

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
1 February 2007 (01.02.2007)

PCT

(10) International Publication Number
WO 2007/014221 A2

(51) International Patent Classification:
G06F 12/00 (2006.01)

(21) International Application Number:
PCT/US2006/028872

(22) International Filing Date: 25 July 2006 (25.07.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/189,323 26 July 2005 (26.07.2005) US

(71) Applicant (for all designated States except US): **INVEN-
SYS SYSTEMS, INC.** [US/US]; 33 Commercial Street,
Foxboro, Massachusetts 02035 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **NARYZHNY,
Yevgeny, V.** [RU/US]; 4810 Park Newport, Newport
Beach, California 92660 (US). **AVERGUN, Mikhail**
[US/US]; 12618 Caminito Radiante, San Diego, California
92130 (US).

(74) Agent: **JOY, Mark**; Leydig, Voit & Mayer, Ltd., Two
Prudential Plaza, Suite 4900, Chicago, Illinois 60601 (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, HN, HR, HU, ID, IL, IN, IS, JP,
KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT,
LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA,
NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC,
SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ,
UA, UG, US, UZ, VC, VN, 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, IS, IT, LT, LU, LV, 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 "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND SYSTEM FOR TIME-WEIGHTED CACHE MANAGEMENT

(57) Abstract: Disclosed is a technique for managing memory items in a cache. An "age-lock" parameter is set to protect the newer memory items. When an incoming memory item (such as a history block header) is to be added to the cache, the amount of free space in the cache is checked. If there is insufficient free space for the incoming memory item, then space is freed up by removing memory items from the cache. No memory items protected by the age-lock parameter are removed. Of the older items, the selection for removal follows any of a number of well know cache management techniques, such as the "least recently used" algorithm. A "maximum size" parameter can be set for the cache. If the cache exceeds this maximum size, then free space is released and memory items are removed to decrease the cache size.



WO 2007/014221 A2

METHOD AND SYSTEM FOR TIME-WEIGHTED CACHE MANAGEMENT

TECHNICAL FIELD

[0001] The present invention generally relates to computing and networked data storage systems, and, more particularly, to techniques for managing (e.g., storing, retrieving, processing, etc.) streams of supervisory process control and manufacturing and production information. Such information is typically rendered and stored in the context of supervising automated industrial processes.

BACKGROUND OF THE INVENTION

[0002] Industry increasingly depends upon highly automated data acquisition and control systems to ensure that industrial processes are run efficiently and reliably while lowering the overall production costs. Data acquisition begins when a number of sensors measure aspects of an industrial process and report their measurements back to a data collection and control system. Such measurements come in a wide variety of forms. By way of example the measurements produced by a sensor include: a temperature, a pressure, a pH, a mass or volume flow of material, a counter of items passing through a particular machine or process, a tallied inventory of packages waiting in a shipping line, cycle completions, a photograph of a room in a factory, etc. Often sophisticated process management and control software examines the incoming data associated with an industrial process, produces status reports and operation summaries, and, in many cases, responds to events and to operator instructions by sending commands to controllers that modify operation of at least a portion of the industrial process. The data produced by the sensors also allow an operator to perform a number of supervisory tasks including: tailoring the process (e.g., specifying new setpoints) in response to varying external conditions (including costs of raw materials), detecting an inefficient/non-optimal operating condition or impending equipment failure, and taking remedial action such as moving equipment into and out of service as required.

[0003] A simple and familiar example of a data acquisition and control system is a thermostat-controlled home heating and air conditioning system. A thermometer measures a current temperature; the measurement is compared with a desired temperature range; and, if necessary, commands are sent to a furnace or cooling unit to achieve a

desired temperature. Furthermore, a user can program or manually set the controller to have particular setpoint temperatures at certain time intervals of the day.

[0004] Typical industrial processes are substantially more complex than the above described simple thermostat example. In fact, it is not unheard of to have thousands or even tens of thousands of sensors and control elements (e.g., valve actuators) monitoring and controlling all aspects of a multi-stage process within an industrial plant. The amount of data sent for each measurement and the frequency of the measurements varies from sensor to sensor in a system. For accuracy and to facilitate quick notice and response of plant events and upset conditions, some of these sensors update and transmit their measurements several times every second. When multiplied by thousands of sensors and control elements, the volume of data generated by a plant's supervisory process control and plant information system can be very large.

[0005] Specialized process control and manufacturing and production information data storage facilities (also referred to as plant historians) have been developed to handle the potentially massive amounts of production information generated by the aforementioned systems. An example of such a system is the WONDERWARE IndustrialSQL Server historian. A data acquisition service associated with the historian collects time-series data from a variety of data sources (e.g., data access servers). The collected data are thereafter deposited with the historian to achieve data access efficiency and querying benefits and capabilities of the historian's relational database. Through its relational database, the historian integrates plant data with event, summary, production, and configuration information.

[0006] Traditionally, plant historians have collected and archived streams of raw data representing process, plant, and production status over the course of time. The status data are of value for purposes of maintaining a record of plant performance and for presenting and recreating the state of a process or plant equipment at a particular point in time. However, individual pieces of data taken at single points in time are often insufficient to discern whether an industrial process is operating properly or optimally. Further

processing of the raw data often renders more useful information for operator decision making.

[0007] Over the years vast improvements have occurred with regard to networks, data storage and processor device capacity, and processing speeds. Notwithstanding such improvements, supervisory process control and manufacturing information system designs encounter a need to either increase system capacity and speed or to forgo saving certain types of information derived from raw data because creating and maintaining the information on a full-time basis draws too heavily from available storage and processor resources. Thus, while valuable, certain types of process information are potentially not available in certain environments. Such choices can arise, for example, in large production systems where processing raw data to render secondary information is potentially of greatest value.

SUMMARY OF THE INVENTION

[0008] In view of the foregoing, the present invention provides a technique for managing memory items in a cache. An "age-lock" parameter is set to protect the newer memory items. When an incoming memory item (such as a history block header) is to be added to the cache, the amount of free space in the cache is checked. If there is insufficient free space for the incoming memory item, then space is freed up by removing memory items from the cache. No memory items protected by the age-lock parameter are removed. Of the older items, the selection for removal follows any of a number of well know cache management techniques, such as the "least recently used" algorithm.

[0009] In some embodiments, a "maximum size" parameter is set for the cache. If the cache exceeds this maximum size, then free space is released and memory items are removed to decrease the cache size. A conflict between this parameter and the age-lock parameter can arise when the combined size of all memory items protected by the age-lock parameter is greater than the maximum size parameter. Some embodiments address this conflict by overriding the age-lock parameter, that is, by removing memory items that would otherwise be protected by that parameter. Other embodiments address the

conflict by overriding the maximum size parameter, that is, by allowing the cache to grow to include room for all the memory items protected by the age-lock parameter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] While the appended claims set forth the features of the present invention with particularity, the invention, together with its objects and advantages, may be best understood from the following detailed description taken in conjunction with the accompanying drawings of which:

[0011] Figure 1 is a schematic diagram of an exemplary networked environment wherein a process control database server embodying the present invention is advantageously incorporated;

[0012] Figure 2 is a schematic drawing of functional and structural aspects of a historian service embodying the present invention;

[0013] Figure 3 is a schematic diagram of an exemplary history block storage mechanism managed by the historian;

[0014] Figure 4 is a flowchart of a technique for managing the storage of history blocks according to the present invention;

[0015] Figure 5 is a schematic diagram of an exemplary cache memory mechanism managed by the historian; and

[0016] Figure 6 is a flowchart of a technique for managing a cache memory according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0017] As noted previously in the background, a plant information historian service maintains a database comprising a wide variety of plant status information. The raw plant status information, when provided to operations managers in its unprocessed form, offers limited comparative information such as how a process or the operation of plant equipment has changed over time. In many cases, performing additional analysis on raw data streams to render secondary information greatly enhances the information value of

the raw data. In embodiments of the invention, such analysis is delayed until a client requests such secondary information from the historian service for a particular timeframe. As such, limited historian memory and processor resources are only allocated to the extent that a client of the historian service has requested the secondary information. In particular, the historian service supports a set of advanced data retrieval operations wherein raw data are processed to render particular types of secondary information “on demand” and in response to “client requests.”

[0018] The terms “client requests” and “on demand” are intended to be broadly defined. The plant historian service embodying the present invention does not distinguish between requests arising from human users and requests originating from automated processes. Thus, a “client request,” unless specifically noted, includes requests initiated by human/machine interface users and requests initiated by automated client processes. The automated client processes potentially include processes running on the same node as the historian service. The automated client processes request the secondary information and thereafter provide the received secondary information, in a service role, to others. Furthermore, the definition of “on demand” is intended to include both providing secondary information in response to specific requests as well as in accordance with a previously established subscription. By performing the calculations to render the secondary information on demand, rather than calculating (and tabling) the information without regard to whether it will ever be requested by a client, the historian system embodying the present invention is better suited to support a very broad and extensible set of secondary information types meeting diverse needs of a broad variety of historian service clients.

[0019] In an embodiment of the present invention, the historian service supports a variety of advanced retrieval operations for calculating and providing, on demand, a variety of secondary information types from raw data previously stored in the historian database. Among others, the historian service specifically includes the following advanced data retrieval operations: “time-in-state,” “counter,” “engineering units-based integral,” and “derivative.” “Time-in-state” calculations render statistical information relating to an amount of time spent in specified states. Such states are represented, for

example, by identified tag/value combinations. By way of example, the time-in-state statistics include, for a specified time span and tagged state value: total amount of time in the state, percentage of time in the state, the shortest time in the state, and the longest time in the state.

[0020] The following description is based on illustrative embodiments of the invention and should not be taken as limiting the invention with regard to alternative embodiments that are not explicitly described herein. Those skilled in the art will readily appreciate that the example of Figure 1 represents a simplified configuration used for illustrative purposes. In many cases, the systems within which the present invention is incorporated are substantially larger. The volume of information handled by a historian in such a system would generally preclude pre-calculating and storing every type of information potentially needed by clients of the historian.

[0021] Figure 1 depicts an illustrative environment wherein a supervisory process control and manufacturing/production information data storage facility (also referred to as a plant historian) 100 embodying the present invention is potentially incorporated. The network environment includes a plant floor network 101 to which a set of process control and manufacturing information data sources 102 are connected either directly or indirectly (via any of a variety of networked devices including concentrators, gateways, integrators, interfaces, etc.). The data sources 102 comprise, for example, programmable logic controllers that are coupled to, communicate with, and control a variety of devices such as plant floor equipment, sensors, and actuators. A set of I/O (input/output) servers 104 (also referred to herein as data access servers) acquire data from the data sources 102 via the plant floor network 101.

[0022] The exemplary network environment includes a production network 110. In the illustrative embodiment the production network 110 comprises a set of human/machine interface (HMI) nodes 112 that execute plant floor visualization applications supported, for example, by Wonderware's INTOUCH visualization application management software. The data driving the visualization applications on the HMI nodes 112 are acquired, by way of example, from the plant historian 100 that also

resides on the production network 110. The historian 100 includes services for maintaining and providing a variety of plant, process, and production information including historical plant status, configuration, event, and summary information.

[0023] A data acquisition service 116, interposed between the I/O servers 104 and the plant historian 100, operates to maintain a continuous, up-to-date, flow of streaming plant data between the data sources 102 and the historian 100 for plant/production supervisors (both human and automated). The data acquisition service 116 acquires and integrates data (potentially in a variety of forms associated with various protocols) from a variety of sources into a plant information database, including time-stamped data entries, incorporated within the historian 100.

[0024] The physical connection between the data acquisition service 116 and the I/O servers 104 can take any of a number of forms. For example, the data acquisition service 116 and the I/O servers can be distinct nodes on the same network (e.g., the plant floor network 110). However, in alternative embodiments the I/O servers 104 communicate with the data acquisition service 116 via a network link that is separate and distinct from the plant floor network 101. In an illustrative example, the physical network links between the I/O servers 104 and the data acquisition service 116 comprise local area network links (e.g., Ethernet) that are generally fast, reliable, and stable and thus do not typically constitute a data-stream bottleneck or source of intermittent network connectivity.

[0025] The connection between the data acquisition service 116 and the historian 100 can also take any of a variety of forms. In an embodiment of the present invention, the physical connection comprises an intermittent or slow connection 118 that is potentially too slow to handle a burst of data, unavailable, or faulty. The data acquisition service 116 therefore includes components and logic for handling the stream of data from components connected to the plant floor network 101. In view of the potential throughput and connectivity limitations of connection 118, to the extent secondary information is to be generated or provided to clients of the historian 100 (e.g., HMI nodes 112), such information should be rendered after the data have traversed the connection 118. In an

embodiment, the secondary information is rendered by advanced data retrieval operations incorporated into the historian 100.

[0026] Figure 2 depicts functional components associated with the historian 100. The historian 100 generally implements a storage interface 200 comprising a set of functions and operations for receiving and tabling data from the data acquisition service 116 via the connection 118. The received data are stored in one or more tables 202 maintained by the historian 100.

[0027] By way of example, the tables 202 include pieces of data received by the historian 100 via a data acquisition interface to a process control and production information network such as the data acquisition service 116 on network 101. In the illustrative embodiment each data piece is stored in the form of a value, a quality, and a timestamp. These three parts to each data piece stored in the tables 202 of the historian 100 are described briefly below.

[0028] Timestamp: The historian 100 tables data received from a variety of “real-time” data sources, including the I/O Servers 104 (via the data acquisition service 116). The historian 100 is also capable of accepting “old” data from sources such as text files. Traditionally, “real-time” data exclude data with timestamps outside of ± 30 seconds of a current time of a clock maintained by a computer node hosting the historian 100. However, real-time data with a timestamp falling outside the 30-second window are addressable by a quality descriptor associated with the received data. Proper implementation of timestamps requires synchronization of the clocks utilized by the historian 100 and data sources.

[0029] Quality: The historian 100 supports two descriptors of data quality: “QualityDetail” and “Quality.” The QualityDetail descriptor is based primarily on the quality of the data presented by the data source, while the Quality descriptor is a simple indicator of “good,” “bad,” or “doubtful,” derived at retrieval time. Alternatively, the historian 100 supports an OPCQuality descriptor that is intended to be used as a sole data quality indicator that is fully compliant with OPC quality standards. In an alternative embodiment, the QualityDetail descriptor is utilized as an internal data quality indicator.

[0030] Value: A value part of a stored piece of data corresponds to a value of a received piece of data. In exceptional cases, the value obtained from a data source is translated into a NULL value at the highest retrieval layer to indicate a special event, such as a data source disconnection. This behavior is closely related to quality, and clients typically leverage knowledge of the rules governing the translation to indicate a lack of data, for example by showing a gap on a trend display.

[0031] The following is a brief description of the manner in which the historian 100 receives data from a real-time data source and stores the data as a timestamp, quality, and value combination in one or more of its tables 202. The historian 100 receives a data point for a particular tag (named data value) via the storage interface 200. The historian compares the timestamp on the received data to (1) a current time specified by a clock on the node that hosts the historian 100 and (2) a timestamp of a previous data point received for the tag. If the timestamp of the received data point is earlier than or equal to the current time on the historian node then:

If the timestamp on the received data point is later than the timestamp of the previous point received for the tag, the received point is tabled with the timestamp provided by the real-time data source.

If the time stamp on the received data point is earlier than the timestamp of the previous point received for the tag (i.e., the point is out of sequence), the received point is tabled with the timestamp of the previously tabled data point "plus 5 milliseconds." A special QualityDetail value is stored with the received point to indicate that it is out of sequence. (The original quality received from the data source is stored in the "quality" descriptor field for the stored data point.)

[0032] On the other hand, if the timestamp of the point is later than the current time on the historian 100's node (i.e., the point is in the future), the point is tabled with a time stamp equal to the current time of the historian 100's node. Furthermore, a special value is assigned to the QualityDetail descriptor for the received and tabled point value to

indicate that its specified time was in the future. (The original quality received from the data source is stored in the “quality” descriptor field for the stored data point.)

[0033] The historian 100 can be configured to provide the timestamp for received data identified by a particular tag. After proper designation, the historian 100 recognizes that the tag identified by a received data point belongs to a set of tags for which the historian 100 supplies a timestamp. Thereafter, the time stamp of the point is replaced by the current time of the historian 100’s node. A special QualityDetail value is stored for the stored point to indicate that it was timestamped by the historian 100. The original quality received from the data source is stored in the “quality” descriptor field for the stored data point.

[0034] It is further noted that the historian 100 supports application of a rate deadband filter to reject new data points for a particular tag where a value associated with the received point has not changed sufficiently from a previous stored value for the tag.

[0035] Having described a data storage interface for the historian 100, attention is directed to retrieving the stored data from the tables 202 of the historian 100. Access, by clients of the historian 100, to the stored contents of the tables 202 is facilitated by a retrieval interface 206. The retrieval interface 206 exposes a set of functions, operations, and methods (including a set of advanced data retrieval operations 204), callable by clients on the network 110 (e.g., HMI clients 112), for querying the contents of the tables 202. The advanced data retrieval operations 204 generate secondary information, on demand, by post-processing raw data stored in the tables 202. In response to receiving a query message identifying one of the advanced data retrieval options carried out by the operations 204, the retrieval interface 206 invokes the identified one of the set of advanced data retrieval operations 204 supported by the historian 100.

[0036] The plant historian 100 implements novel memory management techniques so that it can handle the enormous amount of data coming into it from the data sources 102. (The data coming in may be “indirect,” such as, for example, the headers of history blocks. History blocks directly store data coming from the data sources 102, while the headers are created to manage the history blocks.) Figures 3 through 6 and the

accompanying text illustrate some of these techniques. In this discussion, details of the data themselves are abstracted out in order to focus more particularly on the novel management techniques.

[0037] Figure 3 shows the plant historian 100 with two disk-based memory stores 300 and 302. The illustration is not meant to suggest that these two memory stores 300 and 302 must be located physically within the plant historian server 100, but rather shows that these stores are controlled by the plant historian 100. Some process control environments embody the present invention with only one memory store 300, while others may have more than the two illustrated in Figure 3.

[0038] The two memory stores 300 and 302 contain, to be as general as possible, “memory items” 304. For example, these memory items 304 could be history blocks, one created each day, that capture the evolving status of the process control environment. Each of the memory items 304 contains (in addition to whatever else) creation (or storage) time information, as suggested by the timeline 306 ranking the memory items 304 from older to newer.

[0039] The remaining items in Figure 3 are best illustrated in conjunction with Figure 4. Figure 4 presents the logic of a novel technique usable by the plant historian 100 to manage storage space requirements for the flood of memory items 304 coming into the memory stores 300 and 302. In step 400 of Figure 4, a “free-space size threshold” is set. This threshold is shown by the dashed line 308 of Figure 3. An age parameter is set in step 402 of Figure 4 and is illustrated by the dashed line 312 of Figure 3.

[0040] In step 404 of Figure 4, the amount of free space in the memory store 300 is compared against the free-space size threshold 308. The amount of free space can decrease when, for example, a new memory item 310 is received for storage in the memory store 300. Considering Figure 3, it is clear that if the new memory item 310 is added on “top” of the memory items 304, then the “lowest” (that is, the oldest) memory item 304 will be pushed down below the free-space size threshold 308. In some implementations, this may be acceptable in the short term, but an effort will be made to increase the amount of free space in the memory store 300. This effort is made in step

406 where memory items 304 are removed, oldest first, from the memory store 300 until the amount of free space is greater than the free-space size threshold 308.

[0041] As a complementary memory management technique, in step 408 the plant historian 100 periodically checks the age of its memory items 304 against the age parameter 312. Any memory items 304 older than this (as are two in Figure 3) are removed.

[0042] In those implementations with multiple memory stores, the memory items 304 removed from the first memory store 300 are moved to the second memory store 302. In this “chain” of memory stores, the first memory store 300 can be on a very fast, but expensive and space-limited, disk drive. Additional memory stores can be slower and less expensive. This chain is cheaper and potentially more reliable than having a single, very large memory store, while allowing the first memory store 300 to be as fast as possible to support human interactions and critical control processes.

[0043] The second memory store 302 (and any subsequent ones) can implement the same memory management techniques described above for the first memory store 300 or can implement other techniques.

[0044] Figures 5 and 6 together illustrate another memory management technique usable by the plant historian 100. Rather than the disk-based memory stores 300 and 302 of Figure 3, the cache 500 of Figure 5 is RAM-based. This fact and the fact that all of the memory items 502 in the cache 500 are reproduced elsewhere (such as on disk) present the need and the opportunity for a different memory management technique.

[0045] Instead of being of a fixed size, the cache 500 can grow (“downward” in Figure 5) when necessary. Limiting this growth is a “maximum size parameter” set for the cache 500 in step 600 of Figure 6 and illustrated in Figure 5 by the dashed line 510.

[0046] Because the cache 500 is used to provide fast retrieval of its memory items 502, some of those memory items 502 can be “locked” in the cache. The “age-lock parameter,” set in step 602 of Figure 6 and illustrated by the dashed line 508 of Figure 5, implements this locking as discussed below.

[0047] When a new memory item 506, such as a history block header, is to be added to the cache 500, the size of this new item 506 is compared against the amount of free space in the cache 500 (step 606). If the amount of free space is insufficient, then the two management techniques of steps 608 and 610 come into play. Free space is created in step 608 by removing memory items 502 from the cache 500. Memory items 502 can be selected for removal via the well known "least recently used" algorithm. However, no items newer than the age-lock parameter 508 are allowed to be removed. Alternatively, in step 610, free space is created simply by increasing the size of the cache 500, subject to the maximum size parameter 510.

[0048] It may happen that steps 608 and 610, either separately or in conjunction, cannot provide enough free space for the new memory item 506. For example, the age-lock parameter 508 can be set to lock so many memory items 502 that the cache 500 runs into the maximum size parameter 510. In this case, a compromise is reached by either discarding a memory item 502 otherwise protected by the age-lock parameter 508 or by allowing the cache 500 to expand beyond the maximum size parameter 510.

[0049] In any case, the new memory item 506 is added to the now sufficient free space in step 612 of Figure 6.

[0050] In view of the many possible embodiments to which the principles of the present invention may be applied, it should be recognized that the embodiments described herein with respect to the drawing figures are meant to be illustrative only and should not be taken as limiting the scope of the invention. Those of skill in the art will recognize that some implementation details are determined by specific situations. Although the environment of the invention is described in terms of software modules or components, some processes may be equivalently performed by hardware components. Therefore, the invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.

CLAIMS

We claim:

1. In an industrial process control environment, a method for managing a cache of memory items, the method comprising:
 - setting an age-lock parameter for the cache, the age-lock parameter specifying a duration;
 - loading a memory item that is not in the cache;
 - comparing a size of the loaded memory item with an amount of free space in the cache;
 - if the size of the loaded memory item is greater than the amount of free space in the cache, then selecting and removing memory items already in the cache and marking as free the space previously occupied by the removed memory items until the amount of free space in the cache is greater than or equal to the size of the loaded memory item, wherein selecting memory items to be removed comprises not selecting any memory items newer than the duration specified by the age-lock parameter; and
 - adding the loaded memory item to the cache.
2. The method of claim 1 wherein the loaded memory item comprises a history block header.
3. The method of claim 1 wherein loading a memory item that is not in the cache comprises loading a memory item from a long term storage.
4. The method of claim 1 wherein selecting memory items to be removed comprises applying a "least recently used" algorithm.
5. The method of claim 1 further comprising:
 - setting a maximum size parameter for the cache.

6. The method of claim 5 further comprising:

if a current size of the cache is greater than the set maximum size for the cache, then releasing free space in the cache, if any, and selecting and removing memory items already in the cache and releasing the space previously occupied by the removed memory items until the current size of the cache is less than or equal to the set maximum size for the cache, wherein selecting memory items to be removed comprises not selecting any memory items newer than the duration specified by the age-lock parameter.

7. The method of claim 5 further comprising:

if a total size of memory items in the cache newer than the duration specified by the age-lock parameter is greater than the set maximum size parameter for the cache, then overriding the set maximum size parameter for the cache and retaining the memory items in the cache newer than the duration specified by the age-lock parameter.

8. The method of claim 5 further comprising:

if a total size of memory items in the cache newer than the duration specified by the age-lock parameter is greater than the set maximum size parameter for the cache, then overriding the age-lock parameter and selecting and removing memory items already in the cache and releasing the space previously occupied by the removed memory items until the current size of the cache is less than or equal to the set maximum size for the cache.

9. A computer-readable medium having computer-executable instructions for performing a method for managing a cache of memory items, the method comprising:
 - setting an age-lock parameter for the cache, the age-lock parameter specifying a duration;
 - loading a memory item that is not in the cache;
 - comparing a size of the loaded memory item with an amount of free space in the cache;
 - if the size of the loaded memory item is greater than the amount of free space in the cache, then selecting and removing memory items already in the cache and marking as free the space previously occupied by the removed memory items until the amount of free space in the cache is greater than or equal to the size of the loaded memory item, wherein selecting memory items to be removed comprises not selecting any memory items newer than the duration specified by the age-lock parameter; and
 - adding the loaded memory item to the cache.
10. The computer-readable medium of claim 9 wherein the loaded memory item comprises a history block header.
11. The computer-readable medium of claim 9 wherein loading a memory item that is not in the cache comprises loading a memory item from a long term storage.
12. The computer-readable medium of claim 9 wherein selecting memory items to be removed comprises applying a "least recently used" algorithm.
13. The computer-readable medium of claim 9 wherein the method further comprises:
 - setting a maximum size parameter for the cache.

14. The computer-readable medium of claim 13 wherein the method further comprises:

if a current size of the cache is greater than the set maximum size for the cache, then releasing free space in the cache, if any, and selecting and removing memory items already in the cache and releasing the space previously occupied by the removed memory items until the current size of the cache is less than or equal to the set maximum size for the cache, wherein selecting memory items to be removed comprises not selecting any memory items newer than the duration specified by the age-lock parameter.

15. The computer-readable medium of claim 13 wherein the method further comprises:

if a total size of memory items in the cache newer than the duration specified by the age-lock parameter is greater than the set maximum size parameter for the cache, then overriding the set maximum size parameter for the cache and retaining the memory items in the cache newer than the duration specified by the age-lock parameter.

16. The computer-readable medium of claim 13 wherein the method further comprises:

if a total size of memory items in the cache newer than the duration specified by the age-lock parameter is greater than the set maximum size parameter for the cache, then overriding the age-lock parameter and selecting and removing memory items already in the cache and releasing the space previously occupied by the removed memory items until the current size of the cache is less than or equal to the set maximum size for the cache.

17. In an industrial process control environment, a system for managing memory items, the system comprising:
 - the memory items;
 - a cache configured for storing the memory items; and
 - an age-lock parameter for the cache, the age-lock parameter specifying a duration;wherein the cache is further configured for:
 - loading a memory item that is not in the cache;
 - comparing a size of the loaded memory item with an amount of free space in the cache;
 - if the size of the loaded memory item is greater than the amount of free space in the cache, then selecting and removing memory items already in the cache and marking as free the space previously occupied by the removed memory items until the amount of free space in the cache is greater than or equal to the size of the loaded memory item, wherein selecting memory items to be removed comprises not selecting any memory items newer than the duration specified by the age-lock parameter;
 - and
 - adding the loaded memory item to the cache.
18. The system of claim 17 wherein the loaded memory item comprises a history block header.
19. The system of claim 17 wherein loading a memory item that is not in the cache comprises loading a memory item from a long term storage.
20. The system of claim 17 wherein selecting memory items to be removed comprises applying a "least recently used" algorithm.
21. The system of claim 17 further comprising:
 - a maximum size parameter for the cache.

22. The system of claim 21 wherein the cache is further configured for:
- if a current size of the cache is greater than the set maximum size for the cache, then releasing free space in the cache, if any, and selecting and removing memory items already in the cache and releasing the space previously occupied by the removed memory items until the current size of the cache is less than or equal to the set maximum size for the cache, wherein selecting memory items to be removed comprises not selecting any memory items newer than the duration specified by the age-lock parameter.
23. The system of claim 21 wherein the cache is further configured for:
- if a total size of memory items in the cache newer than the duration specified by the age-lock parameter is greater than the set maximum size parameter for the cache, then overriding the set maximum size parameter for the cache and retaining the memory items in the cache newer than the duration specified by the age-lock parameter.
24. The system of claim 21 wherein the cache is further configured for:
- if a total size of memory items in the cache newer than the duration specified by the age-lock parameter is greater than the set maximum size parameter for the cache, then overriding the age-lock parameter and selecting and removing memory items already in the cache and releasing the space previously occupied by the removed memory items until the current size of the cache is less than or equal to the set maximum size for the cache.




FIG. 1




FIG. 2

FIG. 3





FIG. 4






FIG. 5

FIG. 6

