



(43) International Publication Date  
3 January 2014 (03.01.2014)

- (51) International Patent Classification:  
*G06F 1/12* (2006.01)      *G01R 23/00* (2006.01)
- (21) International Application Number:  
PCT/US2013/048566
- (22) International Filing Date:  
28 June 2013 (28.06.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
61/666,101      29 June 2012 (29.06.2012)      US
- (71) Applicant: **FINITE STATE RESEARCH LLC** [US/US];  
2718 Creeks Edge Parkway, Austin, TX 78733 (US).
- (72) Inventors: **DOUGAN, Cort**; 11304 Cherisse Drive, Austin, TX 78739 (US). **YODAIKEN, Victor**; 2718 Creeks Edge Parkway, Austin, TX 78733 (US).
- (74) Agents: **ZOLTICK, Martin, M.** et al.; ROTHWELL, FIGG, ERNST & MANBECK, P.C., 607 14th Street, N.w., Suite 800, Washington, DC 20005 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Published:**  
— with international search report (Art. 21(3))  
— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))



WO 2014/005016 A1

(54) Title: SYSTEM FOR MAINTAINING ACCURATE IDEAL CLOCK TIME

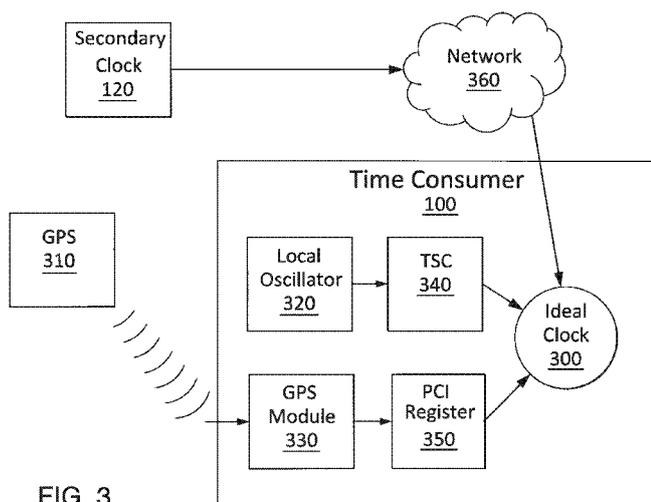


FIG. 3

(57) Abstract: Methods, time consumer systems, and computer program products for maintaining accurate time on an ideal clock (300) of a timing device (100) are disclosed. The method includes receiving (S1100) time information from a local clock (320/340), a reference clock (110, 310/330/350), and one or more secondary clocks (120). The method further includes calculating (S1110) frequencies for the local clock, the reference clock, and the one or more secondary clocks. The method further includes comparing (S1120) the calculated frequencies of the reference clock to the calculated frequencies of the one or more secondary clocks. The method further includes detecting (S1130) a holdover and/or a compromise situation based on the comparison. The method further includes syntonizing (S1140) the ideal clock (300) to one or more of the calculated frequencies.

## SYSTEM FOR MAINTAINING ACCURATE IDEAL CLOCK TIME

### TECHNICAL FIELD

**[001]** The present invention relates generally to a method, time consumer system, and computer program product for maintaining accurate time on a timing device during periods when the reference clock is not accessible and/or is compromised. The field of use includes but is not limited to telecommunications, wireless communication, electric power distribution and transmission, and high-frequency trading.

### BACKGROUND

**[002]** Many electronic control systems and computing devices require an accurate time. One widely used method to accomplish this is to provide an oscillator or other timing device as a "local clock" and to synchronize the local clock to an external "reference clock". This process is analogous to checking a wrist watch (the local clock) against a presumed accurate time announced on the radio (the reference clock). Even if the wrist watch itself gains or loses time over an interval, regular updates from the radio provide enough information to accurately estimate the current time from the time shown by the wristwatch. The time on the wristwatch can be relied on between updates because it can only drift from correct time by a small amount between updates and that drift can be quantified. Technical challenges, however, arise in some fields of use as a result of very high precision requirements.

**[003]** The problem of "holdover" is the problem of maintaining an accurate time even when updates from the reference clock are interrupted for

a relatively extended interval. The problem of compromise detection involves detecting frequency changes in the reference clock. Providing a method for holdover and protecting against compromise is important in telecommunications, industrial control, power transmission and distribution, military equipment, and also in a growing number of commercial computing applications. For example, financial trading systems depend on the distribution of accurate time to compute trading latency and find patterns in trading results. These trading systems are thus vulnerable to interruptions in or compromise of GPS signals.

#### SUMMARY

**[004]** Particular embodiments of the present invention include methods, systems, and computer program products for providing holdover capability and compromise detection by deriving high quality “secondary reference frequency” information from a secondary clock (or collection of clocks) despite inaccuracies in clock signal and/or high signal delay variation. In particular, the methods, systems, and computer program products of the present invention permit the use of a time signal delivered via a network, a network within an enterprise or even the Internet, as a source of precise frequency information. Particular embodiments of the disclosed solution derive the secondary reference frequency and apply this derived frequency information to holdover and compromise detection. The methods, systems, and computer program products of the present invention enable ubiquitous availability and low cost of access to such time sources. The secondary clock

is used to measure both the current frequency of the local clock and the reference clock. If the reference clock frequency changes, the possibility of compromise can be detected. If the reference clock becomes unavailable, the frequency measure derived from the second clock can be used to compensate for frequency changes in the local clock.

**[005]** According to some embodiments, a method for maintaining accurate time on an ideal clock of a timing device is disclosed. The method includes receiving time information from a local clock, a reference clock, and one or more secondary clocks. The method further includes calculating frequencies for the local clock, the reference clock, and the one or more secondary clocks. The method further includes comparing the calculated frequencies of the reference clock to the calculated frequencies of the one or more secondary clocks. The method further includes detecting a holdover and/or a compromise situation based on the comparison. The method further includes syntonizing the ideal clock to one or more of the calculated frequencies.

**[006]** According to some embodiments, a time consumer system for maintaining accurate time on an ideal clock is disclosed. The time consumer system includes a processor and a computer readable memory coupled to the processor. The computer readable memory contains computer readable program instructions that configure the processor to receive time information from a local clock, a reference clock, and one or more secondary clocks. The time consumer is further configured to calculate frequencies for the local clock, the reference clock, and the one or more secondary clocks. The time

consumer is further configured to compare the calculated frequencies of the reference clock to the calculated frequencies of the one or more secondary clocks. The time consumer is further configured to detect a holdover and/or a compromise situation based on the comparison. The time consumer is further configured to syntonize the ideal clock to one or more of the calculated frequencies.

**[007]** According to some embodiments, a computer program product includes a non-transitory computer readable medium having stored thereon executable computer instructions for maintaining accurate time in an ideal clock of a timing device. The computer instructions are configured to cause a computer to receive time information from a local clock, a reference clock, and one or more secondary clocks. The computer instructions are configured to cause a computer to calculate frequencies for the local clock, the reference clock, and the one or more secondary clocks. The computer instructions are configured to cause a computer to compare the calculated frequencies of the reference clock to the calculated frequencies of the one or more secondary clocks. The computer instructions are configured to cause a computer to detect a holdover and/or a compromise situation based on the comparison. The computer instructions are configured to cause a computer to syntonize the ideal clock to one or more of the calculated frequencies.

**[008]** In one embodiment, a time consumer can use the secondary clock to measure a time interval, detect the number of "ticks" counted by the local clock during that time interval and compute local clock frequency. If the local clock counts N ticks during the interval and the secondary clock

measures elapsed time  $T$ , the average duration of a clock tick is  $T/N$  and the average frequency is  $N/T$ .

**[0009]** When the reference clock is operational, the same measurement can be performed and the frequencies of the reference clock and the secondary clock can be compared. If these change with respect to each other, the frequency of one must have changed and thus a possible compromise event has been detected.

**[0010]** If the ideal clock begins in a synchronized state and if the derived secondary reference frequency is sufficiently accurate, holdover can be implemented by using the secondary reference frequency to correct for drift in the local clock or local clocks.

**[0011]** In one embodiment, a computing device connected to a GPS time source as the reference clock may also be able to receive time over the Internet from the National Institute for Standards and Technology (NIST) clocks using NTP. The NIST internet clock, in general, is not suitable for the purpose of synchronizing the ideal clock because packet delay variations can be several milliseconds long on current Internet technology. However, over intervals of time significantly longer than the delay variation, the elapsed time will be accurate. This permits the ideal clock computation to update its frequency estimates for the local clock and accurately compensate for variations in local clock frequency.

**[0012]** In another embodiment, the time consumer may repeatedly query the secondary clock, discarding any responses that take too long or

otherwise do not meet filter criteria to improve the quality of the duration measurement.

**[0013]** In yet another embodiment, the time consumer may collect a series of measurements and then compute a "best fit" frequency or best fit function for the local clock. Those skilled in the art will appreciate that variations on how to fit the frequency estimate to the data can be made without departing from the disclosed methods and devices of the present invention.

**[0014]** In another embodiment, the time consumer may track multiple overlapping intervals from the secondary clock and use these to find patterns in the frequency changes of the local clock to improve frequency compensations. For example, some local clocks show a periodic behavior, speeding up and slowing down possibly in response to fans turning on and off. Once such patterns have been characterized, the time consumer may use them to predict frequency changes.

**[0015]** In one embodiment, the time consumer is a server computer, the reference time is a bus card with a GPS clock module, the local clock is an oscillator connected to a processor "time stamp register" which operates as a counter, and the secondary clock is an internet connection to the National Institute of Standards and Technology (NIST) time via a network device connected to a local network and then out to the Internet via switches and routers.

**[0016]** In another embodiment, the time consumer is a server computer, the reference clock is a PTP grandmaster connected via IRIG to a

GPS receiver, the local clock is an oscillator connected to a processor "time stamp register" which operates as a counter, and the secondary clock is the same NIST connection as in the previous embodiment. The secondary clock can be employed for holdover and compromise detection for any type of primary reference clock.

**[0017]** In yet another embodiment, the time consumer is a base station for a cellular telephone system, the reference clock is a GPS receiver, the local clock is a TCXO oscillator and the secondary clock is a receiver for a time signal encoded in a radio signal from a neighboring base station. As long as a precise frequency can be derived from the secondary clock, the method can be employed.

**[0018]** In another embodiment, a server computer may use multiple secondary clock sources connected to multiple remote clocks over a network. The elapsed time may be computed as a weighted average of the measured elapsed times or via a voting mechanism or some other means of increasing the reliability of the frequency measure from the secondary sources.

**[0019]** In another embodiment, the secondary clock frequency may come from the same physical device as the primary external reference clock. For example, a number of commercial "network time clock" devices that act as PTP master clocks synchronized to a GPS source in one mode, can switch to modes where the quality of the synchronization is degraded or the protocol changes. One such device will switch to a holdover mode using an internal temperature controlled oscillator or rubidium timer when GPS reception is interrupted or of poor quality. However, due to several factors, the quality of

that holdover time may not be as high as the quality of the holdover time that can be produced using the method of the present invention by discarding the synchronization information from the PTP source while using its frequency information as in the example above. Similarly, some devices have a failure mode in which they cannot produce PTP but fall back on a less accurate NTP implementation. The present invention permits the client system to use the NTP source for frequency information during a holdover.

**[0020]** In another embodiment, the time consumer may produce a reliable frequency calculation from a number of unreliable clocks. For example, a time consumer connected to a number of other time consumers may request that each forward its local clock times and then average together this time to produce a secondary clock from which a reliable frequency can be derived. Those skilled in the art will appreciate that more sophisticated aggregation methods than simple averaging can be used without departing from disclosed methods and devices of the present invention.

**[0021]** The above and other aspects and embodiments are described below with reference to the accompanying drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0022]** The accompanying drawings, which are incorporated herein and form part of the specification, illustrate various embodiments of the present disclosure and, together with the description, further serve to explain the principles of the disclosure and to enable a person skilled in the pertinent art

to make and use the embodiments disclosed herein. In the drawings, like reference numbers indicate identical or functionally similar elements.

**[0023]** Fig. 1 is a block diagram showing a time consumer connected to both a reference clock and a secondary clock, in accordance with exemplary embodiments of the present invention.

**[0024]** Fig. 2 is a block diagram showing a time consumer connected to a reference clock and a plurality of secondary clocks, in accordance with exemplary embodiments of the present invention.

**[0025]** Fig. 3 is a block diagram illustrating the calculation of an ideal clock from the local clock, reference clock, and secondary clock, in accordance with exemplary embodiments of the present invention.

**[0026]** Fig. 4 is a timeline illustrating the concept of signal delay variation, in accordance with exemplary embodiments of the present invention.

**[0027]** Fig. 5 is a timeline illustrating the simultaneous recording of secondary time and local frequency information, in accordance with exemplary embodiments of the present invention.

**[0028]** Fig. 6 is a timeline illustrating multiple secondary time and local frequency measurements, in accordance with exemplary embodiments of the present invention.

**[0029]** Fig. 7 is a bus card that contains a GPS receiver which produces a running time value, in accordance with exemplary embodiments of the present invention.

**[0030]** Fig. 8 is a block diagram of a time consumer, in accordance with exemplary embodiments of the present invention.

**[0031]** Fig. 9 is a flow chart illustrating the steps performed by the time consumer to calculate local frequency drift, in accordance with exemplary embodiments of the present invention.

**[0032]** Fig. 10 is a flow chart illustrating the steps performed by the time consumer to perform either holdover or detect a compromise, in accordance with exemplary embodiments of the present invention.

**[0033]** Fig. 11 is a flow chart illustrating a method for maintaining accurate time on an ideal clock of a timing device.

#### DETAILED DESCRIPTION

**[0034]** The methods and devices of the present invention can derive frequency data from a possibly inaccurate source of time and use that data to detect frequency changes in a reference clock and to maintain correct time during failures of the reference clock.

**[0035]** For example, if a radio, which serves as a reference clock, is only intermittently operational, a highly accurate chronometer might be purchased and synchronized with radio updates so that when the radio is unavailable, e.g., during a holdover, a wrist watch may be periodically synchronized to the chronometer. If the accuracy of the chronometer is higher than the wrist watch, the wrist watch time can be maintained close to the reference time until the radio is available again. Additionally, suppose that the clock in the radio station temporarily slows, e.g., a compromise, and

the wrong time is announced. A secondary time source may be used to detect such compromise.

**[0036]** In what follows, a “clock” is a device or combination of devices and software that produces a time signal (that may, for example, be an electrical signal or a value in a register or in a network packet). A “time consumer” is a computing or communications or control device that reads or receives or detects time signals from clocks. A “time consumer” may contain one or more “local clocks” and may communicate with one or more external clocks. A time consumer generates an “ideal clock” that is in some way derived from the signals or values obtained from external clocks in combination with the signals or values obtained from local clocks and perhaps other information. A reference clock is a clock that has been chosen as the authoritative time source. The problem of “synchronization” in this context is the problem of keeping the ideal clock adjusted so that it produces a time as close as possible to the time produced by the reference clock. The problem of “holdover” concerns maintaining synchronization during an interruption or degradation of the signal from the reference clock. The problem of “compromise” involves detecting changes in the frequency of the reference clock.

**[0037]** According to one example,, the time consumer is a computer server, the reference clock is a device following signals received from the Global Positioning System (GPS) satellites, and the local clock is a computer crystal oscillator driving a counter. However, the crystal oscillators provided on commercial computer boards are known to change frequency during

operation, particularly, but not only, as device temperature changes. Additionally, the GPS signal is not always available, and the GPS receiver is not completely reliable. One solution to holdover in such a configuration is to provide an additional more stable secondary oscillator that is temperature controlled (TCXO) or "ovenized" (OCXO), such as an atomic clock of some sort (like a Rubidium clock), and to fall back on this secondary reference when the primary reference clock is not available. Producing a time that involves a combination of a reference clock and a local oscillator is a well-known process.

**[0038]** The problem of compromise of the reference time source involves temporary inaccuracies in the external time source. That is, if the external time source changes frequency for some reason, it would be useful to detect this and then go into holdover mode until the drift ceases or some external source (such as an operator) validates the signal. In particular, the GPS signal has been shown to be susceptible to "spoofing" in which someone deliberately modifies the time signal. The provision of a rubidium clock or temperature controlled or "ovenized" oscillator is one method for cross checking the authoritative time source.

**[0039]** Using the time of a GPS signal as a reference clock is just one example. The external reference time may also be obtained from alternative satellite systems (such as COMPASS and GALILEO), from an atomic clock, over an IRIG signal, a Pulse Per Second (PPS) signal, or via a computer network using IEEE 1588 Precision Time Protocol (PTP), the Network Time Protocol (NTP) or some other messaging protocol. Additionally, synchronous

Ethernet can be used to directly distribute a time frequency that might be used to verify the external time reference or to correct the external time reference. Similarly, the crystal oscillator might be a counter, or a MEMS device, or something else that provides a local time interval.

**[0040]** Computation of time by the time consumer, given a local clock and an external reference clock, usually involves a process of constructing a synthetic clock with an "ideal time." This is done by both *synchronizing* (setting the ideal time to match the external reference time) and *syntonizing* (setting the frequency of the ideal clock to match the frequency of the external time reference). Syntonization involves determining the frequency of the local clock so that intervals computed by that source can be adjusted to match the interval that would have been measured by the external reference. In the illustrative example, the time broadcast on the radio might be written down and the offset from the wristwatch computed. At the next update on the radio, the drift of the wristwatch over the interval could be computed and this would give information on the difference between the frequency of the reference time source and the frequency of the wristwatch. The wristwatch time could be "corrected," at times between radio updates, by using the frequency adjustment. In a more realistic example, the frequency differential between the crystal oscillator and the GPS clock would be computed and used to compute the ideal time.

**[0041]** Signal Delay Variation (SDV) is the time variation of the time it takes for clock data (the signal) to travel from the clock to the time consumer and become available to the time consumer for calculation. When the time is

communicated over a packet network, SDV is called Packet Delay Variation (PDV). Because of SDV, the signal from the clock carries a time that is inaccurate by some variable quantity when it arrives at the time consumer. Generally, SDV is very low for the local clock, and to be usable for synchronization purposes, the reference clock SDV must also be small. If the SDV of the reference clock increases too much, a holdover condition can be triggered.

**[0042]** There are several disadvantages to providing a second reference clock in the form of a rubidium clock or temperature controlled oscillator. The most significant is the additional expense of temperature controlled oscillators and the even more expensive rubidium clocks. A second disadvantage is that these devices are long-term stable but some may be subject to considerable inaccuracies over shorter intervals. These devices are also not immune to device failure, and if the rubidium clock ceases to work, it cannot be used to provide holdover. Finally, there may be mechanical, power supply or other issues involved in providing these devices. Note that in order to minimize SDV, the rubidium clock/TXCO or similar device is generally placed physically close to the time consumer – making it more expensive and more difficult to share among multiple time consumers. Accordingly, exemplary embodiments of the invention provide alternative means for providing holdover capabilities and authoritative source compromise detection.

**[0043]** Fig. 1 is a block diagram showing a time consumer 100 connected to both a reference clock 110 and a secondary clock 120. In exemplary embodiments, time consumer 100 may be connected to a

reference clock 110 and secondary clock 120. In some embodiments, the connections between time consumer 100 and reference clock 110 and secondary clock 120 may be hardwired so that the connection is within the same computing device. In other embodiments, a hardwire connection may be used across different computing device. In some embodiments, one or both of the connections may be across one or more communications networks via a cable or wireless connection.

**[0044]** Fig. 2 is a block diagram showing a time consumer 100 connected to a reference clock 110 and a plurality of secondary clocks 120a-120n. In some embodiments, the "secondary clock" may be an aggregated collection of clocks. The connections between the secondary clocks 120a-120n and the time consumer 100 may consist of a hardwired connection in the same computing device, or a wired or wireless connection across one or more computing devices, or any combination of hardwired, wired, and wireless connections.

**[0045]** Fig. 3 is a block diagram illustrating the calculation of an ideal clock 300 from the local clock 320/340, reference clock 310, and secondary clock 120. Fig. 3 shows one embodiment in which a time consumer 100 includes a GPS module 330 that drives a PCI register 350 and a local oscillator 320 that drives a time stamp counter (TSC) 340. In some embodiments, the GPS module 330 may be a component of a bus card illustrated in Fig. 5, as discussed below, that is configured to receive GPS time information from a GPS device 310. In the illustrated embodiment, GPS device 310, as coupled with the GPS module 330 and PCI Register 350,

serves as a reference clock. Additionally, the local oscillator 320, coupled with the TSC 340, serves as a local clock.

**[0046]** In some embodiments, secondary clock time information is transmitted by a secondary clock 120 to a communications network 360. The ideal clock 300, located in time consumer 100, may receive the secondary time information from secondary clock 120 via communications network 360. In some embodiments, the secondary clock data may be, for example, time information from one or more NIST clocks.

**[0047]** In some embodiments, local clock time information, generated from the local oscillator 320 and TSC 400, may be transmitted to ideal clock 300. In some embodiments, reference clock time information from GPS device 310 is transmitted to ideal clock 300. Frequency data extracted from the secondary clock 120 is used both to crosscheck the GPS 310 time information (for compromise detection) and to maintain the correct time if the GPS clock in GPS 100 fails (holdover).

**[0048]** In a financial trading system, for example, holdover may only be needed until the trading day completes and a new module can be installed (or reception problems can be cleared up). A processor may be connected to a crystal oscillator 320 that drives a running counter called a Time Stamp Counter (TSC) 340. The TSC 340 is the local time source. In this embodiment, a network device connects the time consumer 100, which is a computing device, to a computer network 360 connected to the Internet, and the time consumer 100 periodically requests the time from NIST or another secondary clock 120 using NTP. In normal operation, the data for the ideal

clock 300 is obtained by periodically polling the reference clock, or GPS 310, for reference time, using that information to compute the actual frequency of the local clock and then, between updates, using the TSC interval since the last time read from the reference time, plus current frequency estimate for the local time source. If the system detects that the reference clock is no longer working, the average TSC count over an interval previously computed, for example over a 5 minute interval, is computed against the time reported by the Internet connection to the secondary clock 120, or NIST clock. In practice, these computations may also involve more sophisticated calculations, including the use of smoothing and filtering methods.

**[0049]** In another embodiment, a time consumer 100 contains a network card with an oscillator 320 that is connected over a high speed local network 360 to a device that acts as a IEEE 1588 PTP "master clock." The processor is connected to a crystal oscillator 320 that drives a running counter called a Time Stamp Counter (TSC) 340. The TSC 340 is the local time source. In this embodiment, the same or a second network device connects the time consumer 100 to a computer network 360 connected to the Internet and the device periodically requests the time from a secondary clock 120 using NTP. In normal operation, the ideal clock 300 is obtained by some combination of time provided in PTP packets, delay information from the network card oscillator 320 (used to reduce PDV imprecision on PTP packets), historical data with some smoothing/filtering algorithms, and a computation of the actual frequency of the local clock. Then, between PTP updates, the ideal clock 300 is obtained by using the TSC interval, the last

time read from the reference time, plus current frequency estimate for the local time source. If the system detects that the PTP reference clock 310 is no longer working, the same method used in the previous embodiment is employed to compute the actual frequency of the local time source.

**[0050]** The methodology is applicable to any type of time consumer 100 device that can be configured to communicate to a secondary clock 120 via a network, including, for example, a mobile device such as a wireless telephone, a personal digital assistant, an electronic reader (e.g. an Amazon Kindle), or a personal tablet computer (e.g., an Apple iPad); a server computer; a laptop computer; and a personal computer.

**[0051]** Referring now to Fig. 4, a timeline illustrating the concept of signal delay variation is shown. A signal 400 carrying time information is produced by the reference clock 110 at time  $t_0$  and only arrives at the time consumer 100 at time  $t_3$ , as illustrated by the actual time timeline 410. As a result of signal delay, the point at which the time information, produced at time  $t_0$ , is received at the time consumer 100 at  $t_3$ , the time information is no longer accurate. Furthermore, if the delay between  $t_0$  and  $t_3$  varies, it is difficult for the time consumer to accurately compensate for delay.

**[0052]** Referring now to Fig. 5, a timeline illustrating the simultaneous recording of secondary time and local frequency information is shown. The TSC count  $n_0$  511 could be recorded at time consumer 100 just as a time update from a secondary clock, such as a NIST clock,  $t_0$  501 is received. In accordance with exemplary embodiments, the TSC count  $n_1$  512 may be recorded at the time consumer 100 just as a secondary clock update  $t_1$  502 is

received, indicating the passage of some computed time interval. This information may be used to update the computed frequency of the local clock with the simplest calculation being that the frequency is  $(n_1-n_0)/(t_1-t_0)$ .

**[0053]** Embodiments of the methods and devices of the present invention rely on an extraction of precise frequency data from an imprecise time source. The secondary clock 120 is an imprecise time source because of the variable delay in response to requests for time from the client or time consumer 100. For example, suppose the time consumer 100 requests a time and receives a response 10 milliseconds later. The time value in the response packet, in theory, could have been produced by the remote clock, or secondary clock 120, at almost any time within that 10 millisecond interval. For example, if the request took 1 millisecond to reach the remote clock and the remote clock took 1 millisecond to generate a response, and 8 milliseconds passed before the response was received by the client, the time value in the packet is 8 milliseconds less than the actual time. On the other hand, if the request took 5 milliseconds to reach the clock, the clock took 4 milliseconds to produce the response, and the response took 1 millisecond to reach the client, the time value in the packet is only 1 millisecond less than the actual time.

**[0054]** Although this variation, often described as Packet Delay Variation (PDV), can be reduced via a variety of filtering and smoothing techniques and via hardware assist, such as time stamps from switching and routing and network devices along the route, there are limits to such techniques. However, exemplary embodiments of the present invention rely

on the previously unexploited high accuracy of frequency information that can be extracted from such a secondary clock 120 source. If, for example, we measure a 2 minute interval with such a source and there is a 10 millisecond delay on both ends, the worst case error in the interval is 20 milliseconds, which is 0.000167% a value that can be reduced with longer errors and/or smoothing/filtering.

**[0055]** In one embodiment, the time consumer 100 will use historical data to determine the minimal response time delay that can be expected, and then begin an interval only after some request has produced a response with such a delay. The same method can be used by a time consumer 100 to determine when to end the interval – since the important quality is how exactly the interval is measured, not whether the interval is of some preset duration. So, an interval of 122 seconds plus or minus 5 milliseconds is more precise than one of exactly 120 seconds plus or minus 20 milliseconds. The time consumer 100 can also track multiple intervals and overlapping intervals to better measure actual frequency changes in the local time source. In cases where network quality permits, the time scale can be reduced: given a high quality connection to a secondary time source, the intervals might be milliseconds, or even of shorter duration. However, the intervals should be long enough to permit accurate frequency time calculation.

**[0056]** For purposes of illustration, suppose that the a received NTP time is associated with a delay ranging between  $\frac{1}{2}$  and 10 milliseconds. When a time  $t$  is read from a NTP packet, the time consumer 100 only has information that the real time is between  $t + \frac{1}{2} \text{ millisecond}$  and  $t + 10$

*milliseconds*. This level of delay variation is much too high for synchronization. However, because the secondary clock 120, such as a NIST clock, is very precise, an interval beginning when time  $t$  is read from a packet and ending when time  $t + 2 \text{ minutes}$  is read from a packet is within 20 milliseconds of 2 minutes. That raw error is very small compared to the interval and, if the external clock is reliable, it can be used to reliably measure the frequency of the local clock over that interval. The raw error can be made more precise with filtering and/or smoothing. The exact duration of the appropriate interval, the appropriate smoothing and filtering methods, and other details will depend on the quality of the connection to the clock and the required accuracy and sensitivity of the ideal clock.

**[0057]** Referring now to Fig. 6, a timeline illustrating multiple secondary time and local frequency measurements is displayed. Instead of the simple elapsed time calculation described above, a method for using interpolation or other signal processing methods to produce a function for estimating local clock frequency via data fitting can be used. In one embodiment, the time consumer 100 can record a sequence of matching pairs 600a-h of TSC counts and clock signals from the secondary clock 120, and then use some interpolation or curve fitting algorithm to determine an equation for the frequency of the local clock. A least squares algorithm, for example, has excellent results in certain environments. Variations in interpolation method and/or the application of some filtering to the raw data can be used in different circumstances. There are many well-known methods from signal processing that those skilled in the art will appreciate could be used for this purpose.

**[0058]** Referring now to Fig. 7, a bus card that contains a GPS receiver which produces a running time value is displayed. The reference clock 110 may be a bus card 710 that includes a GPS receiver and module 330, and electronic components including one or more oscillators 720 and control devices 700 that produce a running time value in a register that can be read by the main processor 730 of the time consumer 100. In the current state of practice, the reference time from the bus card can be accurate down to a few tens of nanoseconds.

**[0059]** Referring now to Fig. 8, a block diagram of a time consumer is shown. In some embodiments, time consumer 100 may be a communication device that includes a processor 730, a network interface device 810, a memory 820, and an ideal clock 300. Computer readable program code may be stored in a non-transitory computer readable medium/memory 820, such as, but not limited to, magnetic media (*e.g.*, a hard disk), optical media (*e.g.*, a DVD), memory devices (*e.g.*, random access memory), and the like. In some embodiments, time consumer 100 may be connected to an external or internal database to store, for example, the time and clock frequency information disclose above.

**[0060]** In some embodiments, computer readable program code is configured such that when executed by a processor 730, the code causes the processor 730 to perform steps described above and shown in the figures. In other embodiments, the processor 730 of time consumer 100 is configured to perform the steps described above and below without the need for code. That

is, for example, the processor 730 may consist of one or more ASICs. Hence, the features of the present invention described herein may be implemented in hardware and/or software. For example, in particular embodiments, the ideal clock 300 may be implemented as hardware or software, or by any suitable combination of hardware and/or software.

**[0061]** Now referring to Fig. 9, a flow chart illustrating the steps performed by the time consumer to calculate local frequency drift are displayed. In some embodiments, the time consumer 100 collects time data from reference clock 110 in step S900. In step S901, if enough time data is collected from reference clock 110, then the frequency of the ideal clock 300 is matched to the frequency of reference clock 110 in step S902. If there is not enough time data, then the time consumer 100 continues to collect time data from reference clock 110 in Step S900. Once the frequency is matched in step S902, an updated frequency estimate is calculated in step S903. As explained above, the estimate may be calculated using a variety of methods, such as smoothing and filtering methods

**[0062]** Now referring to Fig. 10, a flow chart illustrating the steps performed by the time consumer to perform either holdover or detect a compromise are displayed. In some embodiments, ideal clock 300 of time consumer 100 will continue to operate in step S1000 based on a certain reference clock 110 frequency. In step S1010, if the reference frequency has not changed, then the ideal clock 300 will continue to operate in step S1000 based on the currently used reference clock 110 frequency. However, if the reference frequency has changed, then in step S1020 it is determined if the

reference clock 110 has failed. For example, as described above, the frequencies of the reference clock and the secondary clock can be compared. If the frequencies change with respect to each other, then the frequency of one has changed and thus a possible compromise or holdover has been detected. If the reference clock 110 has not failed, then there is no compromise and the ideal clock 300 is matched to the changed reference frequency 300, and the ideal clock 300 continues to operate in step S1000. However, if the reference clock 110 has failed, then a holdover procedure is performed in step S1030. In accordance with exemplary embodiments, the holdover procedure is implemented by using the secondary clock 120 reference frequency to correct for drift in the local/ideal clock or clocks 300.

**[0063]** Now referring to Fig. 11, a flow chart illustrating a method for maintaining accurate time on an ideal clock of a timing device is displayed. In some embodiments, time information from a local clock, a reference clock, and one or more secondary clocks may be received in step S1100. In step S1110, in some embodiments, the received time information may be used to calculate the frequencies of the local clock, the reference clock, and the one or more secondary clocks. In step S1120, in some embodiments, the calculated frequencies of the reference clock may be compared to the calculated frequencies of the one or more secondary clocks. In step S1130, in some embodiments, a holdover and/or a compromise situation may be detected based on the comparison. In step S1140, in some embodiments, the ideal clock may be synchronized to one or more of the calculated frequencies.

**[0064]** While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present disclosure should not be limited by any of the above-described exemplary embodiments. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the disclosure unless otherwise indicated herein or otherwise clearly contradicted by context.

**[0065]** While the processes described above and illustrated in the drawings are shown as a sequence of steps, this was done solely for the sake of illustration. Accordingly, it is contemplated that some steps may be added, some steps may be omitted, the order of the steps may be re-arranged, and some steps may be performed in parallel.

CLAIMS

1. A method for maintaining accurate time on an ideal clock (300) of a timing device (100), comprising:

receiving (S1100) time information from a local clock (320/340), a reference clock (110, 310/330/350), and one or more secondary clocks (120);

calculating (S1110) frequencies for the local clock, the reference clock, and the one or more secondary clocks;

comparing (S1120) the calculated frequencies of the reference clock to the calculated frequencies of the one or more secondary clocks;

detecting (S1130) a holdover and/or a compromise situation based on the comparison; and,

syntonizing (S1140) the ideal clock to one or more of the calculated frequencies.

2. The method of claim 1, further comprising:

determining if the calculated frequency of the reference clock has changed with respect to the calculated frequencies of the one or more secondary clocks, wherein a change detects a holdover or compromise situation.

3. The method of claims 1 or 2, further comprising:

syntonizing the ideal clock to the calculated frequencies of the one or more secondary clocks if a holdover is detected.

4. The method of any of claims 1-3, further comprising:
  - syntonizing the ideal clock to the calculated frequency of the reference clock if no holdover is detected.
  
5. The method of any of claims 1-4, wherein the one or more secondary clocks comprise one or more National Institute for Standards and Technology Internet clocks.
  
6. The method of any of claims 1-5, wherein the calculated frequencies of the reference clock, the one or more secondary clocks, and the local clock comprise best fit frequencies based on the received time information from the reference clock, the one or more secondary clocks, and the local clock, respectively.
  
7. The method of any of claims 1-6, further comprising:
  - storing the calculated frequencies.
  
8. A time consumer system (100) for maintaining accurate time on an ideal clock (300) comprising:
  - a processor (730), and
  - a computer readable memory (820) coupled to the processor, wherein the computer readable memory contains computer readable program instructions that configure the processor to:

receive time information from a local clock (320/340), a reference clock (110, 310/330/350), and one or more secondary clocks (120);

calculate frequencies for the local clock, the reference clock, and the one or more secondary clocks;

compare the calculated frequencies of the reference clock to the calculated frequencies of the one or more secondary clocks;

detect a holdover and/or a compromise situation based on the comparison; and,

syntonize the ideal clock to one or more of the calculated frequencies.

9. The time consumer of claim 8, wherein the processor is further configured to:

determine if the calculated frequency of the reference clock has changed with respect to the calculated frequencies of the one or more secondary clocks, wherein a change detects a holdover or compromise situation.

10. The time consumer of claims 8 or 9, wherein the processor is further configured to:

syntonize the ideal clock to the calculated frequencies of the one or more secondary clocks if a holdover is detected.

11. The time consumer of any of claims 8-10, wherein the processor is further configured to:

syntonize the ideal clock to the calculated frequency of the reference clock if no holdover is detected.

12. The time consumer of any of claims 8-11, wherein the time consumer comprises one of: a server computer and a base station for a communications system.

13. The time consumer of any of claims 8-12, wherein the processor is further configured to calculate best fit frequencies for the reference clock, the one or more secondary clocks, and the local clock based on the received time information from the reference clock, the one or more secondary clocks, and the local clock, respectively.

14. The time consumer of any of claims 8-13, further comprising a database, wherein the calculated frequencies are stored in the database.

15. A computer program product comprising a non-transitory computer readable medium having stored thereon executable computer instructions for maintaining accurate time in an ideal clock (300) of a timing device (100), the computer instructions are configured to cause a computer to perform operations comprising the following steps:

receive time information from a local clock (320/340), a reference clock (110, 310/330/350), and one or more secondary clocks (120);

calculate frequencies for the local clock, the reference clock, and the one or more secondary clocks;

compare the calculated frequencies of the reference clock to the calculated frequencies of the one or more secondary clocks;

detect a holdover and/or a compromise situation based on the comparison; and,

syntonize the ideal clock to one or more of the calculated frequencies.

16. The computer program product of claim 15, wherein the computer instructions are further configured to cause a computer to:

determine if the calculated frequency of the reference clock has changed with respect to the calculated frequencies of the one or more secondary clocks, wherein a change detects a holdover or compromise situation.

17. The computer program product of claims 15 or 16, wherein the computer instructions are further configured to cause a computer to:

syntonize the ideal clock to the calculated frequencies of the one or more secondary clocks if a holdover is detected.

18. The computer program product of any of claims 15-17, wherein the computer instructions are further configured to cause a computer to:

syntonize the ideal clock to the calculated frequency of the reference clock if no holdover is detected.

19. The computer program product of any of claims 15-18, wherein the computer instructions are further configured to cause a computer to:

calculate best fit frequencies for the reference clock, the one or more secondary clocks, and the local clock based on the received time information from the reference clock, the one or more secondary clocks, and the local clock, respectively.

20. The computer program product of any of claims 15-19, wherein the computer instructions are further configured to cause a computer to:

store the calculated frequencies.

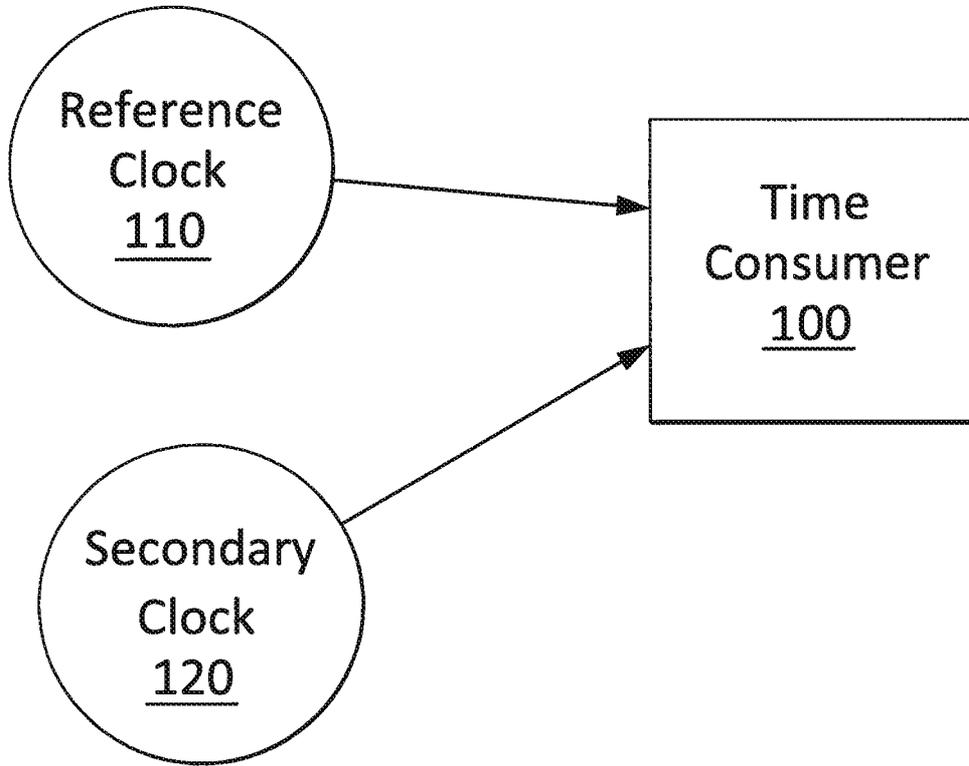


FIG. 1

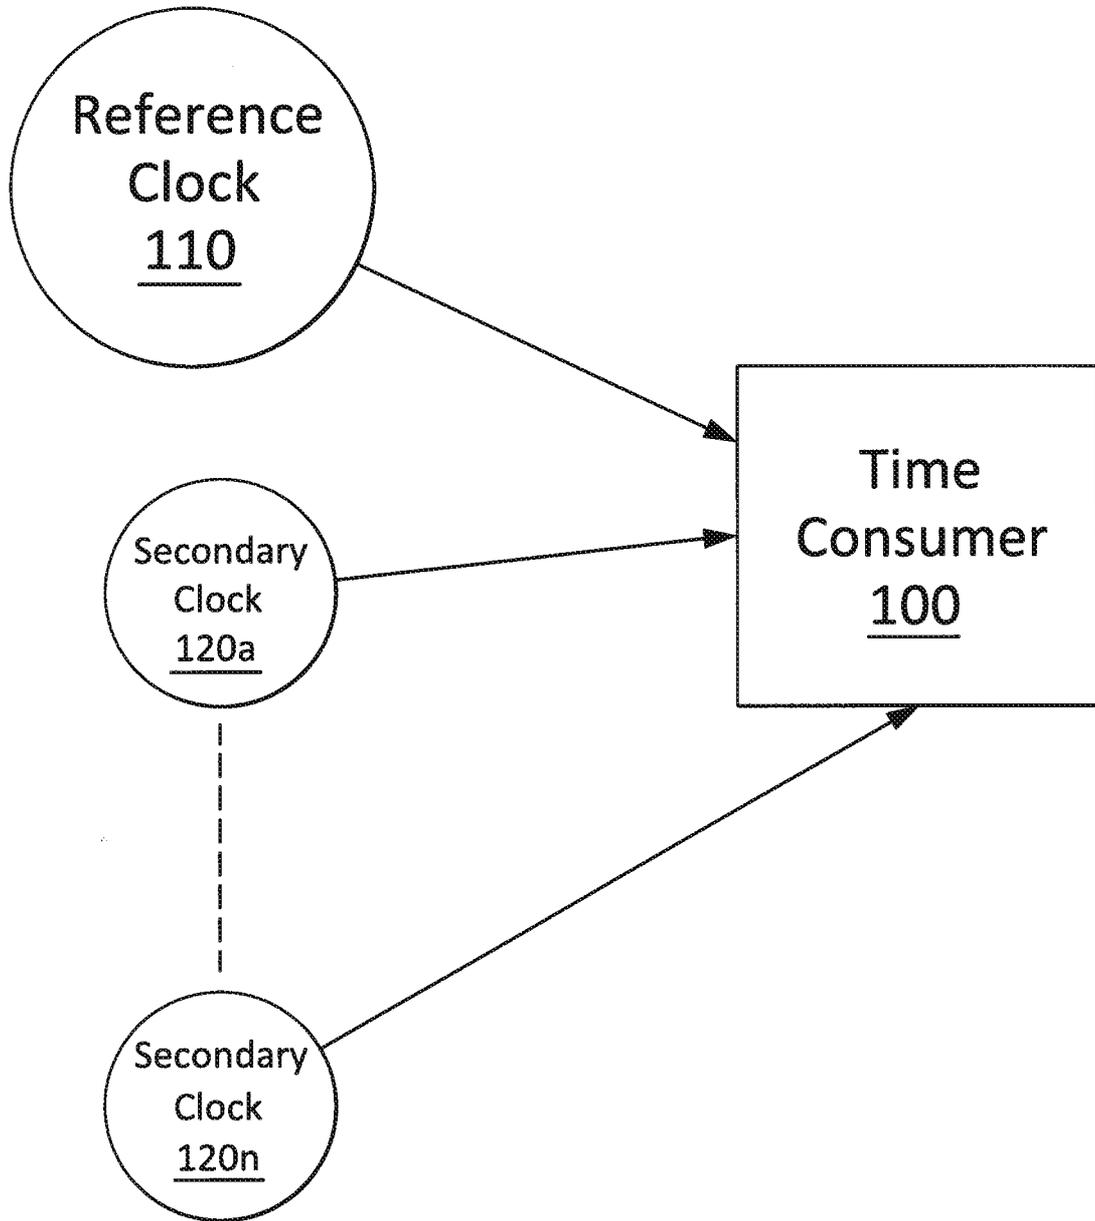


FIG. 2

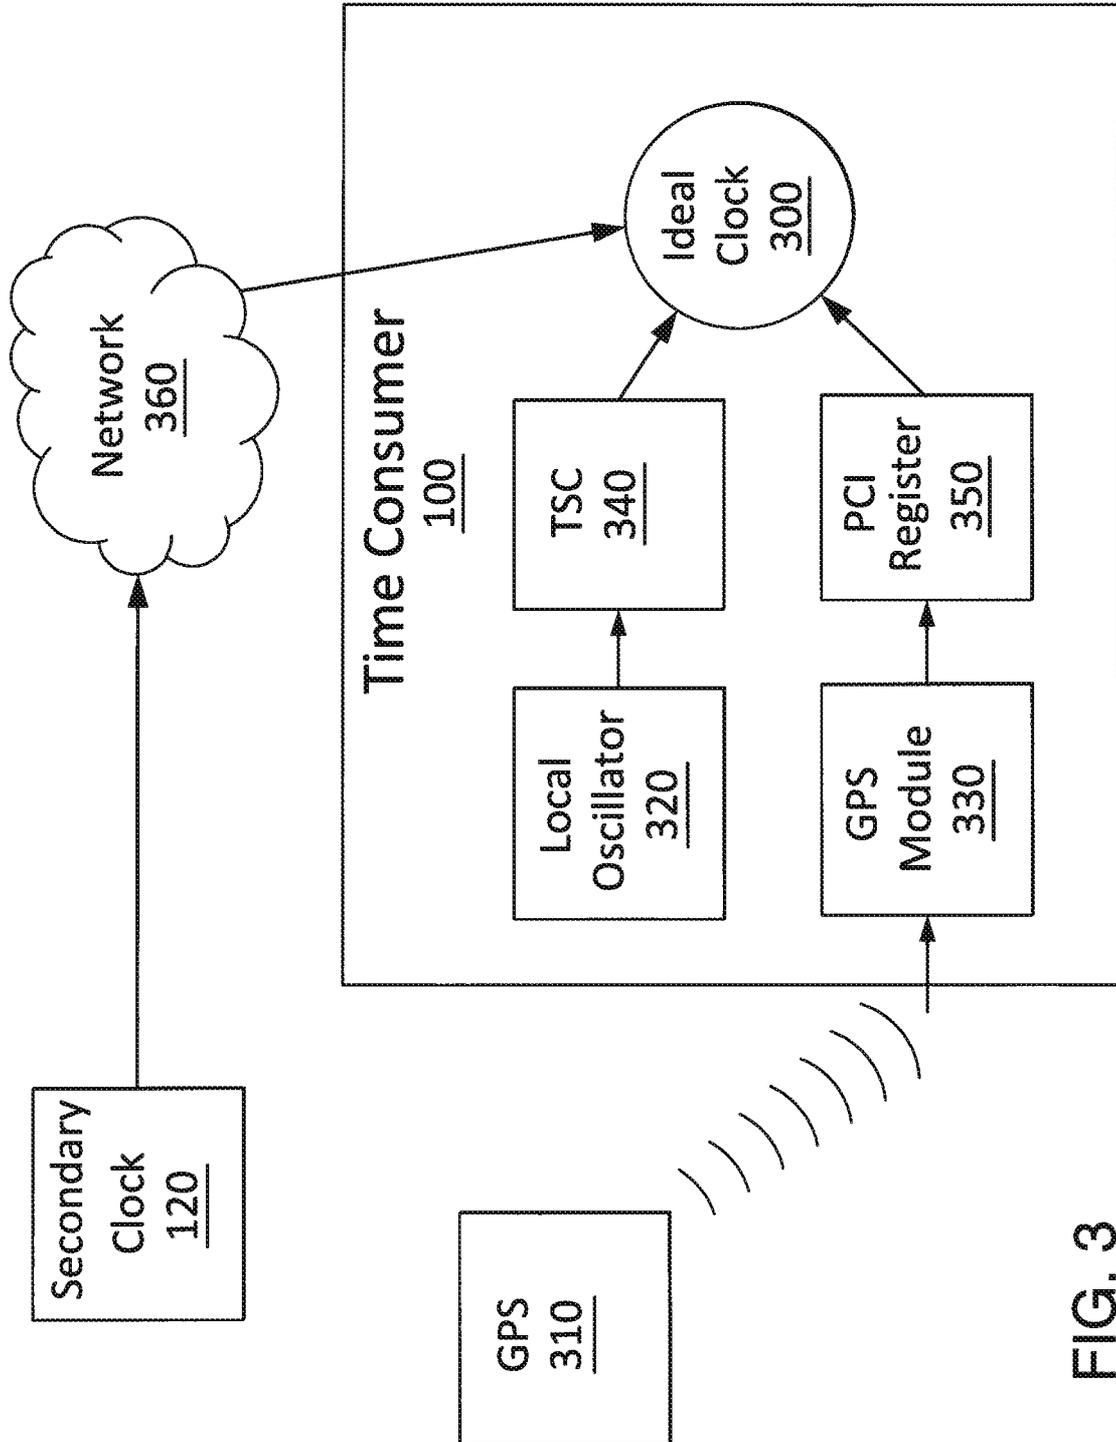


FIG. 3

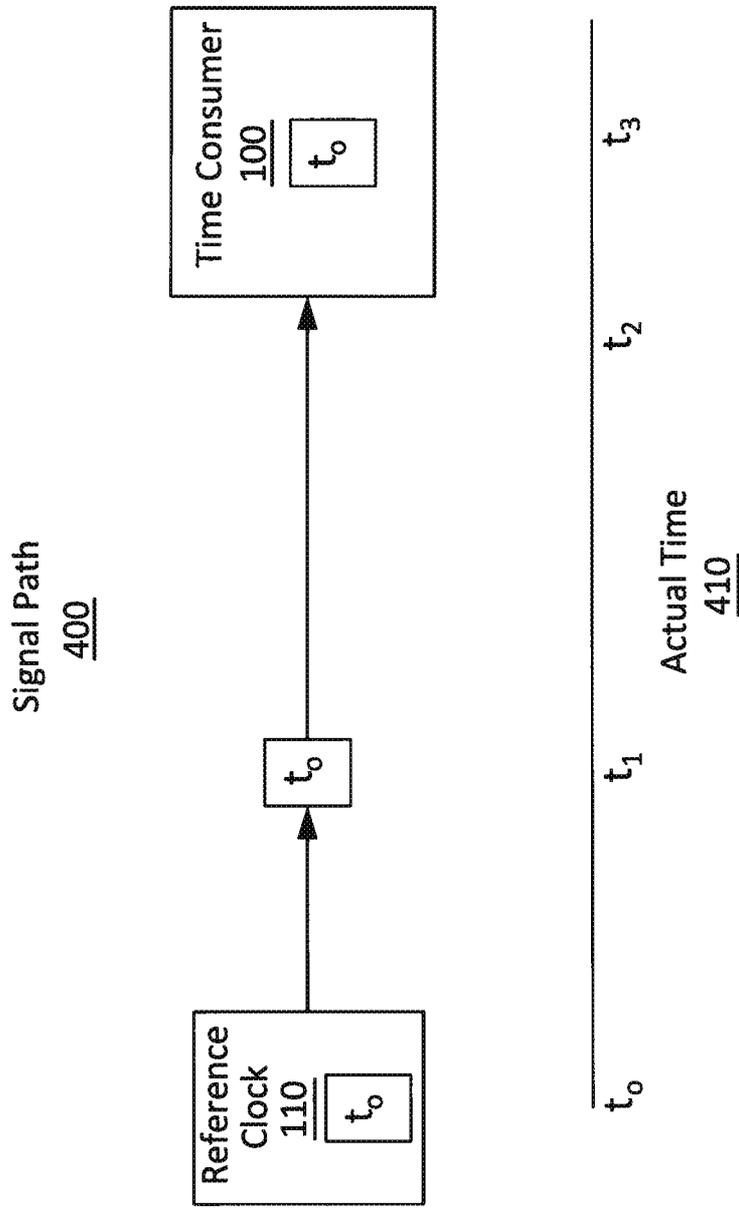


FIG. 4

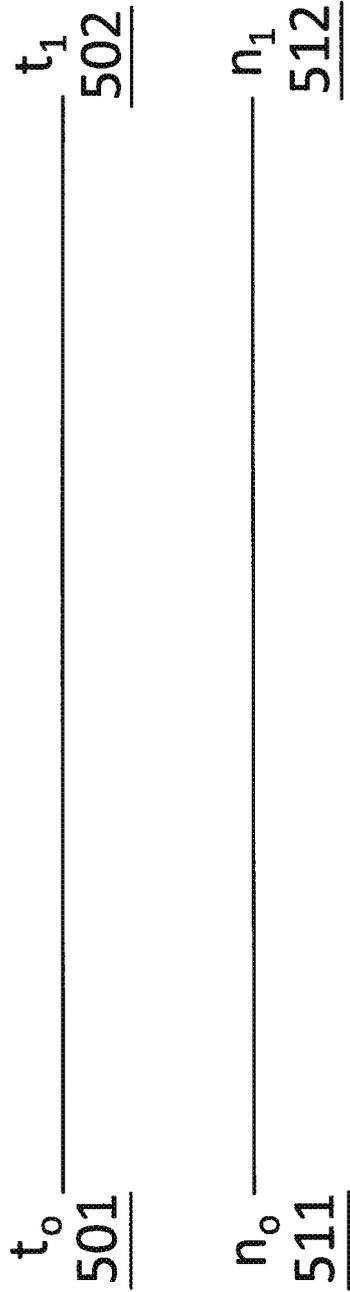
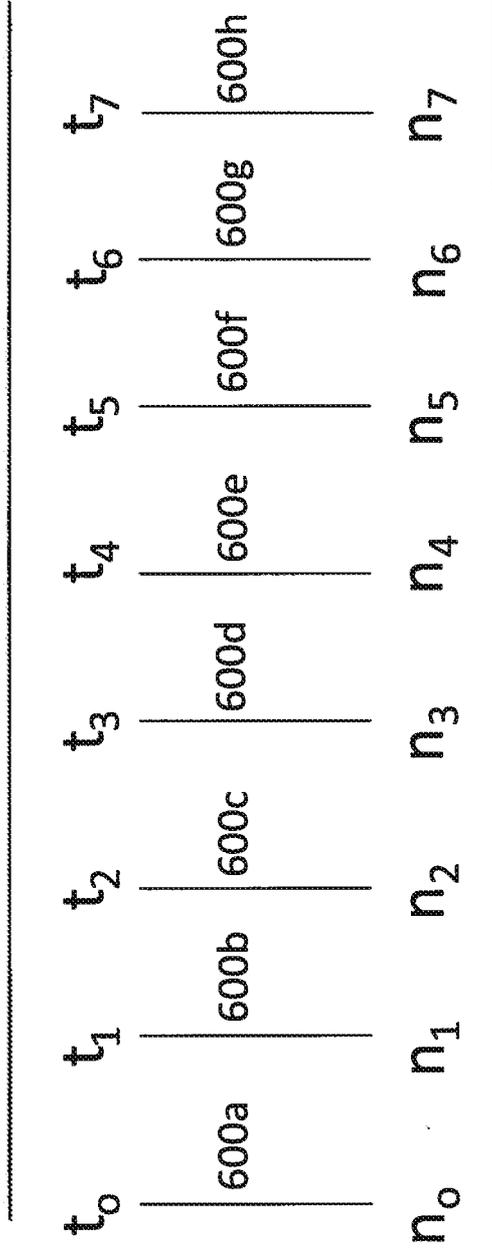


FIG. 5

Times from secondary clock



Tick counts from local clock

FIG. 6

7/11

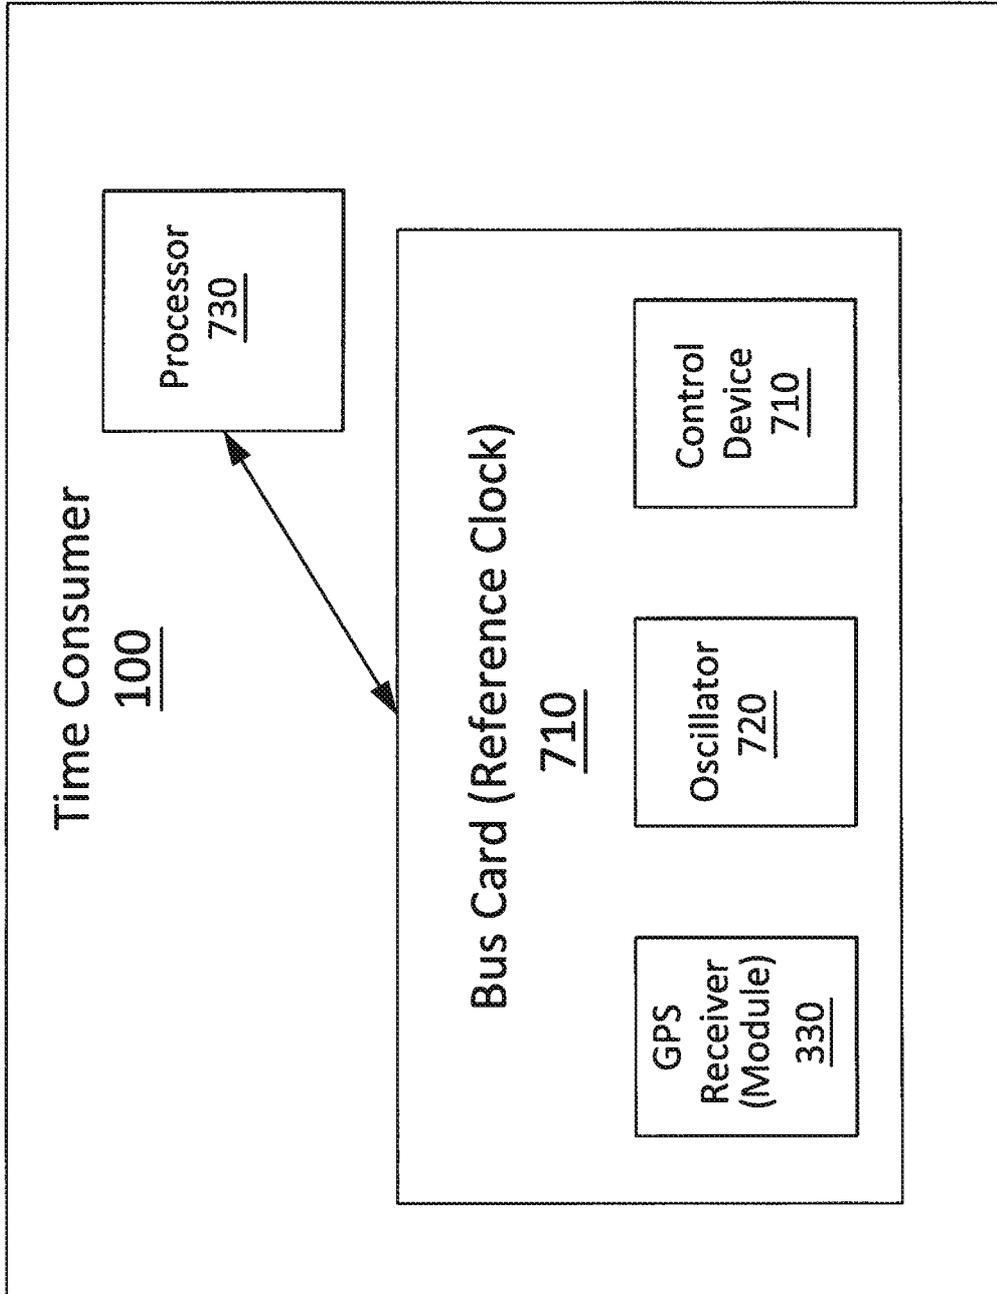


FIG. 7

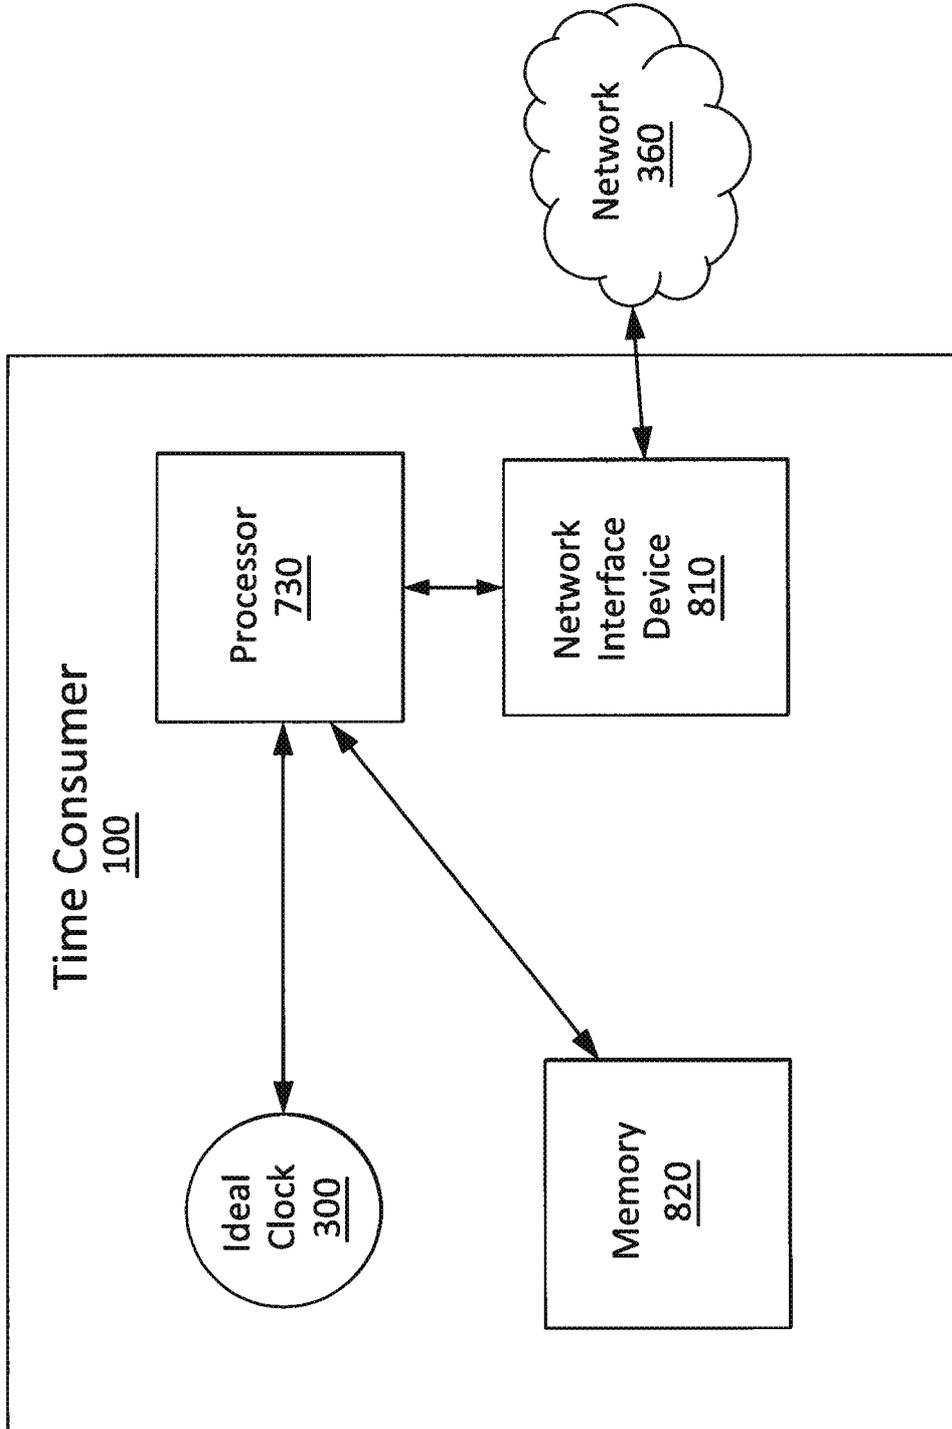


FIG. 8

9/11

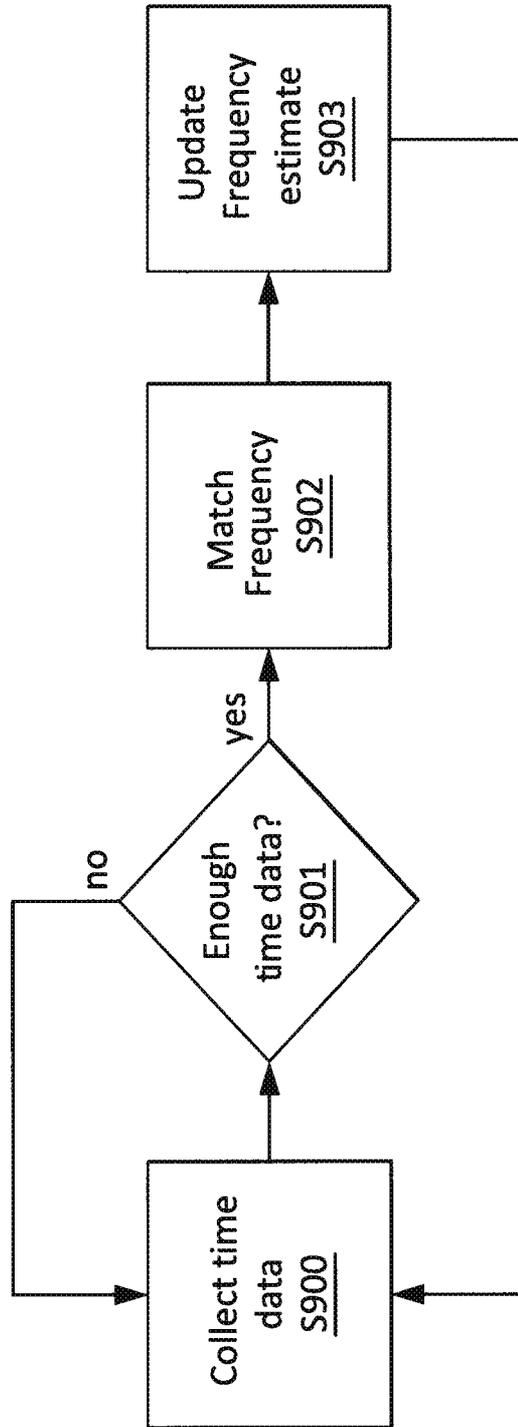


FIG. 9

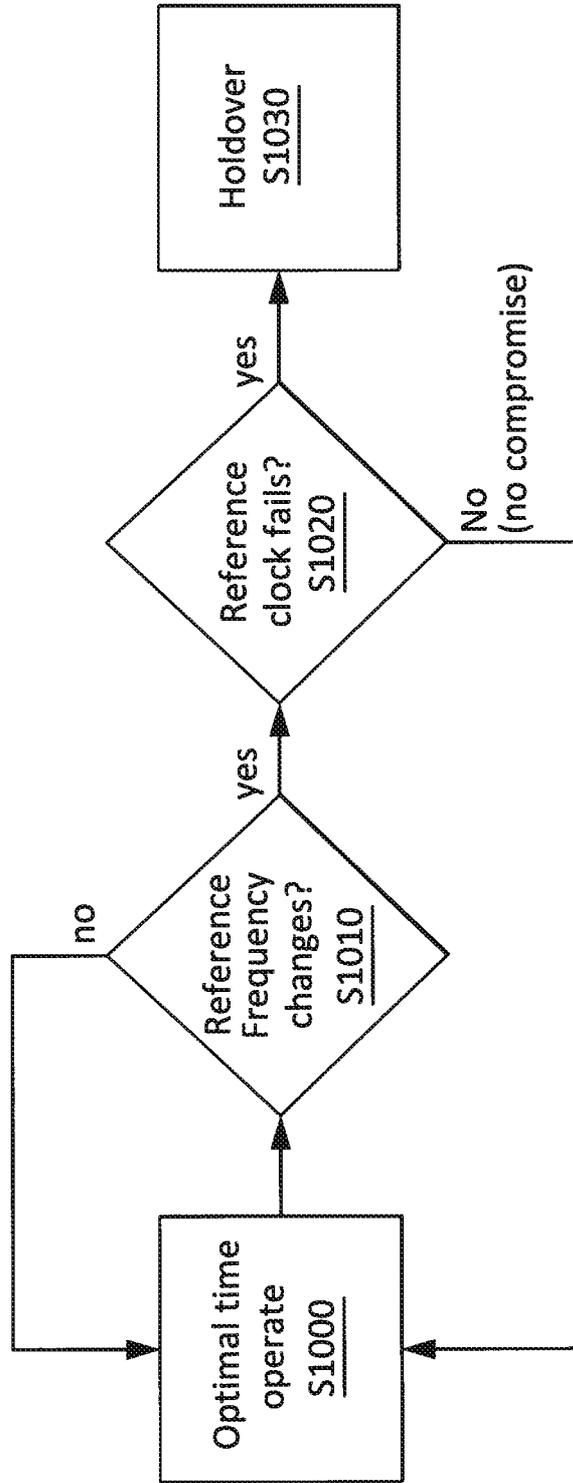


FIG. 10

11/11

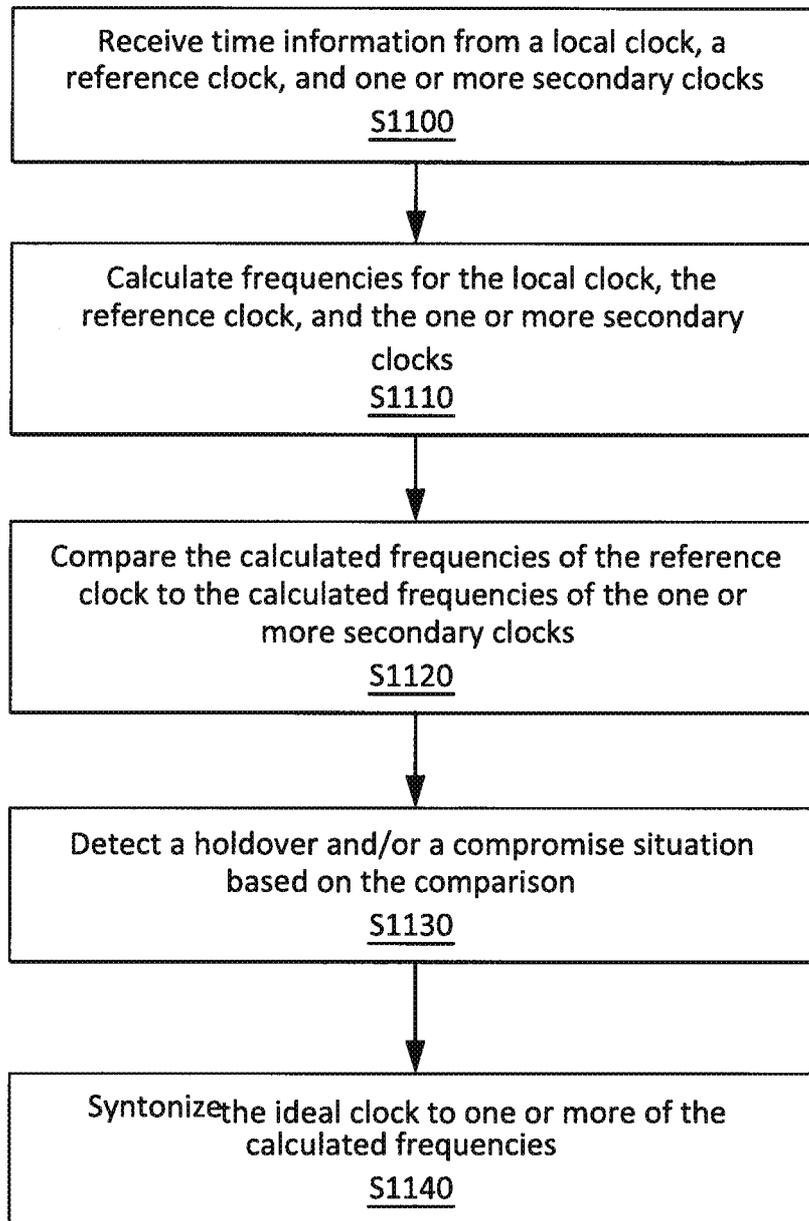


FIG. 11

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US13/48566

<b>A. CLASSIFICATION OF SUBJECT MATTER</b> IPC(8) - G06F 1/12; G01R 23/00 (2013.01) USPC - 368/56; 714/400 According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b> Minimum documentation searched (classification system followed by classification symbols) IPC(8) Classification(s): G06F 1/12; G01R 23/00 (2013.01) USPC Classification(s): 368/46, 52, 56; 714/400 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) MicroPatent (US Granted, US Applications, EP-A, EP-B, WO, JP, DE-G, DE-A, DE-T, DE-U, GB-A, FR-A); DialogPro (Derwent, INSPEC, NTIS, PASCAL, Current Contents Search, Dissertation Abstracts Online, Inside Conferences); IP.com; IEEE.com; Google Scholar, clock syntonizing holdover compromise, multiple reference clocks, time keeping, frequency comparison, accurate time,		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 8,120,432 B2 (SMILEY, R, et al.) February 21, 2012, figures 1-4, column 1, line 1- column 8, line 21; column 9, line 1- column 11, line 41; column 13, line 4- column 15, line 55; column 17, lines 15-24.	1-3, 8-10, 15-17
A	US 2011/0261917 A1 (BEDROSIAN, P.) October 27, 2011, entire document.	1-3, 8-10, 15-17
A	US 2010/0254225 A1 (SCHWEITZER, E, III et al.) October 7, 2010, entire document.	1-3, 8-10, 15-17
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/>		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "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) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "I" 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 "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 "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 "&" document member of the same patent family		
Date of the actual completion of the international search 22 October 2013 (22.10.2013)		Date of mailing of the international search report <b>22 NOV 2013</b>
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201		Authorized officer: Shane Thomas PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US13/48566

**Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1.  Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
  
2.  Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
  
3.  Claims Nos.: 4-7, 11-14, 18-20  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

**Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)**

This International Searching Authority found multiple inventions in this international application, as follows:

1.  As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2.  As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3.  As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
  
4.  No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

**Remark on Protest**

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.