

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 June 2003 (12.06.2003)

PCT

(10) International Publication Number
WO 03/049343 A1

(51) International Patent Classification⁷: **H04J 3/06**

US 60/337,366 (CIP)
Filed on 30 November 2001 (30.11.2001)

(21) International Application Number: PCT/US02/36799

(22) International Filing Date:
14 November 2002 (14.11.2002)

(71) Applicant (for all designated States except US):
SCHLUMBERGERSEMA INC. [US/US]; 100 Mil-
ton Park, 30000 Mill Creek Avenue, Suite 100 - Patents,
Alpharetta, GA 30022 (US).

(25) Filing Language: English

(26) Publication Language: English


(30) Priority Data:
60/337,366 30 November 2001 (30.11.2001) US
10/280,448 25 October 2002 (25.10.2002) US

(72) Inventors; and
(75) Inventors/Applicants (for US only): **SALT, Metin**
[TR/US]; 134 Kia Drive, Alpharetta, GA 30022 (US).
RAGLE, Stephen, J. [US/US]; 470 Carrwick Circle,
Alpharetta, GA 30022 (US). **BENGTSSON, Johan**
[SE/US]; 751 Laurel Street, #110, San Carlos, CA 94070
(US). **PAIK, Ellen** [US/US]; 1103 Cresthaven Drive,
Silver Springs, MD 20903 (US).

(63) Related by continuation (CON) or continuation-in-part
(CIP) to earlier applications:
US 10/280,448 (CIP)
Filed on 25 October 2002 (25.10.2002)

[Continued on next page]

(54) Title: TIME SYNCHRONIZATION USING DYNAMIC THRESHOLDS



(57) Abstract: Methods and apparatus for time synchronization using dynamic thresholds. A method for synchronizing network elements includes receiving at a network element a time synchronization message (102) sent from a master clock element. The network element includes an internal clock to be synchronized with a master clock of the master clock element. The method includes calculating a latency (104) of the time synchronization message. The method includes estimating an adjustment (106) that may be made to the internal clock in response to the time synchronization message. The method includes determining whether the latency calculated is less than the adjustment estimated (108). The method includes adjusting the internal clock (110) when the latency calculated is less than the adjustment estimated.



WO 03/049343 A1



(74) **Agents:** EGAN, William, J. et al.; Fish & Richardson, P.C., 500 Arguello Street, Suite 500, Redwood City, CA 94063 (US).

(81) **Designated States (national):** AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),

Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search 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.

TIME SYNCHRONIZATION USING DYNAMIC THRESHOLDS

BACKGROUND

5 This invention relates to time synchronization over a network.

A network generally includes elements that may need to coordinate their operations. When coordination is required, some of these elements have internal clocks that are synchronized with a master clock. One technique for synchronizing is to periodically send messages from the master clock to the network elements. These
10 messages include a time value that the master clock provides (i.e., a time stamp). An element receiving one of the messages estimates the master clock time by calculating any delay between the time the master clocks stamps the message and when the element processes the message and by adding the calculated delay to the time stamp of the message. The element then can adjust, i.e., synchronize, its internal clock to match the
15 estimated time of the master clock. The mentioned delay is often referred to as latency.

One source of latency is the time it takes for the message to travel from the master clock to the element. One way for an element to calculate this travel time is to send a message on a round trip to the master clock and then measure the time it takes the message to return. The element then splits the round trip time to calculate the time it
20 takes for the message to travel from the master clock to the element.

There are usually thresholds beyond which adjustments to the internal clock of an element being considered may not improve synchronization with the master clock. One such threshold is a margin of error with which a network element can estimate the time of the master clock. The described latency is a significant contributor to this margin of
25 error. Conventional synchronization schemes define a preset set value of latency beyond which the margin of error will be too great to synchronize. An element will thus discard a time synchronization message having latency that is greater than the preset value.

Another threshold is a maximum which a network element can change its internal clock without breaking the network. In this case, the element is able to accurately
30 estimate the time of the master clock but the element's internal clock is so much out of synch with the master clock that adjustment will result in errors in the network. Conventional synchronization schemes also define a preset value for this threshold of maximum adjustment.

SUMMARY

The present invention provides methods and apparatus, including computer program products, for synchronizing a network using dynamic thresholds.

In general, in one aspect, a method for synchronizing network elements
5 includes receiving at a network element a time synchronization message sent from a master clock element. The network element includes an internal clock to be synchronized with a master clock of the master clock element. The method includes calculating a latency of the time synchronization message. The method includes
10 estimating an adjustment that may be made to the internal clock in response to the time synchronization message. The method includes determining whether the latency calculated is less than the adjustment estimated. The method includes adjusting the internal clock when the latency calculated is less than the adjustment estimated.

In general, in another aspect, a computer program product, for synchronizing a first clock of a network with a second clock of the network, includes instructions to
15 cause a processor to calculate a latency between the first clock and the second clock. The product includes instructions to estimate an adjustment of the first clock to synchronize the first clock with the second clock. The product includes instructions to determine whether to adjust the first clock, the determining being based on whether the latency calculated is less than the adjustment estimated. The product is tangibly
20 stored on machine-readable medium,

In general, in another aspect, a computer program product, for synchronizing a first clock of a network with a second clock of the network, includes instructions to cause a processor to calculate a latency between the first clock and the second clock. The product includes instructions to estimate an adjustment of the first clock to
25 synchronize the first clock with the second clock. The product includes instructions to calculate an earliest current time as indicated by the second clock and a latest current time as indicated by the second clock, the calculation assuming that the latency is completely imbalanced. The product includes instructions to determine that the latency calculated is less than the adjustment estimated when a current time as
30 indicated by the first clock is in between the earliest and latest current times. The product is tangibly stored on machine-readable medium

Particular implementations of the invention can provide any combination of the following features. A system in accordance with the invention uses dynamic thresholds to determine whether to adjust an internal clock of a network element

being considered for synchronization. The system need not rely on a preset acceptable error margin and can improve synchronization where conventional systems would discard a time synchronization message. The system can accommodate balanced as well as unbalanced latency. Balanced latency occurs when the travel time of each leg of a round trip is the same. Unbalanced latency occurs when the travel for each leg is different. The system can accommodate latency that is variable, such as latency that occurs in networks having elements that communicate by wireless transmission and reception, such as radio frequency ("RF") transmission and reception. The dynamic thresholds can account for variables, such as the described latency, for each instance of synchronization.

The details of one or more implementations are set forth in the accompanying drawings and the description below. Other features and advantages will be apparent from the description and drawings.

DESCRIPTION OF DRAWINGS

FIG. 1 shows a method for synchronizing one clock of a network with another clock of the network.

FIG. 2 shows an example of how to determine whether the round trip latency between the clocks being synchronized is less than an estimated adjustment to one of the clocks.

FIG. 3 graphically illustrates the concept of the example of FIG. 2.

FIG. 4 shows another example of how to determine whether round the trip latency between the clocks being synchronized is less than the estimated adjustment to one of the clocks.

FIG. 5 graphically illustrates the concept of the example of FIG. 4.

FIG. 6 shows an implementation of the method 100.

Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

Time synchronization techniques in accordance with the invention uses the fact that as long as the latency of a time synchronization message ("sync message") received by a network element is less than an estimated amount of correction of an internal clock of the network element, it is an improvement to adjust the internal clock. Adjustment in such cases brings the time of the internal clock closer in synch

to the time of another clock, such as a master clock of the network, with which the internal clock is being synchronized. The techniques adopt a methodology in which a calculation of the acceptable margin of error is dynamically related to the estimated correction. The acceptable margin of error is not a preset value.

5 FIG. 1 shows a method 100 for synchronizing in accordance with the invention. A network element being synchronized ("sync element") receives a sync message (step 102). The sync element includes an internal clock which is being synchronized with another clock of the network, such as a master clock of the network. The sync message includes a time stamp of, for example, the master clock.
10 The time stamp specifies the time, as indicated by the master clock, when the sync message was sent. The sync message can be sent in response to a request for synchronization ("request") generated by the sync element or, alternatively, in response to a command from a network component such as a controller.

 The sync element calculates the latency of the sync message (step 104). The
15 element can calculate the latency by various methods. One method is to calculate a round trip time for a message to travel between the sync element and a source of the sync message. The source is usually a network element that includes the clock, such as the master clock, with which the internal clock of the sync element is being synchronized. Such as clock will be referred to in this specification as the reference
20 clock. If the source is not a network element that does not include a reference clock, then the latency between the source and reference clock should be calculated and considered in the latency calculation described here. In this round-trip time method, the sync element sends its own message, which can be the above mentioned request for synchronization, to the source of the sync message. The sync element marks the
25 times, as indicated by its internal clock, of when the sync element sends the request and of when the source sends the sync message. The difference between these times is the latency involved in requesting and receiving the particular sync message received in step 102. In this case, the system uses the request to both trigger synchronization and to calculate round trip latency.

30 Alternatively, the sync element can send a message that will be used to only calculate latency and is separate from the request. This type of message will be referred to as a latency calculation message. In this case, the sync element, after it receives the sync message, sends the latency calculation message to the source of the sync message. Upon receiving the latency calculation message, the source returns the

message to the sync element. The sync element marks the times, as indicated by its internal clock, of when the sync element sent and received the latency calculation message. When the sync element performs the described method immediately after receiving the sync message, the difference between the described times is a close
5 approximation of the latency of the sync message.

As an alternative to the round-trip time method, other techniques can be used to determine the latency of the sync message. As discussed, the source is usually a network element having the master clock but need not be.

The sync element estimates the adjustment that will be made to its internal
10 clock (step 106). The sync element can estimate the adjustment by various methods. One method is for the sync element to estimate the current time as indicated by the master clock, i.e., the current time of the master clock, by using the time stamp of the sync message. Then, the sync element calculates the difference between the current time as indicated by its internal clock, i.e., the current time of the internal clock, and
15 the estimated current time of the master clock to derive the estimated adjustment. Alternatively, the sync element can use other methods for estimating the adjustment. For example, instead of using the current time to estimate adjustment, the sync element can use a time of any known event, such as when the sync element receives the sync message. Regardless of whether the sync element calculates the difference
20 between the current time of the master clock and the current time of the internal clock or whether the sync element calculates the difference between the time of a known event, the difference in either case indicates the estimated adjustment to synchronize the internal clock with the master clock.

The sync element determines if the calculated latency, which as described can
25 approximately be the round trip time, is less than the estimated adjustment (decision 108). If it is, then the sync element adjusts its internal clock (step 110). In this case, adjustment will very likely result in an improvement of synchronization. That is, the time difference between the master clock and the internal clock of the element will be reduced. If the latency is not less than the estimated adjustment, then
30 the element does not adjust its internal clock (step 112). In this case, an adjustment will likely not result in an improvement of synchronization because the margin of error in the synchronization calculation, as caused by the latency, is greater than the required estimated adjustment.

Although any network requiring synchronization will benefit from the method 100, this method especially benefits those networks having imbalanced latency such as wireless networks. An example of a wireless network is a telemetry system that includes a system controller ("SC"), one or more master controllers ("MCC"), and one or more terminal units. The system controller includes a master clock and each MCC includes an internal clock. The SC and the MCC must synchronize to, e.g., perform operations at the terminal units. This telemetry system is further described in commonly owned U.S. Patent No. 6,195,018 to Ragle *et al.*, issued on February 27, 2001, and also in commonly owned U.S. Patent Application No. 10/128,928, entitled *Intelligent Two-way Telemetry*, filed on April 22, 2002, which patent and application are hereby incorporated by reference in their entirety.

FIG. 2 shows an example of how to determine whether calculated latency, whether balanced or not, is less than the estimated adjustment. In this example, a sync element, such as one of the described MCCs, sends a message requesting synchronization to an element having a master clock ("master clock element"), such as the described SC. In response, the master clock element sends a sync message back to the network element (i.e., the synch element). The sync element receives the sync message. The sync element calculates the round trip time, which includes the time required for the request to travel from the sync element to the source, for the source to process the request and respond, and for the sync message to travel from the source to the sync element. The round-trip time, in this example, represents the latency.

To determine whether calculated latency is less than the estimated adjustment, the synch element calculates the earliest and latest time values of the current time of the master clock, assuming a completely imbalanced latency (step 202). The earliest current time of the master clock results from one extreme case of imbalanced latency where all of the latency is caused by the time it takes for a request message to travel from the sync element to the master clock element. In this extreme case, the earliest current time of the master clock is the time stamp of the synch message. The latest current time of the master clock results from the other extreme case of imbalanced latency where all of the latency is caused by the time it takes for a synch message to travel from the master clock element to the sync element. In this extreme case, the latest current time of the master clock is the time stamp of the synch message plus the calculated latency.

The sync element determines if the current time, as indicated by its internal clock, is between the calculated earliest and latest time values of the current time of the master clock (decision 204). If the current time as indicated by its internal clock is between the calculated earliest and latest time values of the current time of the master clock, then the calculated latency is less than the estimated adjustment (step 206). Otherwise, the calculated latency is not less than the estimated adjustment (step 208).

Optionally, decision 204 can include additional conditions to preclude adjustment when the estimated adjustment is very small such that any adjustment will likely be detrimental to synchronization. This is the case where the internal clock of the synch element is closely synchronized with the master clock. An example of the conditions is that the time of the synch element is not within a certain band of an estimated master clock time.

FIG. 3 shows a time line that graphically illustrates the concept of the above example. The band of time 302 represents the interval where the internal clock of the synch element is closely synchronized with the master clock and adjustment would likely not result in an improvement of synchronization. The bands of time 304 and 306 represent the interval in which if the current time of the synch element falls, then the round trip latency is less than the estimated adjustment but the internal clock of the synch element is not too closely synchronized with the master clock. When the current time of the synch element falls within this band, adjustment is likely to improve on synchronization. The bands of time 308 and 310 represent time intervals where the round trip latency is greater than estimated adjustment, and adjustment would likely be detrimental to synchronization.

The example described in reference to FIGs. 2 and 3 represents the case where the round trip latency is less than some preset threshold, such as the preset value of latency mentioned in the background section. FIG. 4 shows an example of how to determine whether round trip latency is less than the estimated adjustment when the round trip time is greater than the threshold round trip latency. This is the case where conventional synchronization techniques would discard a sync message.

To determine whether round trip latency is less than the estimated adjustment, the sync element calculates the earliest and latest time values of the current time of the master clock assuming a completely imbalanced latency (step 402). This calculation is similar to the one described for step 202 of FIG. 2.

The sync element determines whether the current time as indicated by its internal clock is within either of two time intervals (decision 404). The first interval lies between the earliest current time of the master clock minus a maximum adjustment band and the earliest current time of the master clock. The second interval

5 lies between the latest current time of the master clock and the latest current time of the master clock plus the maximum adjustment band. The maximum adjustment band delimits adjustment thresholds, beyond which adjustment will very likely cause network errors. That is, adjustment will break the network. If the current time as indicated by the sync element's internal clock lies within these two time intervals,

10 then the round trip value is less than the estimated adjustment and adjustment is likely to improve synchronization (step 406). Otherwise, the round trip value is not less than the estimated adjustment and adjustment is likely to not improve synchronization (step 408). Optionally, the described two time intervals can be expanded to include the interval between them. That is, the two time intervals can be substituted with one

15 time interval delimited by the earliest and the latest times of the two time intervals. With this option, adjustment will still improve synchronization.

FIG. 5 shows a time line that graphically illustrates the concept of the example of FIG. 4. The bands of time 502 and 504 represent the described first and second time interval, respectively. The bands 506 and 508 of time represent the time

20 intervals where synchronization will likely not be improved. The band of time 510 represents the optionally included time interval described above.

Table 1 shows an example of pseudo code for implementing the described methods. The bands of FIGs. 2 through 5 correspond to those defined by the pseudo code. The following definition of terms apply to the pseudo code and to the figures.

25 (The MCC is a synch element and the SC is the master clock element.)

$MSC_{request}$: The time when MCC sends a request to initiate synchronization as indicated by the internal clock of the MCC.

$MSC_{response}$: The time when the MCC receives a response to its request as indicated by the internal clock of the MCC.

30 $SC_{response}$: The time when the SC sends the response to the MCC as indicated by the master clock.

$SC_{estimated}$: This is the estimated current time of SC as indicated by calculated by the MCC.

RTC: Current time as indicated by the MCC's internal clock.

Rt: Round Trip Time; the difference between $MSC_{t_{request}}$ and $MSC_{t_{response}}$.

Δt : Time difference between RTC and estimated SC time

CtsMaxRqstMsgRoundtrip: A preset threshold specifying the maximum
 5 round trip time above which convention synchronization techniques would discard the sync message.

CtsRtcMaxUpdate: A threshold specifying the maximum which a network element can change its internal clock without breaking the network.

```

If (RT < CtsMaxRqstMsgRoundtrip)
{
  if (5 >  $\Delta t$ ) do_nothing(); /*band 302*/
  else if (5 <  $\Delta t$  <= CtsRtcMaxUpdate) RTC =  $SC_{t_{estimated}}$ ; /*bands 304 and
306*/
  else if (CtsRtcMaxUpdate <  $\Delta t$ ) declare (TIME_SYNC_LARGE_ERROR);
/*bands 308 and 310*/
}
else if (smaller_of { |  $SC_{t_{response}} + Rt - RTC$  |, |  $SC_{t_{response}} - RTC$  | } <
CtsRtcMaxUpdate)
{
  if (RTC >  $SC_{t_{response}} + Rt$ ) RTC =  $SC_{t_{response}} + Rt$ ; /*band 504*/
  else if (RTC <  $SC_{t_{response}}$ ) RTC =  $SC_{t_{response}}$ ; /* band 502*/
  else do_nothing(); /*bands 506 and 508*/
}

```

Table 1

10 The logic of the pseudo code is based on the fact that as long as error margin (i.e., $RT/2$) is less than the amount of estimated correction (Δt), it is an improvement to adjust the MCC's internal clock (i.e., update the RTC). However, RTC is updated not to $SC_{t_{estimated}}$ but either to $SC_{t_{response}}$ or to ($SC_{t_{response}} + Rt$). This logic reduces the effect of error margin in the updated RTC and avoids larger fluctuations in subsequent
 15 updates. Alternatively, RTC can be set to any point between $SC_{t_{response}}$ and ($SC_{t_{response}} + Rt$) and would still be an improvement over the current RTC value. Although it is not guaranteed, any point between RTC current value and $SC_{t_{estimated}}$ can also be an improvement.

FIG. 6 shows method 600, which is one implementation of method 100. An MCC of a wireless network sends a request message to a SC of the network (step 602). The MCC includes an internal clock and the SC includes a master clock. The internal clock of the MCC is being synchronized with the master clock of the SC. In response to the request, the SC sends a synch message, which the MCC receives (step 604). As with the sync messages discussed above, the sync message here includes a time stamp that indicates the time, as indicated by the master clock, when the SC sent the synch message. The MCC estimates the adjustment to the internal clock to synchronize it with the master clock (step 606). The MCC calculates the round trip latency (step 606), which includes the time for the request to travel from the MCC to the SC, the time required by the SC to process and respond to the request, and the time for the sync message to travel from the SC to the MCC. The MCC determines whether the latency is less than a threshold, such as a fixed threshold for latency that is typically used by conventional systems (decision 608).

The process flow diverges at this juncture, depending on the result of decision 608. If the round trip latency is less than the threshold, the MCC calculates the earliest and latest time, as indicated by the master clock, when the MCC received the sync message (step 610). The earliest time is the time indicated by the time stamp. This calculation assumes that the time required for the request to travel from the MCC to the SC is zero and, furthermore, that the time required for the sync message to travel from the SC to the MCC is the calculated latency. The latest time is the time indicated by the time stamp plus the calculated latency. This calculation assumes that the time required for the request to travel from the MCC to the SC is the calculated latency and, furthermore, that the time required for the sync message to travel from the SC to the MCC is zero. The MCC determines if the time, as indicated by the internal clock of the MCC, when the MCC receives the sync message (i.e., $MSC_{t_{\text{response}}}$) is between the calculated latest and earliest time (decision 612). If no, then the calculated latency is not less than the estimated adjustment (step 614) and the MCC does not adjust its internal clock (step 616). Otherwise, the MCC determines whether the time, as indicated by the internal clock, when the MCC received the sync message is within five seconds of the time, as indicated by the master clock, of when the SC sent the sync message (decision 620). If the times are within five seconds, then the MCC does not adjust its internal clock (step 616). Otherwise, the MCC adjusts its internal clock.

If the latency is not less than the threshold, the MCC calculates the earliest and latest time, as indicated by the master clock, when the MCC received the sync message (step 624). The calculation is similar to that described in reference to step 610. The MCC determines if the time, as indicated by the internal clock of the MCC, when the MCC receives the sync message (i.e., $MSC_{t_{response}}$) falls within either of two time intervals (decision 626). The time intervals are similar to those described in reference to step 404 of FIG. 4. If no, then the calculated latency is not less than the estimated adjustment (step 628) and the MCC does not adjust its internal clock (step 616). Otherwise, the calculated latency is less than the estimated adjustment (step 630) and the MCC adjusts its internal clock (step 622).

In the above example, the MCC does not estimate adjustment based on the difference between the current times of the MCC and the SC. Rather, the MCC estimates adjustment based on the time when the MCC receives the sync message. With either case, the difference should be the same, or at least very similar, and can be used to adjust the internal clock of the MCC.

The following provides an example scenario. The MCC of the wireless network described above receives a sync message with a time stamp of 988970423 (in seconds) from the SC. The round trip time associated with the sync message is 40 seconds and the maximum predefined latency threshold is 30 seconds. The current time of the MCC's internal clock is 988970603. A conventional system would discard this sync message. Consequently, the MCC would still be 3 minutes off from the master clock of the SC. The systems described herein would however not discard the sync message and, by applying the described techniques, improve synchronization. Assuming a completely imbalanced latency, the SC's current time (as of when the MCC receives the sync message) can be 988970423 or 988970463. Adjustment to either of this time will bring the internal clock closer to the time of the master clock.

The invention can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. Apparatus of the invention can be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a programmable processor; and method steps of the invention can be performed by a programmable processor executing a program of instructions to perform functions of the invention by operating on input data and generating output. The invention can be implemented

advantageously in one or more computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. Each computer program can be implemented in a high-level procedural or object-oriented programming language, or in assembly or machine language if desired; and in any case, the language can be a compiled or interpreted language. Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, a processor will receive instructions and data from a read-only memory and/or a random access memory. The essential elements of a computer are a processor for executing instructions and a memory. Generally, a computer will include one or more mass storage devices for storing data files; such devices include magnetic disks, such as internal hard disks and removable disks; magneto-optical disks; and optical disks. Storage devices suitable for tangibly embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM disks. Any of the foregoing can be supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made. For example, the steps of the described method can be performed in different order and still achieve the same or similar results. Bands 304 and 306 need not be delimited as described but can also be otherwise delimited and still improve synchronization. For example, the bands 304 and 306 of FIG. 3 and as defined in Table 1 can be increased to an area delimited by $(SCT_{estimated} - CtsRtcMaxUpdate)$ and $(SCT_{estimated} + CtsRtcMaxUpdate)$. The clock with which the internal clock of the sync element is being synchronized need not be a master clock but can be any clock of the network. The network refers to a set of elements connect by one or more links. The network can include multiple networks. For example, the network can include two enterprise networks linked by the Internet.

WHAT IS CLAIMED IS:

CLAIMS

1. A method for synchronizing network elements, the method comprising:

receiving at a network element a time synchronization message sent from a master clock element, the network element including an internal clock to be synchronized with a

5 master clock of the master clock element;

calculating a latency of the time synchronization message;

estimating an adjustment that may be made to the internal clock in response to the time synchronization message;

determining whether the latency calculated is less than the adjustment estimated; and

10 adjusting the internal clock when the latency calculated is less than the adjustment estimated.

2. The method of claim 1, wherein determining whether the calculated round trip latency is less than the estimated synchronization adjustment includes:

assuming that the latency is imbalanced.

15 3. The method of claim 2, wherein the time synchronization message includes a time stamp specifying a time, as indicated by the master clock, when the master clock element sent the time synchronization message, and wherein assuming that the latency is imbalanced includes:

20 calculating an earliest current time of the master clock assuming that there is no latency; and

calculating a latest current time of the master clock by adding all of the calculated latency to the time stamp.

4. The method of claim 3, wherein determining whether the latency calculated is less than the adjustment estimated includes:

25 determining whether a current time as indicated by the internal clock is between the earliest and latest current time of the master clock.

5. The method of claim 1, wherein:

calculating latency includes calculating latency that varies from one instance of synchronization to another instance of synchronization; and

30 estimating adjustment includes estimating adjustment that varies from one instance of synchronization to another instance of synchronization.

6. The method of claim 1, wherein:

the time synchronization message includes a time stamp specifying a time, as indicated by the master clock, when the master clock element sent the time synchronization message; and

5 adjusting the internal clock includes adjusting based on the time stamp.

7. The method of claim 1, further comprising:

determining whether the latency calculated exceeds a threshold;

when the latency calculated does not exceed the threshold, applying a first set of criteria to determining whether the latency calculated is less than the adjustment
10 estimated; and

when the latency calculated exceeds the threshold, applying a second set of criteria to determining whether the latency calculated is less than the adjustment estimated.

8. The method of claim 7, wherein applying the first set of criteria includes:

determining whether the master clock and the internal clock are sufficiently
15 synchronized; and

determining whether adjustment to the internal clock will cause an error in the network.

9. The method of claim 8, wherein:

determining whether the master clock and the internal clock are sufficiently
20 synchronized includes determining whether a difference between a current time as indicated by the internal clock and the current time as estimated to be indicated by the master clock is less than five seconds; and

determining whether adjustment to the internal clock will cause an error in the network includes determining whether there will be an error cause by adjusting the
25 internal clock because a difference between the current time as indicated by the internal clock and the current time as estimated to be indicated by the master clock is too large.

10. The method of claim 9, wherein:

calculating a current time includes calculating the current time of when the network element receives the time synchronization message.

30 11. The method of claim 1, further comprising sending a request message from the network element to the master clock element, the request message causing the master clock element to send the time synchronization message, wherein calculating the latency

includes:

recording a time when the network element sent the request message;

recording a time when the network element received the time synchronization message sent from the master clock element; and

5 calculating a difference between the time when the network element sent the request message and the time when the network element received the time synchronization message sent from the master clock element.

12. A computer program product, tangibly stored on machine-readable medium, for synchronizing a first clock of a network with a second clock of the network, the product
10 comprising instructions to cause a processor to:

calculate a latency between the first clock and the second clock;

estimate an adjustment of the first clock to synchronize the first clock with the second clock; and

determine whether to adjust the first clock, the determining being based on whether
15 the latency calculated is less than the adjustment estimated.

13. The product of claim 12, further comprising instructions to:

adjust the first clock when a determination has been made to do so.

14. The product of claim 12, further comprising instructions to:

calculate an earliest current time as indicated by the second clock and a latest current
20 time as indicated by the second clock, the calculation assuming that the latency is completely imbalanced; and

determine that the latency calculated is less than the adjustment estimated when a current time as indicated by the first clock is in between the earliest and latest current times.

25 15. The product of claim 12, further comprising instructions to determine that the first clock is to be adjusted when:

the first clock and the second clock are not synchronized within five seconds; and adjustment does not cause an error in the network..

16. The product of claim 12, further comprising instructions to determine that the latency
30 calculated is less than the adjustment estimated when:

adjustment does not cause an error in the network.

17. The product of claim 12, wherein:

the first clock is include in a first network element and the second clock is included in a second network element; and

calculating a latency between the first clock and the second clock includes calculating the time for a first message to travel from the second network element to the first network element.

18. The product of claim 17, wherein, the product further comprising instructions to:

calculate a second time interval that it takes for a second message to travel from the first network element to the second network element, for the second element to process the second message, and for the second element to generate and send the first message; and

assume the second time interval is the latency.

19. The product of claim 12, wherein:

the network is a wireless network; and

the second clock is a master clock.

20. The product of claim 19, wherein:

the network is a radio frequency network.

21. A computer program product, tangibly stored on machine-readable medium, for synchronizing a first clock of a network with a second clock of the network, the product comprising instructions to cause a processor to:

calculate a latency between the first clock and the second clock;

estimate an adjustment of the first clock to synchronize the first clock with the second clock;

calculate an earliest current time as indicated by the second clock and a latest current time as indicated by the second clock, the calculation assuming that the latency is completely imbalanced; and

determine that the latency calculated is less than the adjustment estimated when a current time as indicated by the first clock is in between the earliest and latest current times.




FIG. 1




FIG. 2




FIG. 3




FIG. 5




FIG. 4




FIG. 6




FIG. 6
(Cont'd)




FIG. 6
(Cont'd)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US02/36799

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) :H04J 3/06

US CL :370/516,517

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/503,518,519516,517; 375/354,371,373

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P	US 6,351,821 B1 (VOTH) 26 February 2002, see entire reference.	1-21
A	US 6,069,887 A (GEIGER ET AL) 30 May 2000, see figs. 1,2.	1-21
A,P	US 6,370,161 B1 (EHRlich ET AL) 09 April 2002, see entire reference.	1-21
A	US 5,784,421 A (DOLEV ET AL) 21 July 1998, see figs.4,5,6.	1-21

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

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

Date of the actual completion of the international search 16 JANUARY 2003	Date of mailing of the international search report 10 FEB 2003
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-9230	Authorized officer MELANIE JAGANNATHAN Telephone No. (703) 305-3900