

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
3 February 2005 (03.02.2005)

PCT

(10) International Publication Number
WO 2005/010685 A2

- (51) International Patent Classification⁷: **G06F**
- (21) International Application Number:
PCT/US2004/022810
- (22) International Filing Date: 16 July 2004 (16.07.2004)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/488,645 18 July 2003 (18.07.2003) US
60/505,640 24 September 2003 (24.09.2003) US
- (71) Applicant (for all designated States except US): **CORE-STREET, LTD.** [US/US]; One Alewife Center, Suite 200, Cambridge, MA 02140 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **LIBIN, Phil** [US/US]; 57 Clay Street, Cambridge, MA 02140 (US). **MICALI, Silvio** [US/US]; 459 Chestnut Hill Avenue, Brookline, MA 02146 (US). **ENGBERG, David** [US/US]; 38 Dover Street, Cambridge, MA 02140 (US).
- (74) Agent: **MUIRHEAD, Donald, W.**; Choate, Hall & Stewart, Exchange Place, 53 State Street, Boston, MA 02109-2804 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— without international search report and to be republished upon receipt of that report
- 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.



WO 2005/010685 A2

(54) Title: CONTROLLING ACCESS TO AN AREA

(57) Abstract: Controlling access includes providing a barrier to access that includes a controller that selectively allows access, at least one administration entity generating credentials/proofs, wherein no valid proofs are determinable given only the credentials and values for expired proofs, the controller receiving the credentials/proofs, the controller determining if access is presently authorized, and, if access is presently authorized, the controller allowing access. The credentials/proofs may be in one part or may be in separate parts. There may be a first administration entity that generates the credentials and other administration entities that generate proofs. The first administration entity may also generate proofs or the first administration entity may not generate proofs. The credentials may correspond to a digital certificate that includes a final value that is a result of applying a one way function to a first one of the proofs.

CONTROLLING ACCESS TO AN AREACross-Reference to Related Applications

This application claims priority to U.S. provisional patent application no. 60/488,645 filed on July 18, 2003, which is incorporated by reference herein, and
5 claims priority to U.S. provisional patent application no. 60/505,640 filed on September 24, 2003, which is incorporated by reference herein, and is a continuation-in-part of U.S. patent application no. 10/876,275 filed on June 24, 2004 (pending) which claims priority to U.S. provisional patent application no. 60/482,179 filed on June 24, 2003, and which is a continuation-in-part of U.S. patent application no. 09/915,180 filed on
10 July 25, 2001 (pending), which is a continuation of U.S. patent application no. 09/483,125 filed January 14, 2000 (now U.S. Patent No. 6,292,893), which is a continuation of U.S. patent application no. 09/356,745 filed July 19, 1999 (abandoned), which is a 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.
15 08/559,533 filed November 16, 1995 (now U.S. Patent No. 5,666,416) which claims priority to U.S. provisional patent application no. 60/006,038 filed on October 24, 1995, and is a continuation-in-part of U.S. patent application no. 10/409,638, filed on April 8, 2003 (pending) which claims priority to U.S. provisional patent application no. 60/370,867, filed April 8, 2002, U.S. provisional patent application no. 60/372,951,
20 filed April 16, 2002, U.S. provisional patent application no. 60/373,218, filed April 17, 2002, U.S. provisional patent application no. 60/374,861, filed April 23, 2002, U.S. provisional patent application no. 60/420,795, filed October 23, 2002, U.S. provisional patent application no. 60/421,197, filed October 25, 2002, U.S. provisional patent application no. 60/421,756, filed October 28, 2002, U.S. provisional patent application
25 no. 60/422,416, filed October 30, 2002, U.S. provisional patent application no. 60/427,504, filed November 19, 2002, U.S. provisional patent application no. 60/443,407, filed January 29, 2003, and U.S. provisional patent application no. 60/446,149, filed February 10, 2003; the teachings of all of which are incorporated herein by reference. And which is a continuation-in-part of U.S. patent application no.
30 10/103,541, filed March 20, 2002 (pending), 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 (pending), and which is a continuation of U.S. patent application no. 09/483,125, filed January 14, 2000, (now U.S. Patent No. 6,292,893), which is a continuation of U.S. patent application no. 09/356,745, filed July 19, 1999, (abandoned), which is a continuation of U.S. patent application no. 08/823,354, filed
5 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 application no. 10/103,541 is also a continuation-in-part of U.S. patent application no. 08/992,897, filed December 18, 1997, (now U.S. Patent
10 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 (abandoned), which is based on U.S. provisional patent application no. 60/004,796, filed October 2, 1995. U.S. patent application no. 08/992,897 is also a continuation-in-part of U.S. patent application no.
15 08/729,619, filed October 11, 1996, (now U.S. Patent No. 6,097,811), which is based on U.S. provisional application no. 60/006,143, filed November 2, 1995. 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 (abandoned), which is a continuation of U.S. patent application no. 08/741,601, filed November 1, 1996 (abandoned), which is based
20 on U.S. provisional patent application no. 60/006,143, filed November 2, 1995. 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 (abandoned), which is a continuation of U.S. patent application no. 08/746,007, filed November 5, 1996 (now U.S. Patent No. 5,793,868), which is based on U.S. provisional patent application no. 60/025,128, filed
25 August 29, 1996. U.S. patent application no. 08/992,897 is also based on U.S. provisional patent application no. 60/035,119, filed February 3, 1997, and is also a continuation-in-part of U.S. patent application no. 08/906,464, filed August 5, 1997 (abandoned), which is a continuation-in-part of U.S. patent application no. 08/763,536, filed December 9, 1996 (now U.S. Patent No. 5,717,758), which is based on U.S.
30 provisional patent application no. 60/024,786, filed September 10, 1996, and is a continuation-in-part of U.S. patent application no. 08/636,854, filed April 23, 1996, (now U.S. Patent No. 5,604,804), and is also based on U.S. provisional patent

application no. 60/025,128, filed August 29, 1996. 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 (abandoned), which is based on U.S. provisional patent application no. 60/025,128, filed August 29, 1996, and is a continuation-in-part of U.S. patent application no. 08/715,712, filed September 19, 1996 (abandoned), and is also a continuation-in-part of 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. 08/992,897 is also a continuation-in-part of U.S. patent application no. 08/752,223, filed November 19, 1996 (now U.S. Patent No. 5,717,757), which is based on U.S. provisional patent application no. 60/025,128, filed August 29, 1996, and is also a continuation-in-part of U.S. patent application no. 08/804,869, filed February 24, 1997 (abandoned), which is a continuation of U.S. patent application no. 08/741,601, filed November 1, 1996 (abandoned), which is based on U.S. provisional patent application no. 60/006,143, filed November 2, 1995. 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 (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. U.S. patent application no. 10/103,541 is also based on U.S. provisional patent application no. 60/277,244, filed March 20, 2001, and U.S. provisional patent application no. 60/300,621, filed June 25, 2001, and U.S. provisional patent application no. 60/344,245, filed December 27, 2001. All of the above are incorporated herein by reference. U.S. patent application no. 10/409,638 is also a continuation-in-part of U.S. patent application no. 09/915,180, filed June 25, 2001 (pending), 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 (now U.S. Patent No. 6,292,893), which is a continuation of U.S. patent application no. 09/356,745, filed July 19, 1999, (abandoned), which is a 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 application no. 60/006,038, filed October 24, 1995. The teachings of all of the above are incorporated herein by

reference. U.S. patent application no. 10/409,638 is also a continuation-in-part of U.S. patent application no. 10/395,017, filed March 21, 2003 (pending), 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 (abandoned), which is a
5 continuation of U.S. 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 (abandoned), which is based on U.S. provisional patent application no. 60/004,796, filed on October 2, 1995,
10 and which is also a continuation-in-part of U.S. patent application no. 08/729,619, filed October 10, 1996 (now U.S. Patent No. 6,097,811), which is based on U.S. provisional patent application no. 60/006,143, filed November 2, 1995, and which is also a continuation-in-part of U.S. patent application no. 08/804,868, filed February 24, 1997 (abandoned), which is a continuation of U.S. patent application no. 08/741,601, filed
15 November 1, 1996 (abandoned), which is based on U.S. provisional patent application no. 60/006,143, filed November 2, 1995, and which is also a continuation-in-part of U.S. patent application no. 08/872,900, filed June 11, 1997 (abandoned), which is a continuation of U.S. patent application no. 08/746,007 filed November 5, 1996 (Now U.S. Patent No. 5,793,868), which is based on U.S. provisional patent application
20 no. 60/025,128, filed August 29, 1996, and which is also based on U.S. provisional patent application no. 60/035,119, filed February 3, 1997, and which is also a continuation-in-part of U.S. patent application no. 08/906,464, filed August 5, 1997(abandoned), which is a continuation of U.S. patent application no. 08/763,536 filed December 9, 1996 (now U.S. Patent No. 5,717,758), which is based on U.S.
25 provisional patent application no. 60/024,786, filed September 10, 1996, and is also a continuation of U.S. patent application no. 08/636,854, filed April 23, 1997, (now U.S. Patent No. 5,604,804), and U.S. provisional patent application no. 60/025,128, filed August 29, 1996, and which is also a continuation-in-part of U.S. patent application no. 08/756,720, filed November 26, 1996(abandoned), which is based on U.S.
30 provisional patent application no. 60/025,128, filed August 29, 1996, and is also a continuation-in-part of U.S. patent application no. 08/715,712, filed September 19, 1996 (abandoned), and is also a continuation-in-part of 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 (now U.S. Patent No. 5,717,757), which is based on U.S. provisional patent application no. 60/025,128, filed August 29, 1996, and is also a
5 continuation-in-part of U.S. patent application no. 08/804,869, filed February 24, 1997 (abandoned), which is a continuation of U.S. patent application no. 08/741,601, filed November 1, 1996 (abandoned), which is based on U.S. provisional patent application no. 60/006,143, filed November 2, 1995, and which is also a continuation-in-part of U.S. patent application no. 08/823,354 filed March 24, 1997 (now U.S. Patent No.
10 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. The teachings of all of the above are incorporated herein by reference.

Background of the Invention

1. Technical Field

This application relates to the field of physical access control, and more particularly to the field of physical access control using processor actuated locks and related data.

2. Description of Related Art

Ensuring that only authorized individuals can access protected areas and devices may be important in many instances, such as in the case of access to an airport, military installation, office building, etc. Traditional doors and walls may be used for protection of sensitive areas, but doors with traditional locks and keys may be cumbersome to manage in a setting with many users. For instance, once an employee is fired, it may be difficult to retrieve the physical keys the former employee was issued while employed. Moreover, there may be a danger that copies of such keys were made and never surrendered.

Smart doors provide access control. In some instances, a smart door may be equipped with a key pad through which a user enters his/her PIN or password. The key pad may have an attached memory and/or elementary processor in which a list of valid PINs/passwords may be stored. Thus, a door may check whether the currently entered PIN belongs to the currently valid list. If so, the door may open. Otherwise, the door may remain locked. Of course, rather than (solely) relying on traditional keys or simple key pads, a more modern smart door may work with cards (such as smart cards and magnetic-strip cards) or contactless devices (e.g., PDA's, cell phones, etc.). Such cards or devices may be used in addition to or instead of traditional keys or electronic key pads. Such magnetic-strip cards, smart cards or contactless devices, designed to be carried by users, may have the capability of storing information that is transmitted to the doors. More advanced cards may also have the ability of computing and communicating. Corresponding devices on the doors may be able to read information

from the cards, and perhaps engage in interactive protocols with the cards, communicate with computers, etc.

An aspect of a door is its connectivity level. A fully connected door is one that is at all times connected with some database (or other computer system). For instance, 5 the database may contain information about the currently valid cards, users, PINs, etc. In some instances, to prevent an enemy from altering the information flowing to the door, such connection is secured (e.g., by running the wire from the door to the database within a steel pipe). On the other hand, a totally disconnected door does not communicate outside of its immediate vicinity. In between these two extremes, there 10 may be doors that have intermittent connectivity (e.g., a wirelessly connected "moving" door that can communicate with the outside only when within range of a ground station, such as the door of an airplane or a truck).

Traditional access control mechanisms suffer from many drawbacks. Fully connected doors may be very expensive. The cost of running a secure pipe to a distant 15 smart door may vastly exceed the cost of the smart door itself. Protecting a wire cryptographically, while possibly cheaper, still has its own costs (e.g., those of protecting and managing cryptographic keys). Moreover, cryptography without steel pipes and security guards cannot prevent a wire from being cut, in which case the no-longer-connected door may be forced to choose between two extreme alternatives: 20 namely, remaining always closed or always open, neither of which may be desirable. In any case, fully connecting a door is often not a viable option. (For instance, the door of a cargo container below sea level in the middle of the Atlantic Ocean is for all practical purposes totally disconnected.)

Disconnected smart doors may be cheaper than connected doors. However, 25 traditional approaches to smart doors have their own problem. Consider, for instance, a disconnected smart door capable of recognizing a PIN. A terminated employee may no longer be authorized to go through that door; yet, if he still remembers his own PIN, he may have no trouble opening such an elementary smart door. Therefore, it would be

necessary to “deprogram” the PINs of terminated employees, which is difficult for disconnected doors. Indeed, such a procedure 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 is
5 terminated may be too impractical.

It is desirable to provide a level of security associated with fully connected doors without incurring the additional costs thereof. As demonstrated, disconnected smart doors and cards do not by themselves guarantee the security, convenience and low cost of the access-control system.

10 Summary of the Invention

According to the present invention, controlling access includes providing a barrier to access that includes a controller that selectively allows access, at least one administration entity generating credentials/proofs, wherein no valid proofs are determinable given only the credentials and values for expired proofs, the controller
15 receiving the credentials/proofs, the controller determining if access is presently authorized, and, if access is presently authorized, the controller allowing access. The credentials/proofs may be in one part or may be in separate parts. There may be a first administration entity that generates the credentials and other administration entities that generate proofs. The first administration entity may also generate proofs or the first
20 administration entity may not generate proofs. The credentials may correspond to a digital certificate that includes a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The digital certificate may include an identifier for the electronic device. The credentials may include a final value that is a
25 result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The credentials may include an identifier for a user requesting access. The credentials/proofs may include a digital signature. The barrier to access may include walls and a door. Controlling access may also include providing a door lock coupled to
30 the controller, wherein the controller allowing access includes the controller actuating

the door lock to allow the door to open. Controlling access may also include providing a reader coupled to the controller, wherein the controller receives credentials/proofs from the reader. The credentials/proofs may be provided on a smart card presented by a user. Controlling access may also include providing an external connection to the controller. The external connection may be intermittent. The controller may receive at least a portion of the credentials/proofs using the external connection or the controller may receive all of the credentials/proofs using the external connection. Controlling access may also include providing a reader coupled to the controller, where the controller receives a remaining portion of the credentials/proofs from the reader. The credentials/proofs may be provided on a smart card presented by a user. The credentials/proofs may include a password entered by a user. The credentials/proofs may include user biometric information. The credentials/proofs may include a handwritten signature. The credentials/proofs may include a secret value provided on a card held by a user. The credentials/proofs may expire at a predetermined time.

According further to the present invention, an entity controlling access of a plurality of users to at least one disconnected door includes mapping the plurality of users to a group, for each time interval d of a sequence of dates, having an authority produce a digital signature $SIGUDd$, indicating that members of the group can access door during time interval d , causing at least one of the members of the group to receive $SIGUDd$ during time interval d for presentation to the door in order to pass therethrough, having the at least one member of the group present $SIGUDd$ to the door D , and having the door open after verifying that (i) $SIGUDd$ is a digital signature of the authority indicating that members of the group can access the door at time interval d , and (ii) that the current time is within time interval d . The at least one member of the group may have a user card and the door may have a card reader coupled to an electromechanical lock, and the at least one member of the group may receive $SIGUDd$ by storing it into the user card, and may present $SIGUDd$ to the door by having the user card read by the card reader. The authority may cause $SIGUDd$ to be received by the at least one member of the group during time interval d by posting $SIGUDd$ into a database accessible by the at least one member of the group. $SIGUDd$ may be a public-key signature, and the door may store the public-key of the authority. The door may

also verify identity information about the at least one member of the group. The identity information about the at least one member of the group may include at least one of: a PIN and the answer to a challenge of the door.

According further to the present invention, controlling physical access also
5 includes assigning real time credentials to a group of users, reviewing the 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 first part and comparing the result to the second part;
10 and allowing physical access to members of the group only if 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. The second part may be provided by an entity other than the authority. The real time credentials may be provided on a smart card. Members of the group may obtain the second part of the real time credentials at a first
15 location. Members of the group may be allowed access to a second location different and separate from the first location. At least a portion of the first part of the real time credentials may represent a one-way hash applied a 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.
20 Controlling physical access may also include controlling access through a door.

According further to the present invention, determining access includes determining if particular credentials/proofs indicate that access is allowed, determining if there is additional data associated with the credentials/proofs, wherein the additional data is separate from the credentials/proofs, and, if the particular credentials/proofs
25 indicate that access is allowed and if there is additional data associated with the particular credentials/proofs, then deciding whether to deny access according to information provided by the additional data. The credentials/proofs may be in one part or in separate parts. There may be a first administration entity that generates the credentials and other administration entities that generate proofs. The first
30 administration entity may also generate proofs or may not generate proofs. The

credentials may correspond to a digital certificate that includes a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The digital certificate may include an identifier for the electronic device. The credentials
5 may include a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The credentials may include an identifier for a user requesting access. The credentials/proofs may include a digital signature. Access may be access to an area enclosed by walls and a door. Determining access may include providing a
10 door lock, wherein the door lock is actuated according to whether access is being denied. Determining access may also include providing a reader that receives credentials/proofs. The credentials/proofs may be provided on a smart card presented by a user. The credentials/proofs may include a password entered by a user. The credentials/proofs may include user biometric information. The credentials/proofs may
15 include a handwritten signature. The credentials/proofs may include a secret value provided on a card held by a user. The credentials/proofs may expire at a predetermined time. The additional data may be digitally signed. The additional data may be a message that is bound to the credentials/proofs. The message may identify the particular credentials/proofs and include an indication of whether the particular
20 credentials/proofs have been revoked. The indication may be the empty string. The additional data may include a date. The additional data may be a message containing information about the particular credentials/proofs and containing information about one or more other credentials/proofs. Determining access may also include storing the additional data. The additional data may include an expiration time indicating how
25 long the additional data is to be saved. The expiration time may correspond to an expiration of the particular credentials/proofs. Determining access may also include storing the additional data for a predetermined amount of time. Credentials/proofs may all expire after the predetermined amount of time. The additional data may be provided using a smart card. The smart card may be presented by a user attempting to gain
30 access to an area. Access to the area may be restricted using walls and at least one door. The additional data may be for a user different from the user attempting to gain access. Determining access may also include providing a communication link and

transmitting the additional data using the communication link. The communication link may be provided the additional data by a smart card. The smart card may require periodic communication with the communication link in order to remain operative. The smart card may be provided with the additional data by another smart card. The additional data may be selectively provided to a subset of smart cards. Determining access may also include providing a priority level to the additional data. The additional data may be selectively provided to a subset of smart cards according to the priority level provided to the additional data. The additional data may be randomly provided to a subset of smart cards.

10 According further to the present invention, issuing and disseminating a data about a credential includes having an entity issue authenticated data indicating that the credential has been revoked, causing the authenticated data to be stored in a first card of a first user, utilizing the first card for transferring the authenticated data to a first door, having the first door store information about the authenticated data, and having the first door rely on information about the authenticated data to deny access to the credential. 15 The authenticated data may be authenticated by a digital signature and the first door may verify the digital signature. The digital signature may be a public-key digital signature. The public key for the digital signature may be associated with the credential. The digital signature may be a private-key digital signature. The credential and the first card may both belong to the first user. The credential may be stored in a second card different from the first card, and the first door may rely on information about the authenticated data by retrieving such information from storage. The credential may belong to a second user different from the first user. The authenticated data may be first stored in at least one other card different from the first card and the authenticated data may be transferred from the at least one other card to the first card. 25 The authenticated data may be transferred from the at least one other card to the first card by first being transferred to at least one other door different from the first door. The entity may cause the authenticated data to be stored in the first card by first causing the authenticated data to be stored on a responder and then having the first card obtain the authenticated data from the responder. The responder may be unprotected. The first door may receive information about the authenticated data from the first card by 30

the authenticated data first being transferred to at least one other card different from the first card. The at least one other card may receive information about the authenticated data from the first card by the authenticated data first being transferred to at least one other door different from the first door. The first door may be totally disconnected or
5 may be intermittently connected.

According further to the present invention, a first door receives authenticated data about a credential of a first user, the process including receiving the authenticated data from a first card belonging to a second user different than the first user, storing information about the authenticated data, receiving the credential, and relying on the
10 stored information about the authenticated data to deny access to the credential. The authenticated data may be authenticated by a digital signature and the first door verifies the digital signature. The digital signature may be a public-key digital signature. The public key for the digital signature may be associated with the credential. The digital signature may be a private-key digital signature. The authenticated data may be stored
15 in the first card by being first stored in at least one other card and then transferred from the at least one other card to the first card. The authenticated data may be transferred from the at least one other card to the first card by first being transferred to at least one door different from the first door. The authenticated data may be stored in the first card by first being stored on a responder and then obtained by the first card from the
20 responder. The responder may be unprotected. The first door may receive information about the authenticated data from the first card by the authenticated data first being transferred to at least one other card different from the first card. The at least one other card may receive information about the authenticated data from the first card by the authenticated data first being transferred to at least one other door different from the
25 first door. The first door may be totally disconnected or may be intermittently connected.

According further to the present invention, assisting in an immediate revocation of access includes receiving authenticated data about a credential, storing information about the authenticated data on a first card, and causing a first door to receive
30 information about the authenticated data. The authenticated data may be authenticated

by a digital signature. The digital signature may be a public-key digital signature. The public key for the digital signature may be associated with the credential. The digital signature may be a private-key digital signature. The credential and the card may both belong to a first user. The first card may become unusable for access if the first card fails to receive a prespecified type of signal in a prespecified amount of time. The credential may belong to an other user different from the first user. The authenticated data may be received by the first card by being first stored in at least one other card different from the first card and then transferred from the at least one other card to the first card. The authenticated data may be transferred from the at least one other card to the first card by first being transferred to at least one other door different from the first door. The first card may obtain the authenticated data from a responder. The responder may be unprotected. The first card may cause the first door to receive information about the authenticated data by first transferring the authenticated data to at least one other card different from the first card. The first card may cause the at least one other card to receive information about the authenticated data by first transferring the authenticated data to at least one other door different from the first door. The first door may be totally disconnected or may be intermittently connected. The first card may eventually remove the stored information about the authenticated data from storage. The credential may have an expiration date, and first card may remove the stored information about the authenticated data from storage after the credential expires. The expiration date of the credential may be inferred from information specified within the credential.

According further to the present invention, logging events associated with accessing an area includes recording an event associated with accessing the area to provide an event recording and authenticating at least the event recording to provide an authenticated recording. Recording an event may include recording a time of the event. Recording an event may include recording a type of event. The event may be an attempt to access the area. Recording an event may include recording credentials/proofs used in connection with the attempt to access the area. Recording an event may include recording a result of the attempt. Recording an event may include recording the existence of data other than the credentials/proofs indicating that access

should be denied. Recording an event may include recording additional data related to the area. Authenticating the recording may include digitally signing the recording. Authenticating at least the event recording may include authenticating the event recording and authenticating other event recordings to provide a single authenticated recording. The single authenticated recording may be stored on a card. The authenticated recording may be stored on a card. The card may have an other authenticated recording stored thereon. The other authenticated recording may be provided by the card in connection with the card being used to access the area. Access may be denied if the other authenticated recording may not be verified. A controller may be provided in connection with accessing the area and where the controller further authenticates the other authenticated recording. The other authenticated recording may be authenticated using a digital certificate. Logging events may also include a user presenting a card to attempt to access the area. Logging events may also include the card further authenticating the authenticated recording in connection with the user attempting to access the area. A controller may be provided in connection with accessing the area and wherein the controller and the card together further authenticate the authenticated recording. Logging events may include providing correlation generation data that indicates the contents of the authenticated recording. The correlation generation data may be bound to the authenticated recording. The correlation generation data may be bound to the authenticated recording and the resulting binding may be authenticated. The resulting binding may be digitally signed. The correlation generation data may be a sequence of numbers and a particular one of the numbers may be assigned to the event. Logging events may also include authenticating a binding of the particular number and the event. Authenticating the binding may include digitally signing the binding. Authenticating the binding may include one way hashing the binding and then digitally signing the result thereof. Correlation generation data for the event may include information identifying an other event. The other event may be a previous event. The other event may be a future event. Logging events may also include associating a first and second random value for the event, associating at least one of the first and second random values with the other event, and binding at least one of the first and second values to the other event. Providing correlation generation data may include using a polynomial to generate the

correlation information. Providing correlation generation data may include using a hash chain to generate the correlation information. The correlation generation data may include information about a plurality of other events. The correlation generation data may include error correction codes. Logging events may also include disseminating the authenticated recording. Disseminating the authenticated recording may include providing the authenticated recording on cards presented by users attempting to access the area. The area may be defined by walls and a door.

According further to the present invention, at least one administration entity controls access to an electronic device by the at least one administration entity generating credentials and a plurality of corresponding proofs for the electronic device, wherein no valid proofs are determinable given only the credentials and values for expired proofs, the electronic device receiving the credentials, if access is authorized at a particular time, the electronic device receiving a proof corresponding to the particular time, and the electronic device confirming the proof using the credentials. The at least one administration entity may generate proofs after generating the credentials. A single administration entity may generate the credentials and generate the proofs. There may be a first administration entity that generates the credentials and other administration entities that generate proofs. The first administration entity may also generate proofs or may not. The credentials may be a digital certificate that includes a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to of a future one of the proofs. The digital certificate may include an identifier for the electronic device. The credentials may include a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The credentials may include an identifier for the electronic device. The electronic device may be a computer, which may boot up only if access is authorized. The electronic device may be a disk drive. At least one administration entity controlling access to an electronic device may include providing proofs using at least one proof distribution entity separate from the at least one administrative entity. There may be a single proof distribution entity or a plurality of proof distribution entities. At least one administration entity controlling access to an electronic device

may include providing proofs using a connection to the electronic device. The connection may be the Internet. At least some of the proofs may be stored locally on the electronic device. At least one administration entity controlling access to an electronic device may include, if the proof corresponding to the time is not available
5 locally, the electronic device requesting the proofs via an external connection. Each of the proofs may be associated with a particular time interval. After a particular time interval associated with a particular one of the proofs has passed, the electronic device may receive a new proof. The time interval may be one day.

According further to the present invention, an electronic device controls access
10 thereto by receiving credentials and at least one of a plurality of corresponding proofs for the electronic device, wherein no valid proofs are determinable given only the credentials and values for expired proofs and testing the at least one of a plurality of proofs using the credentials. The credentials may be a digital certificate that includes a final value that is a result of applying a one way function to a first one of the proofs.
15 Each of the proofs may be a result of applying a one way function to a future one of the proofs. The digital certificate may include an identifier for the electronic device. The credentials may include a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The credentials may include an identifier for the
20 electronic device. The electronic device may be a computer. An electronic device controlling access thereto may also include the computer booting up only if access is authorized. The electronic device may be a disk drive. An electronic device controlling access thereto may also include obtaining proofs using a connection to the electronic device. The connection may be the Internet. At least some of the proofs
25 may be stored locally on the electronic device. An electronic device controlling access thereto may also include, if the proof corresponding to the time is not available locally, the electronic device requesting the proofs via an external connection. Each of the proofs may be associated with a particular time interval. After a particular time interval associated with a particular one of the proofs has passed, the electronic device may
30 receive a new proof. The time interval may be one day.

According further to the present invention, controlling access to an electronic device includes providing credentials to the electronic device and, if access is allowed at a particular time, providing a proof to the electronic device corresponding to the particular time, wherein the proof is not determinable given only the credentials and values for expired proofs. The credentials may be a digital certificate that includes a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The digital certificate may include an identifier for the electronic device. The credentials may include a final value that is a result of applying a one way function to a first one of the proofs. Each of the proofs may be a result of applying a one way function to a future one of the proofs. The credentials may include an identifier for the electronic device. The electronic device may be a computer. Controlling access to an electronic device may include the computer booting up only if access is authorized. The electronic device may be a disk drive. Controlling access to an electronic device may include providing proofs using at least one proof distribution entity separate from the at least one administrative entity. There may be a single proof distribution entity. There may be a plurality of proof distribution entities. Controlling access to an electronic device may include providing proofs using a connection to the electronic device. The connection may be the Internet. At least some of the proofs may be stored locally on the electronic device. Controlling access to an electronic device may include, if the proof corresponding to the time is not available locally, the electronic device requesting the proofs via an external connection. Each of the proofs may be associated with a particular time interval. After a particular time interval associated with a particular one of the proofs has passed, the electronic device may receive a new proof. The time interval may be one day.

Brief Description of Drawings

Figure 1A is a diagram illustrating an embodiment that includes a connection, a plurality of electronic devices, an administration entity, and a proof distribution entity according to the system described herein.

Figure 1B is a diagram illustrating an alternative embodiment that includes a connection, a plurality of electronic devices, an administration entity, and a proof distribution entity according to the system described herein.

5 Figure 1C is a diagram illustrating an alternative embodiment that includes a connection, a plurality of electronic devices, an administration entity, and a proof distribution entity according to the system described herein.

Figure 1D is a diagram illustrating an alternative embodiment that includes a connection, a plurality of electronic devices, an administration entity, and a proof distribution entity according to the system described herein.

10 Figure 2 is a diagram showing an electronic device in more detail according to the system described herein.

Figure 3 is a flow chart illustrating steps performed in connection with an electronic device determining whether to perform validation according to the system described herein.

15 Figure 4 is a flow chart illustrating steps performed in connection with performing validation according to the system described herein.

Figure 5 is a flow chart illustrating steps performed in connection with generating credentials according to the system described herein.

20 Figure 6 is a flow chart illustrating steps performed in connection with checking proofs against credentials according to the system described herein.

Figure 7 is a diagram illustrating a system that includes an area in which physical access thereto is to be restricted according to the system described herein.

Detailed Description of Various Embodiments

Referring to Figure 1A, a diagram 20 illustrates a general connection 22 having a plurality of electronic devices 24-26 coupled thereto. Although the diagram 20 shows three electronic devices 24-26, the system described herein may work with any number of electronic devices. The connection 22 may be implemented by a direct electronic data connection, a connection through telephone lines, a LAN, a WAN, the Internet, a virtual private network, or any other mechanism for providing data communication. The electronic devices 24-26 may represent one or more laptop computers, desktop computers (in an office or at an employees home or other location), PDA's, cellular telephones, disk drives, mass storage devices, or any other electronic devices in which it may be useful to restrict access thereto. In an embodiment herein, the electronic devices 24-26 represent desktop or laptop computers that are used by employees of an organization that wishes to restrict access thereto in case a user/employee leaves the organization and/or one of the computers is lost or stolen. Of course, there may be other reasons to restrict access to one or more of the electronic devices 24-26 and the system described herein may be used in connection with any appropriate implementation.

An administration entity 28 sets a policy for allowing access by users to the electronic devices 24-26. For example, the administration entity 28 may determine that a particular user, U1, may no longer have access to any of the electronic devices 24-26 while another user U2, may access the electronic device 24 but not to the other electronic devices 25, 26. The administrative entity 28 may use any policy for setting user access.

The administrative entity 28 provides a plurality of proofs that are transmitted to the electronic devices 24-26 via the connection 22. The proofs may be provided to the electronic devices 24-26 by other means, which are discussed in more detail below. The electronic devices 24-26 receive the distributed proofs and, using credentials stored internally (described in more detail elsewhere herein), determine if access thereto should be allowed. Optionally, a proof distribution entity 32 may also be coupled to the connection 22 and to the administration entity 28. The proof distribution entity 32

provides proofs to the electronic devices 24-26. In an embodiment herein, a proof would only be effective for one user and one of the electronic devices 24-26 and, optionally, only for a certain date or range of dates.

5 The proofs may be provided using a mechanism like that disclosed in U.S. patent number 5,666,416, which is incorporated by reference herein, where each of the electronic devices 24-26 receives, as credentials, a digital certificate signed by the administrative entity 28 (or other authorized entity) where the digital certificate contains a special value representing an initial value having a one way function applied thereto N times. At each new time interval, the electronic devices may be presented
10 with a proof that consists of a one of the values in the set of N values obtained by the applying the one way function. In such a case, the electronic devices 24-26 may confirm that the proof is legitimate by applying the one way function a number of times to obtain the special value provided in the digital certificate. This and other possible mechanisms are described in more detail elsewhere herein.

15 It is also possible to use one or more of the products provided by CoreStreet, Ltd. of Cambridge, MA to provide the appropriate credentials and proofs as set forth herein or use any other mechanism for generating unique proofs that 1) could only have been generated by an administrative authority (absent an administrative security breach); and 2) can not be used to generate any other proofs. Accordingly, the proofs
20 are such that, given a legitimate proof P1, an unauthorized user may not generate another seemingly legitimate proof P2 for a different purpose (e.g., for a different time interval, different device, etc.). Thus, issued proofs may be stored and distributed in an unsecure manner, which substantially reduces the costs associated with the system. Of course, it is advantageous to maintain proper security for the entity or entities that
25 generate the credentials and/or proofs as well as maintaining appropriate security for any unissued (e.g., future) proofs.

In addition, an unauthorized user in possession of legitimate proofs P1-PN may not generate a new proof PN+1. This is advantageous in a number of instances. For

example, a terminated employee may not himself generate new proofs to provide unauthorized access to his corporate laptop after termination even though he is still in possession of all of the previous legitimate proofs he used for the laptop while he was still employed by the corporation.

5 In an embodiment herein, the electronic devices 24-26 are computers having firmware and/or operating system software that performs the processing described herein where the proofs are used to prevent unauthorized login and/or access thereto. Upon booting up and/or after a sufficient amount of time has passed, the computers would require an appropriate proof in order to operate. In this embodiment,
10 functionality described herein may be integrated with the standard Windows login system (as well as BIOS or PXE environments). The administration entity 28 may be integrated with the normal user-administration tools of corporate Microsoft networks and to allow administrators to set login policies for each user. In many cases, the administration entity 28 may be able to derive all needed information from existing
15 administrative information making this new functionality almost transparent to the administrator and reducing training and adoption costs. The administration entity 28 may run within a corporate network or be hosted as an ASP model by a laptop manufacturer, BIOS maker or other trusted partner. The proof distribution entity 32 may run partially within the corporate network and partially at a global site. Since
20 proofs are not sensitive information, globally-accessible repositories of the proof distribution system may run as web services, thereby making the proofs available to users outside of their corporate networks.

 In an embodiment herein, each of the computers would require a new proof each day. However, it will be appreciated by one of ordinary skill in the art that the
25 time increment may be changed so that, for example, the computers may require a new proof every week or require a new proof every hour.

 In addition, it is also possible to take advantage of a little-used feature of IDE hard drives which allows setting of a password on a drive which must be presented to

the drive before it will spin up and allow access to the contents. If the firmware for the drive were modified to use the system described herein, it is possible that access to a hard drive may be restricted so that, for example, it would not be possible to gain access to a computer hard drive even by placing it in a different computer. This feature
5 may be implemented with other types of hard drives.

In other implementations, the system may be used in connection with accessing data files, physical storage volumes, logical volumes, etc. In some instances, such as restricting access to files, it may be useful to provide appropriate modifications to the corresponding operating system.

10 Referring to Figure 1B, a diagram 20' illustrates an alternative embodiment with a plurality of administrative entities 28a-28c. Although the diagram 20' shows three administrative entities 28a-28c, the system described herein may work with any number of administrative entities. In the embodiment shown by the diagram 20', it is possible for one of the administrative entities 28a-28c (e.g., the administrative entity
15 28a) to generate the credentials while other ones of the administrative entities 28a-28c (e.g., the administrative entities 28b, 28c) generate the proofs or all of the administrative entities 28a-28c generate the proofs. Optionally, the proof distribution entity 32 may be used.

20 Referring to Figure 1C, a diagram 20'' illustrates an alternative embodiment with a plurality of proof distribution entities 32a-32c. Although the diagram 20'' shows three proof distribution entities 32a-32c, the system described herein may work with any number of proof distribution entities. The embodiment shown by the diagram 20'' may be implemented using technology provided by Akamai Technologies Incorporated, of Cambridge, Massachusetts.

25 Referring to Figure 1D, a diagram 20''' illustrates an alternative embodiment with a plurality of administrative entities 28a'-28c' and a plurality of proof distribution entities 32a'-32c'. Although the diagram 20''' shows three administration entities

28a'-28c' and three proof distribution entities 32a'-32c', the system described herein may work with any number of administration entities and proof distribution entities. The embodiment shown by the diagram 20''' combines features of the embodiment illustrated by Figure 1B with features of the embodiment illustrated by Figure 1C.

5 Referring to Figure 2, a diagram illustrates the electronic device 24 in more detail as including a validation unit 42, credential data 44 and proof data 46. The validation unit 42 may be implemented using hardware, software, firmware, or any combination thereof. Upon certain conditions, such as boot up, the validation unit 42 receives a start signal that causes the validation unit 42 to examine the credential data
10 44 and the proof data 46 and, based on the result thereof, generate a pass signal indicating that a legitimate proof has been presented or otherwise generate a fail signal. The output of the validation unit 42 is used by follow on processing/devices such as computer boot up firmware, to determine whether operation can proceed.

In an embodiment herein, the electronic device 24 includes an external interface
15 48 which is controlled by the validation unit 42. As with the validation unit 42, the external interface 48 may be implemented using hardware, software, firmware, or any combination thereof. The external interface 48 is coupled to, for example, the connection 22, and is used to fetch new proofs that may be stored in the proof data 46. Thus, if the validation unit 42 determines that the proofs stored in the proof data 46 are
20 not sufficient (e.g., have expired), the validation unit 42 provides a signal to the external interface 48 to cause the external interface 48 request new proofs via the connection 22. Of course, if the electronic 24 has been lost and/or stolen or if the user is a terminated employee or if there is any other reason not to allow access to the electronic device 24, then the external interface 48 will not be able to obtain a valid
25 proof. In some embodiments, the external interface 48 prompts a user to make an appropriate electronic connection (e.g., connect a laptop to a network).

In an embodiment herein, time data 52 provides information to the validation unit 42 to indicate the last time that a valid proof was presented to the validation unit

42. This information may be used to prevent requesting of proof too frequently and, at the same time prevent waiting too long before requesting a new proof. Interaction and use of the validation unit 42, the external interface 48, the credential data 44, the proof data 46, and the time data 52 is described in more detail elsewhere herein.

5 Referring to Figure 3, a flow chart 70 illustrates steps performed in connection with determining whether to send the start signal to the validation unit 42 to determine if the validation unit 42 should examine the credential data 44 and the proof data 46 to generate a pass or fail signal. Processing begins at a first step 72 where it is determined if a boot up operation is being performed. In an embodiment herein, the proofs are
10 always checked in connection with a boot-up operation. Accordingly, if it is determined at the test step 72 that a boot up is being performed, then control transfers from the step 72 to a step 74 where the start signal is sent to the validation unit 42. Following the step 74 is a step 76 where the process waits predetermined amount of time before cycling again. In an embodiment herein, the predetermined amount of time
15 may be one day, although other amounts of time may also be used. Following step 76, control transfers back to the test step 72, discussed above.

If it is determined at the test step 72 that a boot up operation is not being performed, then control transfers from the test step 72 to a test step 78 where it is determined if the a predetermined amount of time has elapsed since the last running of
20 the validation unit 42. This is determined using the time data element 52 and perhaps the current system time. In an embodiment herein, the predetermined amount of time used at the test step 78 is one day. If it is determined at the test step 78 that the amount of time since the last running of the validation unit 42 is greater than the predetermined amount of time, then control transfers from the test step 78 to the step 74 where the
25 start signal is sent to the validation unit 42. Following the step 74 or following the test step 78 if the amount of time is not greater than the predetermined amount of time, is the step 76, discussed above.

Referring to Figure 4, a flow chart 90 illustrates steps performed in connection with the validation unit 42 determining if a sufficient proof has been received. As discussed elsewhere herein, the validation unit 42 sends either a pass or a fail signal to follow on processing/devices (such as computer boot up firmware or disk drive
5 firmware). Processing begins at a first step 92 where the validation unit 42 determines the necessary proof. The necessary proof is the proof determined by the validation unit 42 sufficient to be able to send a pass signal. The validation unit 42 determines the necessary proof by examining the credential data 44, the proof data 46, the time data 52, and perhaps even the internal/system clock. Following the step 92 is a test step 94
10 which determines if the appropriate proof is available locally (i.e., in the proof data 46) and if the locally provided proof meets the necessary requirements (discussed elsewhere herein). If so, then control transfers from the step 94 to a step 96 where the validation unit 42 issues a pass signal. Following the step 96, processing is complete.

In some embodiments, it may be possible and desirable to obtain and store
15 future proofs in the proof data 46. For example, a user that expects to be without a connection to the administration entity 28 and/or the proof distribution entity 32 may obtain and store future proofs. In these embodiments, the electronic device may automatically poll for future proofs when connected to the administration entity 28 and/or the proof distribution entity 32, which may be provided according to a
20 predefined policy. Alternatively (or in addition), it may be possible for a user and/or electronic device to specifically request future proofs which may or may not be provided according to governing policy.

If it is determined at the test step 94 that the appropriate proof is not locally available (i.e., in the proof data 46), then control transfers from the test step 94 to a test
25 step 98 where the validation unit 42 determines if an appropriate proof is available externally by, for example, providing a signal to cause the external interface 48 to attempt to fetch the proof, as discussed above. If it is determined that the test step 98 that the externally-provided proof meets the necessary requirements (discussed elsewhere here), then control transfers from the test step 98 to the step 96, discussed

above, where the validation unit 42 issues a pass signal. In an embodiment herein, the externally-provided proof is stored in the proof data 46.

If it is determined at the test step 98 that an appropriate proof is not available externally, either because there is no appropriate connection or for some other reason, then control transfers from the test step 98 to a step 102 where the user is prompted to enter an appropriate proof. In an embodiment herein, if a user is at a location without an appropriate electrical connection, the user may call a particular phone number and receive an appropriate proof in the form of a number that may be entered manually into the electronic device in connection with the prompt provided at the step 102. Of course, the user may receive the proof by other means, such as being handwritten or typed or even published in a newspaper (e.g., in the classified section).

Following the step 102 is a test 104 which determines if the user has entered a proof meeting the necessary requirements (as described elsewhere herein). If so, then control transfers from the test step 104 to the step 96, discussed above, where the validation unit 42 issues a pass signal. Otherwise, control transfer from the test step 104 to a step 106 where the validation unit 42 issues a fail signal. Following the step 106, processing is complete.

Referring to Figure 5, a flow chart 120 illustrates steps performed in connection with generating credentials used by the validation unit 42. The steps of the flow chart 120 may be performed by the administration entity 28 which generates the credentials (and a series of proofs) and provides the credentials to the electronic device 24. Other appropriate entities (e.g., entities authorized by the administration entity 28) may generate the credentials. The random value is used in connection with generating the credentials and the proofs and, in an embodiment herein, is generally unpredictable. Following the step 122 is a step 124 where an index variable, I, is set to one. In an embodiment herein, the credentials that are provided are used for an entire year and a new proof is needed each day so that three hundred and sixty five separate proofs may be generated in connection with generating the credentials. The index variable, I, is

used to keep track of the number of proofs that are generated. Following step 124 is a step 126 where the initial proof value, $Y(0)$ is set equal to the random value RV determined at the step 122.

5 Following the step 126 is a test step 128 which determines if the index variable, I , is greater than an ending value, $IEND$. As discussed above, in an embodiment herein, three hundred and sixty five proofs are generated in connection with generating the credentials so that, in this embodiment, $IEND$, is three hundred and sixty five. However, for other embodiments it is possible to set $IEND$ to any number.

10 If it is determined at the test step 128 that the value of I is not greater than $IEND$, then control transfers from the step 128 to a step 132 where $Y(I)$ is set equal to the one way function applied to $Y(I-1)$. The one way function used at the step 132 is such that, given the result of applying the one way function, it is nearly impossible to determine the value that was input to the one way function. Thus, for the one way function used at the step 132, given $Y(I)$, it is very difficult, if not impossible, to
15 ascertain the value of the input (in this case $Y(I-1)$). As used herein, the term one way function includes any function or operation that appropriately provides this property, including, without limitation, conventional one way hash functions and digital signatures. This property of the one way function used at the step 132 is useful in connection with being able to store and distribute issued proofs in an unsecure manner,
20 as discussed elsewhere herein. The credentials and the proofs may be generated at different times or the proofs may be regenerated at a later date by the entity that generated the credentials or by another entity. Note that, for other embodiments, it is possible to have $Y(I)$ not be a function of $Y(I-1)$ or any other Y 's for that matter.

25 Processing begins at a first step 122 where a random value, RV , is generated. Following the step 132 is a step 134 where the index variable, I , is incremented. Following the step 134, control transfers back to the test step 128, discussed above. If it is determined at the test step 128 that I is greater than $IEND$, then control transfers from the test step 128 to a step 136 where a final value, FV , is set equal to $Y(I-1)$. Note

that one is subtracted from I because I was incremented beyond IEND. Following the step 136 is a step 138 where the administration entity 28 (or some other entity that generates the proofs and the credentials) digitally signs the final value, the current date, and other information that is used in connection with the proofs. In an embodiment
5 herein, the other information may be used to identify the particular electronic device (e.g., laptop), the particular user, or any other information that binds the credentials and the proof to a particular electronic device and/or user and/or some other property. Optionally, the date and/or the FV may be combined with the other information. For example, it is possible to use an OCSP-like signed message that simply says, "device
10 #123456 is valid on 1/1/2004" or have a bit in a miniCRL that corresponds to a specific device being on or off. In those case, the credential on the device may authenticate the device (i.e., determine that the device really is device #123456, etc.). OCSP and miniCRL's are know in the art. Following the step 138, processing is complete.

Referring to Figure 6, a flow chart 150 illustrates steps performed by the
15 validation unit 42 in connection with determining the validity of a proof. Processing begins at a first step 152 where the validation unit 42 receives the proof (e.g., by reading the proof from the proof data 44). Following the step 152 is a step 154 where the validation unit 42 receives the credentials (e.g., by reading the credential data 46).

Following step 154 is a test step 156 which determines if the other information
20 that is provided with the credentials is okay. As discussed elsewhere herein, the other information includes, for example, an identification of the electronic device, an identification of the user, or other property identifying information. If it is determined at the test step 156 that the other information associated with the credentials does not match the particular property described by the other information (e.g., the credentials
25 are for a different electronic device or different user), then control transfers from the test step 156 to a step 158 where a fail signal is provided. Following the step 158, processing is complete.

If it is determined at the test step 156 that the other information associated with the credentials is okay, then control transfers from the test step 156 to a step 162 where a variable N is set equal to the current date minus the date associated with the credentials (i.e., the number of days since the credentials were issued). Following the
5 step 162 is a step 164 where the proof value provided at the step 152 has a one way function applied thereto N times. The one way function used at the step 164 corresponds to the one way function used at the step 132, discussed above.

Following step 164 is a test step 166 which determines if the result obtained at the step 164 equals the final value FV that is part of the credentials received at the step
10 154. If so, then control transfers from the test step 166 to a step 168 where a pass signal is provided by the validation unit 42. Otherwise, if it is determined at the test step 166 that the result obtained at the step 164 does not equal the final value FV provided with the credentials at the step 154, then control transfers from the test step 166 to a step 172 where a fail signal is provided by the validation unit 42. Following
15 step 172, processing is complete.

Digital signatures may provide an effective form of Internet authentication. Unlike traditional passwords and PINs, digital signatures may provide authentication that may be universally verifiable and non-repudiable. Digital signatures may be produced via a signing key, SK, and verified via a matching verification key, PK. A
20 user U keeps his own SK secret (so that only U can sign on U's behalf). Fortunately, key 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 could make his own PK as public as possible (so that every one can verify U's signatures). For this reason PK is preferably called the public key. Note that the term "user" may signify a
25 user, an entity, a device, or a collection of users, devices and/or entities.

Public keys may be used also for asymmetric encryption. A public encryption key PK may be generated together with a matching decryption key SK. Again, knowledge of PK does not betray SK. Any message can be easily encrypted with PK,

but the so computed ciphertext may only be easily decrypted via knowledge of the key SK. Therefore, a user U could make his own PK as public as possible (so that every one can encrypt messages for U), but keep SK private (so that only U can read messages encrypted for U).

5 The well-known RSA system provides an example of both digital signatures and asymmetric encryption.

Alphanumeric strings called certificates provide that a given key PK is a public key of a given user U. An entity, often called certification authority (CA), generates and issues a certificate to a user. Certificates expire after a specified amount of time,
10 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 information (including no information), AI. In symbols, $C = \text{SIG}_{CA}(\text{SN}, \text{PK}, \text{U}, \text{D}_1, \text{D}_2, \text{AI})$.

15 Public encryption keys too may provide a means of authentication/identification. For instance, a party knowing that a given public encryption key PK belongs to a given user U (e.g., because the party has verified a corresponding digital certificate for U and PK) and desirous to identify U, may use PK to encrypt a random challenge C, and ask U to respond with the correct decryption.
20 Since only the possessor of SK (and thus U) can do this, if the response to the challenge is correct, U is properly identified.

It is possible to provide a system to control physical access to an area using a smart door (and/or smart virtual door, see description elsewhere herein). A smart door may verify that the person entering is currently authorized to do so. It may be
25 advantageous to provide the door not only with the credential of a given user, but also with a separate proof that the credential/user is still valid in a way that can be securely utilized even by a disconnected door. In an embodiment, such proofs are generated as

follows. Assume that a credential specifies the door(s) a user may enter. Then, for each credential and each time interval (e.g., each day), a proper entity E (e.g., the same entity that decides who is authorized for which door at any point in time, or a second entity working for that entity) computes an authenticated indication (PROOF) that a given
5 credential is valid on the given time interval. (If credentials do not identify the doors users are authorized to enter, a PROOF may also specify the door(s) the credential is good for on the given time interval).

A PROOF of E may consist of a digital signature of E indicating in an authenticated manner that a given credential is valid for a given interval of time, for
10 instance: $SIG_E(ID, Day, Valid, AI)$, where ID is information identifying the credential (e.g., the credential's serial number), Day is an indication of the given time interval (without loss of generality intended, a given day), Valid is an indication that the credential is deemed valid (this indication can be omitted if E never signs a similar data string unless the credential is deemed valid), and AI indicates any additional
15 information (including no information) deemed useful. In some instances, the signature of E may be a public-key signature. (But it could also be a private-key signature, that is, one that may be produced and verified via a single, secret key, known both to the signer and the verifier.) If the credential consists of a digital certificate, one sub-embodiment may consist of a short-lived certificate, that is, a digital signature that re-issues the
20 credential for the desired time interval (e.g., a digital certificate specifying the same public key, the same user U and some other basic information as before, but specifying the start date and the expiration date so to identify the desired ---without loss of generality intended--- day). For instance, letting, without loss of generality intended, a short-lived certificate last for a day, in such sub-embodiment a PROOF may take the
25 form $SIG_E(PK, U, D_1, D_2, AI)$, where start-date D_1 indicates the beginning of a given day D and end-date D_2 the corresponding end of day D, or where $D_1=D_2=D$; or more simply using a single date-information field to identify the day in question, $SIG_E(PK, U, Day, AI)$. If E coincides with the original certification authority, a short-lived-certificate PROOF may take the form $SIG_{CA}(PK, U, D_1, D_2, AI)$ or $SIG_{CA}(PK, U, Day,$
30 AI).

Being authenticated, a user may not manufacture his own PROOF of the day (i.e., the PROOF of the day of his own credential), nor can he change his PROOF of yesterday into his own PROOF of today, nor the PROOF of another user for today into his own for today. Because PROOFS are essentially unforgeable and inalterable, these PROOFS need not be protected. Thus, entity E may make the PROOFS available with negligible cost. For instance, E may post all the PROOFS of a given day on the Internet (e.g., make the PROOFS available via Akamai servers or the equivalent), or send the PROOFS to responders/servers that may be easily reached by the users. For instance, to a server located at the entrance of an airport (or office building) where many of the doors to be correctly accessed are located. This way, an employee coming to work may easily pick up his own PROOF (e.g., by inserting his own card into a card reader coupled with the server) and ---say---store the PROOF onto his own card, together with his own credential. This way, when the user presents his card to a door that his credential authorizes to access, the door can not only verify the credential but also receives and verifies a PROOF of current authorization, without needing to be connected at all! The door verifies the PROOF (e.g., the digital signature of E via E's public key that it may store since installation) and that the time interval specified by the PROOF is proper (e.g., via its own local clock). If all is fine, the door grants access else, the door denies access. In essence, the door may be disconnected and yet its PROOF verification may be both relatively easy (because the door may receive the PROOF by the most available party: the very user demanding access) and relatively secure (though the door receives the PROOF from arguably the most suspicious party: the very user demanding access). In fact, a user demanding access may typically be in physical proximity of the door, and thus can provide the PROOF very easily, without using any connection to a distant site, and thus operate independent of the door's connectivity. At the same time, the user demanding access may be the least trustworthy source of information at that crucial time. Nonetheless, because the user may not manufacture or alter a PROOF of his own current validity in any way, the door may be sure that a properly verified PROOF must be produced by E, and E would have not produced the PROOF if E knew the user to be not authorized for the given time interval. When a user stops being authorized, E will stop issuing PROOFS of authorization for the user, and thus the user can no longer enter even disconnected

doors, because the user will lack the PROOF that a door needs to verify in order to grant access. Thus, by utilizing the user demanding access to prove proper and current authorization, the system described herein dispenses with inconveniences associated with other systems, i.e., the need to dispatch a crew to re-program disconnected doors.

5 This approach also enables one to manage disconnected-door access by “role” (or by “privilege”). That is, rather than having a credential specify the door(s) that its user is authorized to enter, and then issue –e.g., daily-- a PROOF of current validity of a credential (or rather than issuing a PROOF specifying that a given credential authorizes his user to enter some door(s) on a given time interval), disconnected doors
10 may be programmed (e.g., at installation time) to grant access only to users having a given role. For instance, a cockpit door in an airplane may be programmed to grant access only to PILOTS and INSPECTORS. The credentials may be issued to employees primarily to vouch for their identity (which does not change), while each PROOF that E –e.g., daily—issues for a given credential may also specify (e.g., in the
15 AI field) the role(s) of its corresponding user on that day. For instance, PROOF = SIG_E(ID, Day, PILOT, AI) proves on day Day the user corresponding to credential identified by ID is a pilot. This way, employees may “migrate” from one role to the next without having their credential reissued, and without any need to specify within a user credential or in its corresponding daily PROOF which doors the user may access
20 that day. Note that the number of such doors may be huge. Thus, specifying within a user credential all the doors a user may be authorized to access may be cumbersome. Moreover, if new doors are added (e.g., because new airplanes are bought) then the pilot’s credential may have to be reissued, which is cumbersome too, to specify the additional doors.

25 The time intervals appropriate for a given credential may be specified within the credential itself, or may be specified by the credential and the PROOF together. For instance, a credential may specify a given start day and that it needs to be proved valid every day, while the PROOF may specify time interval 244, to mean that the PROOF refers to day 244 after the start day specified in the credential.

The system described herein may also be advantageous relative to more expensive connected-doors systems. For instance, assume that all doors were securely connected to a central database, and that a sudden power outage occurs (e.g., by sabotage). Then the connected doors may be forced to choose between two extreme alternatives: ALWAYS OPEN (good for safety but bad for security, particularly if terrorists caused the outage) and ALWAYS CLOSED (bad for safety but good for security). By contrast, in case of a sudden power outage, the system described herein offers a much more flexible response, some (no longer) connected doors may remain always closed, others always open and others yet may continue to operate as per the disconnected-door access control described herein. That is, the doors, relying on batteries, may open only if the right credential and the right PROOFs are presented. In fact, before the outage occurs it is possible for all employees to receive their expected PROOFs regularly.

Entity E may of course produce PROOFs specifying different time intervals for different credentials. For instance, in an airport facility, police officers and emergency personnel may every day have a PROOF specifying the next two weeks as the relevant time interval, while all regular employees may have daily PROOFs specifying only the day in question. Such a system may provide better control in case of a long and unexpected power outage. Should such a power outage occur, the daily usual distribution of PROOFs may be disrupted and ordinary employees may not receive their daily PROOFs, but policemen and emergency handlers may still carry in their cards the two-week proofs they received the day before and thus may continue to operate all doors they are authorized to enter (e.g., all doors).

It should be realized that the approach described herein encompasses using credentials consisting of a reduced form of certificate, that may be called minimal certificates. A minimal certificate may essentially omit the user name and/or the identifier ID of the certificate (or rather replace the user name and/or the identifier ID with a public key of the certificate, which may be unique for each certificate). For instance, a minimal certificate credential may take the form $C = \text{SIG}_{CA}(\text{PK}, D_1, D_2, \text{AI})$

with the understanding that proper presentation of this credential includes proving knowledge of the secret key SK corresponding to PK (e.g., by a challenge-response method). The door may know beforehand whether (or not) proper presentation of a credential relative to PK (preferably if currently validated) should result in granting access. Alternatively, a minimal credential C may specify (e.g., in AI) whether or not a user who knows the corresponding SK is entitled to enter a given door. A PROOF relative to a minimal certificate whose public key is PK, may be of the form $SIG_E(ID, Day, Valid, AI)$ or $SIG_E(PK, Day, Valid, AI)$, or $SIG_E(ID, Day, AI)$ if it is understood that any similar signature indicates validity by implication. Alternatively, a currency PROOF of a minimal certificate may take the form of the re-issuance of a minimal short-lived certificate: e.g., $SIG_E(PK, D_1, D_2, AI)$, where start date D_1 indicates the beginning of a given day D and D_2 the corresponding end of day D, or $D_1=D_2=D$; or $SIG_E(PK, Day, AI)$; or, letting E coincide with the original certification authority, $SIG_{CA}(PK, D_1, D_2, AI)$ or $SIG_{CA}(PK, Day, AI)$. In general, any method described herein directed to certificates should be understood to apply to minimal certificates as well.

A smart door may verify the validity and currency of a user's credentials which may be accompanied by a corresponding proof. The credentials/proofs used by a user to obtain access to an area may be similar to the credentials/ proofs used in connection with controlling access to electronic devices, as discussed elsewhere herein. The following are examples of credentials/proofs, some of which may be combined with others:

1. a PIN or password, entered at a key pad associated with the door or communicated to the door by a user's card;
2. biometric information, provided by a user via a special reader associated with the door;
3. a traditional (handwritten) signature, provided by a user via a special signature pad associated with the door;
4. a digital certificate for a public key PK (e.g., such a credential can be stored in a user's card and the right user/card may use the corresponding

- secret key SK to authenticate/identify itself to the door --e.g., via a challenge response protocol). For instance, if PK is a signature public key, the door may ask to have signed a given message and the right user --the only one who knows the corresponding secret signing key SK-- may provide the correct requested signature; if PK is a public encryption key, the door may request to have a given challenge ciphertext decrypted, which can be done by the right user, who knows the corresponding secret decryption key SK;
- 5 5. an enhanced digital certificate that includes a daily "validation value" (which assures that the certificate is valid on this particular date), stored
10 in a user's card and communicated to the door;
 6. a digital signature of a central authority confirming that a user's certificate is valid at the current time, communicated to the door by a server or a responder;
 - 15 7. a digital certificate that is stored in a user's card and communicated to the door, as well as a daily "validation value" communicated to the door by a server or a responder;
 8. a secret, stored in a user's card, knowledge of which is proven to the door by an interactive (possibly zero-knowledge) protocol with the door;
 - 20 9. a secret-key signature of an authority, stored in a user's card, indicating that the user is authorized to enter on a particular day.

Thus, in some instances, credentials/proofs are provided in a single part while, in other instances, credentials/proofs are provided in separate parts, the credentials and, separately, the proofs. For example, where the credentials/proofs consists of an
25 enhanced digital certificate that includes a daily validation value which indicates that the certificate is valid on this particular date and is associated with a user and communicated to the door, the credentials (the enhanced digital certificate) may be provided separately (by different means and/or at different times) from the proofs (the daily validation value). Similarly, the credentials and the proofs may be all generated
30 by the same authority or may be generated by different authorities.

Referring to Figure 7, a diagram illustrates a system 200 that includes an area 202 in which physical access thereto is to be restricted. The area 202 is enclosed by a plurality of walls 204-207. The wall 207 has a door 212 therein for providing egress to the area 202. In other embodiments, more than one door may be used. The walls 204-
5 207 and the door 212 provide a barrier to access to the area 202. The door 212 may be locked using an electronic lock 214, which prevents the door 212 from opening unless and until the electronic lock 214 receives an appropriate signal. The electronic lock 214 may be implemented using any appropriate elements that provide the functionality described herein, including, without limitation, using off-the shelf electronic locks.

10 The electronic lock 214 may be coupled to a controller 216, which provides an appropriate signal to the electronic lock 214 to allow the door 212 to be opened. In some embodiments, the electronic lock 214 and the controller 216 may be provided in a single unit. The controller 216 may be coupled to an input unit 218, which may receive a user's credentials and, optionally, also receive a corresponding proof indicating that a
15 user is currently authorized to enter the area 202. The input unit 218 may also receive a hot revocation alert (HRA) indicating that the user is no longer allowed to enter the area 202. HRA's are described in more detail hereinafter. The input unit 218 may be any appropriate input device such as a key pad, a card reader, a biometric unit, etc.

Optionally, the controller 216 may have an external connection 222 that may be
20 used to transmit data to and from the controller 216. The external connection 222 may be secure although, in some embodiments, the external connection 222 may not need to be secure. In addition, the external connection 222 may not be required because the functionality described herein may be provided using stand-alone units having no external connections. In instances where the external connection 222 is provided, the
25 external connection 222 may be used to transmit credentials, proofs, HRA's and/or may be used in connection with logging access to the area 202. Logging access is described in more detail elsewhere herein. Note that the external connection 222 may be intermittent so that, for example, at some times the external connection 222 provides connectivity for the controller 216 while at other times there may be no external
30 connection for the controller 216. In some instances, the external connection 222 may

be used to transmit a portion of the credentials/proofs (e.g., a PKI digital certificate) while a user presents to the input unit 218 a remaining portion of the credentials/proofs (e.g., a daily validation value used in connection with the digital certificate).

In some embodiments, a user may present a card 224 to the input unit. As
5 discussed elsewhere herein, the card 224 may be a smart card, a PDA, etc. that provides data (e.g., credentials/proofs) to the input unit 218. The card 224 may get some or all data from a transponder 226. In other instances, the card 224 may get data from other cards (not shown), from the input unit 218 (or some other mechanism associated with
accessing the area 202), or some other appropriate source.

10 In a first example, credentials and proofs may be maintained using a pin/password with physical protection. In this example, every morning a server generates a new secret password SU for each authorized user U and communicates the new SU to specific doors to which U is allowed to access. The communication may be encrypted to be sent using unsecure lines or may be transmitted to the doors via some
15 other secure means. When U reports to work in the morning, the central server causes the U's card to receive the current secret password SU. The secret password SU is stored in the secure memory of the card, which can be read only when the card is properly authorized (e.g., by the user entering a secret PIN in connection with the card or by connecting with trusted hardware on the server or the doors). Whenever the user
20 attempts to access a door, the card securely communicates SU to the door. The door then checks if the value SU received from the card matches the value received from the server in the morning, and, if so, allows access.

Thus, SU is the user's credential for a day. This system has the advantage that each credential is of limited duration: if an employee is terminated or his card is stolen,
25 his credentials will not be useful the next day. The system, however, requires some connectivity: at least a brief period of connectivity (preferably every morning) is needed to update the door. This transmission should be secured (e.g., physically or cryptographically).

In another example, the user's credentials include secret-key signatures. This example utilizes signatures, either public-key signatures (e.g., RSA signatures) or secret-key signatures (e.g., Message Authentication Codes, or MACs). For instance, an access-control server uses a secret key SK to produce signatures, and the door has
5 means to verify such signatures (e.g., via a corresponding public key or by sharing knowledge of the same SK). When a user U reports to work in the morning on a day D, the server causes the user's card to receive a signature Sig authenticating U's identifying information (e.g., the unique card number, or U's secret password, or biometric information such as U's fingerprints) and the date D. When U attempts to
10 access a door, the card communicates the signature Sig to the door, which verifies its validity possibly in conjunction with identifying information supplied by U, and the date supplied by the door's local clock. If all is correct, the door allows access.

In this technique, the signature Sig may be considered the user's credentials and proof together. This method has its own advantages: the cards need not store secrets,
15 and the doors need not maintain secure connections to a central server, nor a long list of valid credentials.

In another example, the user's credentials include a digital certificate with hash-chain validity proofs similar to those generated in connection with the flow chart 120 of Figure 5. This example utilizes public-key signatures and a one-way hash function H
20 (implementing a special type of digital signature). A central authority has a key pair: a public key PK (known to the doors) and a secret key SK that is not generally known. For a user U, the authority generates a random secret value X0 and a computes values $X_1 = H(X_0)$, $X_2 = H(X_1)$, ..., $X_{365} = H(X_{364})$. Because H is a one-way hash function, each value of X cannot be computed from the next value of X. The authority
25 issues to U a digital certificate Cert, signed using SK and containing the value X365, valid for one year. Then, when U reports to work on day i, the authority causes the user's card to receive the day's validation value Xj, where $j = 365 - i$. When U attempts to access a door, the card communicates the validation value Xj and certificate Cert containing X365 to the door. The door verifies the validity of the Cert with public

key PK of the authority and also checks that H applied i times to X_j produces X_{365} . Note that the "one year" and 365 may be replaced with any other time period.

5 Thus, the user's certificate Cert as well as the validation value X_j make up the user's credentials/proof. This system has many advantages: neither the door nor the card need to store any secrets; the door need not have any secure connections; the certificate can be issued once a year, and thereafter the daily computational load on the central authority is minimal (because the authority just needs to retrieve X_j); the daily validation values can be provided by unsecured (cheap) distributed responders, because they need not be secret.

10 A credential/proof for a user U is often limited in its duration, which is useful in a number of circumstances. For example, if U is an employee of an airport and is terminated, his credentials/proof may expire at the end of the day and he will be no longer able to access the airport's doors. For more precise access control, it may be desirable to have shorter-duration credentials. For example, if the credential/proof for
15 U includes the hour and the minute as well as the date, then U can be locked out of the airport within one minute of being terminated. However, shorter-duration credentials/proof may require more frequent updating, which adds expense to the system. It could be inconvenient if every employee at an airport had to upload new credentials/proof onto his or her card every minute. Thus, there may be an inherent
20 tension between the desires to have short-term credentials and to have a lower-cost system, which may lead to credentials that are sometimes longer than desired. For example, U may need to be locked out of the airport immediately, but his credential won't expire until midnight. It is therefore desirable to provide for immediate revocation of credentials that have not yet expired.

25 Note that, if credentials/proofs are always stored in a secured database that is queried by doors each time access is requested, it is relatively straight-forward to revoke credentials/proofs by, for example, removing the revoked credentials/proofs from the database. However, having a door query a secure database each time is

expensive. First, because this adds a significant delay to the transaction since the user wants access the door right away, but he must wait for the query to be properly completed. Second, because this communication is preferably conducted over a secure channel, which can easily cost \$4,000 per door (or more) or be entirely unavailable in some cases (e.g., for doors of airplanes or cargo containers). Third, because a single secure database may only handle a limited query load, and replicating a secure database is in itself expensive and time consuming (e.g., because the costs of keeping the database secure must be duplicated and the effort to keep these copies synchronized must be added). Therefore, unlike the fully connected approach, disconnected or intermittently connected approaches (such as those in the examples above) require less communication and often store credentials/proofs on non-secured responders or on the cards themselves. In such a case, simply removing credentials/proofs from the database may not suffice. To refer again to the above examples, the password SU, or the authority signature, or the validation value Xj would somehow have to be removed from a user's card or the doors. Moreover, even such a removal may not always guarantee revocation of a credential since a credential stored in an unsecured responder may be available to anyone, including a malicious attacker who could save it and attempt to use it after its removal from the user's card. Thus, even though cost-effective solutions with limited-duration credentials exist, these solutions, by themselves, do not necessarily provide sufficient revocation of a non-expired credentials/proofs.

Revoking credentials/proofs, may be performed using a Hot Revocation Alert (HRA), which is a (preferably authenticated) piece of data transmitted to the door that will prevent the door from granting access to a user with revoked (though possibly unexpired) credentials/proofs. For example, an HRA may consist of a digitally signed message indicating that given credentials/proofs have been revoked. Note, however, that a signature may not always be involved in an HRA. For example, in the case of a securely connected door, just sending an HRA along the protected connection may suffice. However, as mentioned above, securely connected doors may be expensive in some instances and impossible (or nearly so) in others.

It is useful if HRA's are authenticated so that an entity to which an HRA is presented may be relatively certain that the HRA is genuine. Letting ID be an identifier for the revoked credentials/proofs C (in particular, ID may coincide with C itself), then SIG(ID, "REVOKED", AI) may be an HRA, where "REVOKED" stands for any way of signaling that C has been revoked ("REVOKED" may possibly be the empty string if the fact that the credentials/proofs are revoked could be inferred by other means—such as a system-wide convention that such signed messages are not sent except in case of revocation), and AI stands for any additional information (possibly date information—such as the time when the credentials/proofs have been revoked and/or the time when the HRA was produced— or no information). The digital signature SIG may be, in particular, a public-key digital signature, a secret-key digital signature, or a message authentication code. It is also possible to issue an authenticated HRA by properly encrypting the information. For example, an authenticated HRA may take the form ENC(ID, "REVOKED", AI).

Another notable example of an authenticated HRA is described in U.S. Patent 5,666,416, which is incorporated herein by reference. The issuing authority incorporates into a credential/proof C a public key PK (of a digital signature scheme) that is unique to C, so that a digital signature relative to that PK indicates that C is revoked. In a special embodiment of such a scheme, PK may consist of a value Y1 computed as $Y1 = H(Y0)$, where H is a (preferably hashing) one-way function and Y0 is a secret value. When credential/proof C is revoked, the HRA consisting of just Y0 is issued. Such an HRA can be verified by hashing Y0 and checking that the result matches the value Y1 which belongs to the credential/proof C.

Note that a signature may not be required for an HRA. For example, in case of a securely connected door, just sending (ID, "REVOKED", AI) along the protected connection may suffice as an HRA. However, the advantage of authenticated HRA's is that HRA's themselves need not be secret. Authenticated HRAs, once authenticated by the appropriate authority, may be store on one more (possibly geographically dispersed) responders. Furthermore, these responders may be unprotected (unlike the issuing authority), because they are not storing secret information. Greater reliability

may be provided at a lower cost by replicating multiple unprotected responders. Some additional advantages of the authenticated-HRA example of U.S. Patent 5,666,416 are: (1) the HRA is relatively short (can be as short as twenty bytes), (2) is relatively easily computed (simply a look-up of the previously stored Y0) and (3) is relatively easily
5 verified (just one application of one-way hash function).

Authenticated HRAs may be particularly advantageous for efficient broad dissemination, as further described below. When an HRA transits through multiple points on the way the door, there may be multiple possibilities for an incorrect HRA to be inserted into the system. Indeed, an HRA received by the door not directly through
10 or from the issuer via a secure connection may be no more than a mere rumor of particular credential's revocation. If the HRA is authenticated, however, this rumor can be readily confirmed by the door, which can verify its authenticity.

In general, an HRA may be specific to a single credential/proof or may provide revocation information about a multiplicity of credentials/proofs. For instance, if
15 ID1,...,IDk are identifiers for revoked credentials, an HRA may consist of the single digital signature SIG(ID1,...,IDk; "REVOKED"; AI). Consider the case of a door that stores information identifying the credentials/proofs which have the right to access the door. If such a door receives an HRA indicating that one or more credentials/proofs are revoked, the door does not need to store the HRA. It suffices for the door to erase the
20 identified credential(s)/proof(s) from its storage (or mark them "REVOKED" somehow). Then, if a user with a revoked credential/proof attempts access, the door will not allow access because the presented credential/proof is not currently stored therein or, if stored therein, is marked "REVOKED".

Consider now a case of a door that does not store information identifying all
25 allowed credentials/proofs, but rather verifies whether a credential/proof is allowed when presented. When a user presents a credential/proof to such a door, the door may first verify whether the credential/proof is valid, disregarding HRA's. (For instance, if the credential/proof includes a digital signature, the door verifies the signature. In

addition, if the credential/proof includes an expiration time, the door may also verify that the credential/proof is not expired, e.g., using an internal clock.) But even if all the checks are passed, the door may still deny access if the credential/proof is indicated as being revoked by an HRA. Therefore, it is helpful if such a door has information
5 concerning relevant HRAs. One way to achieve this is for the door to save all HRAs presented to the door. On the other hand, in some instances, this may become impractical. Consider a system where many credentials/proofs could be used to go through that door. For example, the Department of Transportation is envisaging a 10,000,000-credential system for a variety of individuals (including pilots, airport staff,
10 airline employees, mechanics, baggage handlers, truck drivers, police, etc.) who may at one time or another be allowed access to a given door. At a modest 10% annual revocation rate, the door may have 1,000,000 HRAs to store by the end of a year, which may be a very costly (if not infeasible) task. Moreover, if the quantity of the HRAs cannot be precisely determined in advance, the designers of a system may have to
15 overestimate the storage size for HRA's in order to be on the safe side, and build even more storage capacity (at even more cost) into the door.

This problem may be addressed by means of removable HRAs. This means having an HRA indicate a time component specifying when the HRA can be safely removed from storage. For instance, in a system with credentials/proofs of limited
20 duration, this can be achieved by (1) having a credential/proof include an expiration time after which the credential/proof should not be accepted by the door as valid for access; (2) having an HRA revoking the credentials/proof include the expiration time and (3) having the door remove from its storage the HRA revoking the credentials/proof after the expiration time. For instance, the expiration time for a
25 credential/proof could be the time at which the credential/proof expires (and the expiration time could be explicitly included and authenticated within the credential/proof or it could be implied by system-wide conventions) Removing such HRA after the expiration time does not harm security. In fact, if the door does not store the HRA that revokes a particular credential/proof, it may be because the door erased
30 the HRA from memory after expiration, at which point the out-of-date credential/proof will be denied access by the door anyway.

Note that step (2) above may be optional in cases where the expiration time can be indicated in an HRA implicitly or indirectly. For instance, the HRA may have the form SIG(C, "REVOKED", AI), and the credentials/proof may include its own expiration date. In addition, step (1) above may be optional since removable HRAs
5 may also be implemented with HRAs that do not indicate the expiration times of the revoked credentials at all. For instance, if all credentials in a particular system are valid for at most one day, then all HRAs may be erased after being stored for a day. (More generally, if the maximum lifetime of a credential/proof may be inferred somehow, then a corresponding HRA may be erased after being stored for that amount of time.)
10 As for another example, when presented with a credential/proof with a particular expiration time, the door may look for an HRA revoking the credential. If one exists and the expiration time has passed already, then the door may safely remove the HRA. Else, the door may store the expiration time in connection with the stored HRA, and remove the HRA after that time.

15 A door may remove HRAs after their expiration in a variety of ways. In some cases, HRA removal may be accomplished efficiently by maintaining a data structure (such as a priority queue) of HRAs based on expiration times. Alternatively, the door may periodically review all HRAs in storage and purge the ones that are no longer needed. As another alternative, the door may erase an HRA if, when encountering the
20 HRA, the door realizes the HRA is no longer relevant. For instance, the HRAs may be stored in a list that is checked each time a credential is presented for verification. Whenever an expired HRA is encountered in such a list, the expired HRA may be removed. As yet another alternative, the door may remove HRAs only as needed, when memory needs to be freed (perhaps for other HRAs).

25 Removable HRAs may significantly reduce the storage required at the door. Using the above example of 10,000,000 users and 10% annual revocation rate, then, if HRAs expire and are removed, on average, in one day, only 2,740 (instead of 1,000,000) HRAs may need to be stored. This reduced storage requirement is a great potential advantage of removable HRAs.

It is useful for HRAs to be made available to the doors as quickly as possible, in order to inform the doors of credentials/proofs that are no longer acceptable. This may be a problem for disconnected doors, but it may also be a problem for fully connected doors. Of course, a fully connected door may be sent an HRA over the connection of the door when the HRA is issued. However, this transmission may still be blocked or jammed by a determined enemy. (e.g., if the connection to the door is secured by cryptographic means, an enemy may just cut the wire, or alter/filter the traveling signals. If the connection to the door is secured by running a wire in a steel pipe, then such jamming and blocking may be harder, but still is not impossible.) Such malicious jamming and blocking of an HRA may be even easier to carry out for doors with intermittent (e.g., wireless) connectivity.

In order to make it harder for an enemy to prevent a door from receiving an HRA, an HRA may be carried by a revoked card itself. For instance, when a card communicates with a database or a connected door (or any door that knows of the relevant HRA), the door may send the HRA to the card, which may store the HRA. In particular, this can be done without any indication to the user, so as to protect against users who may wish to tamper with the card and remove the HRA. This method is more effective if the card carries a tamper-proof hardware component or data (e.g., encrypted data) that is not easily read/removed by the user. When the card is subsequently used in an attempt to gain access to any (even fully disconnected) door, the card may communicate its HRA to the door, which, upon proper verification, may deny access (and, in some instances, store the HRA).

The HRA may be sent over a wireless channel (e.g., via a pager or cellular network or via satellite) to the card itself. This may be accomplished even if the card has limited communication capabilities – for example, by placing a wireless transmitter at a location that each user is likely to pass. For instance, at a building, such a transmitter may be placed at every building entrance, to provide an opportunity for every card to receive the transmission whenever a user of one of the cards enters the building. Alternatively, the transmitter may be placed at the entrances to the parking lot, etc.

To prevent a malicious user from blocking the transmission (by, for example, wrapping the card in material that will be impenetrable by the transmitted signal), the card may in fact *require* that it receive periodic transmissions in order to function properly. For example, the card may expect a signal every five minutes in order to
5 synchronize its clock with that of the system, or may expect to receive another periodic (preferably digitally signed) signal, such as a GPS signal, or just expect appropriate noise at the appropriate frequencies. If such a signal is not received with a reasonable time interval, the card may “lock out” and simply refuse to communicate with any door, this making itself unfit for access. Note that such a system may be more economical
10 and convenient than simply broadcasting all HRAs to all cards, because HRAs are custom and continually changing messages. Thus, broadcasting HRA’s to all cards may require putting up a special purpose satellite or customizing an already existing one. The above method instead takes advantage of already available signals for broad transmissions and installs very local transmitters for the custom messages.

15 Alternatively, a user may be prevented from blocking transmissions to a card if the security policy requires the user to wear the card visibly, as a security badge, or to present it at an appropriate place (within transmission range) to a guard. An additional technique for disseminating an HRA for a particular card/credential/proof may include using OTHER cards to carry the HRA to doors. In a version of this, Card 1 may (e.g.,
20 when picking up its own daily credential/proof, or wirelessly, or when communicating with a connected door, or when making any kind of connection) receive an HRA, HRA2, revoking a credential/proof associated with a different card, Card 2. Card 1 may then store HRA2 and communicate HRA2 to a door, which then also stores HRA2. Card 1 may in fact provide HRA2 to multiple doors, e.g., to all doors or all
25 disconnected doors that access or communicate with Card 2 for a particular period of time (e.g., for an entire day). At this point, any door (even if disconnected) reached by Card 1 may be able to deny access to the holder of Card 2 that contains the revoked credential/proof. Preferably, HRA2 is digitally signed or self authenticating, and any door reached by Card 1 checks the authenticity of HRA2 so as to prevent the malicious
30 dissemination of false HRAs.

This may be enhanced by having a door reached by Card 1 communicate the learned HRA2 to another card, Card 3, that subsequently accesses it or communicates with the door. This is useful because Card 3 may reach doors that Card 1 will not reach or will reach later than Card 3. This process may continue by having these additionally
5 reached doors communicate to other cards, etc. Moreover, it is possible that some doors, even though not fully connected to a central database, may have connections to each other. Such doors thus may exchange available HRAs similarly. If cards have communication capability with each other --e.g., when in proximity-- they may also exchange information about HRAs that they store.

- 10 Note that authenticated HRAs may be particularly advantageous with the HRA dissemination techniques discussed herein. Indeed, sending HRAs through multiple intermediaries (cards and doors) may provide multiple points of failure where HRAs may be modified or false HRAs may be injected by an adversary. In a sense, unauthenticated HRAs may become mere rumors by the time they reach the doors.
15 Authenticated HRAs, on the other hand, may be guaranteed to be correct no matter how they reach the doors.

In instances where resources are not a significant concern, all HRAs could be stored and disseminated in this manner. It may also be possible to adopt some optimizations. For instance, a card may manage HRA storage like a door, and remove
20 expired HRAs to free internal card storage and to prevent unnecessary communication with other doors. Minimizing storage and communication may be useful within such a system, because, even though the number of unexpired revoked credentials may be short, it is possible that some components (e.g., some cards or doors) may not have enough memory or bandwidth to handle all unexpired HRAs.

- 25 Another possibility for minimizing storage and communication includes selecting which HRAs are to be disseminated via which cards. For instance, HRAs may come with priority information, indicating the relative importance of spreading knowledge about a particular credential/proof as quickly as possible. For example,

some HRAs may be labeled “urgent” while others may be labeled “routine.” (A gradation of priorities may be as fine or coarse as appropriate.) Devices with limited bandwidth or memory may record and exchange information about higher-priority HRAs, and only if resources permit, may devote their attention to lower-priority ones.

5 As another example, an HRA that prevents a card to access a given door may be disseminated via cards that are more likely to quickly reach that door (e.g., cards whose credential enables access to that door or doors in its vicinity). Indeed, the card and the door may engage in a communication with the goal of establishing which HRAs to accept for storage and/or further dissemination. Alternatively, HRAs or cards to store

10 them may be selected in a way that involves randomness, or a door may provide an HRA to a certain number of cards (e.g., the first k cards the door “encounters”).

The use of such dissemination techniques may reduce the likelihood that a user with revoked credentials/proofs will be able to gain access since even for a disconnected door a user would have to get to the door before any other user provides

15 an appropriate HRA thereto with an up-to-date card. The exchange of information among cards and doors may help ensure that many cards are quickly informed of a revocation. This approach may also be used as a countermeasure against “jamming” attacks that attempt to disconnect a connected door and prevent the door from receiving the HRA. Even if the jamming attack succeeds and the door never gets informed of the

20 HRA by the central servers or responders, an individual user’s card may likely inform the door of the HRA anyway. It is noted that the actual method of exchanging the HRAs among cards and doors may vary. In case of a few short HRAs, it may be most efficient to exchange and compare all known HRAs. If many HRAs are put together in one list, the list may contain a time indicating when the list was issued by the server.

25 Then the cards and doors may first compare the issue times of their lists of HRAs, and the one with older list may replace it with the newer list. In other cases, more sophisticated algorithms for finding and reconciling differences may be used.

Efficient HRA dissemination may be accomplished by (1) issuing an authenticated HRA; (2) sending the authenticated HRA to one or more cards; (3)

having the cards send the authenticated HRA to other cards and/or doors; (4) having doors store and/or transmit to other cards the received HRAs.

It may be useful to present in detail some sample HRA use:

SEQUENCE 1 (directly from “authority” to door):

- 5 1. Entity E revokes a credential/proof for a user U and issues an HRA A containing the information that the credential/proof has been revoked;
2. A is transmitted via wired or wireless communication to a door D;
3. D verifies the authenticity of A and, if verification succeeds, stores information about A;
- 10 4. When U attempts to access D by presenting the credential/proof, the door D observes that the stored information about A indicates that the credential/proof is revoked and denies access.

SEQUENCE 2 (from “authority” to a user’s card to door):

- 15 1. Entity E revokes a credential/proof for a user U and issues an HRA A containing the information that the credential/proof has been revoked;
2. Another user U’ reports to work and presents his card to E in order to obtain his current credential/proof;
3. Along with the current credential/proof for U’, the HRA A is transmitted to the card of U’; the card stores A (the card may or may not verify the authenticity of A, depending on the card’s capabilities);
- 20 4. When U’ attempts to access a door D, his card transmits his credential/proof along with A to D
5. D verifies the authenticity of A and, if verification succeeds, stores A;
6. When U attempts to access D by presenting his credential/proof, the door D observes A revoking U’s credential/proof and denies access.
- 25

SEQUENCE 3 (from “authority” to another door to a user’s card to door):

1. Entity E revokes a credential/proof for a user U and issues an HRA A containing the information that U's credential/proof has been revoked;
2. A is transmitted via wired or wireless communication to a door D';
3. D' verifies the authenticity of A and, if verification succeeds, stores A;
- 5 4. Another user U' with his own credential/proof presents his card to D' in order to gain access to D'. D', in addition to verifying credentials/proofs of U' and granting access if appropriate, transmits A to the card of U'. The card stores A (the card may or may not verify the authenticity of A, depending on the card's capabilities).
- 10 5. When U' attempts to access a door D, his card transmits his own credential/proof along with A to D
6. D' verifies the authenticity of A and, if verification succeeds, stores A;
7. When U attempts to access D by presenting his credential/proof, the door D observes A revoking U's credential/proof and denies access.

15 SEQUENCE 4 (from "authority" to the user's card to door):

1. Entity E revokes a credential C for a user U and issues an HRA A containing the information that C has been revoked;
2. The user U, carrying his card, passes a transmission point located near the building entrance, which causes his card to receive A; the card stores
- 20 A (the card may or may not verify the authenticity of A, depending on the card's capabilities);
3. When U attempts to access a door D, his card transmits A along with C to D
4. D verifies the authenticity of A and, if verification succeeds, stores A
- 25 and denies access to U;
5. If U again attempts to access D by presenting C, the door D observes the previously stored A revoking C and denies access.

Sometimes it may be useful to establish, after the fact, who attempted to access a particular door, at what time, what credentials/proofs were presented, and whether

30 access was denied or granted. It may also be useful to know if a door's mechanism

became jammed, if a switch or sensor failed, etc. To this end, it may be desirable to maintain event logs of the events that take place. Such logs may be particularly useful if readily available at some central location so that they may be inspected and acted upon. For instance, in case of hardware failure, a repair team may need to be
5 dispatched promptly. There are, however, two major problems with such logs.

First, if a door is connected, it may be easier to collect logs by sending them via the connection. However, collecting event logs may be more difficult for disconnected doors. Of course, one way to collect logs is to send a person to every disconnected door to physically deliver the logs back to the central location, but this approach is
10 costly.

Second, for an event log to be believed, the integrity of the whole system surrounding the generation, collection and storage of the logs should be guaranteed. Else, for instance, an adversary may create false log entries or delete valid ones. Traditional approaches such as physically securing the communication channels and
15 data storage facilities are very costly (and may not be sufficient by themselves).

Conventional logs may vouch that "a certain user went to a certain door" by the mere existence of such a log entry, which must be assumed to be valid. However, this may not be appropriate for a high-security application. Consider a user U accused of damaging some property behind a locked door D. A traditional log entry may provide
20 only weak evidence that U entered D: one would have to trust that no one maliciously falsified the log entry. Thus, it is desirable to have logs that provide much stronger evidence, because they may not be "manufactured" by an enemy. In particular, indisputable logs may prove that door D (possibly with the cooperation of U's card) created the record in the log.

25 The system described herein addresses this in the following manner: Whenever a door receives a credential/proof presented as part of a request for access, the door

may create a log entry (e.g., a data string) containing information about the event, for example:

- time of request;
- type of request (if more than one request is possible – for example, if the request
- 5 is for exit or for entry, or to turn the engine on or off, etc.);
- credential/proof and identity presented (if any);
- whether the credential/proof verified successfully;
- whether the credential/proof had a corresponding HRA;
- whether access was granted or denied.

10 Log entries may also contain operational data or information on any unusual events, such as current or voltage fluctuations, sensor failures, switch positions, etc. One way to produce an indisputable log includes having the door digitally sign event information by means of a secret key (SK). The resulting indisputable log may be represented by $SIG(\text{event}, AI)$, where AI stands for any additional information. The

15 signature method used by door D may be public-key or private-key.

If it is useful to emphasize the public key PK relative to which the signature is valid, or the secret key SK used in producing the signature, or the door that generated the signature, it may be possible to symbolically represent the indisputable log by $SIG_{PK}(\text{event}, AI)$, $SIG_{SK}(\text{event}, AI)$, or $SIG_D(\text{event}, AI)$. Such a log may be

20 indisputable because an enemy may not forge the door's signature without knowing the relevant secret key. On the other hand, the authenticity of the log could be checked by any properly informed verifier (e.g., one that knows the door's PK or the door's SK) without having to trust the integrity of the database storing the log, or that of the system transmitting the log. In general, logs may be made indisputable not only by digitally

25 signing each entry, but also by using a digital authentication step for multiple entries. For instance, the door could authenticate a multiplicity of events E1, E2, ... by means of a digital signature: symbolically, $SIG(E1, \dots, E2, AI)$. As usual, here and anywhere in this application, a digital signature may mean the process of digitally signing the one-

way hash of the data to be authenticated. In particular, stream authentication may be viewed as a special case of digital signature. For instance, each authenticated entry could be used to authenticate the next (or the previous) one. One way to do this consists of having an authenticated entry include the public key (in particular, the public key of a one-time digital signature) used for authenticating the next or other entries.

Logs and indisputable logs may also be made by cards (in particular, a card may make an indisputable log by digitally signing information about an event E : in symbols, $SIG(E, AI)$). All of the log techniques described herein may also be construed to relate to card-made logs.

In addition, other logs and indisputable logs may be obtained by involving both the door and the card. For instance, during a request of door access, the card may provide to the door the card's own (possibly indisputable) log entry to the door. The door may inspect the log entry and grant access only if the door finds the log entry "acceptable." For instance, the door may verify the digital signature of the card authenticating the log entry; or the door may verify that time information included in the card's log entry is correct according to a clock accessible to the door.

Other types of indisputable logs may be obtained by having both the door and the card contribute to the generation and/or authentication of a log entry. For instance, the card may authenticate a log entry and the door may then also authenticate at least part of the log entry information, and vice versa. In a particular embodiment, a card C may give the door its signature, $x = SIG_C(E, AI)$, of the log entry, which the door will countersign -- in symbols, $SIG_D(x, AI')$ —and vice versa. Alternatively, the door and the card may compute a joint digital signature of the event information (e.g., computed by means of a secret signing key split between the door and the card, or by combining the door's signature with that of the card into a single "multi" signature). Several multi-signature schemes may be used, in particular that of Micali, Ohta and Reyzin.

It is possible to include additional information into the logs. It may then be checked if the information as reported by the card and as reported by the door agrees. For instance, both the card and the door may include time information into the log entries, using clocks available to them. In addition, the card (and possibly also the door) may include location information (such as obtained from GPS) into the log entry. Alternatively, if current location is unavailable (e.g., because GPS reception capability is unavailable), information on latest known location (and possibly how long ago it was established) may be included. This way, particularly in the case of a mobile door (such as the door of an airplane), it may be possible to establish where the door and the card were located when the event occurred.

Of course, even an indisputable log entry as above may be maliciously deleted from the database or prevented from reaching the database altogether. To protect against such deletions, it is useful to provide a Deletion-Detectable Log Systems. Such systems may be built by using (1) an authentication scheme (e.g., a digital signature scheme), (2) a correlation-generating scheme and (3) a correlation-detection scheme as follows. Given one log event E (part of a sequence of —possibly past and/or future— events), the correlation-generating scheme may be used to generate correlation information CI, which is then securely bound to E by means of the authentication scheme to generate a deletion-detectable log entry. The correlation-generating scheme may ensure that, even if events themselves are uncorrelated and the existence of one event may not be deduced from the existence of other events, CI is generated in such a way as to guarantee that for missing log entries no properly correlated information is present, something that can be detected using the correlation-detection scheme. In some instances, the system may also guarantee that even if some log entries are missing, others can be guaranteed authentic and/or individually indisputable.

In a first example, the correlation information CI of the log entries may include sequentially numbering the log entries. The corresponding correlation-detection scheme may consist of noticing the presence of a gap in the numbering sequence. But to obtain a deletion-detectable log system, a proper binding between CI and the log entries is found, which may not be easy to do, even if secure digital signatures are used for the

authentication component of the system. For instance, having the i -th log entry consist of $(i, \text{SIG}(\text{event}, \text{AI}))$, is not secure, because an enemy could, after deleting a log entry modify the numbering of subsequent entries so as to hide the gap. In particular, after deleting log entry number 100, the adversary may decrease by one the numbers of log entries 101, 102, etc. The enemy may so hide his deletions because, even though the integrity of the event information is protected by a digital signature, the numbering itself may not be. Moreover, even digitally signing also the numbers may not work. For instance, assume that the i -th log entry consists of $(\text{SIG}(i), \text{SIG}(\text{event}, \text{AI}))$. Then an enemy could: (1) observe and remember $\text{SIG}(100)$, (2) delete entry number 100, (3) substitute $\text{SIG}(100)$ in place of $\text{SIG}(101)$ in original entry 101, while remembering $\text{SIG}(101)$, and so on, so as to hide the deletion completely.

Neither of the above two methods produces the desired secure binding of CI and log entries. Indeed, by securely binding (1) the numbering information together with (2) the event being numbered, we mean that an enemy may not manufacture the binding of some number j together with event information about the i -th event E_i , when j is different than i , even if provided with (a) a secure binding of number i and E_i and (b) a secure binding of number j and E_j . For instance, the i -th log entry may consist of $\text{SIG}(i, E_i, \text{AI})$. This way, the deletion of the i -th log entry will be detected given later log entries. This is so because a later log entry may carry with it a number greater than i , which cannot be removed, modified or switched with another log-entry numbering information by the adversary, because it is securely bound to the log entry. For instance, assume the adversary deletes the log entry number 100: $\text{SIG}(100, E_{100}, \text{AI})$. As long as the adversary cannot delete all subsequent log entries (which would require continual access to the database), to hide his deletion, the adversary would need to create another log entry with the same number 100. However, this may be difficult because (a) the adversary cannot generate a brand new 100th log entry $\text{SIG}(100, E', \text{AI}')$ since he does not have the door's secret signing key; (b) the adversary cannot modify an existing log entry without invalidating the digital signature (e.g., cannot change $\text{SIG}(101, E_{101}, \text{AI}_{101})$ into $\text{SIG}(100, E_{101}, \text{AI}_{101})$ even if he remembers the deleted entry $\text{SIG}(100, E_{100}, \text{AI}_{100})$); (c) the adversary cannot extract a signature of a portion

of the log entry indicating the number 100 and bind it with a digital signature to another log entry.

Such secure binding can also be achieved by means other than digitally signing together the entry number and the event being numbered. For instance, it can be achieved by one-way-hashing the entry number and the event being numbered and then signing the hash, symbolically $SIG(H(i, E_i, AI))$. As for another example, it can be achieved by including the hash of the number into the digital signature of the event or vice versa: e.g., symbolically $SIG(i, H(E_i), AI)$. It can also be achieved by signing the numbering information together with the digital signature of the event information: e.g., symbolically $SIG(i, SIG(E_i), AI)$. As yet another alternative, one can separately sign (1) the numbering information together with a unique string x , and (2) the event information together with the string x , symbolically $(SIG(i, x), SIG(x, E_i, AI))$. (Such string x could be a nonce.)

Deletion-detectable logs may also be achieved by securely binding with the log entry correlation information other than sequential numbering information. For instance, one can include in log entry i some identifying information from a prior log entry, for example, entry $i-1$. Such information may be a collision-resistant hash of entry $i-1$ (or a portion of log entry $i-1$): symbolically, log entry i can be represented as $SIG(H(\text{log entry } i-1), E_i, AI)$. Then if the adversary attempts to remove log entry $i-1$, such removal would be detected when log entry i is received, because the hash of the previously received log entry, $H(\text{log entry } i-2)$, would not match $H(\text{log entry } i-1)$ (because of the collision-resistance of H), whereas $H(\text{log entry } i-1)$, because it is securely bound to log entry i , could not be modified by the adversary without destroying the validity of a digital signature. Here by log entry i we may also mean a subset of its information, such as E_i .

Note that it need not be log entry $i-1$ whose information is bound with entry i : it may be another prior or future entry, or, in fact, a multitude of other entries. Moreover, which log entries to bind with which ones may be chosen with the use of randomness.

Other correlation information may also be used. For instance, each log entry i may have securely bound with it two values (e.g., random values or nonces) x_i and x_{i+1} : symbolically, e.g., $SIG(x_i, x_{i+1}, E_i, AI)$. Then two consecutive log entries may always share one x value: for instance, entries i and $i+1$ will share x_{i+1} . However, if a log entry is deleted, this will no longer hold (because the adversary cannot modify signed log entries without detection unless it knows the secret key for the signature). For instance, if entry number 100 is deleted, the database will contain $SIG(x_{99}, x_{100}, E_{99}, AI)$ and $SIG(x_{101}, x_{102}, E_{101}, AI)$ and one can observe that they are not sharing a common x value. Such correlation information may take other forms: in fact, a log entry may be correlated with multiple other log entries. This can be accomplished, in particular, by use of polynomials to generate correlation information (e.g., two or more log entries may each contain the result of evaluating the same polynomial at different inputs). Such correlation information may also make use of a hash chain: for instance, starting with a value y_1 , let $y_2=H(y_1)$, $y_3=H(y_2)$, ..., etc., and then securely bind y_i with E_i : e.g., the i -th log entry may be symbolically represented as $SIG(y_i, E_i, AI)$. Then consecutive log entries i and $i+1$ may have correlation values y_i and y_{i+1} such that $y_{i+1} = H(y_i)$. If the adversary deletes a log entry, however, this may no longer hold and thus deletion can be detected. For instance, if entry 100 is deleted, the database will contain $SIG(y_{99}, E_{99}, AI)$ and $SIG(y_{101}, E_{101}, AI)$ (which, as before cannot be modified by the adversary without distorting the digital signatures). Then the deletion can be detected because $H(y_{101})$ will not match y_{99} . Use of multiple hash chains, perhaps used non-consecutive entries and in both directions, may also provide such correlation information.

In another embodiment, each log entry may contain an indication of some or all of the previous or even subsequent events, thus making logs not only deletion-detectable, but also reconstructible in case of deletions. Reconstructible log systems may be built by using (1) an authentication scheme (e.g., a digital signature scheme), (2) a reconstruction-information-generating scheme and (3) a reconstructing scheme as follows. Given one log event E (part of a sequence of —possibly past and/or future— events), the reconstruction-information-generating scheme is used to generate reconstruction information RI , which is then securely bound to other log entries by

means of the authentication scheme. The reconstruction-information-generating scheme ensures that, even if the log entry corresponding to event i is lost, other log entries contain sufficient information about E so as to allow reconstruction of E from RI present in other log entries. For instance, the $i+1^{\text{st}}$ entry may contain information
5 about all or some of the previous i events, generated by the reconstruction-information-generating scheme. Therefore, if an enemy succeeded somehow in erasing the j -th log entry from the database, information about the j -th event E_j will show up in one or more subsequent entries, making it possible to reconstruct information E_j even in the absence of the j -th log entry, using the reconstructing scheme. Thus, it would not be
10 enough for an enemy to have temporary access to the database: he would have to monitor the database "all the time" and delete multiple log entries to prevent information about the j -th event from being revealed. Choosing which events to include into a log entry can be done by the reconstruction-information-generating scheme in a randomized fashion, so as to make it harder for an enemy to predict when
15 information about a given event will show up in successive logs. Preferably, the system for reconstructible logs may also be deletion-detectable and indisputable.

Note also that reconstruction information about event j included into another log entry need not be direct. It may consist of a partial entry j , or of its hash value h_j (in particular, computed by the reconstruction-information-generating scheme via a one-
20 way/collision-resistant hash function), or of its digital signature, or of any other indication. In particular, if a one-way collision-resistant hash function H is used, then it is possible to indisputably restore information about the j -th event from a log entry i which contains h_j : symbolically, if the i -th entry is signed, the corresponding indisputable log may take the form $\text{SIG}(h_j, E_i, AI)$. For instance, if one suspects that a
25 particular user entered a particular door at a particular time, one can test if the value h_j matches the hash $H(E_j)$ of a log entry E_j that would have been created in response to such an event. This is indisputable because of the collision-resistance property of H : it is essentially impossible to come up with an entry E'_j different from E_j such that $H(E'_j)=H(E_j)$.

Log entries E_j may be created in a way that would make it easier for one to guess (and hence verify) what the log entry for a given event should be (for instance, by using a standardized format for log entries, using a coarse time granularity, etc.). One-way hash may be particularly useful because of its small size: it may be possible to

5 hash many or even all prior log entries for inclusion into a subsequent entry. For instance, entry $i+1$ can include $h_1=H(E_1)$, $h_2=H(E_2)$, ..., $h_i=H(E_i)$. Alternatively, one can nest (some of) the hashes, thus reducing the amount of space required. For instance, if one nested them all, then the second log entry would include $h_1=H(E_1)$, the third log entry would include $h_2 = H(E_2, h_1)$, Thus, if one can reconstruct or

10 observe log entries 1 through $i-1$ and log entry $i+1$, then one can indisputably reconstruct log entry i . This system may be improved by encrypting (some of) the information in a log entry (e.g., with a key known only to the database), so that the enemy cannot see which information he must destroy in order to compromise reconstructibility of a particular event. Actually, once the log is protected by

15 encryption, such encrypted logs (preferably indisputable encrypted logs) can be shipped to another (second) database without any loss of privacy. This makes deletions even more difficult for an enemy: now he has to gain access to two or more databases to falsify logs.

Reconstructible logs may also be achieved through use of error-correcting

20 codes. In particular, this can be done by generating multiple components ("shares") of each log entry and sending them separately (perhaps with other log entries) in such a way that, when sufficiently many shares have been received, the log entry may be reconstructed by the reconstructing scheme, which may invoke a decoding algorithm for the error-correcting code. These shares can be spread randomly or

25 pseudorandomly, thus making it harder for the adversary to remove sufficiently many of them to prevent reconstruction of a log entry when enough shares eventually arrive.

Event logs (whether created by cards, by doors, or jointly by both) may be carried by cards to facilitate their collection. When a card reaches a connected door, or communicates with a central server, or is otherwise able to communicate with the

30 central database, it can send the logs stored in it. This can be done similarly to the

dissemination of HRAs, except that HRAs may be sent from a central point to a card, whereas logs may be sent from the card to the central point. All the methods of disseminating HRAs, therefore, apply to the collection of event logs. Specifically, a method for disseminating HRAs can be transformed into a method for collecting event logs by (1) substituting a sender for the receiver and vice versa; (2) replacing an HRA with a log entry.

In particular, a card C1 may collect events logs for events unrelated to C1, such as access by another card C2, or a malfunction of a door D. Moreover, event logs for one door D1 may be stored (perhaps temporarily) on another door D2 (perhaps carried there by a card C1). Then, when another card C2 communicated with D2, it may receive some of these log entries and later communicate them to another door or to a central location. This broad dissemination may ensure that event logs reach the central point faster. (Moreover, it is possible that some doors, even though not fully connected to a central database, may have connections to each other. Such doors thus may exchange available event logs similarly. If cards have communication capability with each other —e.g., when in proximity— they may also exchange information about event logs that they store.) In such collecting process, indisputable logs are advantageous, because they do not need to be carried over secured channels, as they cannot be falsified. Therefore, they do not rely on the security of cards or connections between cards and doors. Deletion-detectable logs provide additional advantages by ensuring that, if some log entries are not collected (perhaps because some cards never reach a connected door), this fact may be detected. Reconstructible logs may additionally allow for reconstruction of log entries in case some log entries do not reach a central database (again, perhaps because some cards never reach a connected door).

In some instances, all event logs could be stored and disseminated in this manner. Else, it may be useful to adopt some optimizations. One optimization approach is to have event logs come with priority information, indicating the relative importance of informing a central authority about a particular event. Some log entries may be of more urgent interest than others: for instance, if a door is stuck in an open or closed position, if unauthorized access is attempted, or if unusual access pattern is

detected. In order to speed the delivery of such important information to the location where it can be acted upon, information in access logs may be labeled with tags indicating its importance (or its importance may be deduced from the information itself). For example, some log entries may be labeled "urgent" while others may be
5 labeled "routine;" or they may be labeled by numbers or codewords that indicate their degree of importance. (A gradation of priorities may be as fine or coarse as appropriate.) More effort or higher priority may be devoted to spreading information of higher importance. For instance, higher priority information may be given to more cards and/or doors in order to increase the likelihood that it will reach its destination
10 sooner or more surely. Also, a card or a door, when receiving information of high priority, may make room for it by removing low-priority information from its memory. Likewise, a door may decide to give high-priority information to every card that passes by, whereas low-priority information may be given to only a few cards or may wait until such time when the door is connected.

15 Alternatively or in addition to the above techniques, cards may be selected to store particular log entries in a way that involves randomness, or a door may provide a log entry to a certain number of cards (e.g., the first k cards it "encounters"). The use of such dissemination techniques may significantly reduce the likelihood that an important entry in an event log will be unable to reach the central location where it can
20 be acted upon. In particular, it may be used as an effective countermeasure against "jamming" attacks that attempt to prevent a broken door from communicating its distress. The actual method of exchanging the logs among cards and doors may vary. In case of a few entries, it may be most efficient to exchange and compare all known entries. In other cases, more sophisticated algorithms for finding and reconciling
25 differences may be in order.

It may be useful to present in detail some sample ways in which event logs may be collected. Below, "authority" A includes some central point or database in which event logs are collected.

SEQUENCE 1 (directly from door to authority):

1. Connected door D creates an indisputable log entry E in response to an event.
2. E is transmitted via wired or wireless communication to the authority A.
- 5 3. A verifies the authenticity of E and, if verification succeeds, stores E.

SEQUENCE 2 (from door to a user's card to authority):

1. Door D creates an indisputable log entry E in response to an event.
2. A card C of a user U that is presented for access to D receives and stores E (in addition to access-related communication). The card may or may not verify the authenticity of E.
- 10 3. When U leaves work and presents his card to A at the end of the work day, E is transmitted to A by the card.
4. A verifies the authenticity of E and, if verification succeeds, stores E.

SEQUENCE 3 (from door to a user's card to another (connected) door to authority):

- 15 1. Door D creates an indisputable log entry E in response to an event.
2. A card C of a user U that is presented for access to D receives and stores E (in addition to access-related communication). The card may or may not verify the authenticity of E.
3. Later, U presents his card C for access to another (connected) door D' .
- 20 D', in addition to verifying credentials and granting access if appropriate, receives E_i from C. D' may or may not verify the authenticity of E.
4. E is transmitted by D' via wired or wireless communication to the authority A.
- 25 5. A verifies the authenticity of E and, if verification succeeds, stores E.

Protected areas may be defined by walls and physical doors, such as doors through which a human may enter, or doors of a container, of a safe, of a vehicle, etc. Protected areas may also be defined by virtual doors and walls. For instance, an area

may be protected by a detector that can sense an intrusion, and possibly sound an alarm or send another signal if authorization is not provided. Such an alarm system is an example of a virtual door: in an airport, often entering the gate area through an exit lane will trigger such an alarm, even though no physical doors or walls have been violated.

5 Another example of a virtual door is a toll booth: even though many toll booths contain no physical bars or doors, a given car may or may not be authorized to go through the booth. Such authorization may depend, for instance, on the validity of a car's electronic toll billing token. Yet another example is that of a traffic control area. For instance, to enter the downtown of a given city, or a road leading to a nuclear facility,

10 an army barrack, or another sensitive area, a vehicle must have proper authorization, for purposes such as billing, security or congestion control.

In addition, protection may not be needed only for areas, but also for devices, such as airplane engines or military equipment. For instance, it may be necessary to ensure that only an authorized individual can start the engines of an airplane or of a

15 truck carrying hazardous materials.

There are many ways to use credentials/proofs for access control. Note that, for the disclosure herein, the term "day" below should be understood to mean general time period in a sequence of time periods, and "morning" to mean the beginning of a time period.

20 Throughout this application, "doors" should be construed to include all types of portals (e.g., physical and/or virtual), access-control systems/devices, and monitoring systems/devices. In particular, they include key mechanisms used to start engines and control equipment (so that our invention, in particular, can be used to ensure that only currently authorized users may start a plane, operate an earth-mover or otherwise

25 access and control various valuable and/or dangerous objects, devices and pieces of machinery). Consistently with this convention, we shall refer to "entering" as being granted the desired access (whether physical or virtual).

Similarly, for concreteness but without loss of generality intended, a card may be understood to mean any access device of a user. It should be understood that the notion of a card is sufficiently general to include cellular phones, PDAs, and other wireless and/or advanced devices, and a card may include or operate in conjunction
5 with other security measures, such as PINs, password and biometrics, though some of these may “reside” in the brain or body of the cardholder rather than in the card itself.

In addition, the expression “user” (often referred to as a “he” or “she”) broadly, may be understood to encompass not only users and people, but also devices, entities (and collections of users, devices and entities) including, without limitation, user cards.

10 The system described herein may be implemented using any appropriate combination of hardware and software including, without limitation, software stored in a computer readable medium that is accessed by one or more processors. In addition, the techniques used for encryption, authentication, etc. may be combined and used interchangeably, as appropriate. In that regard, each of the following U.S. patents and
15 applications is incorporated by reference herein:

- U.S. provisional patent application no. 60/004,796, filed October 2, 1995;
- U.S. provisional patent application no. 60/006,038 filed on October 24, 1995;
- U.S. provisional patent application no. 60/006,143, filed November 2, 1995;
- U.S. provisional patent application no. 60/024,786, filed September 10, 1996;
- 20 U.S. provisional patent application no. 60/025,128, filed August 29, 1996;
- U.S. provisional patent application no. 60/033,415, filed December 18, 1996;
- U.S. provisional patent application no. 60/035,119, filed February 3, 1997;
- U.S. provisional patent application no. 60/277,244, filed March 20, 2001;
- U.S. provisional patent application no. 60/300,621, filed June 25, 2001;
- 25 U.S. provisional patent application no. 60/344,245, filed December 27, 2001;
- U.S. provisional patent application no. 60/370,867, filed April 8, 2002;
- U.S. provisional patent application no. 60/372,951, filed April 16, 2002;
- U.S. provisional patent application no. 60/373,218, filed April 17, 2002;

- U.S. provisional patent application no. 60/374,861, filed April 23, 2002;
U.S. provisional patent application no. 60/420,795, filed October 23, 2002;
U.S. provisional patent application no. 60/421,197, filed October 25, 2002;
U.S. provisional patent application no. 60/421,756, filed October 28, 2002
5 U.S. provisional patent application no. 60/422,416, filed October 30, 2002;
U.S. provisional patent application no. 60/427,504, filed November 19, 2002;
U.S. provisional patent application no. 60/443,407, filed January 29, 2003;
U.S. provisional patent application no. 60/446,149, filed February 10, 2003;
U.S. provisional patent application no. 60/482,179 filed on June 24, 2003
10 U.S. provisional patent application no. 60/488,645 filed on July 18, 2003;
U.S. provisional patent application no. 60/505,640 filed on September 24, 2003;
U.S. patent application no. 08/715,712, filed September 19, 1996;
U.S. patent application no. 08/741,601, filed November 1, 1996;
U.S. patent application no. 08/756,720, filed November 26, 1996;
15 U.S. patent application no. 08/804,868, filed February 24, 1997;
U.S. patent application no. 08/804,869, filed February 24, 1997;
U.S. patent application no. 08/872,900, filed June 11, 1997;
U.S. patent application no. 08/906,464, filed August 5, 1997;
U.S. patent application no. 09/915,180, filed July 25, 2001;
20 U.S. patent application no. 10/103,541, filed March 20, 2002;
U.S. patent application no. 10/244,695 filed September 16, 2002;
U.S. patent application no. 10/409,638, filed on April 8, 2003;
U.S. patent application no. 10/876,275 filed on June 24, 2004;
U.S. Patent No. 5,604,804;
25 U.S. Patent No. 5,666,416
U.S. Patent No. 5,717,757;
U.S. Patent No. 5,717,758;
U.S. Patent No. 5,793,868;
U.S. Patent No. 5,960,083;
30 U.S. Patent No. 6,097,811; and
U.S. Patent No. 6,487,658.

While the invention has been disclosed in connection with various embodiments, modifications thereon will be readily apparent to those skilled in the art. Accordingly, the spirit and scope of the invention is set forth in the following claims.

What is claimed is:

1. A method of controlling access, comprising:
 - providing a barrier to access that includes a controller that selectively allows access;
 - 5 at least one administration entity generating credentials/proofs, wherein no valid proofs are determinable given only the credentials and values for expired proofs;
 - the controller receiving the credentials/proofs;
 - the controller determining if access is presently authorized; and
 - if access is presently authorized, the controller allowing access.
- 10 2. A method, according to Claim 1, wherein the credentials/proofs are in one part.
3. A method, according to Claim 1, wherein the credentials/proofs are in separate parts.
4. A method, according to Claim 3, wherein there is a first administration entity that generates the credentials and other administration entities that generate proofs.
5. A method, according to Claim 4, wherein the first administration entity also
15 generates proofs.
6. A method, according to Claim 4, wherein the first administration entity does not generate proofs.
7. A method, according to Claim 1, wherein the credentials correspond to a digital certificate that includes a final value that is a result of applying a one way function to a
20 first one of the proofs.
8. A method, according to Claim 7, wherein each of the proofs is a result of applying a one way function to a future one of the proofs.

9. A method, according to Claim 7, wherein the digital certificate includes an identifier for the electronic device.
10. A method, according to Claim 1, wherein the credentials include a final value that is a result of applying a one way function to a first one of the proofs.
- 5 11. A method, according to Claim 10, wherein each of the proofs is a result of applying a one way function to a future one of the proofs.
12. A method, according to Claim 1, wherein the credentials include an identifier for a user requesting access.
- 10 13. A method, according to Claim 1, wherein the credentials/proofs include a digital signature.
14. A method, according to Claim 1, wherein the barrier to access includes walls and a door.

15. A method, according to Claim 14, further comprising:
providing a door lock coupled to the controller, wherein the controller allowing access includes the controller actuating the door lock to allow the door to open.
16. A method, according to Claim 1, further comprising:
5 providing a reader coupled to the controller, wherein the controller receives credentials/proofs from the reader.
17. A method, according to Claim 16, wherein the credentials/proofs are provided on a smart card presented by a user.
18. A method, according to Claim 1, further comprising:
10 providing an external connection to the controller.
19. A method, according to Claim 18, wherein the external connection is intermittent.
20. A method, according to Claim 18, wherein the controller receives at least a portion of the credentials/proofs using the external connection.
21. A method, according to Claim 20, wherein the controller receives all of the
15 credentials/proofs using the external connection.

22. A method, according to Claim 20, further comprising:
providing a reader coupled to the controller, wherein the controller receives a remaining portion of the credentials/proofs from the reader.
23. A method, according to Claim 22, wherein the credentials/proofs are provided on a
5 smart card presented by a user.
24. A method, according to Claim 1, wherein the credentials/proofs include a password entered by a user.
25. A method, according to Claim 1, wherein the credentials/proofs include user biometric information.
- 10 26. A method, according to Claim 1, wherein the credentials/proofs include a handwritten signature.
27. A method, according to Claim 1, wherein the credentials/proofs include a secret value provided on a card held by a user.
- 15 28. A method, according to Claim 1, wherein the credentials/proofs expire at a predetermined time.




FIGURE 1A




FIGURE 1B

3/10




FIGURE 1C

4/10




FIGURE 1D

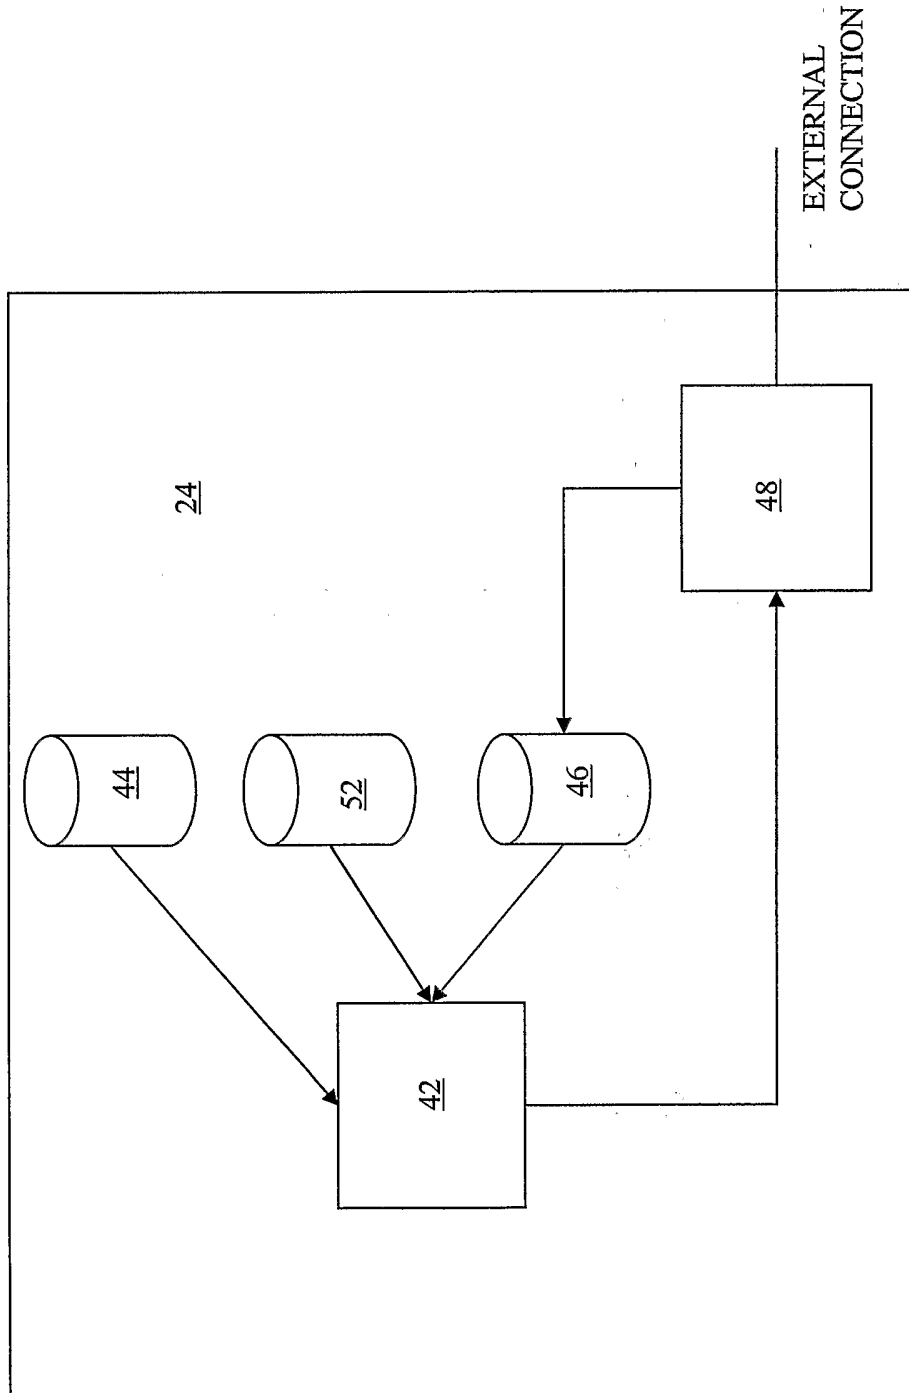


FIGURE 2




FIGURE 3

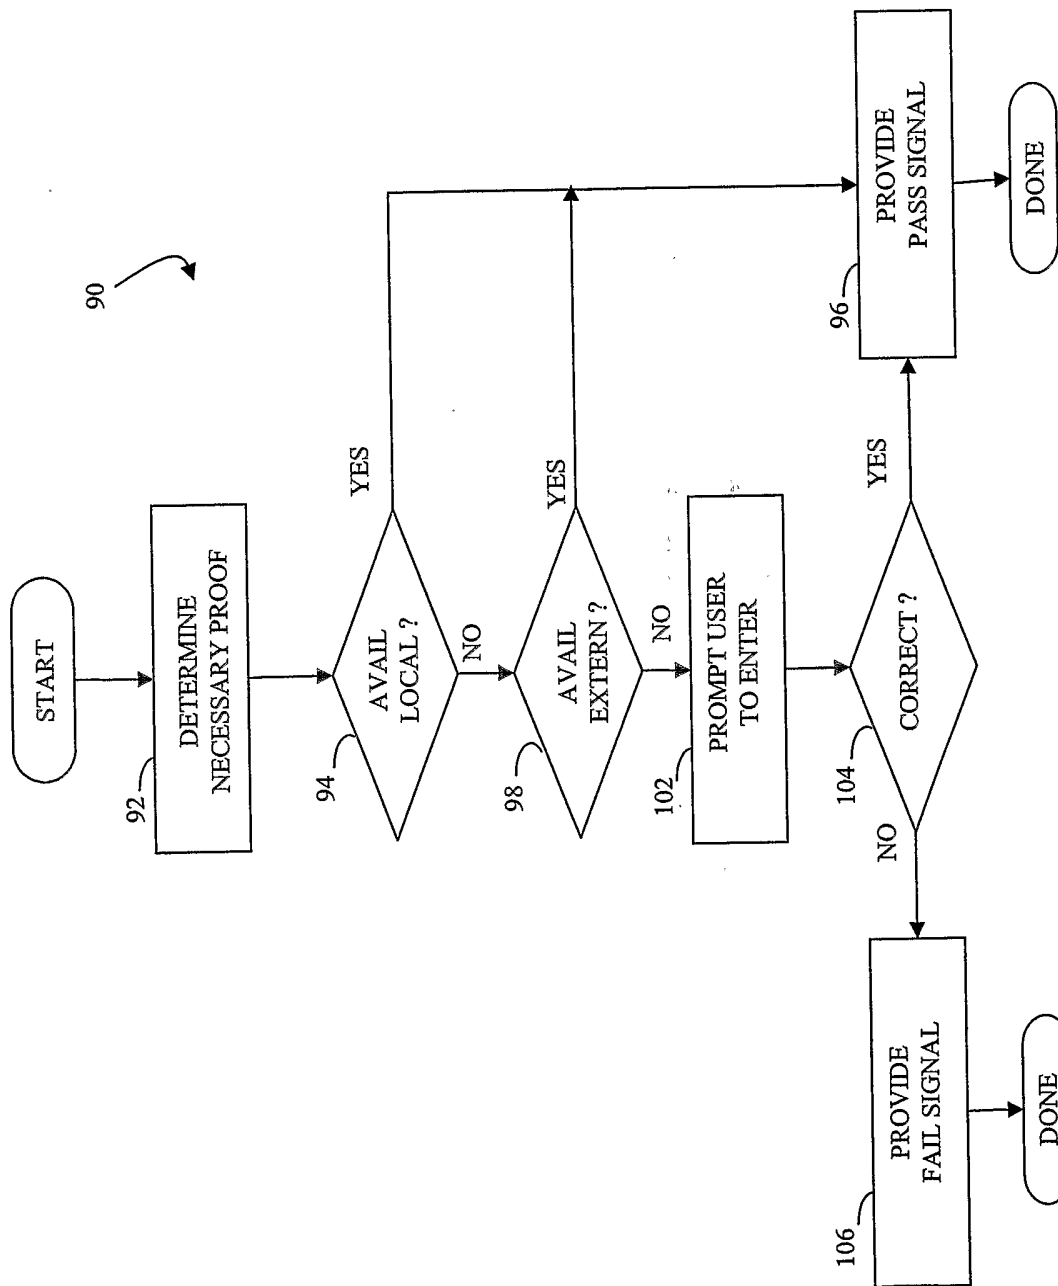


FIGURE 4

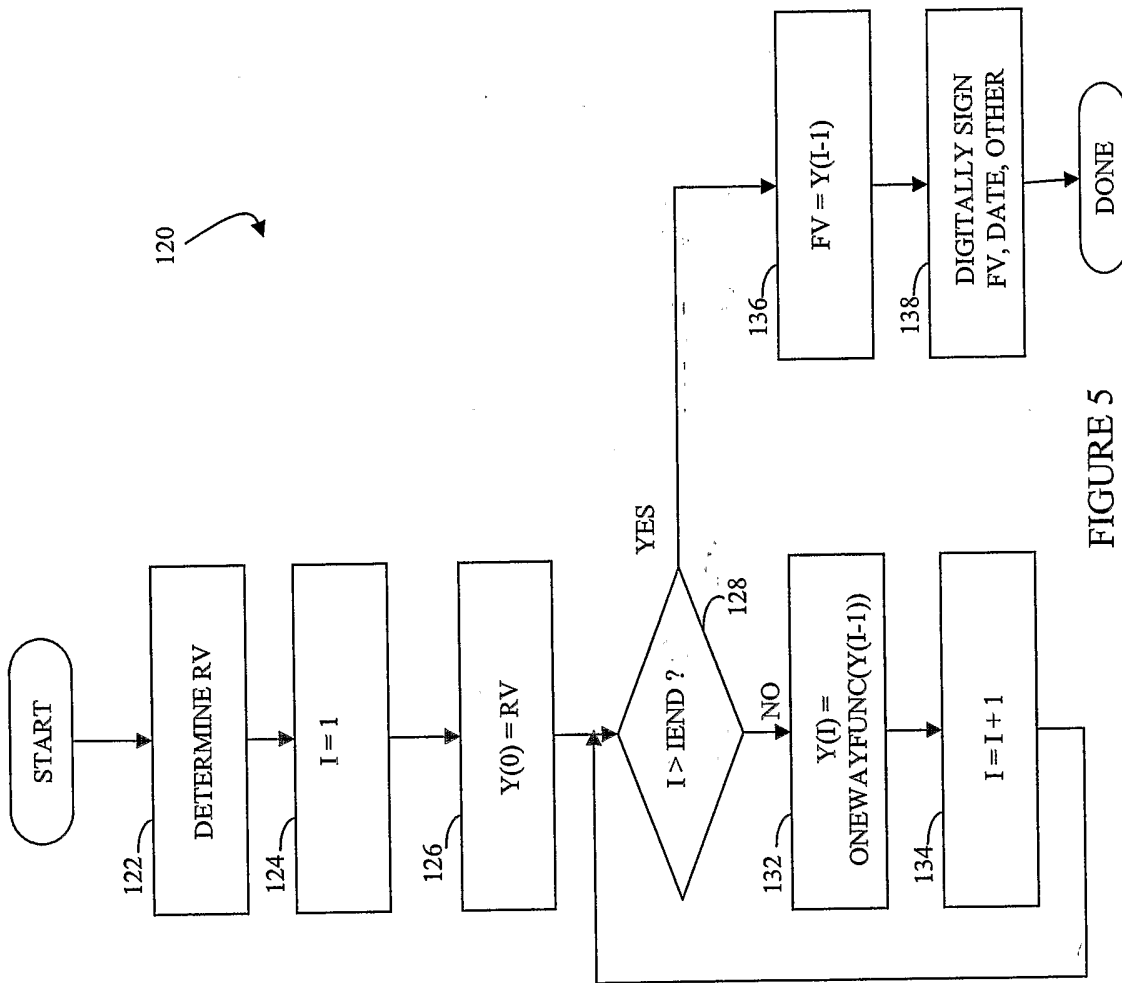


FIGURE 5

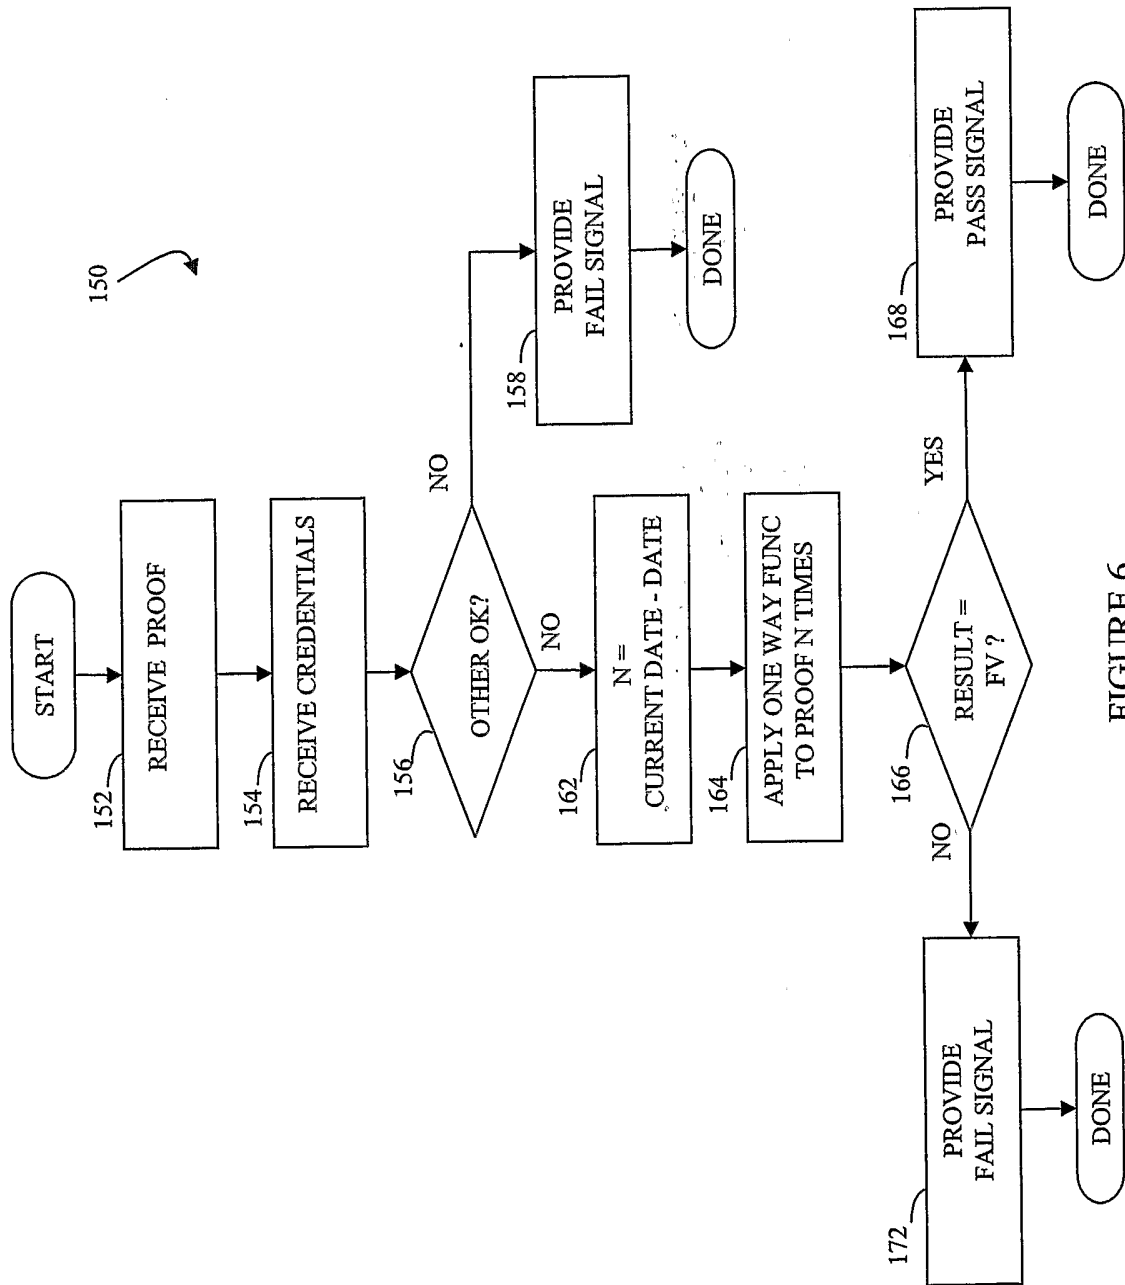


FIGURE 6




FIGURE 7