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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
27 May 2010 (27.05.2010)

PCT

(10) International Publication Number
WO 2010/059331 A1

(51) International Patent Classification:
H03L 7/22 (2006.01) *H04L 7/00* (2006.01)

(21) International Application Number:
PCT/US2009/061847

(22) International Filing Date:
23 October 2009 (23.10.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/116,628 20 November 2008 (20.11.2008) US

(71) Applicant *(for all designated States except US)*: **RAMBUS INC.** [US/US]; 4440 El Camino Real, Los Altos, CA 94022 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **YIP, Tsunwai, Gary** [US/US]; 6363 Wisteria Way, San Jose, CA 95129 (US). **CHANG, Kun-yung** [US/US]; 296 Ramon Drive, Los Altos, CA 94024 (US). **NGUYEN, David** [US/US]; 10925 Miramonte Road, Cupertino, CA 95014 (US).

(74) Agent: **KREISMAN, Lance, M.**; Mahamedi Paradice Kreisman LLP, 550 South Winchester Blvd., Suite 605, San Jose, CA 95128 (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, 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, KM, 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, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *of inventorship (Rule 4.17(iv))*

Published:

— *with international search report (Art. 21(3))*

(54) Title: JITTER CORRECTION IN HIGH SPEED DATA LINK

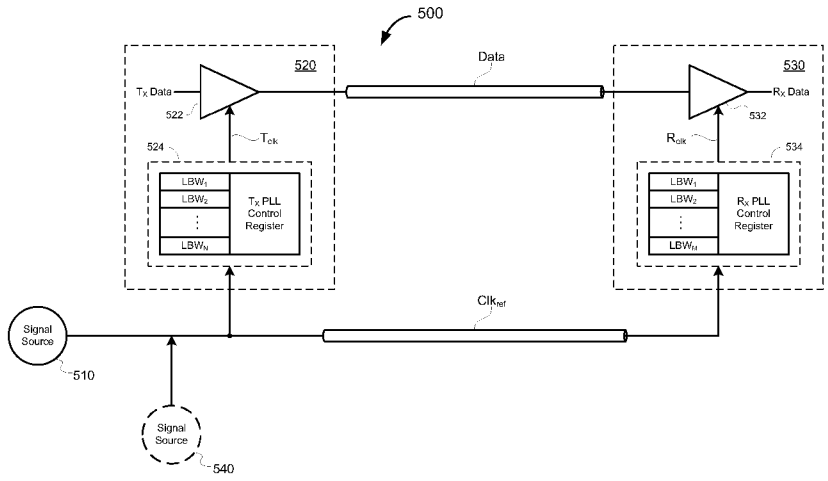

FIG. 5

(57) Abstract: A jitter correction operation, wherein a first timing signal is generated in a locked-loop circuit to time a transmission of a data signal within a corresponding transmit circuit. A second timing signal is generated within another locked-loop circuit to time a sampling of the data signal within a corresponding receive circuit. A loop bandwidth configuration of at least one of the locked-loop circuits is adjusted in order to reduce a difference between loop bandwidths of the respective locked loop circuits.

WO 2010/059331 A1

JITTER CORRECTION IN HIGH SPEED DATA LINK

TECHNICAL FIELD

[0001] The disclosure herein relates to data communications systems generally and more specifically to mitigating the effects of jitter in such systems.

BACKGROUND

[0002] Jitter effects random offsets in the phase of the timing signal, causing an edge of the actual timing signal to be slightly early or late relative to an ideal timing signal edge. During high speed data transfers, even the slightest offsets in timing may significantly reduce the effective margin within which a corresponding data signal may be accurately sampled.

[0003] Figure 1 illustrates a timing relationship where certain lines are shown in bold to illustrate possible effects of jitter on the sampling margin. More specifically, an early data signal and a late timing signal ("Rclk") are highlighted to show the reduced sampling margin which may result from jitter. When the sampling instant is late with respect to the center of the data eye, there is less room for error in the timing of the Rclk and data signals. As a result of this reduced sampling margin, there is an increased likelihood for the data signal to be inaccurately sampled.

[0004] Conventional approaches to mitigating jitter-induced sampling-time error generally involve adding high-precision filters or other circuitry within the on-chip timing circuitry to reduce timing signal jitter. Unfortunately, such circuitry tends to add substantial cost and consume valuable die area.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] The disclosure herein is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

Figure 1 illustrates a timing relationship where certain lines are shown in bold to illustrate possible effects of jitter on sampling margin;

Figure 2 illustrates one embodiment of a jitter error correction operation;

Figure 3 illustrates a timing relationship between a data signal and a receive timing signal where jitter is suppressed using jitter alignment;

Figure 4 illustrates one embodiment of a data communications system that uses jitter error correction;

Figure 5 illustrates an exemplary data communications system with jitter correction;

Figure 6 illustrates a table lookup operation for determining which phase-locked loop circuit (e.g., a transmit or receive circuit) has the greater loop bandwidth configuration;

Figure 7 illustrates a test and evaluation operation for determining which phase-locked loop circuit has the greater loop bandwidth configuration;

Figure 8 illustrates a method of single-step programming of the loop bandwidth configurations;

Figure 9 illustrates a method that relies upon iterative adjustment of the loop bandwidth in order to suppress the effects of jitter;

Figure 10 illustrates another method of iterative adjustment of the loop bandwidth configurations;

Figure 11 illustrates another method of jitter error correction;

Figure 12 illustrates an exemplary data communications system having control circuitry to control jitter correction; and

Figure 13 illustrates an exemplary data communications system having control circuitry to control jitter correction across multiple data communications devices.

DETAILED DESCRIPTION

[0006] In embodiments disclosed herein, a method and apparatus is provided for reducing the effects of jitter on the sampling margin within which to sample a data signal. The jitter of a receive timing signal (e.g., a clock or strobe signal) can be substantially matched to the jitter of a transmit timing signal (used for generating the data signal), such that the receive timing signal effectively tracks jitter in the corresponding data signal. In more specific embodiments, one or more phase-locked loop (PLL) circuits having adjustable loop bandwidth configurations can be adjusted to match the jitter profiles of both the data signal and the receive timing signal. For purposes of discussion, the terms (a) "jitter" and "noise" and (b) "loop bandwidth configuration" (or LBW configuration) and "loop bandwidth setting" (or LBW setting) may each be used interchangeably. Also, although specific reference is made to phase-locked loop circuits, one of ordinary skill in the art may apply the teachings disclosed herein to delay-locked loop (DLL) circuits, or any other locked-loop circuit configurations, with relatively little or no modification.

[0007] Many of the embodiments described in this disclosure involve matching loop bandwidth configurations in both of the transmit (T_X) and receive (R_X) PLL circuits, to effectively match jitter and use generated timing signals in a manner that partially or completely cancels out jitter. In many of the embodiments discussed in this disclosure, the same reference timing signal is provided to both the T_X PLL and the R_X PLL (and these circuits are then adjusted to match jitter, to the extent possible). However, in other embodiments, a different reference timing signal may be provided to each of the T_X PLL and the R_X PLL, individually.

[0008] Each of the T_X and R_X PLLs effects low-pass filtering on the received reference timing signal, the low-pass filtering being defined by a loop bandwidth (e.g., cut-off frequency) of the respective PLL circuit. In other words, the T_X PLL filters a portion of the jitter profile of the reference timing signal according to a specific loop bandwidth (LBW) configuration of the T_X PLL circuit. Similarly, the R_X PLL filters a portion of the jitter profile of the reference timing signal according to a LBW configuration that is specific to the R_X PLL circuit.

[0009] The T_X PLL and the R_X PLL are generally packaged independently of one another (e.g., they are manufactured by different sources), and are thus initially configured with different LBW settings. However, there is very little likelihood of the two PLLs having identical LBW configurations even if both are provided by the same manufacturer. As a result, the jitter provided to each of the Tclk and Rclk signals are typically out of sync, which contributes to a reduction in the data sampling margin, as described above.

[0010] There may also be a fixed offset in signal propagation times between the transmit circuit and the receive circuit. For example, depending on the length and/or type of transmission mediums used for signaling between the T_X and R_X PLLs, different amounts of delay may be applied to each of the data signal and a corresponding reference signal. However, such differences in propagation time are also mitigated while adjusting for jitter according to the embodiments herein.

[0011] Figure 2 illustrates one embodiment of a jitter error correction operation. Jitter is a characteristic of timing signals (e.g., strobe and/or clock signals) which is typically introduced by a timing signal source and its associated distribution circuitry. At 210, a transmit timing (Tclk) signal is first generated within a transmit phase-locked loop (T_X PLL) circuit. For example, this typically involves passing a reference timing signal through a phase detector and a voltage controlled oscillator (VCO), an output of which is provided back to the phase detector in a negative feedback configuration. A frequency divider circuit is typically provided along

the negative feedback path of a PLL circuit. In an alternative embodiment, a delay-locked loop (DLL) circuit may be substituted for the T_X PLL simply by removing (or omitting) this frequency divider component. The T_X PLL circuit scales the reference timing signal (typically operating at a relatively low frequency) to a substantially higher frequency (i.e., the desired frequency of the data signal), thus generating the Tclk signal used for transmitting data. The Tclk signal is a periodic clock signal, but as with each of the embodiments discussed in this disclosure, other forms of timing signals may also be used, such as for example a strobe signal.

[0012] A receive timing (Rclk) signal is generated within a receive (R_X) PLL circuit, at 220. As was mentioned above in the case of the T_X PLL, a DLL or other locked-loop circuit may also be substituted for the R_X PLL. The receive PLL circuit scales a reference timing signal (e.g., to a higher frequency) to match the frequency of the Tclk signal, such that the Rclk signal may be aligned with data transmissions and used by a receiver to sample those transmissions at the appropriate times.

[0013] Step 230 illustrates an optional intermediate step of determining which phase locked-loop circuit has the greater (possible) loop bandwidth configuration. For example, the relative LBWs may be determined by looking up their respective values in a lookup table (e.g., based on the manufacturer's specifications for each of the PLL circuits), as discussed in greater detail below in reference to Figure 6. The relative LBWs may also be determined using a test and evaluation process, as discussed in greater detail below in reference to Figure 7.

[0014] At 240, one or both of the T_X and R_X PLLs are programmed to have nominally matching LBW configurations. In other words, the LBW configurations for the T_X and/or R_X PLLs are chosen such that the effects of jitter on data sampling between the T_X and R_X circuits are substantially reduced (e.g., mitigated). In the case where only one of the PLL circuits is programmable (i.e., has an adjustable LBW configuration), the LBW configuration of the programmable PLL is adjusted at step 240. Assuming both PLL circuits have adjustable LBW

configurations, then each of the T_X and R_X PLLs may be programmed with a (respective) nominal LBW configuration.

[0015] The programming operation, at 240, may be performed in either a single programming step, or via a series of iterative adjustments. As will be described in greater detail below, matching the LBW configurations of the T_X and R_X PLLs effectively synchronizes the jitter in the Tclk signal with the jitter in the Rclk signal. For example, matching the LBW setting of the T_X PLL with the LBW setting of the R_X PLL enables the Rclk signal to substantially track the jitter in the Tclk signal, thus allowing the sampling margin for the data communication to be maintained relatively constant (i.e., near-ideal). Several embodiments of this operation will be presented in the discussion below, including: (a) an embodiment where the PLL with the smaller maximum LBW is set to maximum, and the other PLL programmed to match it as closely as possible; (b) an embodiment where only one PLL is adjustable, and is programmed to have its LBW match the LBW of the non-programmable; and (c) embodiments where different combinations of LBW (irrespective of which is greater or smaller) are experimented with, to determine a combination which gives the least jitter.

[0016] The aforementioned techniques for jitter "alignment" are advantageous for a number of reasons. For example, rather than require very precise (i.e., expensive) low-pass filtering circuitry to remove the jitter from both the Rclk and Tclk signals, low-cost PLL architectures may be leveraged by nominally matching the jitter in the Rclk signal to the jitter in the Tclk signal (or vice-versa). This is further advantageous because it mitigates the relative timing error on data sampling due to jitter.

[0017] Figure 3 illustrates a timing relationship between a data signal and a receive timing (Rclk) signal where jitter is suppressed using jitter alignment. A jitter-corrected Rclk signal is shown alongside a corresponding Data signal exhibiting jitter. The dotted lines illustrate ideal (i.e., jitter-free) waveforms of the respective Data and Rclk signals, to which the jitter-afflicted

Data and Rclk signals are compared (indicated by the solid lines). The ideal sampling margin is illustrated as the interval (W_{ID}) from the midpoint of the ideal data eye to either edge (i.e., transition) of the data eye. As shown, an ideal sampling instant (of the ideal Rclk signal) is aligned with the ideal data eye midpoint. Thus, a leading edge of the ideal Rclk signal typically trails a corresponding leading edge of the Data signal by an interval substantially equal to the ideal sampling margin W_{ID}. Timing instants of the Tclk signal are indicated by corresponding edges of the Data signal.

[0018] At T1, the edge of the actual Data signal is early with respect to a corresponding edge of the ideal Data signal, which arrives (i.e., transitions) at T2, due to jitter in the Tclk signal. At T5, the edge of the actual Data signal is again early with respect to a corresponding edge of the ideal Data signal, which arrives at T6. However, it should be noted that the duration of time between T1 and T2 is greater than the duration of time between T5 and T6. In other words, the leading edge of the Data signal (i.e., at T1) is "more early," relative to the ideal Data signal, than the trailing edge of the Data signal (i.e., at T5). This, in effect, produces a wider-than-ideal data eye, and causes the midpoint of the data eye to be located farther (relative to either edge of the data eye) than the ideal sampling margin W_{ID}.

[0019] At T8, the edge of the actual Data signal is late with respect to a corresponding edge of the ideal Data signal, which arrives at T7, due to jitter. At T12, the edge of the actual Data signal is again late with respect to a corresponding edge of the ideal Data signal, which arrives at T11. This example has been drawn so that the duration of time between T7 and T8 is greater than the duration of time between T11 and T12, i.e., such that the leading edge of the Data signal (i.e., at T8) is "more late," relative to the ideal Data signal, than the trailing edge of the Data signal (i.e., at T12). This effectively produces a narrower-than-ideal data eye, and causes the midpoint of the data eye to be located closer (relative to either edge of the data eye) than the ideal sampling margin W_{ID}.

[0020] Referring now to the Rclk signal, it shown that the sampling instant of the ideal Rclk signal at T4 is noticeably late with respect to the actual midpoint of the corresponding data eye. This putative timing error demonstrates the reduction of data sampling margin which would have occurred, purely as a result of the Data signal arriving early, at T1 and T5. On the other hand, the sampling instant of the ideal Rclk signal at T9 is noticeably early with respect to the actual midpoint of the corresponding data eye. This putative timing error shows the reduction of data sampling margin which would have occurred, purely as a result of the Data signal arriving late, at T8 and T12.

[0021] According to one embodiment, the same (or at least substantially similar) jitter that is represented in the Data signal (i.e., Tclk signal) is effectively applied to the Rclk signal. To implement this objective, the transmit and receive PLLs (used in generating the Tclk and Rclk signals, respectively) may be programmed to reduce differences between their loop bandwidth configurations, e.g., so as to be matched as closely as possible. This goal may be effected by programming the LBW configuration of one of the PLL circuits (i.e., either the transmit PLL or the receive PLL) to nominally match the LBW configuration of the other PLL circuit, or by programming both PLLs if such is supported by the system. This effectively synchronizes the jitter in the Rclk signal to the jitter in the Tclk signal, thus enabling the timing errors in the Rclk signal to substantially offset the timing errors in the Tclk signal.

[0022] As a result, the actual Rclk signal at T3 leads the ideal Rclk signal (at T4) by substantially the same interval with which the actual Data signal at T1 leads the ideal Data signal (at T2). The interval between the actual Rclk signal and the actual Data signal thus remains substantially equal to the ideal sampling margin W_{ID} (i.e., $|T1 - T3| \approx W_{ID}$). It should be noted that the sampling instant of the actual Rclk signal at T3 may still be slightly early relative to the midpoint of the data eye, since the actual data eye is wider than ideal as depicted in FIG. 3. However, the sampling time error (i.e., the timing offset of the sampling instant with

respect to the midpoint of the data eye) is substantially mitigated in comparison to the putative timing error associated with the ideal Rclk signal (at T4).

[0023] After correcting for jitter, the actual Rclk signal at T10 trails the ideal Rclk signal (i.e., at T9) by substantially the same interval with which the actual Data signal at T8 trails the ideal Data signal (at T7). Accordingly, the interval between the actual Rclk signal and the actual Data signal remains substantially equal to the ideal sampling margin W_{ID} (i.e., $|T8 - T10| \approx W_{ID}$). Paralleling the example presented just above, the sampling instant of the actual Rclk signal at T10 may still be slightly late with respect to the midpoint of the data eye, since the actual data eye is depicted as narrower than ideal. However, the sampling time error is substantially mitigated in comparison to the putative timing error associated with the ideal Rclk signal (at T9).

[0024] Synchronizing jitter in the Rclk and Tclk signals offers greater resiliency, in the sampling time error, to random phase fluctuations in both the Rclk and Tclk signals. This, in turn, allows for more robust performance in such data communications systems.

[0025] In another embodiment, jitter adjustment may be performed on only one of the timing signals (e.g., the Rclk signal or the Tclk signal, used for generating the Data signal), or on both the Tclk and Rclk signals. It should be noted that, although the Data and Rclk signals are illustrated as differential and single-ended waveforms, respectively, they are not so limited. In other words, the Data signal may alternatively be a single-ended waveform and/or the Rclk signal may be a differential waveform.

[0026] Figure 4 illustrates one embodiment of a data communications system that uses jitter error correction. The data communications system 400 includes: a signal source 410; a transmit circuit 420, made up of a data transmitter 422 and a phase-locked loop (T_X PLL) circuit 424; and a receive circuit 430, made up of a data receiver 432 and a phase-locked loop (R_X PLL) circuit 434. For purposes of describing an exemplary operation of the data communications

system 400, it is assumed that both the T_X PLL 424 and the R_X PLL 434 have programmable loop bandwidths. The signal source outputs a reference timing signal (Clk_{ref}) to the transmit circuit 420 and the receive circuit 430. For example, the Clk_{ref} signal may either be a periodic clock signal or a strobe signal.

[0027] The signal source 410 typically operates at a relatively low frequency (e.g., to save on power and manufacturing costs). Thus, the T_X PLL 424 receives the low-frequency Clk_{ref} signal and generates a transmit timing (Tclk) signal at a specified frequency for data communication. The data transmitter 422 outputs a data signal to the receive circuit 430 by transmitting the T_X Data in response to the Tclk signal. For example, the data transmitter 422 may be a digital-to-analog converter (DAC).

[0028] The R_X PLL 434 receives the Clk_{ref} signal and generates a receive timing (Rclk) signal at the specified frequency. The data receiver 432 then samples the data signal in response to the Rclk signal to reproduce the T_X Data (e.g., as R_X Data). For example, the data receiver 432 may be an analog-to-digital converter (ADC).

[0029] If desired, the signaling source 410 may be integrated with the transmit circuit 420 or the receive circuit (i.e., the transmit circuit 420 may output the Clk_{ref} signal to the receive circuit 430, or vice-versa). In another embodiment, separate signal sources may be provided for each of the transmit circuit 420 and the receive circuit 430. The PLL circuits (i.e., the T_X PLL 424 and/or the R_X PLL 434) may be implemented as a delay-locked loop (DLL) circuit, or other similar locked-loop circuitry. It should also be noted that, in some embodiments, either the T_X PLL 424 and/or the R_X PLL 434 may have a static or fixed loop bandwidth.

[0030] The graph 460 illustrates the transfer functions for the T_X PLL and the R_X PLL (H_{TX_PLL} and H_{RX_PLL}, respectively). Frequency is plotted on the x-axis and amplitude (e.g., attenuation) is plotted on the y-axis. As illustrated, the T_X PLL and the R_X PLL filter the received Clk_{ref} signal according to the transfer functions H_{TX_PLL} and H_{RX_PLL}, respectively.

The Clk_{ref} signal has an inherent jitter profile which varies with respect to frequency. In this example, LBW_{TX} and LBW_{RX} represent the maximum loop bandwidth settings of the T_X and R_X PLLs, respectively. The mismatch between LBW_{TX} and LBW_{RX} effectively creates a jitter passband, wherein the jitter profile filtered through the T_X PLL is different than the jitter profile filtered through the R_X PLL. This jitter passband typically causes the jitter in the T_X PLL to be out of sync with the jitter in the R_X PLL, thus resulting in a relative timing error between the Rclk signal and the Tclk signal.

[0031] The graph 480 illustrates the transfer functions H_{TX_PLL} and H_{RX_PLL} after loop bandwidth adjustment. As shown, the jitter passband may be substantially reduced by lowering the loop bandwidth of the R_X PLL to match the loop bandwidth of the T_X PLL. In this specific example, loop bandwidth adjustment is performed within the R_X PLL. However, in other embodiments, the loop bandwidth of the T_X PLL may be adjusted to match the loop bandwidth of the R_X PLL (e.g., assuming the maximum LBW_{TX} is greater than the maximum LBW_{RX}). Eliminating the jitter passband allows the jitter in the Rclk signal to be effectively in sync with the jitter in the Tclk signal. As a result, the relative timing error between the Rclk signal and the Tclk signal may be substantially mitigated.

[0032] Figure 5 illustrates an exemplary data communications system with jitter correction. The data communications system 500 includes a signal source 510, a transmit circuit 520, and a receive circuit 530. The signal source 510 outputs a reference timing signal (Clk_{ref}) to the transmit circuit 520 and the receive circuit 530.

[0033] The transmit circuit 520 includes a data transmitter 522 and a phase-locked loop (T_X PLL) circuit 524. The T_X PLL 524 has a programmable loop bandwidth, wherein a selected one of the loop bandwidths LBW_1-LBW_N may be provided to the T_X PLL control register as the loop bandwidth of the T_X PLL 524. For example, the control register may store one of a set number of preconfigured loop bandwidth settings, or alternatively, the control register may be

programmed to include any value within the continuous range of allowable loop bandwidths (i.e., LBW_1 to LBW_N). The T_X PLL 524 receives the Clk_{ref} signal and generates a transmit timing (Tclk) signal at a specified frequency. The data transmitter 522 outputs a data signal to the receive circuit 530 by transmitting the T_X Data in response to the Tclk signal.

[0034] The receive circuit 530 includes a data receiver 532 and a phase-locked loop (R_X PLL) circuit 534. The R_X PLL has a programmable loop bandwidth, wherein a selected one of the loop bandwidths LBW_1-LBW_M may be provided to the R_X PLL control register as the loop bandwidth of the R_X PLL 524. As with the T_X PLL, the control register may store one of a set number of preconfigured loop bandwidth settings, or alternatively, the control register may be programmed to include any value within the continuous range of allowable loop bandwidths (i.e., LBW_1 to LBW_M). The R_X PLL 534 receives the Clk_{ref} signal and generates a receive timing (Rclk) signal at the specified frequency. The data receiver 532 then samples the data signal in response to the Rclk signal to reproduce the T_X Data (e.g., as R_X Data).

[0035] Jitter correction is performed by matching the loop bandwidth of the T_X PLL 524 to that of the R_X PLL 534. In a first example, it is assumed that the maximum loop bandwidth of the R_X PLL 534 is greater than the maximum loop bandwidth of the T_X PLL 524 (e.g., $LBW_M > LBW_N$). Jitter correction in this embodiment may be performed by programming the control register of the T_X PLL 524 with its maximum loop bandwidth setting (e.g., LBW_N), and then programming the control register of the R_X PLL 534 with a loop bandwidth setting (e.g., LBW_1-LBW_M) that is substantially equivalent to the maximum loop bandwidth of the T_X PLL 524. In a second example, it is assumed that the maximum loop bandwidth of the T_X PLL 524 is greater than the maximum loop bandwidths of the R_X PLL 534 (e.g., $LBW_N > LBW_M$). In this case, jitter correction may be performed by programming the control register of the R_X PLL 534 with its maximum loop bandwidth setting (e.g., LBW_M), and then programming the control

register of the T_X PLL 524 with a loop bandwidth setting (e.g., LBW_1-LBW_N) that is substantially equivalent to the maximum loop bandwidth of the R_X PLL 534.

[0036] The data communications system 500 may include an additional signal source 540 for injecting additional noise into the Clk_{ref} signal for calibration and/or testing purposes. For example, it may be desirable to monitor the effects of loop bandwidth adjustment in the T_X PLL 524 and/or the R_X PLL 534, relative to changes in a signal quality metric. Thus, the signal source 540 may be controlled to inject additional noise into the Clk_{ref} signal. In this embodiment, the signal source 540 may be programmable (e.g., having a programmable frequency and/or amplitude). If desired, the signal source 540 may be integrated with either the transmit circuit 520 or the receive circuit 530, or both

[0037] Figure 6 illustrates a table lookup operation which may be used to determine which phase locked-loop circuit has the greater loop bandwidth configuration. Once this is determined, the circuit with the greater LBW may then be selected for adjustment to minimize difference between the LBWs. For purposes of FIG. 6, it should be assumed that the each of the T_X and the R_X PLL are programmable PLL circuits having adjustable (e.g., multiple) LBW configurations.

[0038] At 610, a lookup of the loop bandwidth for the receive PLL is performed. For this example, specific LBW configurations for the R_X PLL may be programmatically stored within a lookup table (or register) in the R_X PLL circuit. In other words, the lookup table may store a number of possible "configuration values" indicating the possible LBW configurations for the corresponding PLL circuit. The LBW configurations may be stored (or indexed) in order of lowest (e.g., minimum) to greatest (e.g., maximum) configuration values, or vice-versa. A system (or device) controller may then interrogate (e.g., read from) the lookup table for the corresponding LBW configuration information. For example, the controller may extract all of the LBW configuration values from the table at once or, alternatively, the controller may

selectively extract only the configuration value associated with a given index (e.g., only the maximum configuration value).

[0039] A similar lookup of the loop bandwidth for the T_X PLL may then be performed, at 620. As described above, this may involve interrogating the programmed register in the T_X PLL for the corresponding configuration values or, alternatively, a separate device controller may be used to extract the LBW configuration information specifically stored in the T_X PLL.

[0040] At 630, the loop bandwidth of the R_X PLL (LB_{WRX}) is compared with the loop bandwidth of the T_X PLL (LBW_{TX}). For example, the maximum LBW configuration of the R_X PLL may be compared with the maximum LBW configuration of the T_X PLL. If LBW_{RX} is greater than LBW_{TX}, at 640, then the R_X PLL is determined to have the greater (possible) LBW configuration. On the other hand, if LBW_{TX} is greater than LBW_{RX}, at 650, then the T_X PLL is determined to have the greater (possible) LBW configuration.

[0041] Figure 7 illustrates a test and evaluation operation for determining which phase-locked loop circuit has the greater loop bandwidth configuration. For purposes of this discussion, it is assumed that both the R_X PLL and the T_X PLL have programmable LBWs and that these PLLs have each been programmed to have their maximum LBW; the LBWs for the two PLLs may then be adjusted to reduce jitter by reducing the LBW of the PLL with the larger LBW downward, so as to minimize the difference between the LBWs, and so that both PLL's are set to have the maximum matching LBWs possible. To perform this process, a designer may wish to implement a protocol (such as exemplified by Figure 7) for determining which PLL has the greater LBW and for responsively directing LBW adjustment efforts on that PLL.

[0042] At 710, the loop bandwidth configurations of both the transmit phase-locked loop (T_X PLL) circuit and the receive phase-locked loop (R_X PLL) circuit are set to their maximum values, respectively.

[0043] At 720, an initial signal quality metric (SQ_I) can be determined, and used as a frame of reference for subsequent jitter adjustment. For example, the signal quality metric SQ_I may correspond to a percentage of the unit interval (UI%) of the data signal. In general, the unit interval of the data eye corresponds to the width of the data eye, and thus the UI% of the data signal indicates the actual width of the data eye relative to the theoretical or ideal data eye width. For example, the initial signal quality metric SQ_I based on the UI% of the data signal may be less than 100% (as shown above, with respect to Figure 3).

[0044] The signal quality metric may alternatively correspond to any determinable characteristic of the data communication. For example, a predetermined sequence of data may be transmitted from the T_X PLL to the R_X PLL, and the number of incorrectly (or correctly) received data bits may be used as the initial signal quality metric SQ_I. In this example, errors or inconsistencies between the sampled data and the transmitted data may correspond to sampling errors resulting from jitter in both the T_X PLL and the R_X PLL. In another example, the signal quality metric may correspond to the overall size of the "passing window," or the duration of time for which a signal is valid (i.e., asserted or deasserted).

[0045] The loop bandwidth setting of the one of the PLLs (e.g., the R_X PLL) is subsequently adjusted, at 730. For example, the LBW setting of the R_X PLL may be adjusted by selecting the next-highest loop bandwidth setting (e.g., by programming LBW_{M-1} into the R_X PLL control register, assuming LBW_M is the maximum LBW setting) or, alternatively, by adjusting (e.g., lowering or reducing) a LBW control parameter of the R_X PLL (e.g., by a predetermined amount). Although, in this particular example, the method first operates on the R_X PLL, it should be noted that this choice is arbitrary, and either the either the T_X PLL and/or the R_X PLL may be adjusted to determine the greater of the two.

[0046] At 740, the adjusted signal quality metric (SQ_A) is determined following adjustment. The same "metric" used to determine the SQ_I value (e.g., UI%, accuracy of data transmission, or passing window size) is typically used in determining the SQ_A value.

[0047] Then, at 750, the adjusted signal quality metric SQ_A is compared with the initial signal quality metric SQ_I to determine whether the quality of the data signal improves in response to the loop bandwidth adjustment. For example, assuming the signal quality metric corresponds to a UI% of the data signal, and SQ_I is less than 100%, an improvement in signal quality may be reflected by an increase in the UI%. In other words, it is generally expected that successful jitter adjustment (i.e., by substantially matching the jitter in both the T_X and R_X PLLs) should result in an increased data eye width, such that the actual width of the data eye more closely resembles the theoretical width of the data eye. Notably, depending on the type of signal quality metric used, an increase in the SQ_A value (i.e., relative to the value of SQ_I) may not necessarily reflect an improvement in the quality of the data signal. For example, if the signal quality metric corresponds to a number of incorrectly received data bits ("bit error"), then an SQ_A value that is greater than the SQ_I value would, in fact, indicate degradation in the quality of the data transfer. In this event, it may be necessary to first evaluate how the particular metric varies in relation to changes in signal quality, in order to ensure proper interpretation of data.

[0048] Assuming that the adjusted signal quality metric improves over the initial signal quality metric (e.g., $SQ_A > SQ_I$), then the adjusted PLL circuit (e.g., the R_X PLL in this particular example) is said to have the greater loop bandwidth configuration, at 760, and it may then be used as the subject of adjustment in order to match the PLLs. Alternatively, if the adjusted signal quality metric does not improve over the initial signal quality metric (e.g., $SQ_A \leq SQ_I$), then the other PLL circuit (e.g., the T_X PLL) is said to have the greater loop bandwidth configuration, at 770, and that PLL circuit may be used as the subject of adjustment.

[0049] Notably, the method depicted in Figure 7 may also be used to determine a combination of PLL LBWs that minimizes error or maximizes data eye width, irrespective of which PLL has a greater LBW. In other words, if there is a combination of PLL LBWs (other than a maximum matching LBW configuration) which yields relatively good signal quality characteristics, then that combination of LBWs may be used. Thus, in at least one embodiment of the teachings discussed herein, one or both PLLs may be cycled through in terms of LBW configuration, with effects on signal quality measured and used to pick appropriate LBW values.

[0050] It should also be mentioned that there has been an increasing trend toward programmable PLLs; thus, the PLL with the greater maximum LBW is generally more likely to be programmable that the PLL with the smaller maximum LBW. For these reasons, it may (depending on implementation) be a sound design choice to identify the PLL have the greater maximum and operate upon it as a first step in attempting to match PLLs, as the PLL circuit determined to have the greater LBW configuration is generally (i.e., more likely to be) a programmable PLL.

[0051] Figure 8 illustrates one method of single-step programming of the loop bandwidths. At 810, the loop bandwidth of the selected phase-locked loop circuit (i.e., either the R_X PLL or T_X PLL) is programmed to nominally match the loop bandwidth of the other PLL circuit (e.g., such that $LBW_{RX} = LBW_{TX}$). Notably, the term "program" as used in this context includes both electrical programming and design choices effected during design or implementation; for example, the LBW configurations for the T_X PLL and/or the R_X PLL may be determined (e.g., chosen) during a design stage (i.e., prior to manufacturing), and the programming may be effected during manufacturing, e.g., by blowing fuses or hardwiring specific circuit elements. Alternatively, the programming step 810 may be performed by a controller, based on an in-situ relative loop bandwidth determination (e.g., step 230 of Figure 2). For example, the LBW of

the T_X PLL may be referenced in a lookup table, and the LBW configuration of the R_X PLL may then be set (by software, firmware or hardware) to be substantially equal to that of the T_X PLL (or vice-versa). Information pertaining to a signal quality metric may also be used in determining how the LBW configuration of the R_X PLL should be adjusted relative to the LBW configuration of the T_X PLL (or vice-versa). The loop bandwidth settings may then be stored (e.g., for future operation of the data communications system) at 820.

[0052] As mentioned above, in at least one embodiment, combinations of LBW configurations may be explored without regard to which PLL has the greater LBW. Figure 9 illustrates a method that involves iterative adjustment of the loop bandwidth configurations that may be used in such an embodiment. For purposes of discussion, it is assumed that both the T_X PLL and the R_X PLL have adjustable LBW configurations.

[0053] At 910, the loop bandwidth settings for the T_X PLL and the R_X PLL are set to maximum.

[0054] At 920, an initial signal quality metric (SQ_I) can be determined, in the manner indicated above. For example, the signal quality metric SQ_I may be chosen to be a percentage of the unit interval (UI%) of the data signal. Alternatively, as indicated earlier, the signal quality metric may correspond to any determinable characteristic of the data communication (e.g., the accuracy of the received data bits or the passing window size).

[0055] At 930, the LBW configuration of a selected PLL circuit is adjusted (e.g., lowered or reduced) relative to the other (i.e., the "non-selected") PLL circuit. For purposes of discussion, it is assumed that the selected PLL circuit corresponds to the PLL circuit determined to have the greater LBW configuration (e.g., as described with respect to Figures 6 and/or 7). However, in the case where only one of the PLL circuits has adjustable LBW settings, then that PLL circuit corresponds to the selected PLL, at 930 (i.e., irrespective of which PLL circuit is determined to have the greater LBW configuration).

[0056] The adjusted signal quality metric SQ_A is then determined at 940, post adjustment, for comparison with the initial signal quality metric SQ_I.

[0057] At 950, the adjusting signal quality metric SQ_A is compared with the initial signal quality metric SQ_I. For example, the signal quality metric SQ_A may be determined using any of the techniques described above (e.g., UI%, bit error, or passing window size).

[0058] If it is determined that the signal quality improves (e.g., $SQ_A > SQ_I$), then the current/adjusted LBW setting for the selected PLL circuit is updated (e.g., stored), at 960. As noted above, because a number of different signal quality metrics may be used, it may be necessary to understand how the selected metric varies with respect to changes in signal quality, in order to ensure proper interpretation of the data.

[0059] If the adjustment results in improved signal quality, the adjusted signal quality metric is set as the new initial signal quality metric (i.e., $SQ_I = SQ_A$), as indicated by reference numeral 970 (otherwise, the adjustment operation is terminated, at 980). In other words, threshold (or standard) for detecting an improvement in signal quality is reset, so that subsequent comparisons may use it as a reference in order to determine the optimal adjustment. The loop bandwidth adjustment operation then returns to step 930, where the LBW configuration of the selected PLL is once again adjusted. For example, the process of reducing the LBW setting of the selected PLL and comparing the adjusted signal quality metric SQ_A with the initial signal quality metric SQI (i.e., steps 930 – 970) may be repeated until it is determined that the signal quality does not improve after a subsequent adjustment of the LBW setting of the selected PLL circuit, at 950 (e.g., $SQ_A < SQ_I$, or $SQ_A \leq SQ_I$).

[0060] If desired, a limit may be placed on the number of iterations, such that the LBW configuration of the selected PLL may be adjusted for at most a predetermined number of iterations. For example, an alternative step may be provided, at 955, for determining whether the LBW setting of the selected PLL has been adjusted through a "maximum" (e.g., a user-

defined limit) number of iterations. If it is determined that, at 955, the specified limit of iterations has been reached, the operation subsequently terminates, at 980. In other words, the LBW settings are continually adjusted (and corresponding improvements in the signal quality monitored) for only so long as the maximum number of iterations has not been reached. This alternative step may be useful in providing a cap on the total amount of time and/or system resources allocated to performing a jitter error correction operation.

[0061] It was mentioned earlier that some implementations of the teachings described in this disclosure may perform in situ adjustment. In one embodiment, the iterative adjustment operation described with respect to Figure 9 is performed during an actual data communication session (i.e., while concurrently communicating data between the transmit circuit and the receive circuit). Thus, the stored LBW settings may be updated in near real-time (e.g., continuously, in regular intervals, or on-demand). The trigger to initiate this iterative adjustment operation may also, if desired, be invoked "on-demand." For example, the communications system may be configured to operate in a calibration or test mode (e.g., during startup, or on a periodic basis), or may be performed prior to runtime, using a set of preconfigured (or user-configured) test data.

[0062] Also, although the embodiment just described called for initially setting each LBW to maximum, at 910, this operation is arbitrary, i.e., the LBW's may be initially set to any level (e.g., minimum) with the system configured to cycle through possible combinations of LBW values to determine which combination yields the best signal metrics.

[0063] Figure 10 illustrates another implementation that relies upon iterative adjustment of the loop bandwidth configurations. For purposes of this discussion, it is again assumed that both the R_X PLL and the T_X PLL have programmable LBWs and that these PLLs have each been programmed to have their maximum LBW; the LBWs for the two PLLs may then be adjusted to reduce jitter by reducing the LBW of the PLL with the larger LBW downward, so as

to minimize the difference between the LBWs. As noted above, to perform this process, a designer may wish to implement a protocol for determining which PLL has the greater LBW and for responsively directing LBW adjustment efforts on that PLL.

[0064] At 1010, the loop bandwidth configuration for the non-selected PLL circuit (PLL_{NS}) is set to maximum.

[0065] At 1020, the loop bandwidth configuration for the selected PLL circuit (PLL_S) is also set to maximum. For purposes of discussion, it is assumed that the selected PLL circuit corresponds to the PLL circuit determined to have the greater loop bandwidth configuration (e.g., as described with respect to Figures 6 and 7). In other words, either the T_X PLL or the R_X PLL will be the selected PLL circuit, while the other of the T_X PLL or the R_X PLL will be the non-selected PLL circuit.

[0066] At 1030, an initial signal quality metric (SQ_I) is determined. For example, the signal quality metric SQ_A may be determined using any of the techniques described above (e.g., UI%, bit error, or passing window size).

[0067] Then at 1040, the LBW configuration of PLL_S is adjusted (e.g., lowered or reduced) relative to that of PLL_{NS}.

[0068] The adjusted signal quality metric SQ_A is then determined, at 1050, to be compared with the initial signal quality metric SQ_I.

[0069] At 1060, the adjusted signal quality metric SQA is compared with the initial signal quality metric SQ_I. As noted above, because a number of different signal quality metrics may be used, it may be necessary to understand how the selected metric varies with respect to changes in signal quality, in order to ensure proper interpretation of the data.

[0070] If at 1060 it is determined that the signal quality improves (e.g., $SQ_A > SQ_I$), then the current/adjusted LBW setting for PLL_S is updated (e.g., stored), at 1062.

[0071] At 1064, the adjusted signal quality metric is set as the new initial signal quality metric (i.e., $SQ_I = SQ_A$). In other words, the threshold (or standard) for detecting an improvement in signal quality during subsequent LBW iterations of PLL_S is reset.

[0072] The loop bandwidth adjustment operation then returns to step 1040, where the LBW configuration of PLL_S is once again adjusted. For example, the process of adjusting the LBW setting of PLL_S and comparing the adjusted signal quality metric SQ_A with the initial signal quality metric SQ_I (i.e., steps 1040 – 1064) may be repeated until it is determined that the signal quality does not improve after a subsequent adjustment of the LBW setting of PLL_S, at 1060 (e.g., $SQ_A < SQ_I$, or $SQ_A \leq SQ_I$).

[0073] If at 1060 it is determined that the signal quality does not improve (e.g. $SQ_A < SQ_I$, or $SQ_A \leq SQ_I$), then the current SQ_I value is stored, at 1070. This SQ_I value represents the best (or greatest) signal quality determined for a respective LBW iteration of PLL_{NS}.

[0074] At 1080, the LBW setting for PLL_{NS} circuit is adjusted (e.g., reduced or lowered). The loop bandwidth adjustment operation then returns to step 1020, where the LBW configuration of PLL_S is once again set to its maximum value. In this manner, a resulting signal quality metric from every possible pairing (or permutation) of LBW configurations for both PLL_S and PLL_{NS} is determined. The stored SQ_I values (associated with the best determined signal quality for every LBW iteration of PLL_{NS}) are then compared with one another to determine which pair of LBW settings (for both the T_X PLL and R_X PLL) yields the greatest overall improvement in signal quality.

[0075] Once all of the LBW settings for PLL_{NS} have been exhausted (e.g., the LBW setting for PLL_{NS} has cycled through its lowest adjustable value), the jitter error correction operation terminates, at block 1080. Finally, the T_X and/or R_X PLL may be programmed with the resulting LBW setting(s) determined at the end of the jitter adjustment operation in order to

mitigate the effects of sampling error, due to jitter, when communicating data through the overall signaling system.

[0076] As indicated above, an optional step may also be used (i.e., at 1061) to determine whether the LBW setting of PLL_S has been adjusted through a maximum number of iterations.

[0077] It should be noted that differences in signal propagation times (e.g., Data and Clk_{ref} signals) between the transmit circuit and the receive circuit may be accounted for (e.g., mitigated) by iteratively adjusting the LBW configuration of the selected PLL relative to changes in the signal quality metric (e.g., as described with respect to the embodiments of Figures 9 and 10).

[0078] Figure 11 illustrates yet another jitter correction operation that may be performed using the teachings provided by this disclosure.

[0079] At 1110, a determination is made as to whether the phase-locked loop circuits (i.e., the T_X PLL and R_X PLL) are programmable. Generally, in order to implement the jitter correction embodiments herein, at least one of the PLL circuits should have a programmable LBW. Thus, if it is determined that neither of the PLL circuits have programmable LBW settings, then the jitter correction operation comes to an end, at 1150. If at least one of the PLL circuits has a programmable LBW, then the jitter correction operation proceeds by generating Tclk and Rclk signals in the T_X and R_X PLLs, respectively, at 1120.

[0080] Step 1130 illustrates an optional intermediate step of determining the relative loop bandwidth configurations of the T_X and R_X PLL circuits. For example, the relative LBWs may be determined by looking up their respective values in a lookup table (e.g., as discussed above in reference to Figure 6). Alternatively, the relative LBWs may be determined using a test and evaluation process (e.g., as discussed above in reference to Figure 7).

[0081] At 1140, the T_X and R_X PLLs are programmed to have nominally matching LBWs. For example, this may be performed in a single programming step (e.g., as discussed above in

reference to Figure 8). Alternatively, this LBW adjustment may be performed via a series of iterative adjustments (e.g., as discussed above in reference to Figures 9 and 10).

[0082] As discussed above, in connection with Figure 9, the jitter correction operation may either be performed offline or during system runtime (e.g., via a startup operation, a periodic calibration operation, or on-demand). The jitter error correction operations described herein may also, depending on embodiment, be carried out manually (e.g., by a user of the communications system and/or devices) or automatically (e.g., by a processor and/or controller). In other words, a user may monitor a transmitted data signal for changes in a signal quality metric and manually adjust the loop bandwidth configurations (e.g., program the control register) for either of the PLL circuits. Alternatively (or in addition), a system controller connected to the data communications system may detect changes in quality of the transmitted data signal and automatically adjust the loop bandwidth configurations of either of the PLL circuits.

[0083] Figure 12 illustrates an exemplary data communications system having control circuitry to control jitter correction. The data communications system 1200 includes a signal source 1210, a transmit device 1220, and a receive device 1230. For purposes of illustration, the signal source 1210 is shown separate from both the transmit device 1220 and the receive device 1230, e.g., as separate integrated circuits. However, depending on implementation, the signal source 1210 may be embodied within a common device, such as a transmit device 1220, configured as a single integrated circuit.

[0084] The transmit device 1220 includes transmit (T_X) circuitry 1222 and control circuitry 1224. The T_X circuitry 1222 is provided for transmitting data to the receive device 1230 based, at least in part, on times indicated by the reference timing (Clk_{ref}) signal. For example, the transmit circuitry 1222 may correspond to the T_X circuit as shown in Figures 4 and 5. The T_X circuitry 1222 may initially receive the data from another device (e.g., external to the transmit

device 1222, e.g., in a situation where, for example, the transmit device is part of a memory controller). Alternatively, the data to be transmitted may reside within the transmit device 1220 (e.g., within a local memory or register). The control circuitry 1224, which may be implemented by a state machine, microcontroller or any other processor or control logic (e.g., implemented in software and/or hardware), connects to the T_X circuitry 1222 via control bus Ctrl_A.

[0085] The receive device 1230 includes receive (R_X) circuitry 1232 to receive the data signal based, at least in part, on times indicated by the Clk_{ref} signal. Although not shown, for purposes of simplicity, the receive device 1230 may include additional circuitry for storing and/or processing the received data. For example, the receive device 1230 may correspond to a memory device, and may thus include one or more registers for storing the received data. The control circuitry 1224 further connects to the R_X circuitry via control bus Ctrl_B.

[0086] Implementing the principles discussed above, the control circuitry 1224 may effect a jitter correction operation for the data communications system 1200. For example, the control circuitry 1224 may adjust the LBW setting of one or more PLL circuits within the T_X circuitry 1222 and/or the R_X circuitry 1232 (e.g., via Ctrl_A and Ctrl_B, respectively) using the steps discussed above. Additionally, the control circuitry 1224 may receive information from the T_X circuitry 1222 and/or the R_X circuitry 1232 for purposes of determining how the corresponding PLLs should be adjusted. For example, the information gathered from the T_X circuitry 1222 and/or the R_X circuitry 1232 may be used to determine a signal quality metric (e.g., UI%, accuracy of data transmission, or passing window size). The control circuitry 1224 may also control (i.e., coupled and/or decouple) an additional signal source (not shown, for simplicity) to inject additional noise into the Clk_{ref} signal for calibration and/or testing purposes

[0087] Figure 13 illustrates an exemplary data communications system having control circuitry to control jitter correction across multiple data communications devices. The

communications system 1300 is made up of a single transmit (T_X) device 1320 connected to a series of receive (R_X) devices 1330_1-1330_N. The T_X device 1320 may, in one embodiment, be implemented as part of a memory controller, and the R_X devices 1330_1-1330_N may be implemented as a number of memory devices. Importantly, although these elements are labeled T_X and R_X, respectively, it should be understood that the data path may be unidirectional or bidirectional, and that more than one communication path may be used (e.g., two unidirectional paths). Only one data bus is illustrated in FIG. 13, for purposes of simplifying discussion, but it should be understood that the principles discussed herein may be applied to each such bus.

[0088] The T_X device transmits data to the R_X devices 1330_1-1330_N, along with a corresponding reference timing (Clk_{ref}) signal. Depending on implementation, one or more of the receive devices 1330_1-1330_N may receive (i.e., sample) the transmitted data based, at least in part, on times indicated by the Clk_{ref} signal. The T_X device 1320 may then collect information regarding the quality of the transmitted data signal (e.g., via the control bus Ctrl) and adjust the PLLs of at least one of the T_X device 1320 and/or R_X devices 1330_1-1330_N, accordingly.

[0089] Each of the R_X devices 1330_1-1330_N may include different PLL circuitry for timing reception of the transmitted data. In other words, the LBW settings for each of the R_X 1330_1-1330_N is likely to vary from device to device. Furthermore, some of the R_X devices 1330_1-1330_N may have adjustable PLLs, while others do not. Thus, it may be difficult (if not impossible) to determine a LBW setting which yields the best possible adjusted signal quality metric for data communications between the T_X device 1320 and each of the R_X devices 1330_1-1330_N.

[0090] Applying the principles discussed earlier, the control circuitry within the T_X device adjusts the PLLs of at least one of the T_X device 1320 and/or R_X devices 1330_1-1330_N to reduce the difference between their LBW settings, matching them as closely as practical or possible.

A number of methods may be employed in order to select LBW settings for each device. For example, in one embodiment, the system may select a LBW for the T_X device 1320, and select an "ideal value" for the R_X devices 1330_1-1330_N, setting their LBWs as closely as possible to this ideal value. Alternatively, the LBW for each R_X device may be individually programmed, or some other metric or algorithm may be used to set the R_X LBWs. For example, as mentioned earlier, the control circuitry may cycle through every possible combination of LBW settings for the T_X device 1320 and each R_X device 1330_1-1330_N, one by one. The control circuitry may perform a statistical analysis on the results of the jitter correction operations to determine individual LBW settings for each T_X and R_X device pair which yields the greatest overall improvement in signal quality metric among the devices (e.g., based on the mean, median, or mode of the signal quality metrics). Other methodologies will also no doubt occur to those having skill in digital systems design.

[0091] It should be noted that the various integrated circuits, dice and packages disclosed herein may be described using computer aided design tools and expressed (or represented), as data and/or instructions embodied in various computer-readable media, in terms of their behavioral, register transfer, logic component, transistor layout geometries, and/or other characteristics.

[0092] When received within a computer system via one or more computer-readable media, such data and/or instruction-based expressions of the above described circuits may be processed by a processing entity (e.g., one or more processors) within the computer system in conjunction with execution of one or more other computer programs including, without limitation, net-list generation programs, place and route programs and the like, to generate a representation or image of a physical manifestation of such circuits. Such representation or image may thereafter be used in device fabrication, for example, by enabling generation of one or more masks that are used to form various components of the circuits in a device fabrication process.

[0093] In the foregoing description and in the accompanying drawings, specific terminology and drawing symbols have been set forth to provide a thorough understanding of the present invention. In some instances, the terminology and symbols may imply specific details that are not required to practice the invention. For example, any of the specific numbers of bits, signal path widths, signaling or operating frequencies, component circuits or devices and the like may be different from those described above in alternative embodiments. In other instances, well-known circuits and devices are shown in block diagram form to avoid obscuring the present invention unnecessarily. Additionally, the interconnection between circuit elements or blocks may be shown as buses or as single signal lines. Each of the buses may alternatively be a single signal line, and each of the single signal lines may alternatively be buses. Signals and signaling paths shown or described as being single-ended may also be differential, and vice-versa. A signal driving circuit is said to "output" a signal to a signal receiving circuit when the signal driving circuit asserts (or deasserts, if explicitly stated or indicated by context) the signal on a signal line coupled between the signal driving and signal receiving circuits. The term "coupled" is used herein to express a direct connection as well as a connection through one or more intervening circuits or structures. Integrated circuit device "programming" may include, for example and without limitation, loading a control value into a register or other storage circuit within the device in response to a host instruction and thus controlling an operational aspect of the device, establishing a device configuration or controlling an operational aspect of the device through a one-time programming operation (e.g., blowing fuses within a configuration circuit during device production), and/or connecting one or more selected pins or other contact structures of the device to reference voltage lines (also referred to as strapping) to establish a particular device configuration or operation aspect of the device. The terms "exemplary" and "embodiment" are used to express an example, not a preference or requirement.

[0094] While the invention has been described with reference to specific embodiments thereof, it will be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope. For example, features or aspects of any of the embodiments may be applied, at least where practicable, in combination with any other of the embodiments or in place of counterpart features or aspects thereof. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A method of operation within a signaling system, the method comprising:
generating, in a first locked-loop circuit, a first timing signal to time transmission of a data signal in a transmit circuit;
generating, in a second locked-loop circuit, a second timing signal to time sampling of the data signal in a receive circuit; and
adjusting a loop bandwidth configuration of at least one of the first or second locked-loop circuits to reduce a difference between loop bandwidths of the first and second locked-loop circuits.

2. The method of claim 1 wherein adjusting a loop bandwidth configuration comprises adjusting the loop bandwidth configuration to align a phase of the first timing signal with a phase of the second timing signal.

3. The method of claim 1 wherein adjusting a loop bandwidth configuration of one of the first or second locked-loop circuits comprises:
determining which of the first or second locked-loop circuits has a greater loop bandwidth configuration; and
programming the loop bandwidth of a selected one of the first or second locked-loop circuits to be substantially equal to the loop bandwidth of the other locked-loop circuit based, at least in part, on the determination.

4. The method of claim 3 wherein programming the loop bandwidth of a selected one of the first or second locked-loop circuits comprises:

setting the loop bandwidth configuration of each of the first and second locked-loop circuits to a maximum value; and

adjusting the loop bandwidth configuration of the locked-loop circuit having the greater loop bandwidth configuration to be substantially equal to the loop bandwidth configuration of the other locked-loop circuit.

5. The method of claim 3 wherein determining which of the first or second locked-loop circuits has the greater loop bandwidth configuration comprises referencing the loop bandwidth configuration of at least one of the first or second locked-loop circuits in a table.

6. The method of claim 3 wherein determining which of the first or second locked-loop circuits has the greater loop bandwidth configuration comprises:

adjusting the loop bandwidth configuration of one of the first or second locked-loop circuits; and

identifying a change in a quality of the data signal in response to the adjustment.

7. The method of claim 3 wherein the selected locked-loop circuit corresponds to the locked-loop circuit having the lesser loop bandwidth configuration.

8. The method of claim 6 wherein the quality of the data signal is determined based on at least one of: (i) a unit interval percentage, (ii) a sampling bit error, or (iii) a passing window size of the data signal.

9. The method of claim 6 wherein selecting one of the first or second locked-loop circuits comprises selecting the locked-loop circuit which, upon adjusting the loop bandwidth configuration, yields an improvement in the quality of the data signal.

10. The method of claim 1 wherein adjusting a loop bandwidth configuration of at least one of the first or second locked-loop circuits comprises:

monitoring the data signal for changes in quality; and

iteratively adjusting the loop bandwidth configuration of one of the first or second locked-loop circuits based on the changes in quality of the data signal.

11. The method of claim 10 wherein iteratively adjusting the loop bandwidth configuration of one of the first or second locked-loop circuits comprises iteratively adjusting the loop bandwidth configuration of one of the first or second locked-loop circuits until no improvement is detected in the quality of the data signal.

12. The method of claim 10 wherein iteratively adjusting the loop bandwidth configuration of one of the first or second locked-loop circuits comprises iteratively adjusting the loop bandwidth configuration of one of the first or second locked-loop circuits for a predetermined number of iterations.

13. The method of claim 1 further comprising, prior to adjusting a loop bandwidth configuration of at least one of the first or second locked-loop, introducing a substantial amount of jitter into at least one of the first timing signal or the second timing signal.

14. A method of operation within a signaling system, the method comprising:

generating, in a first locked-loop circuit, a first timing signal to time transmission of a data signal in a transmit circuit;

generating, in a second locked-loop circuit, a second timing signal to time sampling of the data signal in a receive circuit; and

adjusting a first loop bandwidth configuration of at least one of the first or second locked-

loop circuits to reduce jitter associated with the timing of sampling of the data of the first and second timing signals.

15. The method of claim 14 wherein adjusting a first loop bandwidth comprises adjusting the first loop bandwidth configuration such that a phase of the first timing signal is substantially aligned with a phase of the second timing signal.

16. The method of claim 14 wherein adjusting the first loop bandwidth configuration of at least one of the first or second locked-loop circuits comprises:

determining the difference between loop bandwidths of the first and second locked-loop circuits; and

programming the loop bandwidth of a selected one of the first or second locked-loop circuits in dependence upon the loop bandwidth of the other locked-loop circuit.

17. The method of claim 16 wherein determining the difference between loop bandwidths of the first and second locked-loop circuits comprises:

referencing the loop bandwidth of at least one of the first or second locked-loop circuits in a lookup table; and

selecting one of the first or second locked-loop circuits based, at least in part, on the loop bandwidth of the at least one of the first or second locked-loop circuits, respectively.

18. The method of claim 17 wherein determining the difference between loop bandwidths of the first or second locked-loop circuits comprises:

adjusting a second loop bandwidth configuration of at least one of the first or second locked-loop circuits;

identifying a change in a quality of the data signal in response to adjusting the loop

bandwidth of at least one of the first or second locked-loop circuits; and

selecting one of the first or second locked-loop circuits based, at least in part, on the change in the signal quality.

19. The method of claim 18 wherein selecting one of the first or second locked-loop circuits comprises selecting the locked-loop circuit which, upon adjusting the second loop bandwidth configuration, yields an improvement in the quality of the data signal.

20. The method of claim 16 wherein programming the loop bandwidth of a selected one of the first or second locked-loop circuits in dependence upon the a loop bandwidth of the other locked-loop circuit comprises:

monitoring the data signal for changes in quality; and

iteratively reducing the loop bandwidth of the selected locked-loop circuit based on identified changes in the quality of the data signals.

21. The method of claim 14 wherein: adjusting includes programming the loop bandwidth of at least one of the locked-loop circuits, to substantially match their loop bandwidths; and storing at least one of the adjusted loop bandwidth configurations in an integrated circuit device.

22. The method of claim 14 further comprising, prior to adjusting a loop bandwidth configuration of at least one of the first or second locked-loop, introducing additional jitter into at least one of the first timing signal or the second timing signal.

23. A controller for operating a data communications system, the controller comprising measurement circuitry for determining a difference between loop bandwidths of a first locked-loop circuit and a second locked-loop circuit; and

programming circuitry for adjusting a loop bandwidth configuration of at least one of the first and second locked-loop circuits to reduce the difference between loop bandwidths of the first and second locked-loop circuits.

24. The controller of claims 23 further comprising measurement circuitry for determining a difference between loop bandwidths of a first locked-loop circuit and loop bandwidths of a second locked loop circuitry.

25. The controller of claim 23 further comprising a storage circuit for storing the one or more adjusted loop bandwidth configurations.

26. The controller of claim 23, wherein the programming circuitry is adapted to program a locked-loop circuit in an integrated circuit device distinct from the controller.

27. A data communications system comprising:

a first locked-loop circuit for generating a first timing signal to time transmission of a data signal;

a second locked-loop circuit for generating a second timing signal to time sampling of the data; and

a controller for adjusting a loop bandwidth configuration of at least one of the first and second locked-loop circuits to reduce a difference between loop bandwidths of the first and second locked-loop circuits.

28. The system of claim 27, wherein the controller comprises circuitry to determine the difference between loop bandwidths of the first and second locked-loop circuits.

29. The system of claim 27 further comprising a storage circuit for storing the one or more

adjusted loop bandwidths.

30. The system of claim 27 wherein at least one of the first and second locked-loop circuits is programmable.

31. The system of claim 27 further comprising a signal source coupled to provide a third timing signal to at least one of the first or second locked-loop circuits, wherein the third timing signal includes a substantial amount of jitter.

32. A data communications system comprising:

means for generating, in a first locked-loop circuit, a first timing signal to time transmission of a data signal in a transmit circuit;

means for generating, in a second locked-loop circuit, a second timing signal to time sampling of the data signal in a receive circuit; and

means for adjusting a loop bandwidth configuration of at least one of the first or second locked-loop circuits to reduce a difference between loop bandwidths of the first and second locked-loop circuits.

33. Computer-readable media having information embodied therein that includes a description of an integrated-circuit memory device, the information including descriptions of:

a first locked-loop circuit for generating a first timing signal to time transmission of a data signal;

a second locked-loop circuit for generating a second timing signal to time sampling of the data; and

a controller for adjusting a loop bandwidth configuration of at least one of the first or second locked-loop circuits to reduce a difference between loop bandwidths of the

9 first and second locked-loop circuits.

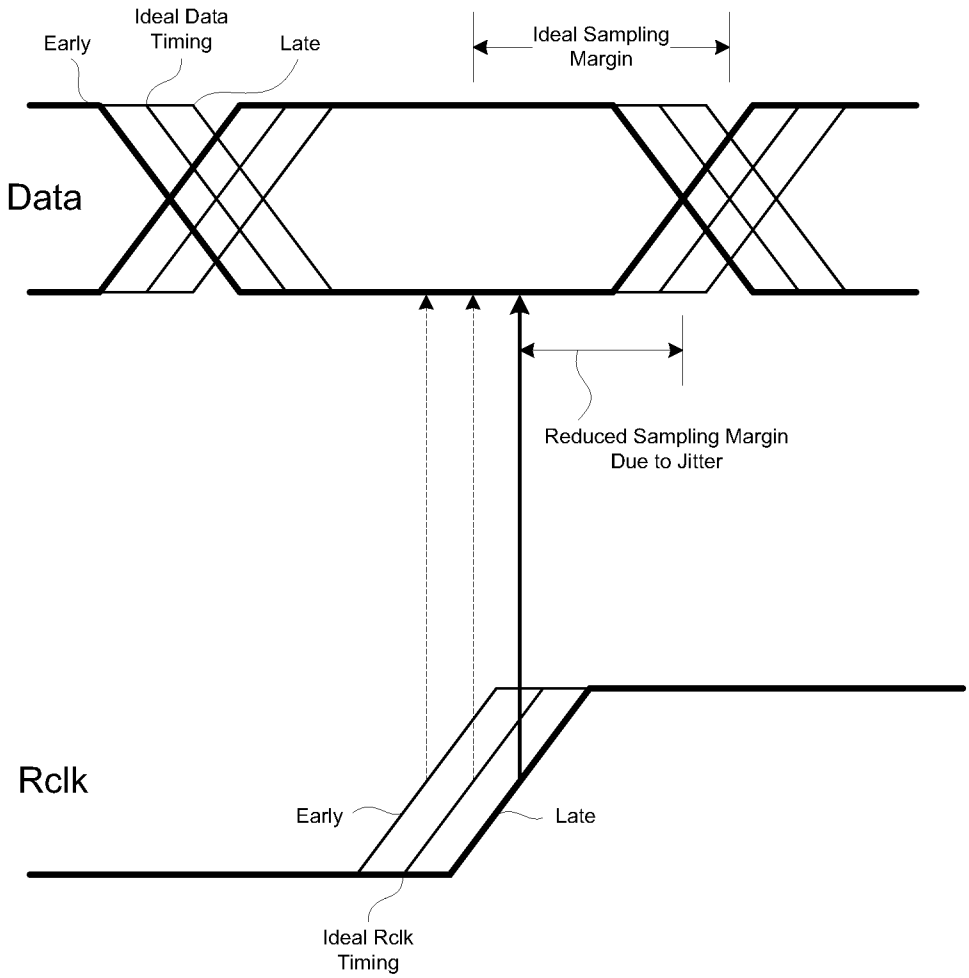
Early
Ideal Data Timing
Late
Ideal Sampling Margin
Data
Reduced Sampling Margin Due to Jitter
Rclk
Early
Late
Ideal Rclk Timing

FIG. 1

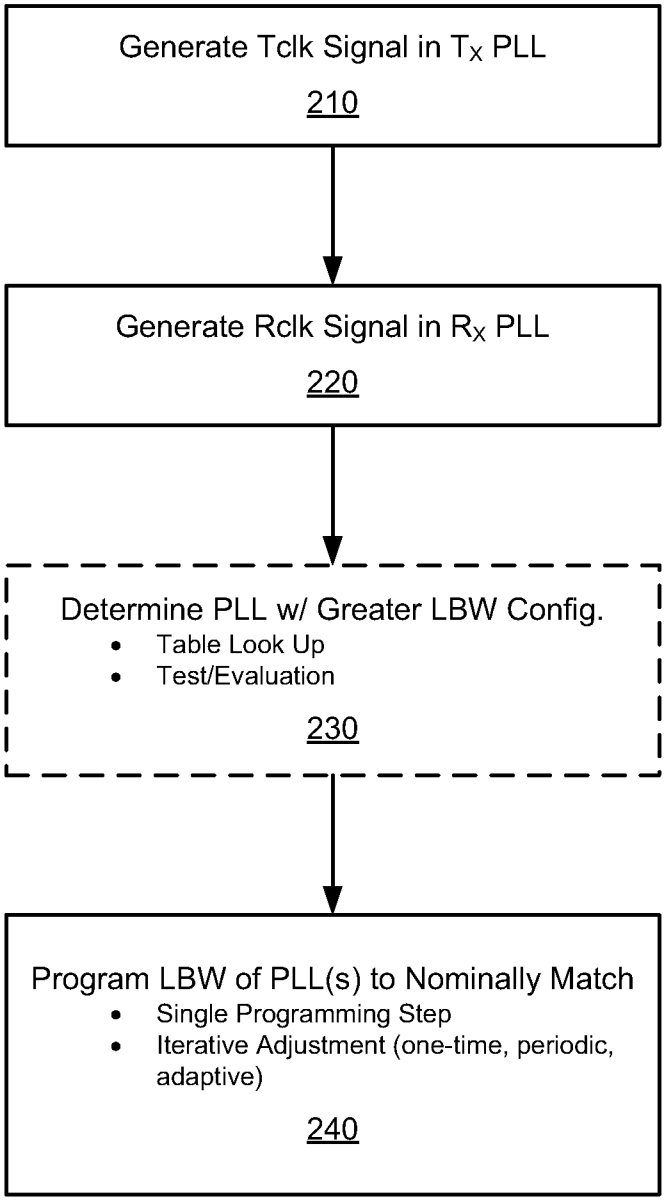
Generate Tclk Signal in T_X PLL
210
Generate Rclk Signal in R_X PLL
220
Determine PLL w/ Greater LBW Config.
• Table Look Up
• Test/Evaluation
230
Program LBW of PLL(s) to Nominally Match
• Single Programming Step
• Iterative Adjustment (one-time, periodic, adaptive)
240

FIG. 2

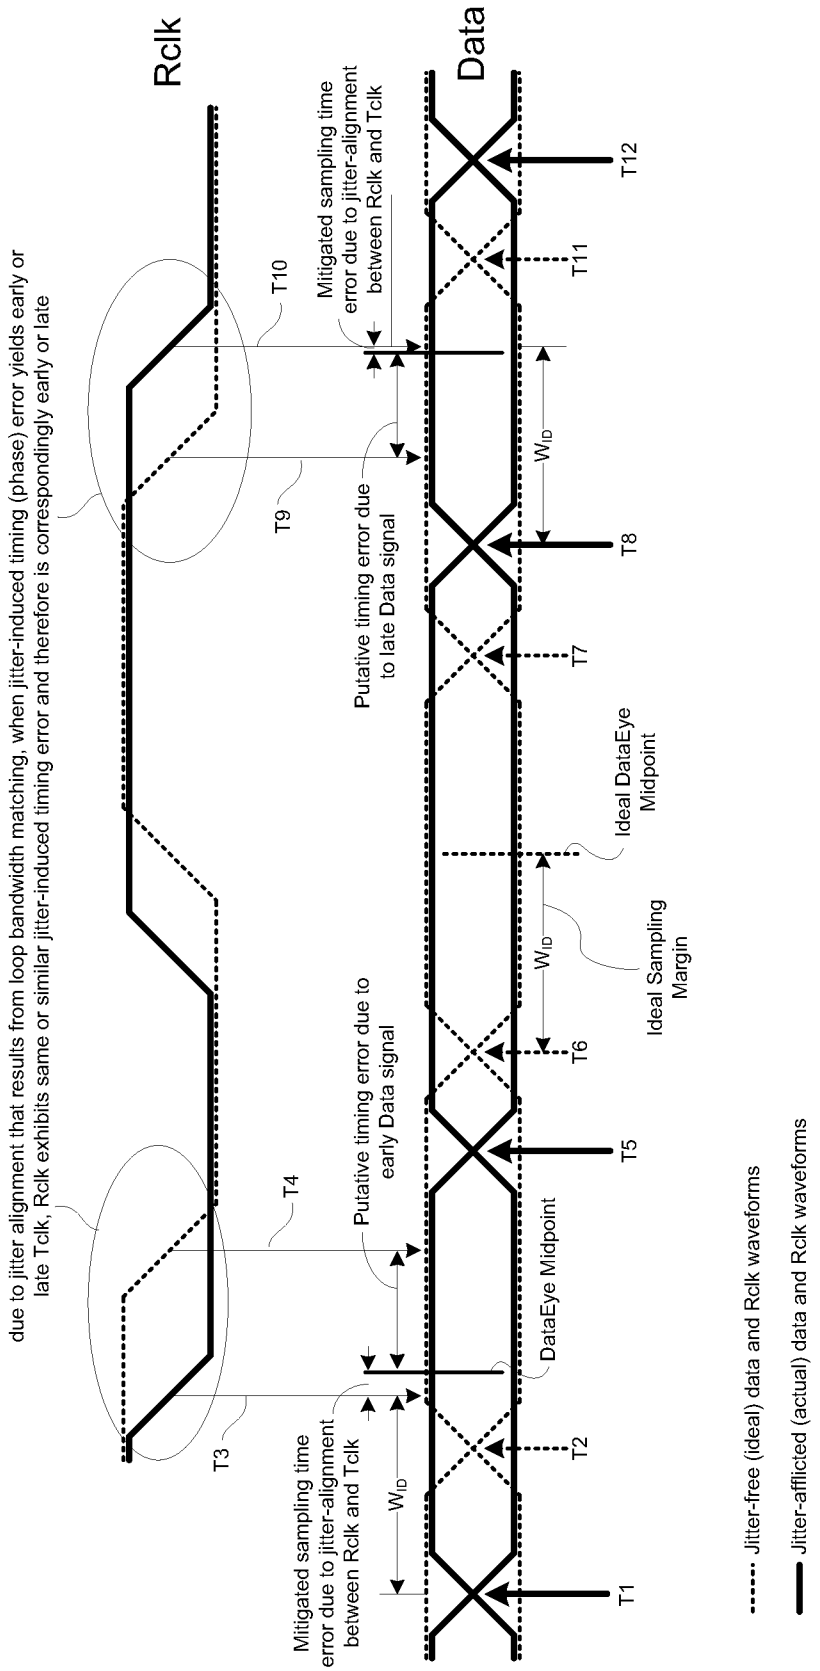
Rclk
Data
due to jitter alignment that results from loop bandwidth matching, when jitter-induced timing (phase) error yields early or late Tclk, Rclk exhibits same or similar jitter-induced timing error and therefore is correspondingly early or late
Mitigated sampling time error due to jitter-alignment between Rclk and Tclk
Putative timing error due to early Data signal
DataEye Midpoint
Putative timing error due to late Data signal
Mitigated sampling time error due to jitter-alignment between Rclk and Tclk
Ideal DataEye Midpoint
Ideal Sampling Margin
W_{ID}
T1
T2
T3
T4
T5
T6
T7
T8
T9
T10
T11
T12
Jitter-free (ideal) data and Rclk waveforms
Jitter-afflicted (actual) data and Rclk waveforms

FIG. 3

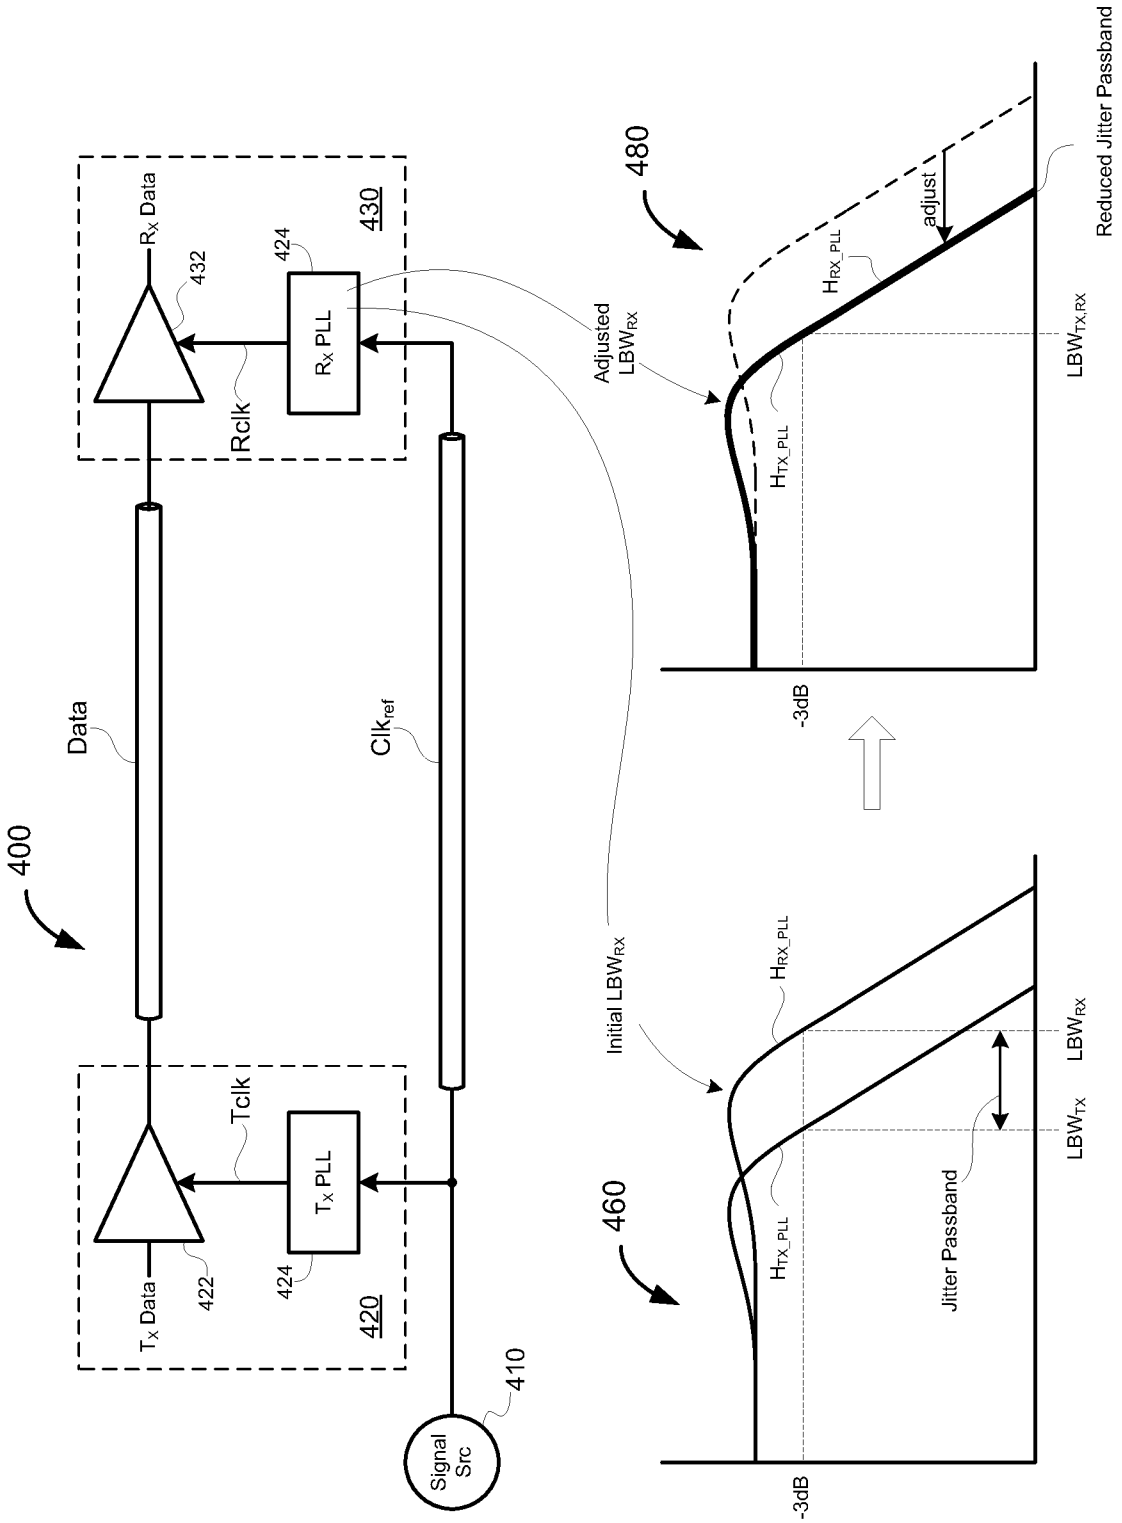
400
Data
Clk$_{ref}$
Rclk
Tclk
R_X Data
T_X Data
R_X PLL
T_X PLL
430
420
424
432
422
410
Signal Src
480
460
Adjusted LBW$_{RX}$
Initial LBW$_{RX}$
H$_{RX_PLL}$
H$_{TX_PLL}$
adjust
Reduced Jitter Passband
LBW$_{TX,RX}$
LBW$_{RX}$
LBW$_{TX}$
Jitter Passband
-3dB

FIG. 4

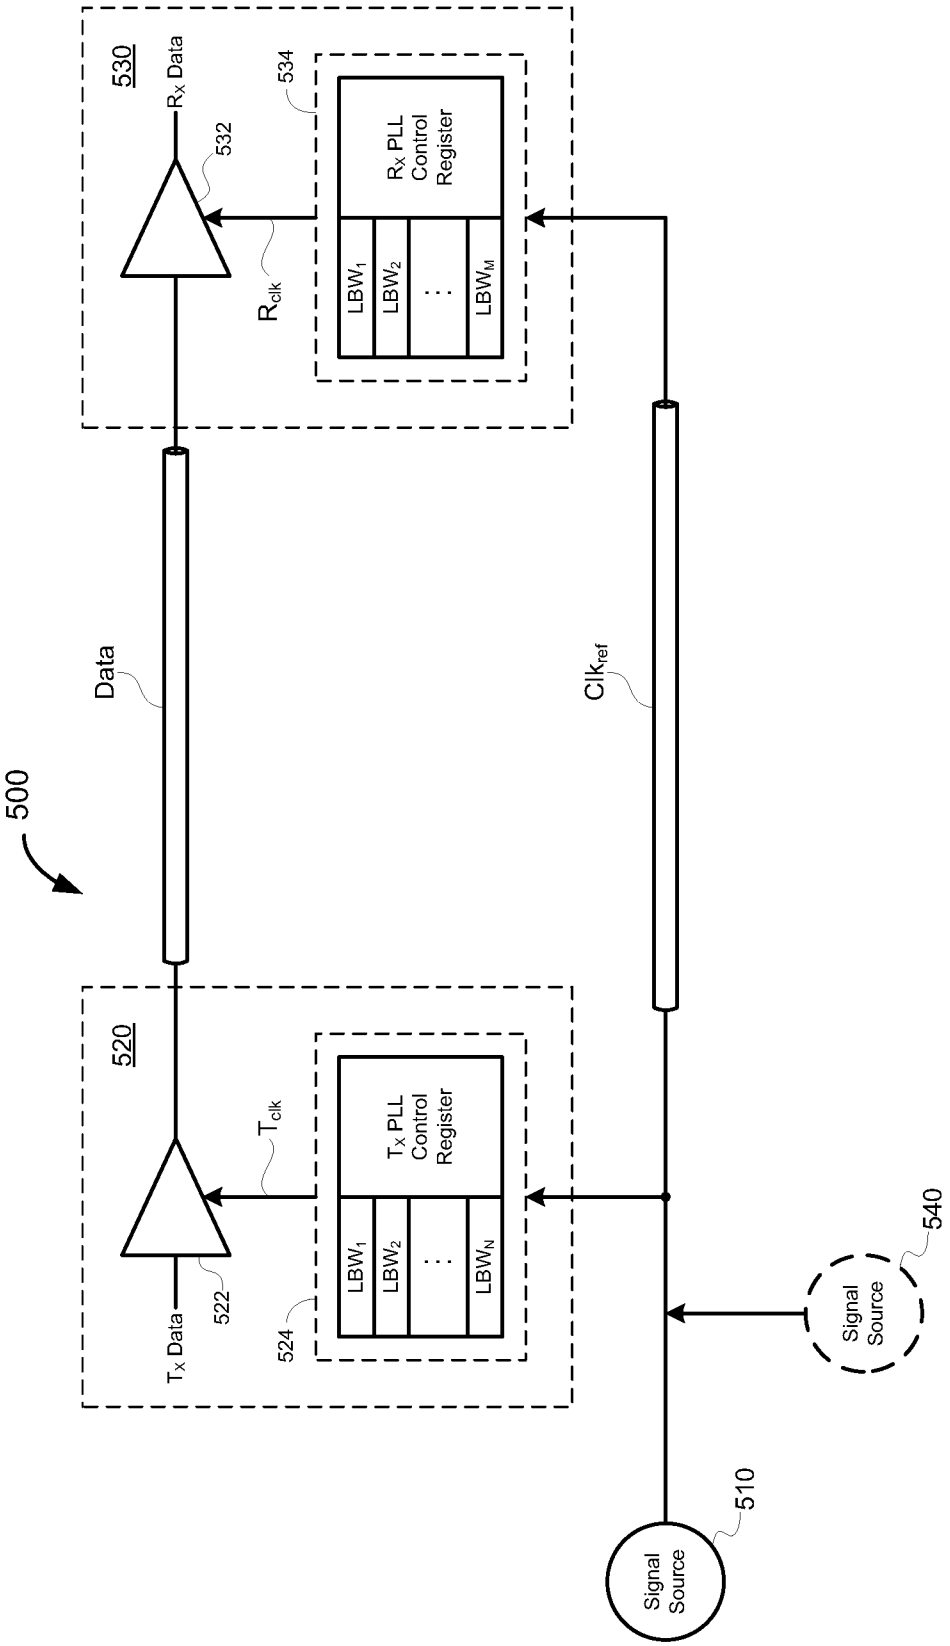
500
530
R_X Data
532
534
R_X PLL Control Register
LBW_1
LBW_2
LBW_M
R_{clk}
Data
Clk_{ref}
520
T_{clk}
T_X PLL Control Register
LBW_1
LBW_2
LBW_N
T_X Data
522
524
Signal Source
540
Signal Source
510

FIG. 5

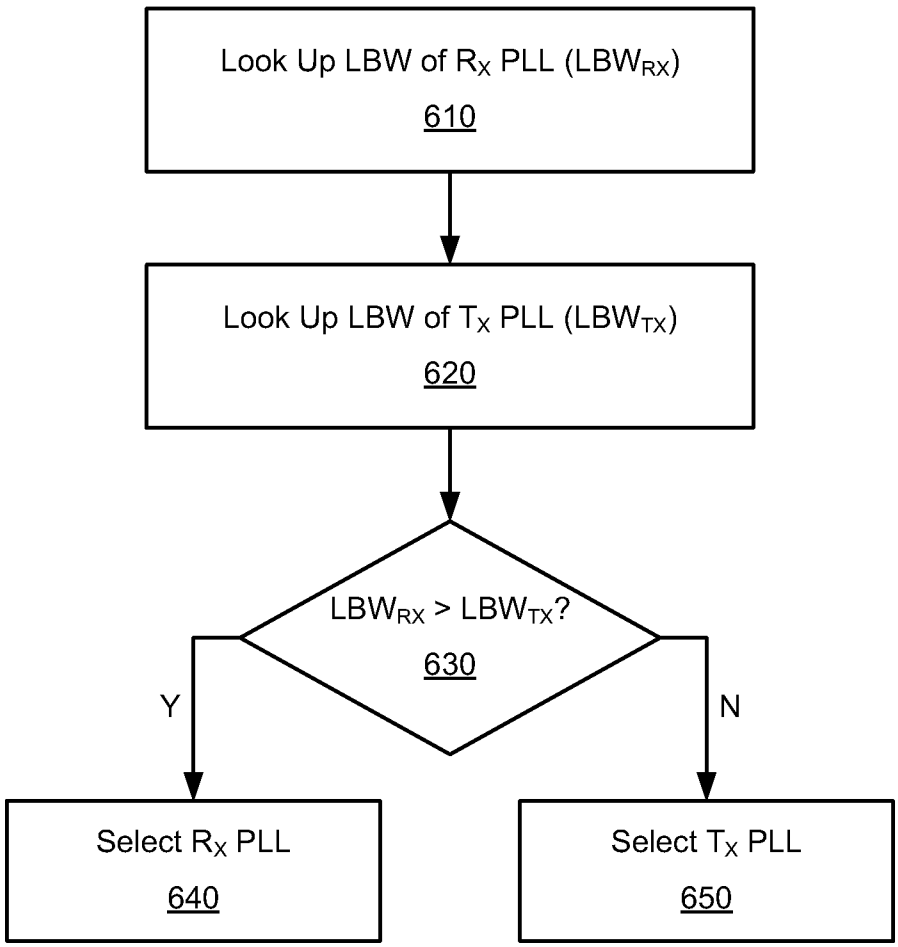
Look Up LBW of R_X PLL (LBW_{RX})
610
Look Up LBW of T_X PLL (LBW_{TX})
620
$LBW_{RX} > LBW_{TX}$?
630
Y
N
Select R_X PLL
640
Select T_X PLL
650

FIG. 6

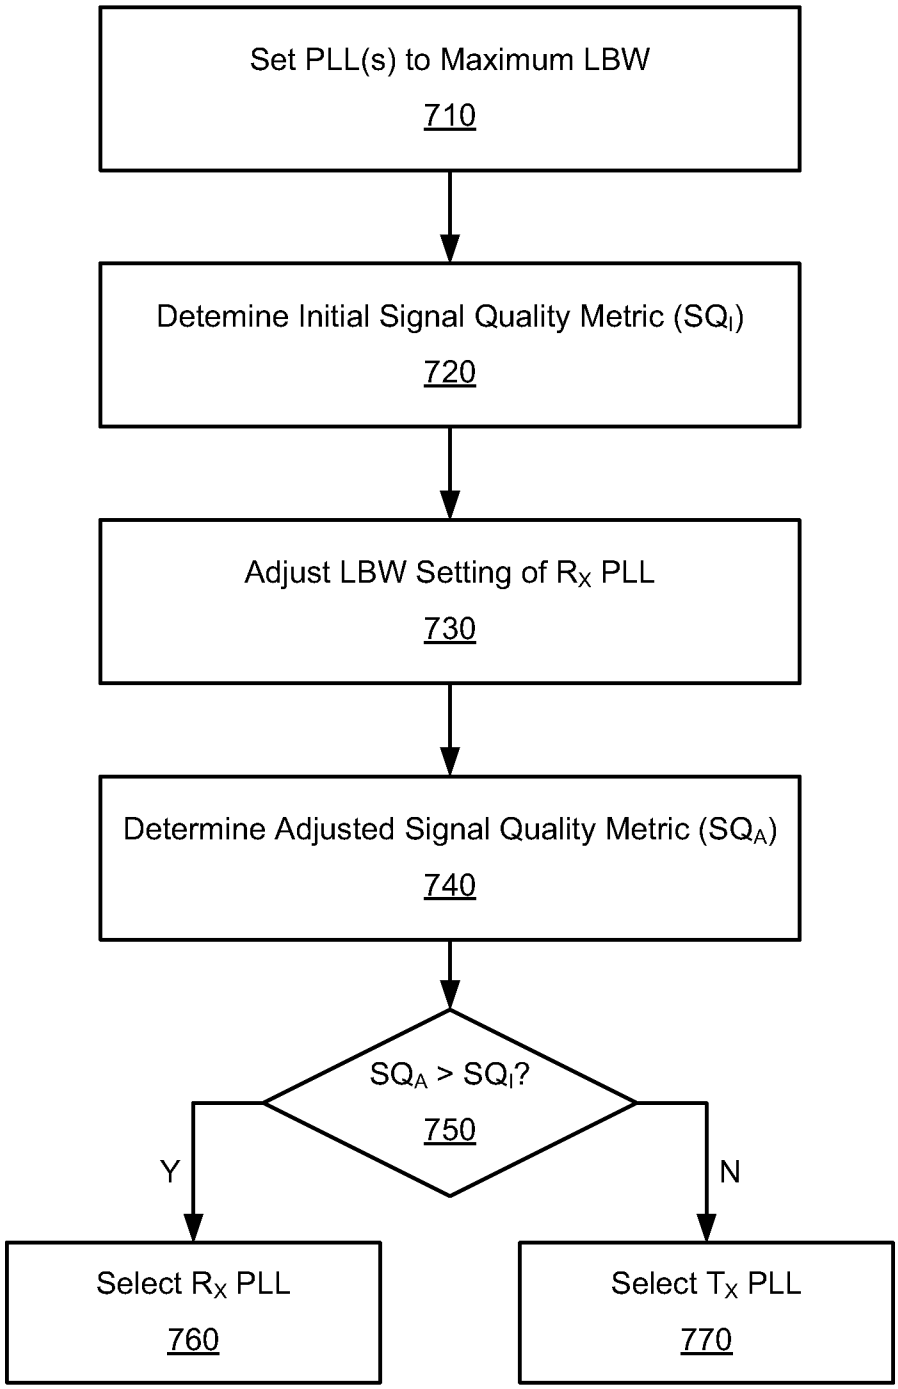
Set PLL(s) to Maximum LBW
710
Detemine Initial Signal Quality Metric (SQ_I)
720
Adjust LBW Setting of R_X PLL
730
Determine Adjusted Signal Quality Metric (SQ_A)
740
$SQ_A > SQ_I$?
750
Y
N
Select R_X PLL
760
Select T_X PLL
770

FIG. 7

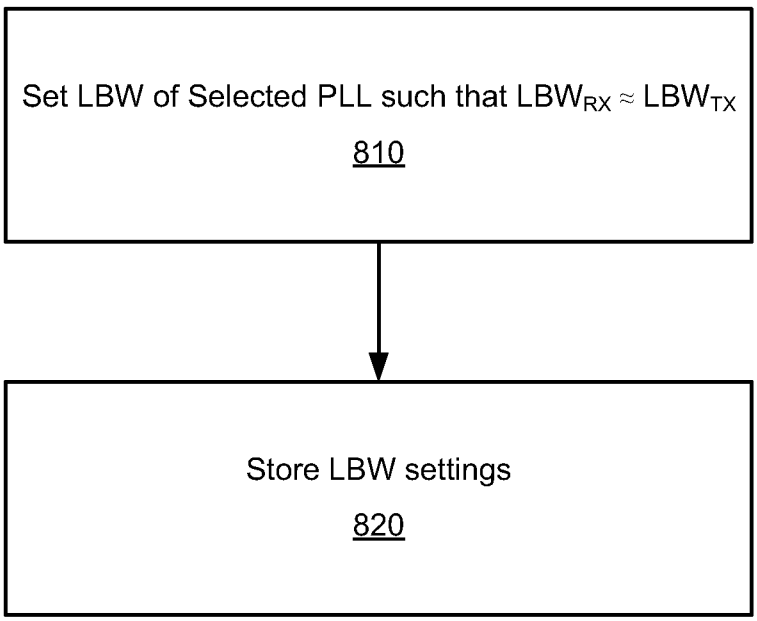
Set LBW of Selected PLL such that $LBW_{RX} \approx LBW_{TX}$
810
Store LBW settings
820

FIG. 8

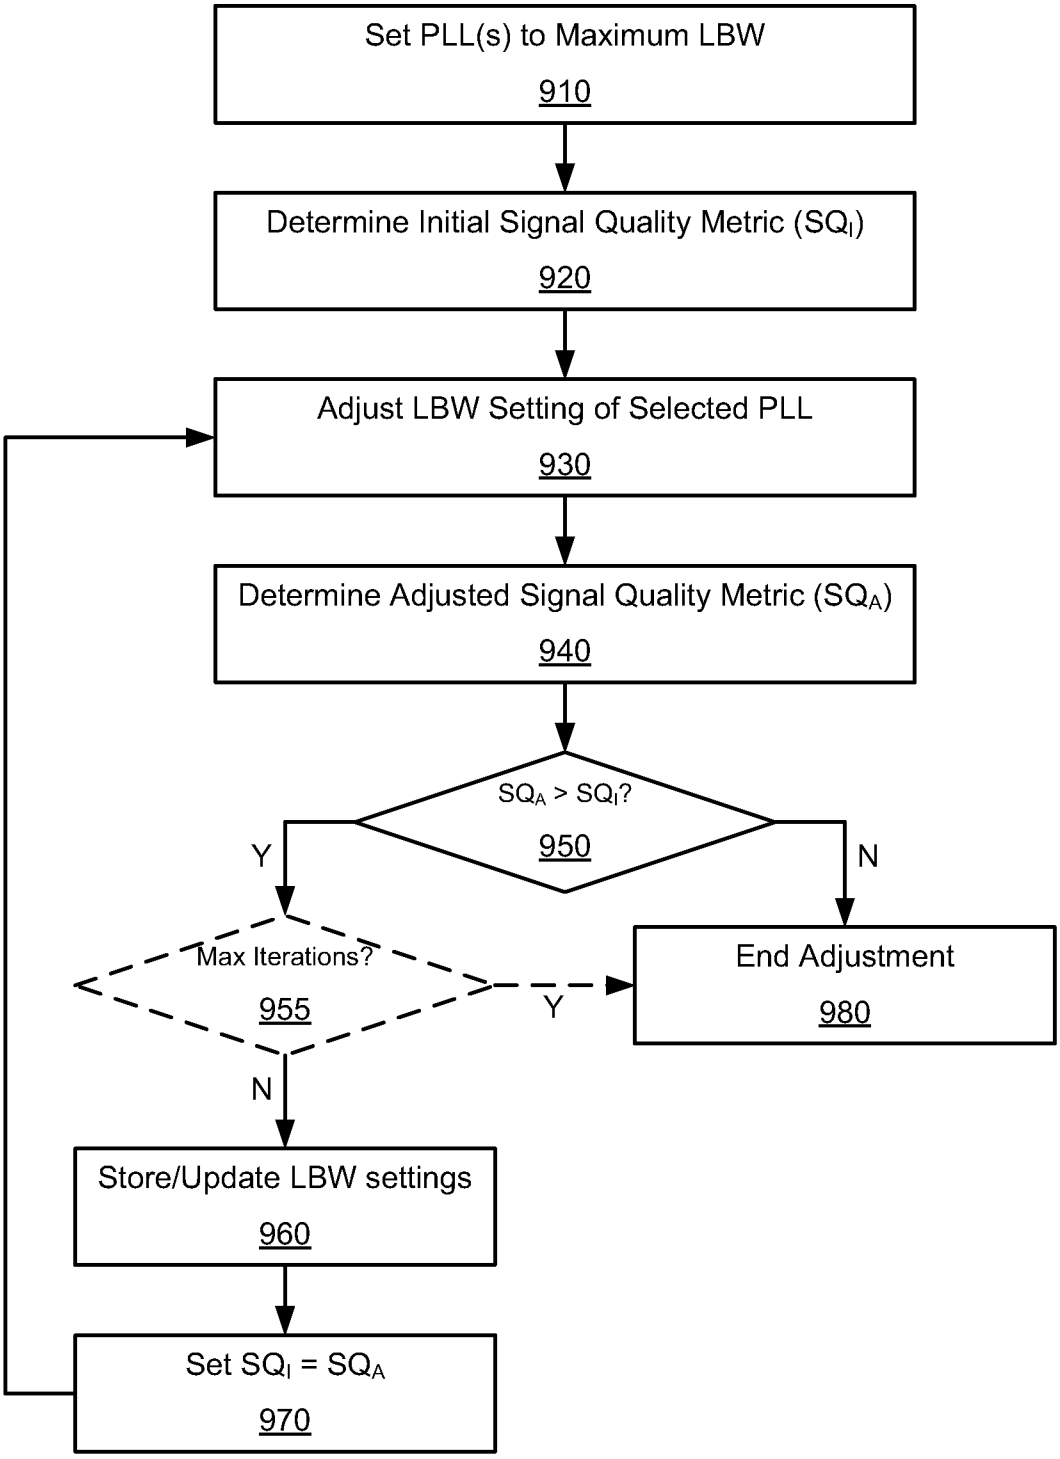
Set PLL(s) to Maximum LBW
910
Determine Initial Signal Quality Metric (SQ_I)
920
Adjust LBW Setting of Selected PLL
930
Determine Adjusted Signal Quality Metric (SQ_A)
940
$SQ_A > SQ_I$?
950
Y
N
Max Iterations?
955
Y
End Adjustment
980
N
Store/Update LBW settings
960
Set $SQ_I = SQ_A$
970

FIG. 9

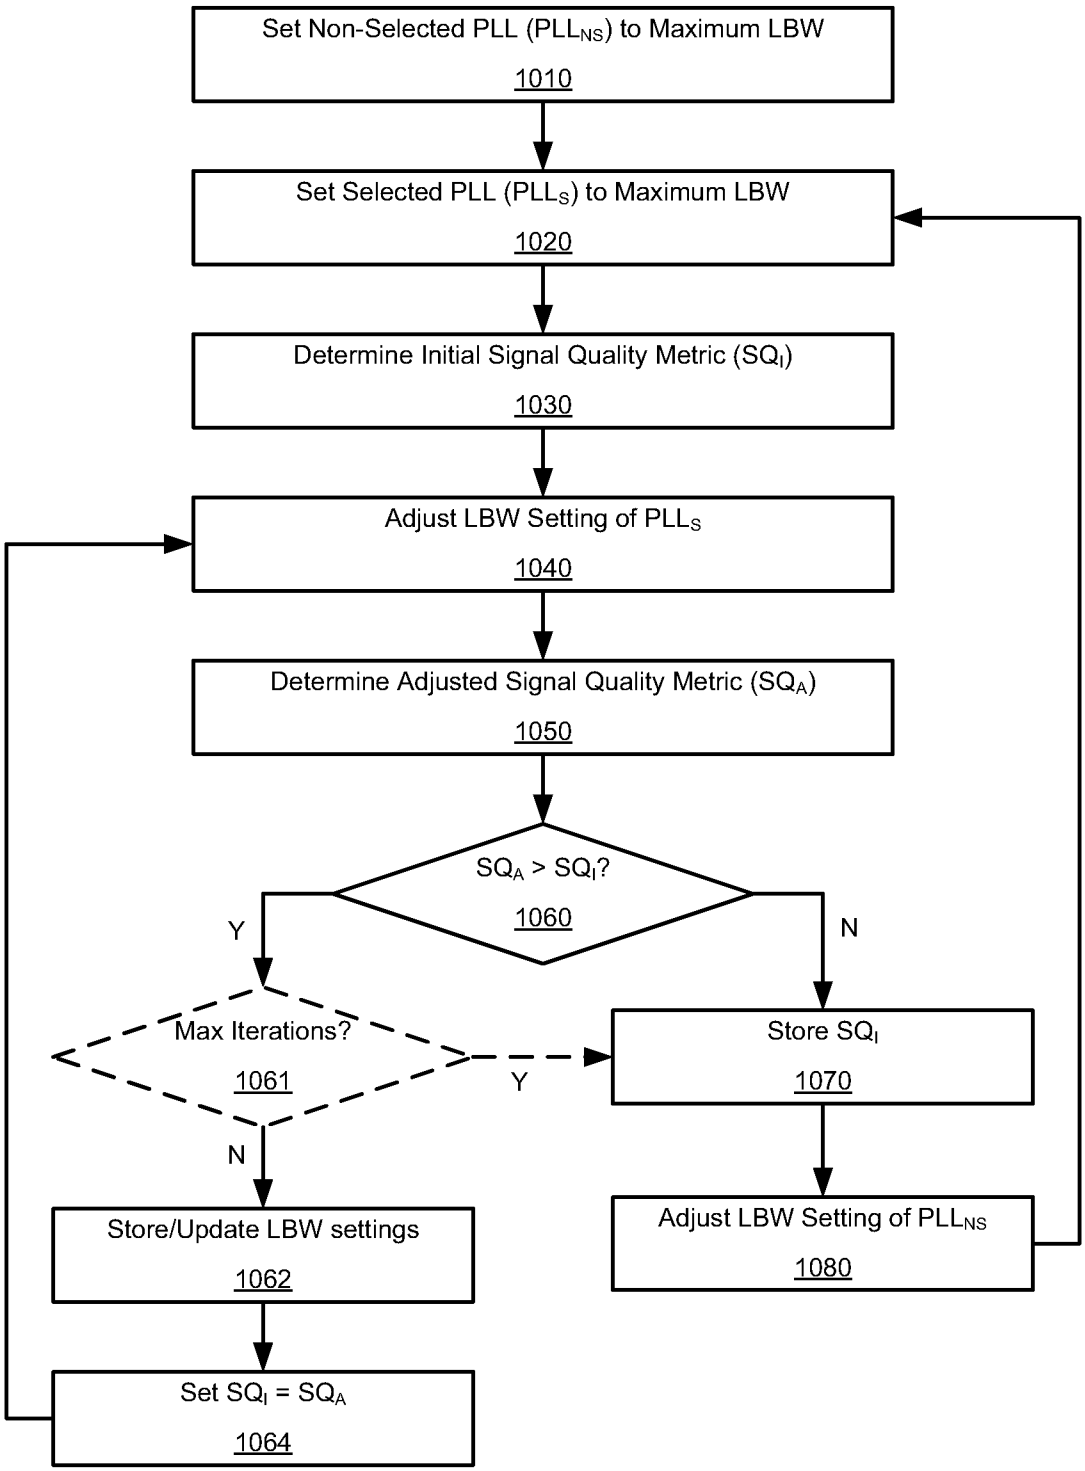
Set Non-Selected PLL (PLL_{NS}) to Maximum LBW
1010
Set Selected PLL (PLL_S) to Maximum LBW
1020
Determine Initial Signal Quality Metric (SQ_I)
1030
Adjust LBW Setting of PLL_S
1040
Determine Adjusted Signal Quality Metric (SQ_A)
1050
$SQ_A > SQ_I$?
1060
Y
N
Max Iterations?
1061
Y
N
Store SQ_I
1070
Adjust LBW Setting of PLL_{NS}
1080
Store/Update LBW settings
1062
Set $SQ_I = SQ_A$
1064

FIG. 10

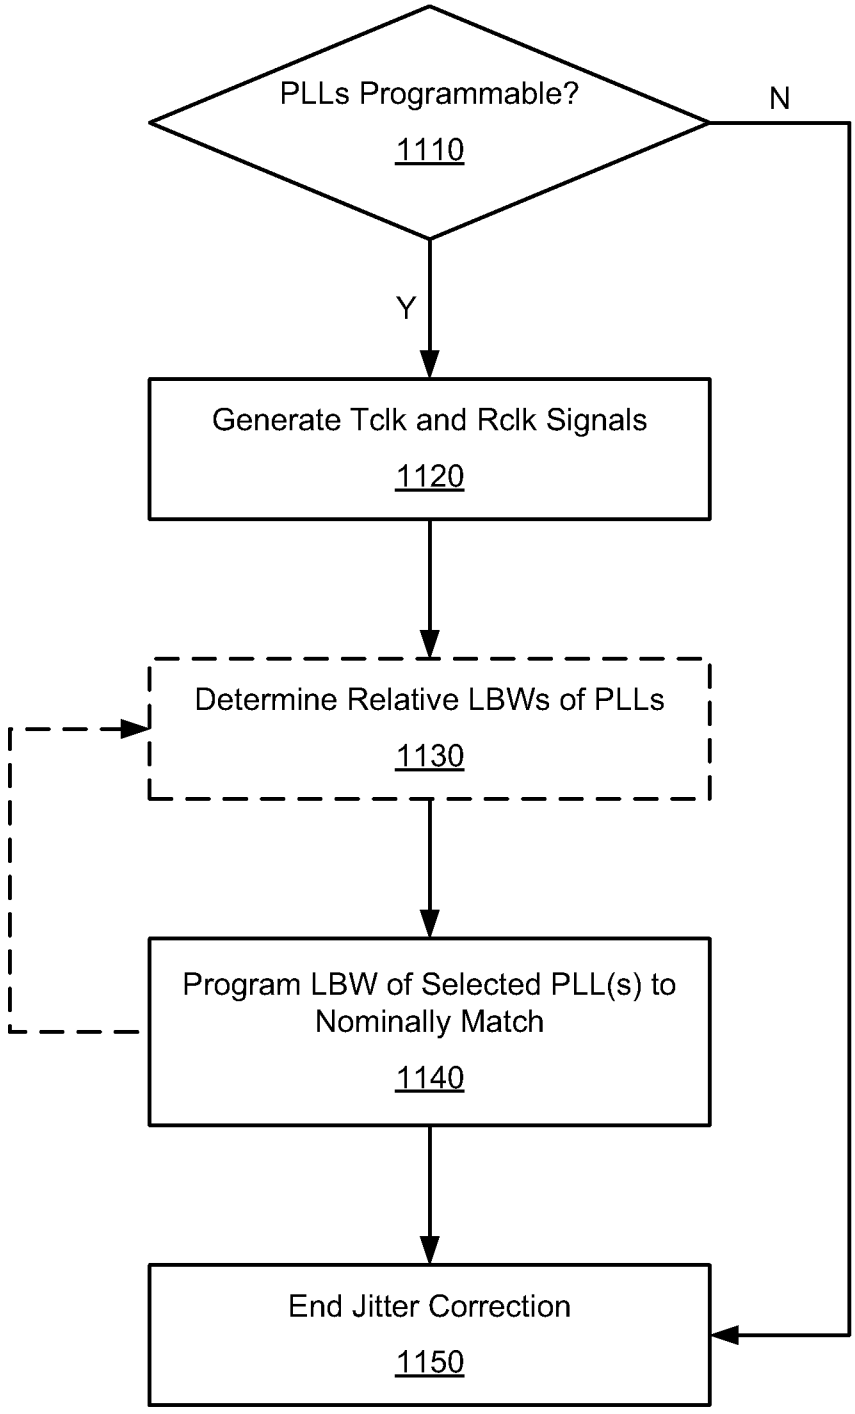
PLLs Programmable?
1110
N
Y
Generate Tclk and Rclk Signals
1120
Determine Relative LBWs of PLLs
1130
Program LBW of Selected PLL(s) to Nominally Match
1140
End Jitter Correction
1150

FIG. 11

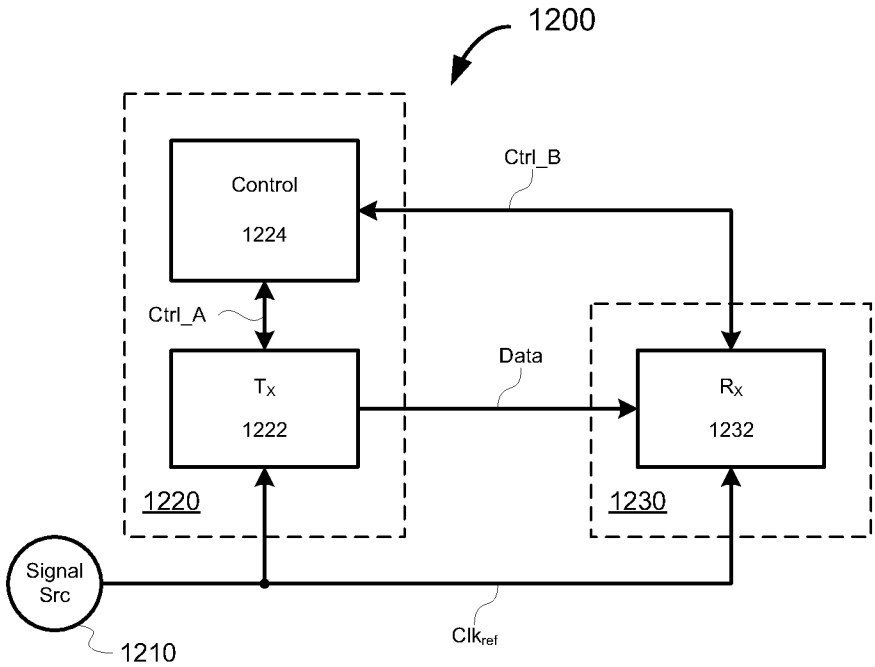
1200
Control
1224
Ctrl_B
Ctrl_A
Data
T_X
1222
R_X
1232
1220
1230
Signal
Src
1210
Clk_{ref}

FIG. 12

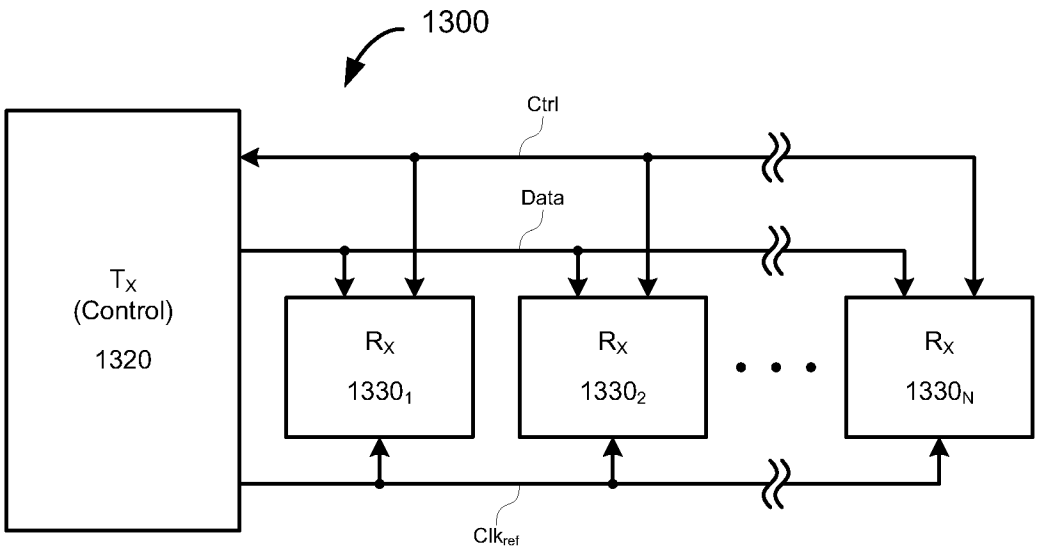
1300
Ctrl
Data
T_X
(Control)
1320
R_X
1330_1
R_X
1330_2
R_X
1330_N
Clk_{ref}

FIG. 13

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2009/061847

A. CLASSIFICATION OF SUBJECT MATTER
INV. H03L7/22 H04L7/00

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)
H03L H04L

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 practical, search terms used)
EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/022303 A1 (TONIETTO DAVIDE [US] ET AL) 5 February 2004 (2004-02-05) paragraph [0020] paragraph [0077] paragraph [0095] - paragraph [0102] claims 1,3,12,13	1-33
X	US 2004/131058 A1 (GHIASI ALI [US]) 8 July 2004 (2004-07-08) paragraph [0104] - paragraph [0111]	1-33
A	WO 2008/115968 A (RAMBUS INC [US]; LEE HAE-CHANG [US]; KIM JAEHA [US]; LEIBOWITZ BRIAN [) 25 September 2008 (2008-09-25) abstract	1-33

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* 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 document 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

"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

"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	Date of mailing of the international search report
19 January 2010	27/01/2010

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31–70) 340–2040, Fax: (+31–70) 340–3016	Kahn, Klaus-Dieter

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2009/061847

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004022303 A1	05-02-2004	NONE	
US 2004131058 A1	08-07-2004	US 2007274350 A1	29-11-2007
WO 2008115968 A	25-09-2008	EP 2130055 A1	09-12-2009