009

- 30 -

X_0 is stored during certificate issuance. Surprisingly more efficient strategies are discussed in the next section.)

Wireless Environment

5 Embodiments are ideal for wireless implementations. Its scalability is enormous: it could accommodate billions of certs with great ease. The bandwidth it requires is negligible, essentially a 30-bit serial number for the query and 20-byte for the response. The computation it requires is negligible, because a certificate-status query is answered by a single table look-up and is immediately verified. Of course, great scalability, minimum
10 bandwidth and trivial computation make the present invention a technology of choice in a wireless environment.

But there is another use that provides an additional advantage in wireless applications. Namely, every morning --e.g., at midnight-- a wireless user may receive a
15 20-byte proof of the validity of his certificate for the remainder of the day. (This 20-byte value can be obtained upon request of the user, or pushed to the user's cellular device automatically --e.g., by means of a SMS message or other control message.) Due to its trivial length, this proof can be easily stored in most cellular telephones and PDAs. Then, whenever the user wants to transact on that day, the user simply sends its own certificate
20 *together with* the cert's 20-byte proof of validity for that day. Because the proof of validity is universally verifiable, the verifier of the cert and proof need not call any CA or any responder. The verifier can work totally off-line. In the cellular environment, in which any call translates into money and time costs, this off-line capability is of great value.

25

COMPARING TO OCSP

Described embodiments and OCSP are both on-demand systems: namely, a user sends a query about the current validity of a certificate and gets back an unforgeable and universally verifiable proof as a response. But there are differences in: Time accuracy;
30 Bandwidth; CA efficiency; Security; and Operating costs.

2003228468 03 Aug 2009

- 31 -

Time Accuracy: In principle, an OCSP response may specify time with unbounded accuracy, while a response according to described embodiments specifies time with a predetermined accuracy: one day, one hour, one minute, etc. In low-value applications, one-day validity is plenty acceptable. For most financial applications, Digital Signature Trust considers a 4-hour accuracy sufficient. (Perhaps this is less surprising than it seems: for most financial transactions, orders received in the morning are executed in the afternoon and orders received in the afternoon are executed the next business day.) In any event, time is not specified by a real number with infinitely many digits. In an on-demand validation system, a time accuracy of less than one minute is seldom meaningful, because the clocks of the querying and answering parties may not be that synchronized. Indeed, in such a system, a time accuracy of 15 seconds is *de facto* real time.

To handle such an extreme accuracy, some embodiments compute hash chains that are roughly 1M long (i.e., needs to compute validity fields of the type X_{1M}), because there are at most 527,040 minutes in a year. If chains so long could be handled efficiently, such embodiments would *de facto* be real time. Computing 1M hashings is not problematic at certificate issuance: 1M hashings can be performed in less than 1 second even using very reasonable platforms, and a certificate is typically issued only once a year, and not under tremendous time pressure. Similarly, 1 second of computation is not problematic for the verifier of a cert validity proof (e.g., a merchant relying on the certificate) considering that he generally focuses just on an individual transaction, and has more time at hand. Computing 1M hashings per certificate-status request would, however, affect the performance of the server producing validity proofs, because it typically handles many transactions at a time. Fortunately, this server needs not to compute all these hashings on-line starting with X_0 , but by table look up - capitalizing on having in storage the full hash-chain of every certificate. Nonetheless, storing 1M-long hash-chains may be a problem in applications with huge numbers of certificates. But, fortunately, as we shall mention

later on, even ordinary servers can, using better algorithms, re-compute 1M-long hash chains with surprising efficiency.

Bandwidth: The preferred embodiment of the present invention has an
5 obvious bandwidth advantage over OCSP. The former uses 20-byte answers, while the latter typically uses 256 bytes.

CA Efficiency: A validity query is answered by a (complex) digital signature
in the OCSP case, and by a (trivial) table look-up in the case of the present
10 invention, as long as the CA stores the entire X-chain for each certificate.

Note that, with a population of 1 million certificates, the CA can afford to store the entire X-chain for each certificate when the time accuracy is one day or one hour. (In the first case, the CA would have to store 365 20-bytes values; that is, 7.3K
15 bytes per cert, and thus 7.3B bytes overall. In the second case, 175.2B bytes overall.) If the time accuracy were 15 seconds, then each hash chain would consist of 1M 20-byte values, and for the entire system the overall storage requirement would be around 10.5 tera-bytes: a sizable storage.

20 To dramatically decrease this storage requirement, the CA may store just a single 20-byte value (i.e., X_0) for each cert, and re-compute from it each X_i value by at most 1M hashings. Alternatively, Jacobsson [5] has found a surprising time/storage tradeoff. Namely, the CA may re-compute all n X_i values, in the right order, by storing $\log(n)$ hash values and performing $\log(n)$ hashings each time. If n
25 were 1M, this implies just storing 20 hash values per cert and performing only 20 hashings each time the cert needs validation. Other non-trivial tradeoffs are possible. In particular, for our 1M-chain case, Reyzin [R] has shown that a CA can compute all X_i values ($i=1M$ down to 1) by storing only 3 hash values and performing at most 100 hashings each time.

30

In sum, even in a de facto real-time application (i.e., using a 15-second time accuracy) the preferred embodiment of the present invention can, by just storing 60

bytes per certificate, replace a complex digital signature operation with a trivial 100-hash operation.

Security and Operating Costs: The last two differences are better discussed
 5 after specifying the type of implementation of the preferred embodiment of the present invention and OCSP under consideration.

Centralized Implementation: Security Analysis

Whenever proving certificate validity relies on the secrecy of a given key, a
 10 secure *vault* ought to protect that key, so as to guarantee the integrity of the entire system. By a centralized implementation of the present invention or OCSP, we mean one in which a single vault answers all validity queries. Centralized implementations are preferable if the number of deployed certificates is small (e.g., no more than 100K), so that the vault could handle the query volumes generated even if almost all
 15 certificates are used in a small time interval, triggering almost simultaneous validity queries. In such implementations, the preferred embodiment is preferable to OCSP in the following respects.

Doomsday Protection: In the traditional OCSP, if (despite vaults and
 20 armored guards) an enemy succeeds in penetrating the vault and compromises the secret signing key, then he can both "resurrect" a previously revoked certificate and "revoke" a still valid one. (Similarly, if the CRL signing key is compromised in a CRL system.) By contrast, in the preferred embodiment of the present invention, penetrating the secure vault *does not help an adversary to forge the validity of any*
 25 *previously revoked certificate.* In fact, when a certificate becomes revoked at day i , not only is its revocation proof Y_0 made public, but, simultaneously, all its X_i values (or at least the values X_0 through X_{365-i}) are deleted. Therefore, after a successful compromise, an enemy finds nothing that enables him to "extend the validity" of a revoked certificate. To do so, he should succeed in inverting the one-way hash H on
 30 X_{365-i} without any help, which he is welcome to try (and can indeed try without entering any secure vault). The worst an enemy can do in a system according to the present invention after a successful compromise is to fake the revocation of valid

2003228468 03 Aug 2009

- 34 -

certificates, thus preventing honest users from authenticating legitimate transactions. Of course, this would be bad, but *not as bad as enabling dishonest users to authenticate illegitimate transactions.*

5 Distributed Implementation: Security and Operating-Cost Analysis

Centralized implementations require all queries about certificate validity to be routed to the same vault. This easily results in long delays and denial of service in applications with millions of active certificates. To protect against such congestion, delays, and denial of service, one might spread the load of answering validity queries across several, geographically dispersed, responder servers. However, in the case of the OCSP each additional responder needs to have a secret signing key, and thus needs to be hosted in a vault, making the cost of ownership of an OCSP system very onerous. A high-grade vault meeting the requirements of financial institutions costs at least \$1M to build and \$1M to run. (A good vault would involve armored concrete, steel doors, back-up power generators, protected fuel depot to run the generator for potentially a long time, etc. Operating it would involve a minimum of 4 different teams for 24X7X365 operations, plus managerial supervision, etc.) In an application requiring 10 such vaults to guarantee reasonably fast response at peak traffic, the cost of ownership of the OCSP system would be \$10M of initial investment and an ongoing budget of \$10M/year. Even if less secure vaults and operations were used, millions of dollars in initial and ongoing costs would still be necessary.

In the case of some embodiments, however, a distributed implementation can be achieved with a single vault (which a CA would have anyway) and an arbitrary number of "un-trusted responders" (i.e., ordinary servers). Let us see the exact details of a distributed system according to described embodiments assuming, to be concrete, that (a) there are 10M certs; (b) there are 1,000 servers, strategically located around the globe so as to minimize response time; and (3) the time granularity is one-day.

CA Operations (Initialization Cost): Every morning, starting with the smallest serial number, compile a 10M-entry array F as follows: For each certificate C having serial number j, store C's 20-byte validity/revocation proof in location j. Then, date and sign F and send it to each of the 1,000 servers.

5

User Operations (Query Cost): To learn the status of a certificate C, send C's serial number, j, (and CA ID if necessary) to a server S.

Server Operations (Answer Cost) : Every morning, if a properly dated and signed array F is received, replace the old array with the new one.

10

At any time: answer a query about serial number j by returning the 20-byte value in location j of the current F.

15 *Operations of the Preferred Embodiment*

1. *Preparing Array F is instantaneous.*

If the whole hash chain is stored for each cert, then each entry is computed by a mere table look-up operation. In an alternative embodiment, it can also be computed on the spot.

20

2. *F contains no secrets.*

It consists of the accurate and full account of which certificates are still valid and which revoked. (The CA's goal is indeed making this non-secret information as public as possible in the most efficient manner)

25

3. *Transferring F to the servers is straightforward.*

This is so because F contains no secrets, requires no encryption, and poses no security risks. Though 10M certs are a lot, sending a 200M-byte file to 1000 servers at regular intervals is very doable.

30

4. *Each server answer is 20-byte long.*

Again, each answer requires no encryption, signature or time stamp.

2003228468 03 Aug 2009

- 36 -

5. *No honest denial of service*

Because each value sent is just 20-byte long, because each such a value is immediately computed (by a table look up), and because the traffic can be spread across 1000 servers, no denial of service should occur, at least during legitimate use of the system.

6. *Servers need not be trusted.*

They only forward 20-byte proofs received by the CA. Being self-authenticating, these proofs cannot be altered and still hash to the relevant targets.

10

Distributed implementations of described embodiments continue to enjoy the same doomsday protection of its centralized counterpart: namely, an enemy successfully entering the vault cannot revive a revoked certificate. Sophisticated adversaries, however, refrain from drilling holes in a vault, and prefer software attacks whenever possible. Fortunately, software attacks, though possible against the distributed/centralized OCSP, cannot be mounted against distributed implementations.

In the OCSP, in fact, the CA is required to receive outside queries from untrusted parties, and to answer them by a digital signature, and thus by means of its precious secret key. Therefore, the possibility exists that OCSP's required "window on the outside world" may be maliciously exploited for exposing the secret signing key.

By contrast, in distributed implementations of described embodiments, there are no such "windows:" the CA is in the vault and never receives or answers any queries from the outside; it only outputs non-secret data at periodic intervals. Indeed, every day (or hour) it outputs a file F consisting of public information. (The CA may receive revocations requests from its RAs, but these come from fewer trusted entities via authenticated channels ---e.g., using secure smart cards.) The untrusted responders do receive queries from untrusted parties, but they answer those queries by means of their file F, and thus by public data. Therefore, in a software attack

against the preferred embodiment of the present invention ordinary responders may only “expose” public information.

SIMPLIFIED PKI MANAGEMENT

5 PKI management is not trivial. (See, for example, *Internet Public Key Infrastructure, Part III: Certificate Management Protocols*; by S. Farrell, A. Adams, and W. Ford; Internet Draft, 1996; *Privacy Enhancement for Internet Electronic Mail – Part II: Certificate-Based Key Management*; by S. Kent and J. Linn; 1989). The preferred embodiment of the present invention may improve PKI management
10 in many applications by: (1) reducing the number of issued certs; (2) enabling privilege management on the cert; and (3) sharing the registration function with multiple independent CAs.

Let us informally explain these improvements in PKI management in a series
15 of specific examples. (Note that features and techniques used in one example can be easily embedded in another. We do not explicitly do this to avoid discussing an endless number of possible variations.)

Turning a Certificate ON/OFF (and Suspending It)

20 Example 7: MUSIC DOWNLOADING

Assume an Internet music vendor wishes to let users download any songs they want, from any of its 1000 servers, for a \$1/day fee. This can be effectively accomplished with digital certificates. However, in this example, U may be quite sure that he will download music a few days of the year, yet he cannot predict *which*
25 or *how many* these days will be. Thus the Music Center will need to issue for U a different one-day certificate whenever U so requests: U requests such a certificate and, after payment or promise of payment, he receives it and then uses with any of the 1000 music servers on that day. Issuing a one-day cert, however, has non-trivial management costs both for the vendor and the user. And these costs must be
30 duplicated each time the user wishes to enjoy another “music day.”

In a preferred embodiment, the present invention can alleviate these costs as follows. The first time that U contacts the vendor, he may be issued a certificate C with issue date $D_1=0$, expiration date $D_2=365$, and a validity field X_{365} , a revocation target Y_1 , and a suspension field Z_{365} . (The vendor's CA builds the suspension field very much as a validity field: by starting with a random 20-byte value Z_0 and then hashing it 365 times, in case of one-day granularity. It then stores the entire hash chain, or just Z_0 , or uses a proper time/storage method to be able to generate any desired Z_i .) At day $i=1, \dots, 365$, if U requests "a day of music" for that day, then the vendor simply releases the 20-byte value X_{365-i} to indicate that the certificate is valid. Else, it releases Z_{365-i} to indicate that the certificate is "suspended." Else, it releases Y_0 to indicate that the certificate is revoked. Optionally, if U and the music vendor agree to --say-- a "week of music starting at day i," then either the 20-byte values for those 7 days are released at the proper time, or the single 20-byte value $X_{365-i-7}$ is released at day i.

That is, rather than giving U a new single-day certificate whenever U wishes to download music, the vendor gives U a *single, yearly* certificate. At any time, this single certificate can be turned ON for a day, by just releasing the proper 20-byte value. Thus, for instance, the preferred embodiment of the present invention replaces issuing (and embedding in the user's browser) 10 single-day certificates by issuing a single yearly cert that, as it may happen, will be turned ON for 10 out of the 365 days of the year. The vendor could also use the method above to issue a cert that specifies a priori the number of days for which it can be turned ON (e.g., a 10-day-out-of 365 cert). Because it has a more predictable cost, such certs are more suitable for a gift.

Turning ON/OFF Many Certificates For The Same User

Example 8: SECURITY-CLEARANCE MANAGEMENT

Digital certificates work really well in guaranteeing that only proper users access certain resources. In principle, privileges could be specified on the cert itself. For instance, the State Department may have 10 different security-clearance levels,

L1,...L10, and signify that it has granted security level 5 to a user U by issuing a certificate C like:

$$C = \text{SIG}_{SD}(\text{SN}, \text{PK}, \text{U}, \text{L5}, \text{D}_1, \text{D}_2, \dots)$$

where again D_1 and D_2 represent the issue and expiration dates.

5

However, specifying privileges on the cert itself may cause a certificate-management nightmare: whenever its privileges change, the cert needs to be revoked. Indeed, the security level of an employee may vary with his/her assignment, which often changes within the same year. For instance, should U's security-clearance level be temporarily upgraded to 3, then the State Department should revoke the original C and issue a new cert C'. This task could be simplified somewhat by having U and thus C' retain the same public key (and expiration date) as before; for instance, by having:

10

$$C' = \text{SIG}_{SD}(\text{SN}', \text{PK}, \text{U}, \text{L3}, \text{D}_1', \text{D}_2', \dots).$$

15

However, U still faces the task of "inserting" the new C' into his browser in a variety of places: his desk-top PC, his lat-top, his cell phone, his PDA, etc. Now, having the CA take an action to re-issue a certificate in a slightly different form is one thing, but counting on users to take action is a totally different thing!

20

This management problem is only exacerbated if short-lived certificates (e.g. certificates expiring one day after issuance) are used. In the context of the present example, single-day certs may enable a State Department employee or user U to attend a meeting where a higher security level is needed. (If U had such a cert in a proper cellular device, smart card or even mag stripe card, he could, for instance, use it to open the door leading to the meeting that day.) The use of short-lived certificates is much broader, and has been advocated because it dispenses with the difficulty of revocation to a large extent (no point revoking a cert that will expire in 24 hours, at least in most applications). However, issuing short-lived certs so that they reside in all pertinent users' browsers still is a management cost.

30

These management costs can be alleviated with use of the preferred embodiment of the present invention as follows. Assuming that one-day time accuracy is enough, the State Department issues to a user U a certificate containing 10 validity fields and 1 revocation field: e.g.,

5 $C = \text{SIG}_{SD}(\text{SN}, \text{PK}, \text{U}, \text{D}_1, \text{D}_2, \mathbf{A}_{365}, \mathbf{B}_{365}, \mathbf{C}_{365}, \mathbf{D}_{365}, \mathbf{E}_{365}, \mathbf{F}_{365}, \mathbf{G}_{365}, \mathbf{H}_{365}, \mathbf{I}_{365}, \mathbf{J}_{365}, \mathbf{Y}_1,)$

where the first validity field, $\mathbf{A}_{365}$, corresponds to security-clearance level 1 ... and the 10th validity field, $\mathbf{J}_{365}$, corresponds to security-clearance level 10, while, as
 10 usual, $\mathbf{Y}_1$ is C's revocation field. Cert C is used as follows. If, on day n, U is in good standing (i.e., cert C is still valid), and U's security-clearance level is 5, then the State Department publicizes (e.g., sends to all its responders in a distributed NOVOMODO implementation) the 20-byte validity proof $\mathbf{E}_{365-n}$. If, on day m, U's security-clearance level becomes 2, then the State Department publicizes $\mathbf{B}_{365-m}$.
 15 And so on. As soon as C becomes invalid (e.g., because U is terminated as an employee or because U's secret key is compromised), then the State Department publicizes $\mathbf{Y}_0$ (and erases "future" A, B, C, D, E, F, G, H, I, and J values from its storage).

20 This way, cert C, though internally specifying its own privileges, needs *not* be revoked when these privileges change in a normal way, and users need *not* load new certs in their browsers. In essence, the preferred embodiment of the present invention has such minimal footprint, that a CA (rather than issuing, revoking, and re-issuing many related certs) can issue with great simplicity a single cert, having a
 25 much higher probability of not being revoked (because changes of security-clearance level do not translate into revocation). As a result, fewer certs will end up been issued or revoked in this application, resulting in simpler PKI management.

30 In sum, the preferred embodiment of the present invention replaces the complex certificate management relative to a set of dynamically changing properties or attributes by a single certificate (with minimum extra length) and a single 20-byte value for attribute.

Telecom companies may use a method similar to that of Example 2 to switch a given wireless device from one rate plan to another, or for roaming purposes.

5 Landlord CAs and Tenant CAs

A main PKI cost is associated to the RA function. Indeed, identifying a user U may require an expensive personal interview and verifying that indeed U knows the right secret key (corresponding to the to-be-certified public key PK). It would be nice if this RA function could be shared across many CAs, while enabling them to
10 retain total independent control over their own certs.

Example 9: ORGANIZATION CERTIFICATES

The Government and big organizations consist of both parallel and hierarchical sub-organizations: departments, business units, etc. An employee may
15 be affiliated with two or more sub-organizations. For instance, in the U.S. Government, he may work for NIST and the Department of Commerce. Issuing a digital certificate for each such affiliation results in a high total number of certificates and a complex PKI management: every time an employee drops/adds one of his/her affiliations, it is best to revoke the corresponding cert/issue a new one.
20 Ideally, two opposites should be reconciled: (1) The Organization issues only one cert per employee, and (2) Each Sub-Organization issues and controls a separate cert for each of its affiliates.

These two opposites can be reconciled by the preferred embodiment of the
25 present invention as follows. To begin with, notice that the preferred embodiment of the present invention is compatible with de-coupling the process of certification from that of validation, the first process being controlled by a CA and the second by a validation authority (VA). For instance, assuming a one-day time accuracy, once a CA is ready to issue a certificate C with serial number SN, it sends SN to a VA, who
30 selects Y_0 and X_0 , secretly stores the triplet (SN, Y_0, X_0) , computes as usual Y_1 and X_{365} , and then returns Y_1 and X_{365} to the CA, who includes them within C. This way, the CA need not bother validating C: the CA is solely responsible for identifying the

user and properly issuing C, while the VA is the only one who can prove C valid or revoked. This de-coupling may be exploited in a variety of ways in order to have *organization certificates* that flexibly reflect internal sub-organization dynamics. The following is just one of these ways, and uses Government and Departments as running examples. The Government as a whole will have its own CA, and so will each Department.

Envisaging k different Departments with corresponding CAs, $CA^1 \dots CA^k$, and one-day time accuracy, a Government certificate C has the following form:

10

$$C = \text{SIG}_{\text{Gov}}(\text{SN}, \text{PK}, \text{U}, \text{D}_1, \text{D}_2, \text{X}_{365}, \text{Y}_1, [\text{X}_{365}^1, \text{Z}_{365}^1], \dots, [\text{X}_{365}^k, \text{Z}_{365}^k])$$

where, as usual, SN is the cert's serial number, PK the public key of the user, U the user's identity, D_1 the issue date, D_2 the expiration date, X_{365} the validity field, Y_1 the revocation field, and where X_{365}^j is the validation field of CA^j ; and Z_{365}^j is the suspension field of CA^j .

Such a certificate is generated by the Government CA with input from the Department CAs. After identifying the user U and choosing a unique serial number SN, the issue date D_1 , and the expiration date D_2 , the Government CA sends SN, PK, U, D_1 , D_2 (preferably in authenticated form) to each of the Department CAs. The jth such CA then: chooses two secret 20-byte values X_0^j and Z_0^j ; locally stores (SN, PK, U, D_1 , D_2 , X_0^j , Z_0^j) or, more simply, (SN, X_0^j , Z_0^j); and returns $[X_{365}^j, Z_{365}^j]$ for incorporation in the Government certificate in position j (or with "label" j).

25

This certificate C is managed with Distributed implementations of the present invention as follows, so as to work as a 1-cert, a 2-cert, ..., a k-cert; that is, as k independent certs, one per Department. On day n, envisaging 100 responders: the Government CA sends all 100 responders the 20-byte value X_{365-n} if C is still valid, and Y_0 otherwise. Then, the jth Department CA sends all 100 responders the 20-byte value X_{365-n}^j to signify that C can be relied upon as a j-cert and Z_{365-n}^j otherwise.

30

Therefore, the Government CA is solely responsible for identifying the user and issuing the certificate, but each of the Department CAs can *independently* manage what de facto is *its own* certificate. (This is absolutely crucial. If CA¹ were the Justice Department and CA² the DOD, then, despite some overlapping interests, it is best that each acts independently of the other). The resulting certificate system is very economical to run. First, the number of certs is greatly reduced (*in principle*, there may be just one cert for employee). Second, a given employee can leave and join different Departments without the need of revoking old certs or issuing new ones. Third, different Department CAs may share the same responders. (In fact, whenever the mere fact that a given user is affiliated with a given Department is not a secret –something that will be true for most departments-- the servers essentially contain only “publishable information”.) Thus a query about the status of C as a j-certificate is answered with two 20-byte values: one as a Government cert and one as a j-cert. This enables one to more nimbly revoke C at a “central level” (e.g., should U lose the secret key corresponding to PK).

Example 10:

In the above example, certificate C was only revocable in a central way, but it could easily be arranged that the responsibility of revocation is push down to individual Departments. For instance, to enable the jth Department CA, in full autonomy, to revoke as well as suspend C as a j-certificate, C may take the following form:

$$C = \text{SIG}_{\text{Gov}}(\text{SN}, \text{PK}, \text{U}, D_1, D_2, [X_{N1}^1, Y_1^1, Z_{N1}^1], \dots, [X_{Nk}^k, Y_1^1, Z_{Nk}^k]).$$

Also, different departments may have different time accuracies for their own certs. This too can be easily accomplished by having C of the following format,

$$C = \text{SIG}_{\text{Gov}}(\text{SN}, \text{PK}, \text{U}, D_1, D_2, [TA^1, X_{N1}^1, Y_1^1, Z_{N1}^1], \dots, [TA^k, X_{Nk}^k, Y_1^1, Z_{Nk}^k])$$

where TA^j is the time accuracy of the j th CA; and N_j is the number of time units between D_1 and D_2 . (E.g., if TA^j is one day and $D_1 - D_2 = 1$ year, then $X_{N_j}^j = X_{365}^j$.)

Within a single organization, one major advantage of issuing certs structured
 5 and managed as above consists in enabling the cert to stay alive though the user
 moves from one sub-organization to another. It should be realized, however, that the
 above techniques are also applicable outside a single-organization domain. Indeed,
 the Government CA can be viewed as a *landlord CA*, the k Department CAs as
tenant CAs servicing *unrelated* organizations (rather than sub-organizations), and
 10 the certificate can be viewed as a *leased cert*. This terminology is borrowed from a
 more familiar example where the advantages of "joint construction and independent
 control" apply. Leased certs are in fact analogous to spec buildings having the
 identical floor footprints.

15 Rather than building just his own apartment, a builder is better off
 constructing a 20-floor building, setting himself up in the penthouse apartment and
 renting or selling out right the other floors. Each of the 20 tenants then acts as a
 single owner. He decides in full autonomy and with no liability to the builder whom
 to let into his flat, and whom to give the keys. A 20-story building is of course less
 20 expensive than 20 times a single-story one: it may very well cost 10 times that. This
 economy of scale is even more pronounced in a leased cert. Indeed, the cost of
 issuing a regular cert and that of issuing a leased one is pretty much the same. Thus
 issuing leased certs could be very profitable to a landlord CA, or at least repay it
 completely of the costs incurred for its own certs. On the other hand, tenant CAs
 25 have their advantage too, in fact: *they save on issuance costs*: they share the cost of
 issuing a cert k ways; and *they save on infrastructure costs*: they share the same
 responders (since they contain only public data).

Natural candidates to act as landlord CAs for external tenant CAs are: credit
 30 card companies; large financial institutions, and again the Government (e.g., via the
 USPS or the IRS). In many cases, in fact, they have long and close relationships
 with millions of "users" and may more easily issue them a digital cert without

investing too many resources for user identification (e.g., a credit card company has been sending bills for years to its customers, and can leverage this knowledge). A credit card company may like the idea of issuing certificates as a landlord CA in order to run more effectively its own affinity program (having hotel chains, airlines etc. as their tenants). The IRS may have already decided to use digital certificates, and leased certs may later on provide them with a revenue stream that will repay of the costs incurred for setting up a faster and better service.

Example 11:

10 So far, the way we have described landlord and tenant CAs requires that the landlord CA cooperates with its own tenant CAs during the issuance process, and thus that it has already identified its tenant CAs beforehand. It is actually possible, however, for a landlord CA to issue rental certs envisioning –say—20 tenant CAs, without having identified all or any of these tenants. Rather, future tenant CAs will
15 be able to rent space in already issued certs. This capability is ideal for new cert-enabled applications. Rather than undergoing the expenses necessary to issue certs to millions of customers, a company offering a new certificate-enabled product may approach a landlord CA having issued millions of certs, rent space in them *after the facts*, and then sign on as customers a large portion of the landlord-CA users by
20 *turning ON* all their corresponding certs *overnight* (without any customer identification and other issuing costs) and then starting managing them according to its own criteria. We shall describe various techniques for enabling this functionality in a forthcoming paper.

25

2003228468 03 Aug 2009

- 46 -

ADDITIONAL SYSTEMS

DEVICE VALIDATION SYSTEM

Let us now see how we can adapt the technology of described embodiments to devices (e.g., cell phones, PDA, Radio Frequency Identification tokens, PCs, Laptops, 5 VCRs, network devices, routers, firewalls, set-top boxes, CD players, game players, DVD devices, etc.).

Consider, for example, the very capacity of turning such devices ON, or letting them continue to operate. If a device is stolen, for instance, it is desired that it no longer 10 operate. On the other hand, if it is not stolen, then it should continue to operate normally. Similarly, if a user "rents" the device, or pays a subscription fee, or uses the device on behalf of a company (e.g., the device is a company laptop), if he no longer pays the rent, or the subscription fee, or no longer works for the company, then the device needs to be 15 turned OFF/disabled. Else, the devices should function properly. Also these devices could be turned ON, OFF, and ON again in a dynamic fashion.

Of course, these functionalities may be accomplished by means of a system according to some described embodiments. In essence, assuming again, for concreteness but without any limitation intended, a daily granularity, the device may carry a digital 20 certificate C, specifying a validity field X, and the device may work on a given day only if it has the daily proof of validity relative to X. The device may have a trusted/protected clock to avoid being fooled. The device (especially if a cellular device) may be "pushed" its own daily validity proof. Alternatively, the device may request to a second entity its own validity proof for the day. For instance, the device may provide its serial number and 25 receive in response the proof of validity of that day.

This works because the integrity of the validity field is guaranteed by a certificate and thus by a CA's digital signature of X (together with other information, such as date information). However, we may protect the integrity of X in the following alternative 30 way: namely, by "burning it in" the device in an

unalterable way: for instance, by writing it in a read-only memory in, say chips (smart-card/PDAs/telephone/laptop etc. chip sets). In this manner, the user of the device cannot alter X in any way. The proof verification algorithm can also be burned in. So that, once an alleged proof of validity P for the given day is presented,
5 then P is hashed the proper number of times and then compared with the burned in X. More generally, here one can use a one-way function F, rather than a one-way hash function. So that the whole process, including manufacturing, looks like this:

A first entity generates an initial value IV, and iterates a one-way function F
10 on V a given number of times so as to obtain a final value FV. A second entity (possibly equal to the first) burns X into a device D. Device D has means to iterate the function F. Device D later receives an alleged n-th proof value PV, where n is a positive integer, and verifies PV by iterating the function F on PV n times and checking that the resulting value equals the burnt-in value X.

15

Device D may consult its own clock to ensure that the n-th proof value corresponds to the current date. The current date may in fact be the n-th date in a sequence of dates starting from a fixed date. The fixed date may be burnt-in the device as well to protect its integrity as well.

20

At each iteration, function F may receive as input (not only the previously computed value but also) additional input. For instance, D's identifier may be an input at each iteration. Such additional inputs may be different at each different iteration as well. For instance, integer k may be an input at iteration k.

25

Also there may not be a single one-way function F. Indeed there may be a sequence of one-way functions, and F_k may be the function applied at iteration k.

The validity field X (being essentially unique to D) could also be used as D's
30 identifier (or part of it), so as to spare dealing with D's serial number and validity field separately.

The described system can be used so far to turn a given device D ON or OFF altogether. But it can also be used to turn ON or OFF a given just a single functionality, or a single functionality of out several possible functionalities. For instance X may be a validity field for functionality FX, Z a validity field for
5 functionality FZ and so on. In this case receiving a validity proof relative to X (Z) means that functionality FX (FZ) is turned ON for that day on device D. Such additional validity fields Z,... can also be burned-in the device D. Also a description/identifier of which functionality is associated to X/Z/... can also be burnt-in.

10

If the number of possible functionalities (and thus the number of validity fields) is large, then the validity fields can be Merkle hashed and then the root value of the Merkle tree may be burnt-in. In this case, to turn ON functionality FX (on a given day), one may provide the device with a proper proof of validity relative to X
15 (for that day), together with the authentication path from X to the root in the Merkle tree. The Merkle authenticating path algorithm may also be burnt-in.

Clock-less Device Validation

As we have seen, the technology of the preferred embodiment can be used to
20 validate devices and by turning them ON or OFF so as to prevent their misuse. Often the security of this application lies in the fact that the device has a clock not controllable by an enemy, possibly the very owner of the device (e.g., a fired employee who, after being fired, wishes to access company data from his company laptop that still lies at home). In fact, even if the company no longer issues a proof of
25 validity for day j, and even if in absence of such a proof of validity the device will not work on day j, an enemy can re-wind the clock of the device, so as to induce the device to believe that the current day is $d < j$, and then play back to the device a proof of validity correctly issued for day d, thereby fooling the device into functioning at day j.

30

The preferred embodiment provides technology that performs device validation even for devices which clock-less, that is, having no clocks, or no secure clocks.

5 The technology envisages a validator, an entity deciding whether a given device should be validated —i.e., turned ON or OFF—at a given date in a sequence of dates. For concreteness, but without limitation intended, let us assume that a given date is a given day in a sequence of days. The device preferably has a secure memory portion and a clock. Though not secure, the device can tell whether the
10 given clock has been reset at least while turned on. For instance, the device can tell, as long as it remains running, that 24hours have passed. The validation software preferably is protected in the device (e.g., running in a protected memory portion, or is burnt in, or resides in firmware), so as to avoid being altered in any way. Notice that some smart cards work in similar way. For instance they have a protected
15 memory portion, they may have a minimum power for keeping in (e.g., secure) storage a given value, and have a clock, but not a battery capable of having the clock running for any significant length of time. Thus, once inserted in a card reader, the smart card's clock becomes active, and the card may monitor the passage of time accurately (e.g., because the clock also is in secure memory), but once the card is
20 taken out of the reader the clock no longer works, though a small value may still be preserved in secure memory.

Example 12:

In this method, the Validator and the device share a secret key K. Key K
25 preferably resides in a secure memory portion of the device. From this key K, both the device and the validator are capable of producing a sequence of values unpredictable (to third parties not possessing K) corresponding to the sequences of dates. For instance, for each day 1,2,... the sequence of values consists of $V1=H(K,1)$, $V2=H(K,2)$,... where H is a one-way hash function, or an encryption
30 function that encrypts 1, 2, ... with key K each time. If, on day j, the validator wishes the device to be active for one more day, it publishes (e.g., it sends to a responder) the value $Vj=H(K,j)$. Assume now the device is turned on day j after

been active on day d and then switched off until day j. Then the device has retained in memory the value $V_d = H(K, d)$ or an indicator (e.g., d) of the latest day in which it was active. The device will not be functioning again until it gets a proof of validity after day d. Alternatively, the device keeps on storing –e.g., in a single variable—

5 the amount of time it has worked, in its own mind, during day d. When the device is switched off, therefore, it may remember not only d but also –say—6 hours and 10 minutes. Thus, when it is switched on again, it will continue to work for 17 hours and 50 more minutes. After that, it will require a proof of validity for a day subsequent to d. Assume now that the device really gets switched on again on day

10 $j > d$. Then the device gets a (alleged) proof of validity V_j for day j (e.g., it is pushed one such proof or it receives such proof after a request to a responder). The device then tries to see whether V_j is a proof of validity subsequent to the proof V_d currently in memory (or relative to a day subsequent to the day d in memory). For instance, the device keeps on producing V_{d+1} , V_{d+2} , ... using its secret key K until

15 the value V_j is produced (or until a given total number of day is exceeded –e.g., one may envisage that we no longer care about the device working at all after 10,000 days). If this is so, then it turns itself ON for another 24 hours (i.e., keeps in memory the new V_j or j, and properly operates and monitors the clock so that, after 24 hours of being continually on are reached, a new value V_{j+1} or V_k for $k > j$ is

20 needed.

The device can be turned OFF by not publishing or feeding it with future proof of validity, or can be turned OFF for good by publishing or causing it to receive a special value such as $H(K, \text{NO MORE})$, or a special value V_{nomore} stored

25 in memory, etc. The device can be suspended on a given day by publishing or causing it to receiving a special value, eg., $H(K, \text{suspend}, j)$. The keys for the validity, revocation and suspension proofs can be the same or different.

This offers a good deal of protection already. Assume that a device is

30 properly used on day j-1 and then it is stolen, and no proof of validity for day j is ever published or otherwise made available to the device. Then, whether or not the device was switched off prior to being stolen, it will stop working starting on day j.

In fact, if it was switched off, when revived it will need a proof of validity for a day after j-1 to turn itself ON properly, and no such proof is forthcoming. If it stolen while being switched on, after 24 hours at most it will stop working anyway.

5 At worse it may happen that the device was switched on (for instance at day j-3), and thus entered in possession of a validity proof V_{j-3} , and then switched off. Assuming that the device is stolen at this point but that its loss is not noticed until day j-1, or that the device is stolen at day j-1 and that an enemy records the values V_{j-2} and V_{j-1} that the device could have seen. Then such an enemy could at most
10 feed the device these two values and make it work for two more days at most.

Example 13:

 This method works essentially as the method disclosed in Example 11, using a sequence of unpredictable values, published or otherwise made available to the
15 device at each of a sequence of dates (e.g., without limitation, days), a clock not secure, etc. but does not use a secret key in the device. For instance, the device stores X_k , the result of iterating one (or more) one-way function F k times on an initial value X_0 as discussed above and with the same variants. Then X_k is written in
firmware (e.g., in a non-alterable way) or stored in a protected portion of memory.
20 The proof of validity for day j simply is X_{k-j} as in the basic scheme of the present invention. Again suspension and revocation can occur in similar ways.

RTC PHYSICAL ACCESS CONFIGURATIONS

MULTIPLE PRIVILEGE MANAGEMENT IN MIXED ENVIRONMENTS

25

 A robust access control system must answer two questions for every user. The first question addresses authentication or identification: "Are you who you say you are?" This question is typically addressed directly or indirectly through
identification badges, keys biometrics, or passcodes. These provide reasonable
30 answers for long-lasting user identification, but don't address the more time critical question of validation: "Are you currently allowed to do what you are trying to do?"

For example, an identification badge can tell you that Alice was hired as an employee some time in the last decade, but cannot independently determine whether she is still an employee with access permissions for the computer server room.

5 For physical access control, a secure lock must determine identity through authentication, and then perform validation to determine whether a user's current privileges allow entry. Some locks perform this validation through wired network connections to a central trusted permissions authority. A physical access solution based entirely on network wired locks has two significant limitations. The cost of
10 each wired lock includes the costs of secure wiring, field control panels, and labor, totaling several thousand dollars per door. The reach of a wired configuration is limited to locks that can be easily accessed by permanent networking. This prevents the use of robust access control for mobile or hard to reach locks such as those on vehicles, storage containers, utility cabinets, etc.

15 The Real Time Credentials technology according to a preferred embodiment of the present invention provides a secure way to perform efficient validation for physical access on both wired and disconnected locks. This allows intelligent door locks to validate current user privileges and permissions without requiring expensive
20 network connections to each lock.

The present disclosure describes several configurations that can be used to provide disconnected validation based on large numbers of independent user privileges. Each configuration offers interoperability with existing access control
25 hardware and software for use in heterogeneous installations. For each configuration, this paper will describe how Real Time Credentials offer increased flexibility while dramatically lowering the total cost of high security.

All four configurations described, below, feature an identical RTC validation
30 process. The primary difference between these schemes is the process of authenticating the user, which impacts price and compatibility with existing access solutions.

Contactless ID/Memory

The first RTC validation configuration is an access control environment based on contactless ID cards with read/write memory access. This is described
 5 using the common MIFARE™ Standard contactless card as an example, but the validation solution would be identical with any memory ID card.

When a MIFARE ID card is used in current networked physical access environments, the lock reads the ID from a card and transmits it to a nearby panel or
 10 server that checks privileges and performs validation. The authentication process is the determination of the card ID, and the validation process is handled remotely based on this ID.

The physical access solution of the present invention can maintain
 15 compatibility with this model for wired doors, but adds support for disconnected doors by using the card's read/write memory to store a digitally signed "validation proof" for that card. This proof is periodically written to the card at any networked reader, and then it can be read at any disconnected lock to establish the current validity and permissions of the user.

20

The following table shows the logical contents of the RTC validation proof that is stored onto the card, along with the approximate storage requirements for each component:

25	Card ID:	#123456	4 bytes
	Status:	card valid	1 byte
	Start time:	8/4/03 09:00	4 bytes
	End time:	8/5/04 08:59	4 bytes
	Authority:	ACME Inc.	20 bytes
30	Privileges:	R&D labs	1 bit to 10 bytes
		Parking	1 bit to 10 bytes
		Locker 53	1 bit to 10 bytes

	Terminal B	1 bit to 10 bytes
Digital Signature		42 bytes
		Total Size: ~ 100 bytes

5 When a user enters a facility through a wired door, the door retrieves the user's complete validation proof in the above format and places it into the memory area on the card. Once the proof is loaded on the card, a disconnected lock can validate the user's permissions through the following steps:

- (1) Perform standard authentication by retrieving the user's card ID;
- 10 (2) Retrieve the RTC validation proof from memory;
- (3) Verify the digital signature matches the known public key of the trusted authority;
- (4) Verify that the proof is current (using the start and end times);
- (5) Verify that the card is valid;
- 15 (6) Check arbitrary access control requirements based on privileges from the proof.

The disconnected lock is configured with a set of access control rules based on privileges, rather than individual user ID. For example, a lock may be configured
 20 to only admit users with the "Parking" privilege, and only during business hours. Since the individual user privileges can be changed through the RTC validation proofs, the locks themselves do not need to be changed as new users are added and removed to change access permissions. In addition, the locks do not need to store any secret keys or data, which means that an individual lock can be disassembled
 25 without any reduction in overall system security.

The RTC validation proofs according to a preferred embodiment of the present invention have certain characteristics that make them uniquely powerful for physical access control environments. Since the proofs are digitally signed, they are
 30 unforgeable and tamper-proof. Since the proofs do not contain any secret keys, they can be public, and transmitted without security risk. The proofs are small enough to be stored on a low-end memory card.

These characteristics allow the RTC validation proofs to be used in cards like MIFARE Standard, while still offering high security cryptographic validation with thousands of independent user privileges per card.

5

Cost. MIFARE 1k Standard cards are available for between \$1 and \$5, depending on manufacturer and volume. A disconnected lock based on MIFARE cards and RTC validation technology could be manufactured for under \$500 per door. With installation, a single door or container could be secured for under \$1000.

10

Security. Simple ID authentication offers weak protection against duplication and forgery. Second and third factor authentication combined with PKI protections can be used to increase authentication security. Credential validation is protected by strong PKI encryption, preventing permission forgery or modification.

15

Contactless Shared Secrets

RTC Credential validation can also be used with identification cards such as HID's iClass that perform validation using secret information that is directly or indirectly shared with all readers. A lock will perform authentication to a card using a randomized challenge/response protocol which proves that the card knows the secret correspondence to its ID.

20

The RCT validation for a shared secret card is identical to the validation for a simple ID card. When a user enters a wired door, the lock will write the current RTC validation proof onto the user's card. This proof is later retrieved by disconnected readers for offline validation.

25

Cost. Contactless shared secret cards with memory are available for between \$5 and \$10, depending on manufacturer and volume. A disconnected lock based on shared secret cards and RTC validation technology could be manufactured for under \$500 per door. With installation, a single door or container could be secured for under \$1000.

30

Security. Shared secret authentication reduces the chance for duplication of individual cards, but compromise of a single offline reader may allow duplication of many cards. Credential validation is protected by strong PKI encryption, preventing
5 permission forgery or modification.

Contactless PKI

Cards capable of performing public key digital signatures offer the highest level of authentication security. This includes cards based on MIFARE PRO X
10 chips as well as many high end JavaCards. Locks may authenticate a card based on a challenge/response protocol without requiring any sensitive information in the locks. This significantly reduces the risk of key duplication.

The RTC validation for a public key card is identical to validation for a
15 simple ID card. When a user enters a wired door, the lock will write the current RTC validation proof onto the user's card, and this proof will be retrieved by disconnected readers for offline validation.

The card's public key will typically be represented by a digital certificate,
20 which can be used for alternate applications such as computer access and email security. High-end public key cards may support additional applications such as information security or stored value, which helps reduce the total cost for each application.

Cost. Contactless PKI cards are available for between \$10 and \$20,
25 depending on manufacturer and volume. A disconnected lock based on MIFARE cards and RTC validation technology could be manufactured for under \$500 per door. With installation, a single door or container could be secured for under \$1000.

30

Security. PKI cards are able to provide strong cryptographic authentication to locks with low risk of key compromise or card duplication. Credential validation

is protected by strong PKI encryption, preventing permission forgery or modification.

TECHNIQUES FOR TRAVERSING HASH SEQUENCES

5

Let H be a one-way hash function. A hash chain of length n is a collection of values

$x_0, x_1, \dots, x_n$ such that $H(x_i) = x_{i-1}$. While x_{i-1} is easy to compute from x_i , computation in the opposite direction is infeasible, due to one-wayness of H .

10

The following is a representation of a hash chain:

$$x_0 \quad (H) \quad x_1 \quad (H) \quad \dots \quad (H) \quad x_{n-1} \quad (H) \quad x_n$$

In many applications (such as, for example, document validation and privilege management services) it is necessary to be able to traverse the hash chain, i.e., to generate the values $x_0, x_1, \dots, x_n$ in order (from left to right in the above chain), over a certain period of time (for example, to output one value a day for a year). Note that the left-to-right order makes this problem difficult, because of one-wayness of H . While it is easy to generate and output, in order, $x_0, x_1, \dots, x_n$, by simply repeatedly applying H , the opposite order requires more time and/or memory.

The two obvious approaches are:

Store just one value, x_n , and, in order to output x_i , compute $H^{n-i}(x_n)$;

25

Store all the values, $x_0, x_1, \dots, x_n$, erasing them as they are being output.

The first approach requires storage of two hash values (one for x_n and the other for the computation of x_i) and $n(n+1)/2$ evaluations of H total, or, on average, $n/2$ evaluations per value output. The second approach requires storage of $n+1$ hash values and n evaluations of H total, or, on average, 1 evaluation per value output.

30

We are interested in intermediate solutions: ones that offer other tradeoffs of memory (the number of hash values stored) versus time (the number of evaluations of H needed).

5 An algorithm has been proposed in the prior art that resulted in the following tradeoff: $\lceil \log_2 n \rceil$ hash values stored and at most $\lfloor \log_2 n \rfloor$ computations of H per hash value output. (See Don Coppersmith and Maruks Jakobsson, Almost Optimal Hash Sequence Traversal, in Matt Blaze, editor, *Financial Cryptography: Sixth International Conference (FC '02)*, Southampton, Bermuda, 11-14, March 2002).

10 Novel Algorithms with Constant Storage

Jakobsson's method requires storage of about $\log_2 n$ hash values, and cannot be used when less storage is available. Note that for a hash chain of length 365, this means that 9 values need to be stored, and for a hash chain of length 1,000,000, this
15 means that 20 values need to be stored. We would like to have an algorithm with lower storage requirements. Moreover, we would like to be able to specify storage requirements that are *independent* of the hash chain length. This way, the same amount of memory would be needed to manage short chains and long chains; thus, one would not need to acquire new memory if hash chain lengths change.

20 For convenience of reasoning about the algorithms, let's call a value x_j that the algorithm stores a *pebble at position j* . Then a pebble is "allowed" the following: (i) to move to a position where another pebble is located (this corresponds to copying a value), or (ii) to move one step left of its current position
25 (this corresponds to evaluating H). Initially, pebbles may start out in arbitrary positions on the hash chain.

Note that the number of pebbles corresponds to the number of hash values stored, and the number of times a pebble takes a step to the left corresponds to the
30 number of evaluations of H . Our goal, then, is to come up with algorithms that reduce the number of pebbles steps (what we will call "cost") given a particular number of pebbles.

Two Pebbles

It is clear that one always needs a pebble at n ---if x_n is not stored, there is no way to recover it and thus no way to output it when it is needed at the end of the traversal. It is also clear that one always needs a pebble at the current position i , in order to be able to output x_i . Thus, at least two pebbles are necessary.

If only two pebbles are used, then one of them must always stay at x_n , and the other has no choice but to start at x_n and move to x_i each time. Thus, the best algorithm for two pebbles takes $n(n+1)/2$ total steps, or $n/2$ steps per output on average. For example, for a hash chain of length 1,000,000, the average number of steps is 500,000 per value output.

Three Pebbles

If we add just one pebble to the two that are absolutely necessary, it turns out that we can dramatically improve on the number of steps.

We will proceed as follows: divide the hash chain up into intervals of length s , where $s = \lceil \sqrt{n} \rceil$ (note that there will be $n/s \approx \sqrt{n}$ intervals). Place pebble number 3 at x_n , and pebble number 2 at x_s . Then, using the algorithm for two pebbles described above, use pebble number 1 to traverse points $x_0 \dots x_s$ (starting each time at x_s). Then place pebble number 2 at x_{2s} (by starting at x_n and moving left), and again use algorithm for two pebbles to traverse $x_{s+1} \dots x_{2s}$. Continue in this manner, each time using the two-pebble algorithm for an interval of length s .

The total number of steps of this algorithm can be computed as follows: to traverse each interval using two pebbles, we need $s(s+1)/2$ steps. In addition, to move pebble number 2 to the beginning of each interval before traversing it, we need $(n-s) + (n-2s) + \dots + s + 0 = (n/s)(n/2)$ steps. Recall that $s = \lceil \sqrt{n} \rceil$, so the average number of steps per output value is $s/2 + (n/s)/2 = \lceil \sqrt{n} \rceil$.

Thus, adding a third pebble to the bare minimum of two allows us to decrease time per output value from $n/2$ to $\sqrt{n}$. This decrease is indeed dramatic: for example, for a hash chain of length 1,000,000, the average number of steps is 1,000 per value output (as opposed to 500,000 needed with two pebbles).

5

Four Pebbles

If we have yet another pebble available, we can again divide the hash chain into intervals. This time, we will set $s = \lceil \sqrt[3]{n} \rceil$, and divide the entire chain into n/s intervals of length s .

10

We will then place pebble number 4 at n , and use it as a starting point for pebble number 3, which will move to the beginning point of each interval of size s , in order from left to right. On each interval, we will use the three-pebble traversal algorithm described above. That is, we will further subdivide each interval into subintervals of size $\lceil \sqrt{s} \rceil$, and place pebble number 2 at the beginning of each subinterval, in order from left to right (pebble number 2 will start, each time, and pebble number 3). Then pebble number 1 will traverse the subinterval, each time starting at pebble number 2.

15

20

Thus, the cost of traversing each interval will be $\sqrt{s}$, or $\lceil n^{1/3} \rceil$ per value output. To that, we have to add the cost of moving pebble number 3 to the beginning of each interval. Pebble number 3 will be moved n/s times: n/s steps at first, $n-2s$ steps next, and so on, giving the average cost of $(n/s)/2 = n^{1/3}/2$ per value output.

25

Thus, the average number of steps per value output is $\lceil 1.5n^{1/3} \rceil$. Using, once again, the example of a chain of length 1,000,000, the average number of steps is 150 per value output.

30

Generalizing to More Pebbles

The general technique that emerges from the above examples is as follows. Given c pebbles, divide the hash chain into $n^{1/(c-1)}$ intervals of length $n^{(c-2)/(c-1)}$

each. Use the technique for $c-1$ pebbles on each of these intervals. The average cost per output value will be $((c-1)/2) n^{(1/(c-1))}$.

5 This generalization can be considered not only for a constant number of pebbles, but also, for example, for $c=1+\log_2 n$. In that case, using the equation $n^{1/\log_2 n} = 2$, we compute that the average cost per output value will be $\log_2 n$ using our algorithms.

10

Improving Worst-Case Cost

Even though the above techniques achieve good average-case cost per output value, some output values will take longer to compute than others.

15

Take, for example, the case of three pebbles. Every time we traverse s pebbles, we have to relocate pebble number 2. Thus, the output value at the leftmost end of an interval will take much longer to compute; for example to compute x_{s+1} , we will need to make $n-(s+1)$ steps. On the other hand, all other pebbles within an interval will take at most s steps.

20

This, of course, may present serious problems in some applications: the computing equipment involved would have to be fast enough to handle these "bad" cases. But if it is already that fast, then there seems to be no point in having good "average" case: we would still need powerful computing equipment, which would simply sit idle on average.

25

In order to prevent this problem, we need to make the cost of the worst-case output value close to the cost of the average-case output value. In the case of three pebbles, this can be accomplished by adding only one extra pebble. Call that pebble "2a." Its job will be to move in advance to where pebble 2 should be next. For example, when pebble 2 is positioned at point s , pebble 2a will start at point n

30

moving toward point $2s$. It will reach point $2s$ exactly when pebble 2 needs to be there---by the time the value s is output.

Thus, while any given interval of size s is being traversed, pebble 2a will
 5 start at position n and move left to the beginning of the next interval. Note that pebble 2a needs to take fewer than n steps in order to get to its destination. The obvious approach would be for pebble 2a to take at most n/s steps for each output value in the interval. This would result in a worst-case cost of $s+n/s \lceil 2\sqrt{n} \rceil$ steps per output value. Note, however, that one can do better: because pebble 1 will
 10 need to take more steps for values at the left end of the interval than values at the right end of the interval, in order to reduce the worst-case cost, pebble 2a should start out "slowly" and then "speed up." This way, the total number of steps taken by pebbles 1 and 2a will stay constant. Specifically, pebble 2a should take $(n/s)/2$ steps at first, $(n/s)/2+1$ steps the next time, and so on, up to $3(n/s)/2$ steps when the last
 15 value of the interval is being output. This will reduce the worst-case cost further to $\lceil 1.5\sqrt{n} \rceil$.

Note that the total number of steps, and thus the average cost per output value, do not increase with the addition of this extra pebble. This is so because the
 20 extra pebble is not doing any extra work, but rather doing work slightly in advance. Thus, with a hash chain of length 1,000,000, the worst-case cost would be 1,500, while the average-case cost would be 1,000 per output value.

This approach extends to more pebbles. If we take the solution with four-
 25 pebbles, and add pebbles 2a and 3a that move in advance into the appropriate positions for pebbles 2 and 3, respectively, we will reduce the worst-case cost to $\lceil 2n^{1/3} \rceil$. Taking again the example of the chain of length 1,000,000, the worst-case cost would be 200, while the average-case cost would 150 per output value.

30 Therefore, in general, with $2c-2$ pebbles, we can traverse the hash chain at the average cost of $((c-1)/2) n^{1/(c-1)}$ per output value, and worst-case cost of $(c/2) n^{1/(c-1)}$ for any given output value.

Again, this generalization can be considered not only for a constant number of pebbles, but also, for example, for $c=1 + \log_2 n$. In this case, using $2\log_2 n$ pebbles, our algorithms will traverse the hash chain with the average cost per output value of $\log_2 n$ and worst-case cost of $1 + \log_2 n$.

The Optimal Solution

Below we describe a method for obtaining an algorithm with provably optimal total (and thus average per output value) computational cost, given any number c of pebbles. Note, however, that for a small values of c , this provably optimal solution will reduce the number of steps only slightly compared to the solutions above.

Suppose we have c pebbles. We must store x_n , which occupies 1 pebble. Then one more pebble will be moved to x_k (for some k to be determined below), by applying H $n-k$ times to x_n . Then, recursively, use the optimal solution for $c-1$ pebbles in order to output $x_0, x_1, \dots, x_k$, in order. Note this amounts to traversing a shorter chain---one of length k , because the value x_k is stored. Then recursively use the optimal solution for c pebbles to output the values $x_{k+1}, \dots, x_n$, in order. This again amounts to traversing a shorter chain---one of length $n-k$, because the first k values are already traversed.

Now define $F(c, n)$ as the number of steps necessary to traverse a hash chain of length n while storing no more than c pebbles at any given time. Clearly, $F(c, 0)=0$ for any $c \geq 1$, and $F(0, n) = \infty$ for any n . Then, in our above method, $F(k, n)=\min_k F(c-1, k)+F(c, n-k-1)+n-k$, and k should be chosen to minimize $F(c, n)$.

It is a simple matter of recursion with memoization (a.k.a. dynamic programming) to find the optimal point k for particular c and n . We present the C code that accomplishes this task. Such optimal points can be easily found in advance and then integrated into the hash traversal code.

Our Implementation of the Optimal Solution for Any Amount of Memory

```

#include "stdio.h"
int **table;
int **ktable;
5
int f(int r, int n)
{
    int k, t_min=-2, t, k_min=-2, t1, t2; // -2 Stands for infinity;
                                     // -1 stands for uninitialized
10
    if( table[r][n]!=-1)
        return table[r][n];
    if (n==0 && r>0) {
        table[r][n] = 0;
15        ktable[r][n] = 0;
        return 0;
    }
    if (r==0) {
        table[r][n]=-2;
20        ktable[r][n]=-2;
        return -2;
    }

    for (k=0; k<n;k++) {
25        t1=f(r-1, k);
        if (t1==-2)
            continue;
        t2=f(r, n-k-1);
        if (t2==-2)
30        continue;
        t=t1+t2+n-k;
        if (t<t_min || t_min==-2) {

```

- 64 -

```

        t_min=t;
        k_min = k;
    }
}
5   table[r][n]=t_min;
   ktable[r][n]=k_min;
   return table[r][n];
}
void main()
10  {
    int max_r, max_n, i, j;

    printf("max balls: ");
    scanf("%d", &max_r);
15  printf("chain length: ");
    scanf("%d", &max_n);

    table = (int**) malloc((max_r+1)*sizeof(int));
    ktable = (int**) malloc((max_r+1)*sizeof(int));
20
    if (table==NULL || ktable==NULL) {
        printf("Out of memory!\n");
        return;
    }
25
    for (i=0; i<=max_r; i++) {
        table[i]=(int*)malloc((max_n+1)*sizeof(int));
        ktable[i]=(int*)malloc((max_n+1)*sizeof(int));
        if (table[i]==NULL || ktable[i]==NULL) {
30         printf("Out of memory!\n");
            return;
        }
    }

```

```

        for(j=0; j<=max_n; j++)
            ktable[i][j]=table[i][j] = -1;
    }

5

    for(i=0; i<=max_r; i++)
        for (j=0; j<=max_n; j++)
            f(i,j);

10

    printf("\nTable for F(r, n) -- the number of steps needed:\n  n\\r");
    for(i=0; i<=max_r; i++)
        printf("%6d", i);
    printf("\n");
15
    for (j=0; j<=max_n; j++) {
        printf("%6d:", j);
        for (i=0; i<=max_r; i++)
            printf("%6d", table[i][j]);
        printf("\n");
20
    }

    printf("\nTable for k -- the optimal position to put the first pebble:\n  n\\r");
    for(i=0; i<=max_r; i++)
        printf("%6d", i);
25
    printf("\n");
    for (j=0; j<=max_n; j++) {
        printf("%6d:", j);
        for (i=0; i<=max_r; i++)
            printf("%6d", ktable[i][j]);
30
        printf("\n");
    }

```

```

    }
  }
\end{verbatim}

```

5

PRIVATE KEY SECURE PHYSICAL ACCESS (REAL TIME CREDENTIALS IN KERBEROS-LIKE SETTINGS)

In general, the scenarios may include multiple doors, and multiple users.

10 Moreover, the access might be controlled by multiple authorities (each authority controlling access through some doors, the sets of doors for different authorities possibly overlapping). On the most general level, the access is controlled by having the users presenting credentials to the doors (verification of such a credential may require interaction between the user and the door, such as PIN entry, as well as an

15 exchange of messages between the door and the user's card). In the case of the doors, it is especially important to support the security of the access with the least cost and even without connectivity of the door to a network or any specific server.

One important observation is that whatever credentials we use, our RTC

20 technology allows to derive important security, infrastructure and cost benefits. RTCs can be utilized in conjunction with either public key cryptography methods (certificates, public key signatures, PKI) as well as the private key cryptographic tools (symmetric or private key signatures and encryption, Kerberos-like systems, etc.)

25

Access control for disconnected doors using public-key technology has been addressed. Here we describe how to adapt those ideas to private-key technology.

Basic Primitives

30 Encryption, signatures, pseudo-random functions

In particular, private-key encryption, private-key signatures (aka MACs), private key random functions, are typical private-key primitives that we shall be

using. For many of our purposes, these primitives could be used interchangeably. For instance, deterministic private-key signature schemes (between two entities who share a secret signing key SK), and random functions F_s (whose seed s is shared between two entities) can actually be considered equivalent. Both produce outputs
 5 that are unpredictable to third parties who might know the corresponding inputs, but not SK or s . For instance, the functions $FSK(x)$ that returns the digital signature of x with secret key SK can, in practice, be considered a good enough pseudo-random function with seed SK. On the other hand, the function $F_s(x)$, that on input x returns the value at x of pseudo-random function F with seed s , could be considered a
 10 private-key signature algorithm with secret key s .

One-way and one-way hash functions

We shall also use another basic primitive: one-way functions F and one-way hash functions H . In essence a function F is one-way if (1) given an input X , one can
 15 efficiently compute $F(X)$, while (2) given $F(X)$, where X has preferably been chosen sufficiently at random so as to be sufficiently unpredictable, computing X is practically impossible (e.g., because too many values for X would have to be tried in principle, and no efficient method exists to narrow the number of possible candidates). A function H is a one-way hash function if it is one-way and (though
 20 preferably mapping longer inputs to shorter ones or arbitrarily long inputs to —say—160-bit ones) it is hard to find two distinct inputs X and Y such that $H(X)=H(Y)$.

In practice, we can use a one-way hash function H to construct other
 25 primitives. For instance, private-key signatures can be constructed in the following simple way. To sign a message M with secret key SK, one computes $H(SK,M)$; that is, one properly combines SK and M —eg, concatenates them—and then hashes the result. Of course, to sign and date M , one can add a date d to this combination and thus compute $H(SK,M,d)$ instead. Similarly, pseudo-random functions can be
 30 constructed as follows. On input x , to produce the output of a pseudo-random function F with seed s , one may compute $H(s,x)$; that is, one may properly combine s and x , and then apply a one-way hash function to the result.

Secure Physical Access

We focus on just the novel aspects introduced by the private-key setting, skipping those common aspects that could be adapted to the new scenario naturally (e.g., the daily/regular computation aspects etc.) We start with a simple scenario.

Single Organization

Let D be a door (with the said mechanism), A an organization that wishes to control access to D, and U a user (possibly working for A), again having a card, CU, with proper identifiers, etc. Then A may control access to D by sharing a secret key SK with D. If A wishes to grant U access to D on day d (time interval d), it computes a proof PUDd, that it is hard for anyone other than A (and possibly D) to compute but easy for D to verify. Let us see how this can be done, both using private-key encryption and private-key signatures.

Private-key encryption solution (with possible proof of identity)

For instance, PUDd may be the encryption, EUDd, of a message specifying U, possibly D as well, and d with the private encryption key SK according to some established private-key encryption algorithm such as DES. Upon receiving EUDd from U's card, D decrypts it with key SK, and if the result specifies both U and the current day (time interval) d, then the door opens. The door may use its own clock to determine whether its own time falls within time interval d.

Here, like elsewhere, U is intended to denote both the user as well as a proper identifier for U. If user U has a card (preferably securely) associated with him, then U may be such card or a proper identifier of it. In the latter case, for instance, the door's card reader may get U from the card and also get EUDd, then it decrypts EUDd with key SK and compares the decrypted U with that furnished by the card, to ensure that they are equal.

Notice that EUDd proves to the door D that user U is authorized to enter through it on time interval d, but this does not prove to D that it is indeed dealing

with user U. Thus, we may augment the basic scheme with a way for U to prove his own identity to the door. This can be done in a variety of ways. In particular, authority A may provide EUDd only to U's card, and U's card is provided with a key pad, and can transfer EUDd to the door D only if the right PIN is entered on its

5 key pad (and the card may self-destroy, or erase its relevant volatile memory content if the wrong PIN is entered more than a given number of times). This way, whenever the door receives EUDd, it knows that it is receiving from U's card (because A only transfers EUDd to U's card) and it knows that the "user behind the card" must be U (as opposed to a malicious user having stolen U's card) because U's card would not

10 work or transfer EUDd to D unless U's PIN has been entered on its key pad. A second way for U to prove his identity to D consists of having U provide his own PIN directly to D. For instance, door D may have its own key pad, and U uses it to enter his own PIN, PINu. The door may have an internal way (e.g., a table) that maps PINu to U, and thus can realize that it is indeed dealing with U. If there are

15 many doors in the system, however, providing and updating (e.g., because of new users joining the systems) a table for each door may be impractical. It is thus preferable to have U's identifier may directly be PINu. For instance, EUDd might be EPINuDd. When user U approaches door D, he enters PINu into D's key pad and his card transfers EPINuDd to the door. The door then checks whether the PIN entered

20 equals that specified in EPINuDd, and in this case it is dealing with the right user and that this same user is authorized by A to go through door D without using any PIN-user table: indeed, the key pad tells D that a user knowing PINu is in front of it, and EPINuDd tells D that the user knowing PINu is currently authorized to go through D. In a third way, rather than directly appearing into EUDd, the user PIN

25 may be securely coupled with EUDd. For instance, A may give EUDd to U's card encrypted with key PINu or with a key K reconstructable from PINu (e.g., $k=H(\text{PINu})$ or $K=H(\text{PINu},d)$ or $K=H(D,\text{PINu},d)$ etc.). In this case, door D will check that the PIN is securely bound to the user's authorization for time interval d. For instance, it uses PINu to decrypt EUDd and checks that EUDd is a proper

30 authorization using the key SK it shares with authority A.

Using Responders

But: how can A easily and securely transfer EUDd to U's card? We propose using responders. These are devices (such as servers or computer terminals/card readers capable of being linked to a server). Preferably these responders need not be vaulted or protected. Such protection could add so much cost and inconvenience to the system that it is crucial to have the system work securely without securing the responders! Ideally, authority A performs an update at every date d of a series of dates. Each date preferably specifies a time interval (e.g., a day). For instance d may be day d or the beginning of day d. During update d, A decides which user U should be granted access to/through D, and computes a proof verifiable by D of this fact. For instance, in an encryption-based shared-key system, this proof may be the string EUDd discussed above and can be verified because A shares with D the key SK that A used to compute EUDd. All these proofs are then sent to the responders. These responders are preferably located in convenient locations. For instance, in an airport system, responders may be located at the airport main entrances. User U then (e.g., when arriving at work) picks up from a responder his own authorization to go through door D. Preferably, U's card may authenticate itself to the responder in order to receive EUDd. This is very convenient, because without wireless or other expensive systems, a user picks up all his daily authorizations for all the doors he is entitled to go through on a given day from the front entrance (through which he may have to go through anyway) and using a traditional mechanism like inserting his own card in a card reader (e.g., to prove that he has shown up at work). After that, he is free to go around the airport and can go easily through all the protected doors D he is entitled to using the authorizations EUDd that he has picked up. But because of this convenience and the fact that the responders are preferably insecure, a malicious user may also pick up a honest user's authorization. It is thus necessary (1) to prevent this from happening without securing the responders and/or (2) ensuring that the authorizations for an honest user cannot be used by anyone else. The latter case can be sufficiently enforced by having users also enter a PIN at the door, as already discussed, preferably securely bound to the authorization released by the card. Thus a malicious user V picking up U's authorization EUDd from a responder cannot impersonate U at the door because it does not know U's PIN. The former protection

can be enforced by having authority A send a responder authorization EUDd after encrypting it with a key SKCU inside U's card CU and known to A. This way, A essentially posts in the responder an encrypted authorization EUDd' that can be turned into an authorization EUDd only by U's card, making it useless for a
 5 malicious V to download someone else's authorization for the day. Even if V manufactures his own card any way he wants, V still would not know SKCU.

It is further possible to have A share a secret key SKD with door D and secret key SKU with the user U. Then PUDd can be a value EUDdk, consisting of
 10 indications of user U, door D and day d, as well as some random secret k, all encrypted (by A) with the secret key SKD. (Note that, in this case, U cannot decrypt EUDdk). In addition, U would receive Ek – namely, k encrypted with SKU. (D and d might be known to U, or could be communicated to U --e.g., by the same responders at the main door.) This way, because U knows SKU, U obtains secret k
 15 as well. In order to enter the door D, card U would send EUDdk to D. D would respond with a random value q, and card U would then send Eq, i.e., q encrypted using secret k. The door D would decrypt Eq, verify that the same q was used, and U is the same as that specified in EUDdk, and that the date d is current and if all the checks are confirmed, will let U through. This mechanism could incorporate PIN
 20 mechanism as above, making it even more secure. Alternative Challenge-Response methods based on k are possible. (In particular, D can compute and send Eq and ask U to send back the correct decryption q.) Such mechanisms provide security even if the attacker monitors the communication between the card and the door.

25 However, an enemy who sees the PIN entered by the user at the door, could after stealing U's card impersonate U, at least during time interval d if U's card has EUDd within it. After that, if U reports stolen his card, A will not any longer make EUDd available to U's card.

30 Private-key signature solution

For instance, PUDd may be the private-key digital signature of a message specifying both U and d (and possibly D as well) with private key SK, known to

both A and D, according to some established private-key signature algorithm. In particular, letting H be a one-way hash function, then $PUDd = H(SK, U, d)$. Upon receiving U from the card, the door's reader may sign U and d with its own private key SK and compare whether the result of this computation matches the string PUDd obtained from the card. Notice that the door reader, carrying a clock, may know what is the current day d, and thus needs not to receive it from the card. This works as long as A grants access for full days at a time. Else, the card also sends d (or the chosen time interval) to the reader, and then the reader digitally signs with SK the obtained U and d, checks that the result indeed equals PUDd, and then that the current time (according to the door's clock) is within d. If so it opens.

Again U may be asked to enter a PIN as part of the transaction. In which case the PIN may also be used as part of U. For instance, U may consist of u and PIN, where u is a string identifying the user, and PIN a password known to the user. In which case, the card transfers to the door reader u, and PUDd (and possibly D or d and additional quantities), the user enter PIN to the door control coupled with the reader, or to the reader itself, and then the reader reconstructs $U = (u \text{ PIN})$, and then signs U with SK to check that PUDd is obtained. Again, if d is card-supplied, it also checks that the current time is within d. This method makes couples a user and his card in a tighter way, so that a enemy that steals the card would have hard time using it without the proper PIN.

Of course, the same SK could be used for a set of doors, in which case by granting access to U for one of them A automatically grants him access to all of them. To allow the greatest granularity of access, each door D may have secret key SKD.

Combining the two approaches

As an example of combining the two approaches U may receive from A (e.g., using mechanisms discussed above, in particular, utilizing encryption) a secret key SKUd for the day d. He may then "prove" to the door D his identity and/or authorization using private-key signatures. Namely, the door D would send to the

card U a random message m ; in response card U would send the signature of m : $H(m, SKU_d)$. Note: computation of this signature may require the PIN_u . The door D then verifies the signature. This may require that the door D knows SKU_d (e.g., having received it from A directly, or compute it from some other information: e.g., $H(SK_d, d, U)$, etc.) Alternatively, A may encrypt SKU_d with a key A shares with D, obtaining $ESKU_d$. Then $ESKU_d$ can be given to U (e.g., as described above), and then U can send it to D together with the signature.

Multiple Organizations

As we have seen it suffices for an organization/authority A to share a secret key SK_d with a door D in order to control which users U may access D in a given time interval d . This process can be extended so as to enable multiple organizations, A, B, C, ..., to independently control access through a door D or set of doors, D1, D2, D3, ... Each organization X shares a secret key SK_{XD} with door D, and then use on the solutions described above. For instance, each organization X may choose SK_{XD} and insert it into D's reader. Each organization X may have to send a team of one or more employee/hired workers/contractors/subcontractors from door to door. But to do so in a facility with lots of doors may be impractical or wasteful, since other organizations may have done so already. Also, if there are or there will be many authorities, then the reader may have difficulty in storing all these keys. In addition, proper precautions should be taken. Else, nothing would prevent an enemy from inserting his own secret key into a door's reader, and then, knowing it, it could use any of the above methods to grant access to himself or his accomplices to that door. For these reasons, we put forward the following solutions. Notice, the same methods could be applied to a single solution as well.

First Solution

As we have seen, a user can go through a secure door if he or his card share a secret key for a given time interval. In a way, therefore, the user and the door share a session key. Kerberos and Needham-Schroeder protocols provide a mechanism for ensuring that pairs of entities share secret session keys, and could be applied here within the overall system. However, these protocols are based on a key-distribution

center that is on-line and must be contacted whenever a shared session key is needed. Thus, we wish to put forward additional and more convenient methods. To begin with, even for implementing a Kerberos/Needham-Schroeder based system, we need a way for a central authority to distribute keys to doors (which may be
5 harder than distributing keys to other authorities).

We envisage a special authority SA (for instance, at an airport, the Airport Authority) to securely distribute keys to door readers. Preferably, SA may be the only entity that can do so. For instance, the door reader is delivered with no secret
10 keys inside, and is manufactured so that once the first set of secret keys (possibly a set of a single key) is inserted, then the reader stores it for a long time, and accepts no other keys for storage in the future. This way, by being the first one to insert any key into the door reader (before, during, or soon after installation), SA ensures that no one else can install secret keys into the door. Alternatively, a control PIN or key
15 is needed for storing other secret keys into a door reader. The door reader is delivered without any control PINs or keys, and is manufactured so that once the first control PIN or key (or possibly a set of them) is inserted, then the reader stores it for a long time, and accepts no other control PINs or keys in the future. However, provided the right control PIN/key is input, then any new key could be inserted and
20 stored into the reader. This way, by being the first one to insert any control PIN/key into the door reader (before, during, or soon after installation), SA ensures that no one else can insert and store a secret key into a door reader.

At this point the SA knows all secret keys of the reader of a door D: for
25 instance, SKAD, SKBD, SKCD, etc. Rather than implementing Kerberos, it might be simpler that SA now gives SKAD to authority A, SKBD to authority B, etc. At this point, authority A/B/... can control users U access to D by either a private-key encryption method or a private-key signature method. Notice that these authorities may operate independently different sets of doors. For instance, assume that
30 1. door D1 has secret key SKXD1 inside its reader, and SA gives SKXD1 to authority X;

2. door D2 has secret key SKXD2 inside its reader, and SA gives SKYD2 to authority y; while

3. SA gives no key of door D1 to Y and no key of door D2 to X.

Then, authority X may control access to door D1 and authority Y may control door

5 D2 in a totally independent manner.

A Better Solution

But even with the above features available we can improve systems such as above in some important respects. Namely:

10

Key-storage size. While it is preferable that a door reader stores different keys for each different organization controlling it, this drives up the number of keys that a reader should securely store.

15

Adding new control. New control issues may come up when a new authority or a new door is introduced in the system. If a door D does not store a key for organization X, and later on it is desired that X gains control over D, then SA must insert a key for X into D's reader. For instance, if a new organization comes up, then the SA must dispatch a team of workers to insert SKXD into every door D that should be control by the new organization. Such physical "tours," however, may be inconvenient. To avoid them, the SA may pre-install additional keys into a door D's reader, and then bind them to new organizations that arise, or to organization that later on must control access through D. This strategy, however, only exacerbates the point described in the first bullet. Furthermore, if a new door is introduced, to be controlled by some already existing authorities, then the SA will have to insert new keys in the door reader, and then deliver the proper secret keys to the already existing authorities that must control it. Though doable, delivering secret keys always is problematic.

30

Taking back control. Once a secret key SKXD is stored in door D and known to organization X, then X will continue to control access through D, even though at a certain point control over D should be exclusively given to different organizations.

To avoid this, SA should again engage into a physical tour and remove SKXD from door D (e.g., by means of a control PIN/key mechanism).

Let us now describe how to bring about these additional improvements.

5 Basic system outline

To begin with, we can have the system work with a single key per door. For instance, the SA stores in door D the single key SKD (and of course keeps track of this information). Such key could potentially be computed by SA deterministically from D's identifier and a secret seed s known only to SA: for instance,
 10 $SKD = H(s, D)$. The SA then gives control over D to authority X by giving X a key SKXD chosen deterministically from SKD and X, for instance as a pseudo-random function with seed SKD evaluated at X (for simplicity we assume throughout that an entity coincides with a proper identifier of it). In particular, we can have $SKXD = H(SKD, X)$. Authority X then uses SKXD to grant user U access to D for a
 15 time interval (e.g., day) d as before. In particular, by using SKXD as the signing key of a private-key signature scheme: for instance, by computing $SKXD \cup d = H(SKXD, U, d)$ and then causing $SKXD \cup d$ to be stored into U's card. When U's card communicates with D's reader, then the card provides the reader with (a) X and (b) $SKXD \cup d$ and possibly other information, such as d (as well as
 20 information about the user U). Upon receiving this information, the reader computes $H(SKD, X)$ and then uses the result (allegedly equal to SKXD) as the signing key of the same private-key signature scheme and signs (U, d) ---in the example above by hashing (U, d) after combining it with SKXD. If the result matches the value tended by the card (allegedly, $SKXD \cup d$), if the time interval is right relative to the reader
 25 clock (and if U entered the right PIN, if PINs are properly used within the above system), then the door opens.

Key Storage, Adding Control

Notice that this single-key-per-door system not only minimizes the key-storage requirements, but also vastly simplifies the problem of adding control. Any
 30 time that an authority X needs to gain for the first time control over a door D, the SA needs not physically reach D and insert (or facilitate X's inserting) a new D-X key

into D's reader. Rather, if D has a key SKD known to the SA, then the SA simply computes the D-X key from SKD (e.g., $SKXD = H(SKD, X)$) and the delivers such D-X key (e.g., electronically) to X.

Taking back control

- 5 For each door D and authority X which is entitled to control D for a time interval (e.g., day) d', the SA computes and makes available its signature of this fact. For instance, this signature may be a private-key signature relative to a key SKD that SA shares with door D. In particular, this signature could be the value $H(SKD, \text{valid}, X, d')$. Notice that even if though being a private-key signature, the
- 10 signature itself can be made public without worries. Indeed, using the H-based implementation of a private-key signature described above, if H is a secure one-way hash function, then computing SKD from $H(SKD, \text{valid}, X, d')$ is very hard. Thus, when user U picks up in his card the right door-control permissions of the day, he may pick up for door D not only SKXD_{Ud}, but also $H(SKD, \text{valid}, X, d')$. The reader
- 15 of door D may then verify SKXD_{Ud} as before, and additionally ascertain that X has indeed control over D for interval d' by hashing together SKD, valid, X and d' and check that the same value tended by the card is obtained, and check that according to its clock the current time is within d'. In fact, only SA (and D) know the secret signing key SKD: authority X only knows $H(SKD, X)$ and computing SKD from
- 20 $H(SKD, X)$ and $H(SKD, \text{valid}, X, d')$ is very hard. Notice that time intervals d and d' may not be the same. For instance, SA may be satisfied to grant control over D to X on a weekly basis, while X may grant access through D to users on a daily basis. Alternatively, the system may replace use of SKXD as above with a time-dependent version of that key: e.g., $SKXDd = H(SKD, X, d)$. Then SA will have to deliver
- 25 SKXDd to each authority X before the time period d. To take back control, SA simply stops sending SKXDd for the periods d for which SA decides to deny X control over door D.

- Notice too, that the system currently allows for some privacy, in that SA needs not
- 30 know which users U are given access by X to D, nor their number. The scheme can be, of course remove this privacy (e.g., reporting or by using a Kerberos system).

Example 14:

Let us now outline our preferred implementation for achieving secure physical access in a systems with a super authority SA, a multiplicity of (preferably disconnected) doors D, a multiplicity of organizations X, a multiplicity of users U.

- 5 The preferred embodiment minimizes key storage and makes it very easy to add and take back control of a door D to organization X.

In the preferred embodiment, SA grants organization X control over door D for a given time interval. During that time interval, X may itself grant a user U
10 access to D.

We envisage SA (and possibly other players) to take action at each of a sequence of dates d corresponding to a sequence of time intervals. For instance, d could be the beginning of a given day and the corresponding time interval the given
15 day. For simplicity, we may use d to mean both the date and the corresponding time interval. (It should be understood, however, that this is not a limitation: for instance, a date could be a given day, and the time interval corresponding to the date the following day.) For concreteness, but without limitation intended, we may assume that each date/time interval is a day.

20

We describe the preferred embodiment using a private-key digital signature. This is without any limitation intended. Our preferred embodiment should be considered implemented with any other private key system as described above. To be more concrete, we assume that the private-key signature is implemented using a
25 one-way hash function H. This is without limitation intended: $H(SK, DATA)$ should always be considered the digital signature with key SK of DATA.

We assume that SA shares a secret key SKD with door D. SA may also share a secret key SKX with organization X. (SKD could be generated by A via a master
30 secret key SK. Similarly for SKX. For instance, SKD could equal $H(SK, D)$ and SKX could equal $H(SK, X)$. SA may then privately —or via encryption—provide D with SKD. Similarly for X.)

At each day d , if the SA wishes to grant organization X access to door D , it computes and causes X to receive secret key $SKXDd$, that is a key securely bound to X, D , and day d that is verifiable by D (e.g., on inputs X and d).

5

For instance, $SKXDd = H(SKD, X, d)$, that is, SA signs X, d with key SKD . SA then causes X to receive $SKXDd$. SA may cause X to receive $SKXDd$ by sending $SKXDd$ to X , preferably after encrypting it with a secret key SKX shared with X . Preferably yet, SA sends the so encrypted $SKXDd$ to X by causing it to be stored in a responder, from which X then downloads it.

10

If X wishes to grant user U access to D in time interval t within day d , X computes and causes U to receive a secret key $SKXDdUt$, that is a key securely bound to X, D, U and t that is verifiable by D .

15

For instance, $SKXDdUt = H(SKXDd, U, t)$, that is, X signs U, t with key $SKXDd$. X then causes U to receive $SKXDdUt$. X may cause U to receive $SKXDdUt$ by sending $SKXDdUt$ to U , preferably after encrypting it with a secret key SKU shared with U . Preferably yet, X sends the so encrypted $SKXDdUt$ to U by causing it to be stored in a responder, from which U then downloads it.

20

If U wishes to access D at time interval t , U causes D to receive X, U, t (e.g., U 's card transfers the to D 's reader).

25

If D receives X, U, t at day d , it computes $SKXDd$ from its secret key SKD and then computes $SKXDdUt$ from $SKXDd$. D then verifies that time interval t is indeed within day d , and using its own clock that indeed the current time is within time interval t . Further, D verifies that it is dealing with U/U 's card by a challenge-response mechanism using key $SKXDdUt$. If these verifications are passed, D opens.

30

For instance, D may compute SKXDd from its secret key SKD by computing $H(SKD, X, d)$, and then compute SKXDdUt from SKXDd by computing $H(SKXDd, U, t)$. For instance, the challenge-response mechanism using key SKXDdUt may consist of having D send a random string q and receive back the
5 encryption of q with key SKXDdUt, or a digital signature of q with key SKXDdUt. Alternatively, D may send Eq, the encryption of q with key SKXDdUt, and must receive back q.

Notice that the preferred scheme should be understood to include using a
10 PIN in conjunction with the above. In particular, any PIN use described in prior sections may be used within the preferred scheme. Notice that the preferred system provides lot of flexibility in that d and t may differ. For instance SA may provide control over D to X for a week d, while X may grant U access to D for a day t within week d. However, we may have $d=t$, in which case t needs not be specified or
15 separately used within the preferred system.

Kerberos approach

Using Kerberos approach directly would not work very well in our secure access application. It is most natural to implement all the doors and the SA as one
20 realm (with SA acting as a Ticket Granting Service, TGS, for that realm). Each organization and its employees would then be a separate realm. The authority for each organization would then act as the Authentication Service, AS, for that realm (as well as possibly its own TGS). According to the Kerberos protocols, each user would then authenticate to the respective authority/AS obtaining a ticket-granting
25 ticket, TGT. This ticket TGT would then be sent by the user to the SA/TGS, along with the request for a service granting ticket for each of the doors the user is entitled to. The SA/TGS would then have to verify the user's eligibility and, if the user –if all is correct– provide these service-granting tickets. This protocol is obviously quite laborious, and places much of the burden on the SA. In particular, it will be SA's
30 responsibility to verify which doors the particular user is entitled to and issue the respective tickets. Moreover, it demands that SA be on-line and engage in the

protocols in real-time. Having the users a channel to the SA presents an extra security threat as well.

Kerberos tickets without protocols

In principle, we could “abandon” the Kerberos protocols and only use the tickets. Namely, all the tickets would be pre-ordered and pre-computed in advance, and the users would pick them up at the time of the main door entry, without engaging in the appropriate Kerberos protocols.

However, many of the above problems would remain – in particular, it would be natural for SA to delegate the control of certain doors to the particular authorities (but in such a way that this control could be easily taken back, possibly to be re-instated at a later point).

Utilizing RTCs within Kerberos

One way to help address this problem is to utilize *Real Time Credentials, RTCs*. For example, we could use the tickets as in the above approach. However, in this approach we may not generate the tickets on a daily basis. Instead, we may use long-range tickets, managing the short-range access controls via RTCs passed in the *Authorization-Data* field of the ticket.

The RTCs could work in this case exactly the same way as in the case of the public key certificates. However, some optimizations are possible here as well.

Utilizing RTCs as above brings a number of possible benefits. These include (but are not limited to):

1. Ease of management.
 - a. Now, SA must be involved relatively infrequently
 - b. Instead of relatively larger tickets, the users will need to pick up much smaller RTCs
 - c. Generating the RTCs can be delegated to the corresponding authority

- d. Taking control back is easy: This can be done in at least two ways. First, simpler and cruder – the tickets may not be renewed by the SA when they do run out. A more refined mechanism will utilize two kinds of RTCs: those issued by SA and those issued by the other authorities. Then each day SA would need to issue a single RTC per each authority, which remains (alternatively, it may have to issue an RTC for each Authority-Door pair, where the Authority is entitled to open the Door). Each authority will also issue an RTC per each user (alternatively, per each User-Door pair, where the User is entitled to open the Door). Note: more traditional Kerberos approach would require even more tickets to be generated and passed around in the on-line protocols.
- e. RTCs allow a clear separation of roles, facilitating many aspects of management and infrastructure.
2. Efficiency.
- a. Space: an RTC is much smaller than a corresponding ticket.
- b. Time: Because they are much shorter (and there are fewer of them and fewer numbers of communication rounds) the communication would be much faster, enabling the users to move through the doors while picking up the RTCs at a reasonable pace.
- c. Load distribution: RTCs can be distributed by non-secured responders. Replication of RTCs would also be neither expensive, nor dangerous.
3. Security.
- a. RTCs are not security-sensitive, once they are generated, and can be managed with greater ease (e.g., by unsecured responders) and without any threat to security.
- b. The separation of tickets and authorizations (via RTC) allows for a greater security in key management (when the keys/tickets are actually generated and communicated)
- c. SA isolation: SA never really needs to have a direct communication line with any of the users.

Beyond Kerberos

It can be observed, that the mechanisms above benefit fairly little from the core Kerberos features (this is largely due to the fact that Kerberos was designed for different applications). So, here we explore how we can utilize RTC-based mechanisms, which are not directly related to Kerberos. These mechanisms could be similar to the private key encryption and private key signature mechanisms above.

In these mechanisms, the special authority SA would share a secret with each organization A (B,C, ...) and with each door D. This can be done, for example, using methods as above so that SA needs to store only a single secret s. The secret shared between SA and A would then be $SKA = \text{Hash}(s, A)$. Similarly, a secret shared between SA and D is $SKD = \text{Hash}(s, D)$. Note, that both A and D also need to store only one secret: SKA or SKD, respectively. In addition, to each organization-door pair (A,D), corresponds an additional secret $SKAD = \text{Hash}(SKD, A)$. This secret can be easily computed by both SK and D. Giving SKAD to A can be necessary but possibly not sufficient for A to control access to the door. In addition, A may need to receive from SA (or from another party) an RTC for the current time period d. This RTC, termed RTCA_d, need not be secret and may certify that A is still in good standing with SA.

Each user U employed by A and entitled to enter the door D may then receive a key $SKAUD = \text{Hash}(SKAD, U)$ from A. Notice that SKAUD can be easily computed by both A and D without any additional secrets. Giving SKAUD to U may be necessary but possibly not sufficient for U to be able to open the door D. In addition, U may need a separate RTC for the current time period d: RTCA_{Ud}.

Notice that this approach has already dramatically simplified the information flows: in the beginning of each time period d, SA sends a single RTCA_d for each organization A. And each organization A sends a single RTCA_{Ud} for each user-door pair. All of these RTCs can be picked up by the employees upon entering the

main gate. Assuming, that a user U is entitled to entering up to 100 doors within the facility, the RTCAUdd's for all the doors could require less than 2KB – an amount manageable even by slow connections (typically, it would take a fraction of a second).

5

To open the door D the user U may need to present the RTCAd and RTCAUdd, as well as perform the authentication based on the secret SKAUD (this authentication may be of the challenge-response type to protect the secret). Notice: since a relatively small number of RTCAd credentials is likely to be present in the system, the validation of these credentials may not need to be done on a per user
10 basis. Instead, each door may validate each RTCAd it receives and cache the result, to be used for other users' validation.

The special authority SA may wish to exercise a finer grain control of the organizations' access to the doors. To achieve such, instead of the per organization
15 credential RTCAd, SA may issue an RTC per each organization door pair (A,D): RTCAUdd. Then it would be possible for SA to grant and take back control over each door by each organization on a daily basis. Note that this may at most double the amount of RTC data that each user would need to receive (still keeping the required
20 transition time for the above example at a fraction of a second).

Aggregate RTCs

One may observe that often the access control rights do not change dramatically from day to day. So, much of the power of the above mechanisms is not utilized. We propose an RTC aggregation mechanism, which can be utilized in such
25 relatively stable environments to increase efficiency even further.

Example 15:

Consider as an example, a case of 100 organizations each having access to 1,000 doors. Therefore, there are a 100,000 of organization-door pairs, and thus,
30 RTCAUdd credentials to be issued and distributed by SA every day. Moreover, if each organization employs around 1,000 people, this would lead to 100,000,000 RTCAUdd credentials to be issued and distributed by all the organizations.

Let us divide all the organization-user-door triplets AUD's into hierarchically arranged groups. It may be easy to visualize these for example as follows. Let all the AUD's correspond to the leaves of a balanced binary tree (ordered in a preferred fashion). Then each node n of the tree corresponds to a set of all the AUD's corresponding to the leaves in the subtree of n . To each such node and a time period d , let there correspond also a credential $RTCnd$. Then the validity of AUD triplet in the period d can be certified by any of the credentials $RTCnd$, for any of the AUD ancestors n . Thus, if all the AUD triplets remain valid on day d , then a single credential $RTCr$, where r is the root of the tree, is sufficient for the whole system.

In general, if there are 100 AUD triplets that become invalid, then at most 1,500 credentials are sufficient to certify the whole system (that is instead of 100,000,000). More generally, at most $k(26 - \lg k)$ credentials are needed for certification of the whole system if k triplets are invalid.

This method can lead to dramatic improvements even if the aggregate RTC's require more values to be stored in the doors and/or users: in the above example, such an overhead may result in at most a factor of 26 overhead in the storage, while saving orders of magnitude (four or five in the example above) in communication. More generally, if a set of all entities to be authorized (in our examples, these were AUD triplets) contains N elements, and k of these are to be excluded, then at most $k(\lg N - \lg k)$ credentials are needed to certify the whole system, while the overhead for the aggregation may be at most $\lg N$. Even more efficient representations of the groups exist in the literature (e.g., while the above is known as subset cover method, we may use also the subset-difference cover and some of the recent results on it)

so, validation of such aggregate credentials may be optimized, e.g., by caching the results at least for the larger groups.

RTC implementations and optimizations

Many different implementations for the real-time credentials are possible. These implementations of RTC's also allow many different optimizations. For example, a real-time credential can be implemented as follows: Let x_0 be a random value, e.g., 20 bytes long. Let x_i be defined as $x_i = \text{Hash}(x_i)$. Let x_n be a public value fixed in some way (e.g., communicated securely by from SA to door D. Then, x_{n-d} would be the real-time credential RTCd for the time period d. It can be verified by applying Hash() to x_{n-d} d times and verify that the result is equal to x_n . This is essentially how RTC's are implemented in the case of public key certificates – for example, there x_n can be included as part of the certificate.

It is possible to use essentially the same implementation here as well. Instead of including x_n inside the certificate, here we may include it as a part of the Kerberos ticket. Or, we might communicate it by some other secure way, such as encrypted with the secret key SKD for the door D, etc.

Another possible implementation of RTCd is simply to set it equal to $\text{Hash}(\text{SKD}, \text{RTC}, d)$, where RTC refers to the credential ID. For example, in order to enable organization A to have control over door D on day d, the credential RTCADd would be used, where RTCADd could be set to $\text{RTCADd} = \text{Hash}(\text{SKAD}, d)$. A credential for user U to access door D on day d, as issued by the organization A may be $\text{RTCAUDd} = \text{Hash}(\text{SKAD}, U, d)$. Such a method allows the credentials to be pre-issued for specific dates well in advance, and without granting access on any days outside the desired time periods (even if these are non-contiguous).

The validation of the above credentials is straightforward. Note, that the above credentials are essentially symmetric signatures with the appropriate keys. In all the above, encryption may be used in place of the Hash.

Notice that we have made the system ore and more efficient at each step. Consider an airport with 1,000 doors, 100 authorities, and 10,000 possible workers, and assume for simplicity that control is given on a daily basis. Then a

Kerberos/Needham-Schroeder system in which a central authority is involved in computing each door-user key must be involved in 100 Million secret keys per day. A system as outlined above, would require SA to generate and deliver to all the authorities less than 100,000 secret keys per day.

5

REAL TIME CREDENTIALS OVER OCSP

We now describe the use of a preferred embodiment of the present invention for Real Time Credential validation technology within an environment that uses the
10 Open Certificate Status Protocol (OCSP) for digital certificate validation. This shows how the inventive technology maintains compatibility with the OCSP standards while offering qualitatively superior security and scalability than traditional OCSP implementations.

15 Traditional OCSP Implementation

CRLs may grow big because they provide proofs of revocation (and thus, indirectly, of validity) about many certificates lumped together. By contrast, the OCSP provides proofs of validity of for individual certificate. OCSP services are typically implemented by OCSP *Responders*. Such a responder is a server that, upon
20 receiving a question from a client (aka Relying Party) about the validity of a given certificate issued by a given CA, provides a digitally signed answer indicating both the status of the certificate and the time of the answer. For doing this, it is necessary for the OCSP responder to know the status of all of the CA's certificates, since it is the CA that can revoke its own certificates. If the OCSP responder were the CA
25 itself, such knowledge is trivially acquired. Else, some other form of keeping the OCSP responder updated about the status of the CA's certificates must be employed. For instance (cfr. U.S. Patent No. 5,717,758, *Witness-Based Certificate Revocation System*), the CA may send the responder its most recent CRL, and the responder may consult that signed document to deduce whether the certificate of interest is
30 currently valid or revoked and so say in its signed response, also indicating the time, as well the time of the next update. (Here it is natural for this update time to coincide

with the date of the next CRL of the CA, since it is that CRL that may trigger a different response.)

Of course, a malicious responder may provide arbitrary signed answers about the certificates of a given CA, with or without consulting the latter's CRLs. For the
5 relying party to securely rely on the digitally signed answer of a OCSF responder about the certificates of a given CA, the OCSF envisages that the CA providing the responder with a responder certificate, a special digital certificate --signed by the CA-- that essentially proves to other parties that the CA trusts the responder to
10 provide accurate proofs about its certificates.

Notice that for this process to work, each OCSF responder (as well as every CA) must have a secret signing key, and this key must be protected (ideally by placing it or the server using it in a vault).

15

FIG. 2 shows this sequence of transactions in a trivial OCSF environment. The fact that secret signing keys are protected is graphically emphasized by putting them with some thick "borders." In case of a signed data, the name of the signer is indicated immediately below. This figure shows the various PKI-sensitive elements
20 of this transaction as shaded boxes. The Certificate Authority itself has a private key, SK1, that must be kept secure to prevent the unauthorized issuance and revocation of certificates. This key is used to sign the CRL that is published to the OCSF Responders. The secret key of responder 1A must also be kept secure, and is used for signing the OCSF responses of responder 1A.

25

Drawback of OCSF

DRAWBACK 1: Computation

Digital signatures are computationally intensive operations. The digital
30 signature created by the Responder on each response is generated at the time of the request, and is by far the most computationally intensive part of the validation

operation: it can easily add anywhere from 50 milliseconds to 1 second to the transaction time.

Even if a responder cached its digital signature about a digital certificate C
5 and then sent the same signature when asked about C until the next update, still the answer to the first user asking about C will be significantly delayed.

DRAWBACK 2: Communication (with centralized implementations)

10 Assume a single validation server implements the OCSP in a centralized manner. Then, all certificate-validity queries would have, eventually, to be routed to it, and the server will be a major "network bottleneck" causing considerable congestion and delays, as shown in FIG. 3. If huge numbers of honest users suddenly query the server, a disrupting "denial of service" will probably ensue.

15

DRAWBACK 3: Security (if distributed implementations)

To prevent the bottleneck problems that centralized OCSP implementations may cause, a CA may consider distributing the request load generated by its
20 certificates by distributing it across several OCSP servers (that it properly certifies). In general, distributing the load of a single server across several (e.g., 100) servers, strategically located around the world, alleviates network congestion. In the OCSP case, however, load distribution introduces worse problems than those it solves. In order to sign its responses to the certificate queries it receives, each of the 100
25 servers should have its own secret signing key. Thus, compromising any of the 100 servers would effectively compromise the entire system.

If a traditional OCSP Responder were compromised, an attacker could do one of three things. First, it could prevent the Responder from issuing any responses.
30 This type of attack is detectable at the Relying Party, and thus not too severe. Second, it could use the discovered secret signing key to sign responses indicating that legitimate certificates are revoked. Third, and most disruptively, it could make

the Responder generate signed responses indicating that a revoked certificate is still valid. This type of false-positive response could allow a terminated employee to regain access to systems, etc.

5 The best way to prevent that a responder could be compromised is to run it from a secure vault, with 24X7 surveillance, etc. Unfortunately, this is a costly option. A truly secure vault, meeting all the requirements needed for a financial CA, may cost over \$1M to build and \$1M/year to operate. Even if one were willing to pick up such expenses, vaults cannot be built overnight: armored concrete does not
10 scale! If a CA needed a few more vaults to lessen the load of its current responders, it may have to wait months before a new one could be constructed.

Moreover, even if several expensive vaults were in place, they may still not be secure. This is so because the OCSP mechanism requires that a responder receive
15 requests coming from un-trusted sources (the clients on the field) and then service them using its secret signing key. The possibility thus exists that a malicious agents prefer to exploit any weakness in the underlying operating system and thus expose the secret signing key to drilling holes night time through an armored-concrete wall. In sum, if no vaults or a sufficiently expensive perimeter protected a responder, the
20 probability of a compromise is very high, but even if a truly secure building housed a responder, a responder would still be vulnerable to a software attack: to a sophisticated digital enemy, the OCSP mechanism makes a vault look much like a bunker with a "window."

25 DRAWBACK 4: *Trust Flow*

OCSP has difficulties in servicing certificate validity requests originating from different security domains. In the scenario shown in FIG. 4, the Responder run by organization #1 is able to provide responses about the status of certificates from
30 CA #1, but Responders run by another organization may not have enough information to provide responses about the "foreign" certificates. For instance

Responder 2A, run by certification authority CA 2, does not know how to answer requests about CA 1's certificates.

5 This problem, deriving from lack of specific knowledge, could be addressed in one of two ways.

First, the Relying Parties from organization #2 could find the Responders from organization #1 to ask them about the status of certificates from CA #1. This limits performance however, since the Responders from organization #1 may be geographically distant from Relying Parties interested in organization #2, so network times may greatly slow overall validation processing.

15 The second alternative is to allow Responders from organization #2 to make responses about certificates from organization #1, by having CA #1 forward its CRLs also to "foreign" responders. This indeed poses no security threats, because CRLs are digitally signed, and because a CA wishes to inform the largest possible audience about the validity of its own certificates. This provides sufficient information to a Responder of organization #2 for answering a request from a Relying party about a certificate of CA 1. But for the Relying Party to take Responder 2A's digitally signed answer really seriously, CA 1 should also certify Responder 2A as trustworthy for answering validity queries about its own certificates. The whole process is illustrated by FIG. 5.

25 This approach provides better scalability and performance, but it muddies the security and trust flow between the two organizations. In the example above, Responder #2A is making an authoritative response to the Relying Party that the certificate #321 of CA #1 is still good. Making an incorrect response for any reason (misconfiguration, hostile attack, or straightforward dishonesty), Responder 2A may cause adverse consequences for users from organization #1. By allowing Responder #2A to make authoritative claims about its own certificates, organization #1 is relinquishing some of the trust that it previously held.

As an example, consider the case where the organizations are credit card issuers. Bank #1 revokes the card certificate for user #321, and it pays to ensure that its Responders are secure and reliable. The Responders from Bank #2 are misconfigured, so when a merchant Relying Party asks about the validity of user #321, they incorrectly respond that the user is valid. The merchant accepts this answer and allows a transaction to proceed for the revoked user.

This type of delegation-of-trust between organizations may be acceptable in some cases, but it is not a generally useful alternative for any large-scale heterogeneous deployment of traditional OCSP.

Real Time Credentials over OCSP

In light of the above problems, we wish to put forward an alternative certificate validation system, Real Time Credentials (RTC), that while keeping compatibility with current OCSP standards, solves all the described drawbacks of traditional OCSP. RTC technology differs from traditional OCSP in that:

1. It *does not delegate trust* to foreign Responders;
2. It *centralizes all validation trust* into a single authority (*the RTC Authority*); yet,
3. It *distributes the query load* from this single authority across an arbitrary number of *unprotected* responders;
4. It *does not decrease security* even in distributed implementations relying on thousands of Responders (and even though these responders are unprotected!);
5. It *improves dramatically the response time* to a query.

This provides a radical improvement over traditional OCSP in terms of security, performance, scalability, and heterogeneity.

The RTC System comprises the following steps:

The CA certifies the RTCA: The new system is centered around the RTC authority (RTCA). This is an entity that may or may not coincide with the CA of a given organization. Preferably, each CA provides its own RTC with a special certificate, the RTCA certificate. The CA preferably digitally signs this certificate, indicating that it trusts and indeed empowers the RTCA to provide certificate validity information about its own certificates. Such a certificate may bind a given verification key PK (for which the RTCA possesses a corresponding secret signing key) to the RTC authority (e.g., identified by a given identifier, OID number) and specify in some fashion that the certificate essentially confers RTC status, and may include other traditional certificate information and formats. In case the two entities coincide, it may still be advantageous for them to have distinct signing keys, so that, in effect, in any case the CA only *issues* certificates and the RTC authority only *manages* them (i.e., proves them valid or revoked). This being the case, even if the CA and the RTCA coincide, an RTCA certificate may still be employed. Preferably each CA has only one RTC, though for redundancy purposes, it may be advantageous to have more than one, whether or not using the same signing key.

The RTCA protects its signing key: The RTCA must protect its signing key, for instance by means of a vault or secure facility. (As we shall see, however, there is no need of additional vaults for certificate validation purposes!) The RTCA may host in the same protected facility more than one server embedding its secret signing key, or securely store (e.g., in Banks' safe security boxes) copies of the key, or host more than one server each having a secret signing key properly certified by the CA.

The CA informs the RTCA of the status of its certificates. For instance, it keeps it apprised on any change in certificate validity in an on-line/real-time fashion (such as sending a message informing the RTCA of a change in certificate status as soon as it occurs). Alternatively, it may send the RTCA its CRLs when produced.

The RTCA individually signs the validity status of each certificate for a given interval of time, independent of any request: Preferably periodically (or at any date

of a sequence of dates), the RTCA, based on its current validation knowledge (e.g., based on the latest CRL of the CA) and *independent* of any Relying Party request, processes each outstanding certificate of its CA, and digitally signs a declaration stating the status of that certificate. The result therefore carries a time component
5 indicating the next update for that certificate. If the period of the RTC depends on the CA-issued CRLs, the update time may be that of the next CRL. The time component may also indicate the issuance time of the CRL used in the processing. In essence, therefore, the RTCA pre-computes a digital signature indicating the status of each certificate for a given time interval T (e.g., from the date of the latest
10 CRL ---or from a date sufficiently close to it--- to the date of the next CRL ---or to a date sufficiently close to it, in either case so as to allow time sufficient from processing all the necessary information). Such pre-computation is performed independent of any relying party request about the certificates. Indeed, preferably the RTCA pre-computes all such signed declaration of certificate status before any
15 queries about certificate status in that time interval are made, or before that time interval altogether. In particular, the RTCA may pre-compute all its signed declarations about time interval T one minute before T starts. The fact that by so doing it is not going to be "synchronized" with the CRL (in case it is used) is not too serious. The CRL itself is not real time, and information about certificate revocation
20 and indeed the very reason for which a certificate has been revoked may take considerably more time. For instance, a user may realize that his secret key has been compromised and thus request that his own certificate be revoked one day after the fact. Thus in any case the certificate has been revoked with a one day delay. Preferably, the RTCA signed declarations of certificate validity are in standard
25 OCSF format. That is, in essence, the RTCA preferably pre-computing OCSF-compliant responses to OCSF requests that have not yet been generated. This is important because OCSF software is already in place, and it would be very convenient to take advantage of the RTC system without having to modify any of the existing relying party software.

30

The RTCA sends his pre-computed signatures of validity status to unprotected responder: After pre-computing such a signature, the RTCA makes it

available (e.g., sends it to) to other parties, including relying parties (e.g., in response to requests of theirs), but, in particular, to responders. These responders need not be protected. In fact they handle RTCA-signed messages, and these cannot in essence be fraudulently modified or altered in an undetectable way. Indeed, the

5 RTCA may easily send them to foreign responders (responders belonging to other organizations) with any jeopardizing security. The RTCA may facilitate the responder processing of its signatures by presenting them to the responder in a suitably organized fashion. For instance, it may present its signed certificate validity status ordered accordingly to the certificate serial number, or in an array, or ensuring

10 that each signed piece of data has the same or suitably closed length, etc. To ensure that all the relevant pre-computed responses have been received, the RTCA may sign and date the totality of its responses (e.g., all those relative to the same time interval and CA).

15 In addition, an RTCA preferably sends to its responders its own RTCA certificate. This transmission needs not occur at every update. In particular can be performed only initially.

The responders store the RTCA-pre-computed signatures: A responder

20 stores the received pre-computed signatures of the RTCA for a sufficient time. Preferably, if these signatures relate to a given time interval T, they store them at least until the end of T. Preferably too, the responders (especially those belonging to the same organization as the RTCA) may be pro-active and check that they received the proper RTCA signatures correctly and on time. For instance, a responder may:

- 25 (1) Verify that the pre-computed responses about a time interval T are received by the beginning of T (or other suitable time related T);
- (2) Verify the received RTCA signatures (and possibly also the proper RTCA certificate);
- (3) Verify whether it has received all signatures (e.g., less than the
- 30 expected number of signatures, less signatures than at last transmission, etc.)

(4) Verify whether it has received a RTCA-signed declaration of validity for a certificate that was previously declared revoked; etc.

If any problem is detected, it may inform the RTCA or another proper entity.

5 *Relying parties ask responders for validity status information:* Relying parties ask responders for the validity status of certificates. Preferably, they do so using the OCSF format for their requests.

10 *Responders answer queries with pre-computed responses:* When asked about the validity of a given certificate, the responder fetches from memory the RTCA pre-computed answer for that certificate and returns it.

A responder may also forward the proper certificate for the RTCA that has signed the pre-computed response

15 *Relying parties verify the pre-computed answers (and RTCA certificates):* Relying parties process the receive responses to ascertain the validity status of the certificate of interest. Preferably, if the response is in OCSF format, they use OCSF software for such processing. Preferably too they verify the proper RTCA
20 certificates.

Throughout this application, it is understood that certificates may be hierarchical certificates and that proofs of the currently validity of CA certificates and CRTA certificates may be added and verified whenever needed.

25 FIG. 6 illustrates the RTC System.

Advantages of the RTC System

30 The RTCA periodically generates digitally signed validity declarations (proofs, since such declarations cannot be forged) for all current certificates of the CA, and then distributes them to any interested responders. (Each proof is preferably

structured as a syntactically correct OCSP response, signed by the RTCA private key.) When a relying party asks about the status of a certificate, the RTC responder is able to return the corresponding pre-generated response which it has cached. The relying party can verify the signature of the RTCA. (In addition, it can also verify
5 the RTCA's certificate, to ensure tat it is dealing with an authentic RTC authority for the given CA. Of course, this like all other certificates can be hierarchical.)

ADVANTAGE 1: *Computation*

Digital signatures are computationally intensive operations. But the RTC
10 system concentrates this difficulty on a single server (entity): the RTCA. It is therefore very easy and relatively inexpensive to equip with this single entity with a computer sufficiently powerful to handle all required digital signatures. By contrast, the RTC responders perform only trivial computations. They essentially (1) store the RTCA signatures and (2) perform just fetch-and-forward operations in response to
15 relying parties queries. Therefore they can be implemented with very inexpensive hardware. As a result, the total RTC cost may be significantly lower than that of the OCSP! At the same time, response time is much quicker. Indeed, the time for a very inexpensive RTC responder for fetching and sending a pre-computed RTCA response is negligible relative to that taken by an OCSP responder which must
20 perform a digital signature in response to a relying party request.

ADVANTAGE 2: *Communication*

In the RTC system, responders may employ trivial hardware and do not need to be secure. Consequently RTC responders are very cheap indeed, and can be
25 deployed in great numbers. That is, one can always afford distributed implementations of RTC system. Therefore, even if enormously many certificate-validity requests are generated in a short amount of time, this load can always be spread across many RTC responders, eliminating the risk of congestion and benign denial of service without incurring much cost. (Notice that the amount of work of
30 the RTCA solely depends on the number of certificates and is not affected by the number of validity-status requests. Thus a single RCA can be used even if millions and millions of validity requested are expected.)

ADVANTAGE 3: *Security*

In the RTC system only the RTCA (besides the CA, if it is a different/differently located entity) be protected. In fact the responders do not store
5 any secret key: they only store the digital signatures of the RTCA, but for all security purposes may be made totally public after being computed by the RTCA. By contrast, each OSCP responder has a secret signing key, compromising which one may compromise the entire system. Therefore defending a single site is preferable and easier than defending many and equally important sites.

10

Moreover, unlike in the OSCP, relying parties cannot easily mount software attacks. In fact, the RTC responders service relying parties' requests with non-secret information. In fact they do not have any secret keys themselves and need only store pre-computed digital signatures. Thus, even if a relying party succeeded in
15 embedding in its query some kind of Trojan horse, it would be able to expose nothing. At most it can expose all a RTC responder knows, and that is the full and accurate account of which certificates are valid and which are revoked in a given time interval. And this not only is non-secret information, but it is even information that a certification authority would like to be universally known, so that no one may
20 rely incorrectly on one of its certificates!

Finally, notice that software attacks cannot be easily mounted against the RTCA either. In fact, though possessing a secret signing key, *the RTCA does not process requests of untrusted sources*. This is so because the RTCA does not
25 answers any untrusted requests: it simply receives inputs from the CA (a very trusted source!) and periodically *outputs* data (signed validity statements). Therefore the very ability to inject a Trojan horse is missing in the RTC system! In other words, not only a single vault may be sufficient in the RTC system, but this vault has no "windows" whatsoever.

30

ADVANTAGE 4: *Trust Flow*

In addition to these advantages, the RTC-over-OCSP approach enables significant flexibility within heterogeneous PKI deployments involving multiple organizations. The following diagram shows how RTC-over-OCSP would be deployed in a cross-CA environment.

5

FIG. 7 shows how a responder from organization #2 can relay (preferably, OCSP compliant) responses from organization #1 without needing to transfer any trust from organization #1 to responders of organization #2. Since RTC responders are simple, non-trusted relays of information, they can be widely distributed and mirrored without reducing overall system security. A relying party queries a responder of organization 2 (Responder 2B) about the validity of a certificate of organization #1. Notice that the (preferably OCSP compliant) response that it gets back is convincing because it is digitally signed by an RTCA of organization #1 (RTCA1). Further, the direct digital signature from the right organization (which is best positioned to know which of its own certificates are still valid, and which has the greatest interest in not making mistakes) is preferably corroborated by the fact that the relying party also gets RTCA1's certificate (preferably signed by CA1) that vouches that RTCA1 is indeed a proper RTC authority of organization 1.

20

In sum, organization #1 enables the responders of organization #2 to provide convincing proofs of validity for organization #1's certificates without relinquishing any amount of control over the validity status of its own certificates. That is, *in the RTC system trust may flow from one organization to another with any associated loss of neither security nor control.*

25

ADVANTAGE 5: *Secure heterogeneity*

FIG. 7 shows the extreme case, where Responders are treated as transparent network infrastructure rather than hardened trust points. It shows how the extreme case of RTC enabling the secure construction of a heterogeneous cloud of Responders that are capable of servicing requests about the status of certificates from many sources. This is similar to the service cloud offered by the Internet's

30

DNS infrastructure, in that it allows for a heterogeneous collection of name servers that transparently interoperate to discover and cache valid responses for queries.

This heterogeneity is a significant advantage of the RTC system over
 5 traditional OCSP. It allows a wide variety of organizations to interoperate so that relying parties from different organizations can cross-validate certificates from other organizations in a secure, reliable, efficient manner.

Real Time Credentials (RTC) is a cost-effective, secure, scalable, and overall
 10 efficient certificate validation system. RTC can (1) provide an alternative to the Open Certificate Status Protocol (OCSP), as well as (2) work within and enhance the OCSP. RTC systems, in fact, even when exercising the option of maintaining compatibility with the OCSP standards, provide significant advantages over the OCSP, so as to offer qualitatively superior security and scalability.

15

RTC Optimizations

2-Party Versus 3-Party Certificate Validation

Let U be a party having a certificate Cu. As part of a transaction with a party
 V, U may send Cu to V (unless V already has it), and possibly perform additional
 20 tasks (such as exhibiting a digital signature relative to a public verification key certified in Cu to belong to U, or being identified by decrypting a random challenge encrypted by V using a public encryption key certified in Cu to belong to U). For the transaction to be secure, V might ascertain the current validity of Cu and make a validity query to a RTC responder. The responder would answer this query by
 25 fetching and returning the most current RTCA-signed declaration about Cu. However, querying an RTC responder *makes 3-party a transaction that would otherwise be 2-party*, increasing the time of the desired U-V transaction.

Thanks to its predictable time intervals, RTC may significantly help.
 30 Namely, party U may, at the beginning of each time interval T (or during it anyway), receive an RTCA-signed declaration Du that Cu is valid throughout T. U can receive Du in response to a request to his (e.g., by making an ordinary relying-party

request) or may be pushed Du (e.g., by an RTC responder or by an RTCA at every update on an automatic basis). In either case, transacting with V during interval T, U may forward Du to V, in addition to all other steps or tasks the transaction entails. Therefore, the U-V transaction is significantly sped up, since V needs not call any
5 third party in order to ascertain the current validity of U's certificate.

Though, in some sense, the "overall time," which includes U obtaining Du, may not be sped up, the U-V transaction will be. Notice that speeding up only the U-V transaction without saving in overall time, may still be quite valuable. In fact,
10 assume RTCA declarations are computed at midnight and specify an entire day as their time interval. Then, U may obtain Du early in the day (when no real pressure exists) and then forward it to V during a time sensitive U-V transaction conducted during working hours, when saving time could be essential. Further efficiency is gained, if U, after obtaining and caching DU, forwards it throughout the day when
15 transacting with several (e.g., 100) parties. This way, for instance, a single relying-party query (that of U itself, possibly made at a relaxed time) successfully replaces 100 relying-party requests (possibly at times of pressure).

Notice that this optimization can also be achieved by the parties V. Namely,
20 after obtaining a response Du from a RTC responder in return to a query about the validity of a certificate Cu of party U, party V can give to U, or make Du available for others to use.

This optimization too applies to the preferred, OCSF-compliant
25 implementations of RTC. Actually, we suggest applying a similar optimization also to traditional OCSF implementations. Namely, a user requests and obtains an OCSF response about his own certificate, and then forwards this OCSF response as part of his transactions to the other parties of the transactions for the appropriate time interval. Alternatively, when asked for the first time by a relying party about the
30 validity of a certificate Cu of party U, an OCSF responder computes its response Ru, returns it to the querying relying party, but also forwards it to U, so that U can cache it and, at least for a while, can forward it as part of its transactions based on Cu.

Certificate-Helped Validation

Notice that the RTC system may be implemented using data found in the individual certificates, thereby saving additional certificates and/or response length.

- 5 As we have seen, the CA may issue an RTCA certificate that empowers a given RTCA to provide authoritative answers about the validity of its own certificates. Such an RTCA certificate ideally specifies the public key that must be used for verifying the RTCA-signed responses. The CA may however, embed this RTCA public key within its own certificates. That is, the CA (with proper format, OID, etc.) may include in a certificate Cu also the public key PK that should be used for
10 verifying the digitally signed responses about Cu's validity. This way, a relying party needs not receive a separate RTCA certificate. When asking an RTC responder for the latest proof of validity for Cu, it may just obtain (e.g., because it so asks) only the RTCA-signed response. In fact, Cu specifies within itself the public
15 verification key that a relying party may use for verifying a proof of validity for Cu. This may yield significant savings in transmission (since the RTC responder may not need to send a separate RTCA certificate, which may be much longer than an RTCA response) and in storage (since the relying party may not need to store the RTCA certificate alongside with the RTCA response, as protection against future
20 claims for having relied on Cu).

Similarly, a certificate Cu may specify its own time intervals. In this case, an RTCA response may not need to specify both the beginning and end of an interval T. In fact, the beginning of T alone (or other simpler specification) may pin-down T.
25 For instance, if Cu specifies daily updates, then any time within a given day suffices to specify the entire day to which a response refers. Alternatively, if it is clear (e.g., from the CA's general policies) that the certificates have validity intervals consisting of a full day, then there is no need for this information to be specified within a certificate, and yet the same savings in RTCA responses apply.

30

Separate Revocation

While an RTC proof of validity or suspension for a given certificate C should specify the time interval to which it refers, a proof of revocation needs not specify any time interval: it suffices for it to specify a single point in time (e.g., the actual time of revocation). Unlike validity and suspension, in fact, revocation traditionally
5 is an irrevocable process. Thus a single revocation time t_r may suffice for proving a certificate revoked. And t_r needs not be the beginning of any time interval T (e.g., it could be any time in "the middle of T"). In case of permanent revocation, therefore, the RTCA needs not send C's revocation proof at all updates dates (e.g., D1, D2, etc.). In principle, a revocation proof could be sent only once (or a few times for
10 redundancy) and then cached by an RTC responder and then returned whenever a relying-party query about C is made.

Notice also that the RTCA may be informed right away that a certificate C has been revoked; for instance, in the middle of a time interval T for which the
15 RTCA has already produced and forwarded a proof of validity for C to the RTC responders. Of course, by the next update, no such proof of validity will be computed for C. But for the time being (i.e., until the end of T) an incorrect proof of validity is out there. Thus, a good counter-measure consists of having proofs of revocation take precedence over proofs of validity. That is, an honest relying party
20 that sees both a proof of validity for C for some time interval T and a proof of revocation for C (at whatever time t), should regard C as revoked (after time t). However, some relying parties may have never seen such a proof of revocation, and thus C may considered by some still valid until the end of T. As we have seen, such problems are somewhat unavoidable, in the sense that even in the traditional OCSP,
25 the news of the revocation of C may take some time to reach the responder, and it may take even longer to realize that C should be revoked. Nonetheless, these problems can be lessened by having the RTCA compute and send all RTC responders a proof of C's revocation (independent of the scheduled dates D1, D2, etc. or D1', D2', etc.) as soon as it learns that it has been revoked (e.g., directly from
30 the CA without waiting the next CRL update). All properly functioning RTC responders will then erase from memory any proof of C's validity and substitute it

with the newly received proof of revocation. This way, from that time on, they will provide relying parties with accurate proofs about C's validity.

System Generality

5 A CA/RTCA/ responder/party/user may be any entity (e.g., person, organization, server, device, computer program, computer file) or a collection of entities.

10 Certificates should be construed to include all kinds of certificates, and in particular hierarchical certificates and flat certificates (cfr. U.S. Patent No. 5,420,927 herein incorporated by reference). Validity status and proofs of validity status may include validity status and proofs of validity status for hierarchical certificates (e.g., validity status and proofs of validity status of all certificates in a chain of certificates). Verifying the validity of a certificate C may include verifying
15 the validity of the CA certificate for the CA having issued C, as well as the validity of the CRTA certificate for the RTCA that providing a signed response about the validity status of C.

20 Though certificates traditionally are digitally signed document binding given keys to given users, following U.S. Patent No. 5,666,416 (herein incorporated by reference), certificates should include all kinds of digitally signed documents. For instance, a vendor, acting as a CA, may certify a price lists of its by digitally signing it (possibly together with date information). Validity status for such certificates is also very crucial. For instance, a vendor may want to prove the current validity of a
25 price list (and refuse honor a given price in a price lists, unless a proof of its currently validity is shown). Thus a customer may wish to ascertain the current validity of a price list document. In particular, the RTC system is ideal (for its scalability and off-line processing) for proving the current validity of web pages. Indeed, the RTCA generated proofs of current validity may be stored next (or in association with) the pages themselves. (In this case, then, a party can be considered
30 a computer file.)

Sending a piece of data D (to party X) should be construed to include making D available (or causing X to receive D).

Three-Factor Authentication With Real-Time Validation

5

The following is an efficient three-factor authentication with real-time validation and revocation performed with no connecting infrastructure at the relying party. This can work for physical access applications such as a door or logical applications such as file or application access. A physical access scenario is described below. Other applications are easy to generalize from this model for those skilled in the art.

Example 16:

1. The user has a credential stored on a wireless device (physical token). This token preferably has the capability of securely storing a digital certificate and private key. Preferably too, the token has a method of long-range (WAN) connectivity (such as GPRS, SMS, pager, CDMA, GSM, etc.) and a method of short-range (PAN) connectivity (such as Bluetooth, IR, RF, etc.) The token may also have one or more additional authentication factors (such as a keypad for a PIN or a biometric reader). This example assumes that the token is a bluetooth cell phone.

2. The door has a control panel with a small CPU capable of performing standard PKI operations and a method of short-range (PAN) connectivity to talk to the physical token. This example assumes a bluetooth-enabled computer similar to our standard demo doors.

3. The user is prompted to enter a PIN number into her cell phone (or enter his own biometric info if a biometric reader is available). This prompt can happen once a day, the first time the user tries to go through a door, every few hours, randomly, upon receipt of a special SMS message, etc. The PIN (or biometric) serves as a second factor of authentication (first being the certificate on the phone) and "unlocks" the phone for use in the physical access application.

4. Once the user comes within range of the door (30 ft for bluetooth), the phone and the door recognize each other and begin the initial authentication and validation sequence:

5

4.1 (optional) The door validates itself to the phone by sending the door's certificate via bluetooth to the phone. The phone checks the certificate and validates the door using any of our standard methods (min-CRL of all doors periodically sent down to the phone is a good approach.) This solves the problem of "rogue readers" and makes sure that the door is a legitimate reader before the phone discloses any information.

4.2 The phone sends the door the user's certificate which contains the user's biometric minutiae. The phone also sends an RTC proof (preferably, either Validation-token ---i.e., a 20-byte proof of validity--- or a Distributed-OCSP proof) to prove its current validity. The proof had been previously received via the WAN in the normal CoreStreet manner, such as that described in U.S. Patent No.5,666,416, Issued September 9, 1997, entitled "Certificate Revocation System".

4.3 The door authenticates and validates the user's certificate in the normal RTC fashion. The door may do this for multiple (or even all) phones currently within range (multiple employees may be near the door).

5. By the time the user reaches the door, the previous steps have been completed. The user scans her finger (or other biometric) on a reader mounted on or near the door (perhaps in the actual doorknob). The door matches the biometric minutiae against the data stored in all validated certificates within range. If the biometric matches, the door opens. Otherwise, the door remains closed.

30 This has the following benefits:

1. Strong authentication (3-factor in this example, more are possible)

2. Transparent to the user (just walk up to the door and open it, no cards or PIN numbers to remember)
3. Real time revocation and validation
4. No connecting infrastructure required at any door - do this at 30,000 feet or in the middle of the ocean
5. Can be built with standard hardware and software components

Step 4.1 is an independent invention of independent interest, since it solves a known problem (eg. identified by the Department of Defense) for which there is no currently known solution. The scheme may be augmented by having "revocation proofs or access logs travel to and/or from other people's cards/phones to disconnected doors".

PROTECTING MOBILE COMPUTING RESOURCES

15

A preferred embodiment of the present invention is based on 20 byte, unforgeable, public "proofs". 20-byte proofs are cryptographically protected using a one-way function called hashing. The process is simple, does not need encryption and does not use digital signatures. These properties make this technology ideal for: large scale deployments (scales to 100s millions); bandwidth limited applications (e.g. wireless applications); offline validation (i.e., network connection not required).

Laptop theft is a serious problem that imposes replacement costs, loss of productivity, loss of unrecoverable (unbacked-up) data, loss of control over sensitive/confidential data (e.g. sensitive operational info, proposals to clients, email, calendar, contacts, pending mergers, new product IP, strategies, and launch plans, financial operating results, private compensation info.), and loss of network and infrastructure details (e.g. user names & passwords, dial-in numbers, IP addressing schemes, DNS naming conventions, and primary mail serves).

2003228468 03 Aug 2009

- 109 -

Some embodiments provide for leases, that is licenses to use for a specified period of time wherein the duration of the lease is a configurable parameter. Such technology forces presence of valid "leases". Leases are 20 byte, unforgeable, "public tokens": valid token, suspension token, and revocation token. New leases are received automatically. A computer may be temporarily disabled and a system administrator or user can unsuspend a laptop. A computer may be permanent disabled with possible recovery by the System Administrator. FIG. 8 is a schematic illustration of the system operation according to some embodiments.

As long as the device is still authorized, a valid lease token is produced once a day (hour, week etc.) by the central authority. Getting a valid lease token onto the protected device can be accomplished in many different ways and is completely transparent to the end user. If the device is stolen, two things happen: valid lease tokens cease to be generated (no way to extend use past the current day); revocation token is propagated to the network (any connection renders device immediately unusable). Stolen devices are turned off within: seconds (best case, if push capability is present); hours (average cast, as soon as any network connection is made); one day (worst case, no connection possible).

The system protects against random theft as well as theft by insiders. Stealing a device makes no sense, since: the hardware is unusable; the software is unusable; and the data is unreadable. Similar to some card radio brands, unusable if stolen and therefore deters theft.

Validity tokens are delivered by the following methods: wired network; wireless network; SMS wireless "push"; pager system; handheld telephone/PDA via infrared port; Bluetooth device; Manually typed in as received via alternate channel (e.g. "7G9L TC77 U8QL S2PS QK2Q EN9V PXXH XPUL"), such as via fax, e-mail, telephone call. FIG. 9 is a schematic illustration of a stolen computer timeline.

Alternative protection methods may be used including: physical anchor for prevention; asset tracking service for recover and as a deterrent; motion sensor and alarm

03 Aug 2009

2003228468

- 110 -

as deterrent; access keys as a deterrent and access control; tracking software for recover and as a deterrent; and data encryption which protects data only.

Potential attacks and results include:

- 5 Removing/circumventing software: Possible if have "administrative privileges" but extremely difficult after revocation. Optional BIOS/hardware countermeasures offer nearly 100% protection.

Replacement/reformat hard drive: All secure data lost and optional BIPS/hardware hooks to prevent drive replacement.

- 10 Move hard drive to another machine to read data: Data can be encrypted.

Prevent Receipt of revocation token: Prolongs laptop operation until lease expires only (worst case).

- 15 Other embodiments of the invention will be apparent to those skilled in the art from a consideration of the specification or practice of the invention disclosed herein. It is intended that the specification and examples be considered as exemplary only and non-limiting.

- 20 Throughout this specification and claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

- 25 The reference in this specification to any prior publication (or information derived from it), or to any matter which is known, is not, and should not be taken as an acknowledgment or admission or any form of suggestion that that prior publication (or information derived from it) or known matter forms part of the common general knowledge in the field of endeavour to which this specification relates.

20032228468 03 Aug 2009

- 111 -

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method for an entity A to control the access of at least one user U to at least one disconnected door D having means to determine the current time, comprising the steps of:
 - for each time interval d of a sequence of dates, having A produce a digital signature SIGUDd, indicating that user U can access door D during time interval d;
 - causing a card of the user U to receive SIGUDd during time interval d for presentation to door D in order to pass through D;
 - having U present the card with SIGUDd to door D; and
 - having D open after verifying that (i) SIGUDd indeed is a digital signature of A indicating that U can access door D at time interval d, and (ii) that the current time is indeed within time interval d, wherein the door D stores, on the card of the user, information corresponding to an access attempt of a different user.
2. The method of claim 1, wherein D has a card reader coupled with an electromechanical lock, and wherein U presents SIGUDd to D by having his card read by D's card reader.
3. The method of claim 1 or claim 2, wherein A causes SIGUDd to be received by U during time interval d by posting SIGUDd into a database accessible by U.
4. The method of any one of claims 1 to 3, wherein SIGUDd is a public-key signature, and wherein D stores the public-key of A.
5. The method of any one of claims 1 to 4, wherein D also verifies identity information about U.
6. The method of claim 5, wherein the identity information about U includes at least one of: a PIN and the answer to a challenge of D.

2003228468 03 Aug 2009

- 112 -

7. The method of any one of claims 1 to 6, further comprising:
receiving the information stored on the card at a database disconnected from the door D.
8. A method for controlling the access to at least one disconnected door, comprising:
for each time interval of a sequence of dates, causing an entity to produce a digital signature, wherein the digital signature indicates that at least one user can access the disconnected door during the time interval;
causing a card of a first user to receive the digital signature during the time interval for presentation to the disconnected door in order to pass through the disconnected door; and
after the first user presents the card with the digital signature to the disconnected door, causing the disconnected door to open after verifying that: (i) the digital signature is a digital signature of the entity indicating that the first user can access the disconnected door at the time interval, and (ii) that a current time is within the time interval, wherein the disconnected door stores, on the card of the first user, information that corresponds to a second user different from the first user.
9. The method of claim 8, wherein the disconnected door has a card reader coupled with an electromechanical lock, and wherein the first user presents the digital signature to the disconnected door by having the card of the first user read by the card reader.
10. The method of claim 8 or claim 9, wherein the entity causes the digital signature to be received by the first user during the time interval by posting the digital signature into a database accessible by the first user.
11. The method of any one of claims 8 to 10, wherein the digital signature is a public-key signature, and wherein the disconnected door stores the public-key of the entity.

20032228468 03 Aug 2009

- 113 -

12. The method of any one of claims 8 to 11, further comprising:
receiving the information stored on the card at a database disconnected from the door.
13. The method of any one of claims 8 to 12, wherein the disconnected door also verifies identity information about the first user.
14. The method of claim 13, wherein the identity information about the first user includes at least one of: a PIN and the answer to a challenge of the disconnected door.
15. A method substantially as hereinbefore described with reference to the accompanying drawings.
16. Computer-readable storage storing executable instructions to enable performance of the method of any one of claims 1 to 15.
17. Apparatus comprising means for performing the method of any one of claims 1 to 15.
18. A system comprising means for performing the method of any one of claims 1 to 15.
19. Apparatus substantially as hereinbefore described with reference to the accompanying drawings.
20. A system substantially as hereinbefore described with reference to the accompanying drawings.




FIG. 1

2 / 9




FIG. 2

3 / 9




FIG. 3




FIG. 4

5/9




FIG. 5




FIG. 6




FIG. 7




FIG. 8

