

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
22 May 2009 (22.05.2009)

PCT

(10) International Publication Number
WO 2009/064439 A2

(51) International Patent Classification:
H04L 12/28 (2006.01) H04L 12/10 (2006.01)

(74) Agent: MILLER, Chad, W.; Weide & Miller, Ltd., 7251
W. Lake Mead Blvd., Suite 530, Las Vegas, NV 89128
(US).

(21) International Application Number:
PCT/US2008/012749

(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, 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, 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.

(22) International Filing Date:
12 November 2008 (12.11.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/987,327 12 November 2007 (12.11.2007) US

(71) Applicant (for all designated States except US): SO-
LARFLARE COMMUNICATIONS, INC. [US/US];
9501 Jeronimo Road, Suite 250, Irvine, CA 92618 (US).

(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, MT, NL,
NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG,
CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(72) Inventor; and
(75) Inventor/Applicant (for US only): ZIMMERMAN,
George, A. [US/US]; 19 Ranchview Road, Rolling Hills
Estates, CA 90274 (US).

[Continued on next page]

(54) Title: ACTIVE IDLE COMMUNICATION SYSTEM




Fig. 3

(57) Abstract: To reduce power consumption and heat generation, an active idle system is proposed that monitors for an idle period and then, after a predetermined time, initiates a silent period. During the silent period data and idle frames are not transmitted. During the silent period, one or more transceiver components may be turned off or forced into some other power saving mode. The predetermined time may be any amount of time and is selected to balance network usage and power savings. Periodically during the silent period, such as at predetermined times, one or more sync or idle frames are transmitted. Received sync or idle frames are processed to maintain receiver settings, synchronization or equalizer adaptation. Restoring active data communication may occur by monitoring the channel during silent periods for a request or only during the predetermined times when sync or idle frames are sent.

WO 2009/064439 A2



Published:

- *without international search report and to be republished upon receipt of that report*

ACTIVE IDLE COMMUNICATION SYSTEM**INVENTOR**

GEORGE A. ZIMMERMAN

5

1. Priority Claim

[001] This patent application claims the benefit of and priority to U.S. Provisional Patent Application Number 60/987,327 which was filed on November 12, 2007 and entitled
10 Active Idle Communication System.

2. Field of the Invention.

[002] The invention relates to data communication and in particular to a method and apparatus for entering a power saving idle state in data communication systems.

3. Related Art.

15 [003] Numerous prior art communications systems, such as current 10G-BaseT systems enable communication at high rates of speed over substantial distances. Solarflare Communications, located in Irvine, California is the leading design and product development entity for 10 Gigabit Ethernet communication systems. While
20 such communication systems operate as specified and desired, there is a global desire to reduce power consumption, such as electric power consumption, without degrading performance.

[004] In prior art Ethernet systems, such as 10 Gigabit systems, the communication system is always active, such that either data or idle frames are always transmitted, received, and processed by the receiver. The idle frames may comprise scrambled

sequences of 1 and 0 bits which are arranged into special codes that indicate that the frames are idle frames, and not data. Active operation of the transmitter, receiver, and receiver processing elements consumes substantial power and as such, is the primary source of power consumption in a communication system.

5 [005] As can be appreciated, during periods when a 10 Gbit/second communication system is not exchanging data, idle frames are still being generated, transmitted, received and processed at 10 Gbit/second rate. This consumes substantial power even when data is not being exchanged, hence.

[006] One proposed solution is to reduce the transmit rate to reduce power
10 consumption. For example, in a 10GBase-T system, it has been proposed to reduce the rate during idle states to save power. The amount of power saved is proportional to the reduction in rate. This proposed solution suffers from several drawbacks. One such drawback is that power savings are less than the solution proposed herein. Another drawback is that to reactivate communication at the full data rate, significant
15 reconfiguration to the settings of the communication system must occur which requires an undesirably long delay. This delay, referred to as the transition time or transition latency, degrades performance and hinders the user's and network's ability to rapidly exchange data.

[007] As a result, there is a need in the art for a method and apparatus to reduce
20 power consumption of a communication system while concurrently minimally affecting performance.

SUMMARY

[008] To overcome the drawbacks of the prior art and to provide additional benefits, a method and apparatus is disclosed for monitoring idle periods and entering the communication system into a power down state. During the power down state, periodic communication and signal processing occurs to maintain the transceiver's communications settings, such as synchronization and equalizer coefficients. Re-activation occurs in response to active data or a reactivation request.

[009] In one embodiment, the innovation comprises a method for reducing power consumption in an Ethernet communication device. This exemplary method comprises monitoring for an idle frame and upon detecting a predetermined number of idle frames, entering a reduce power consumption state in which idle frames are not processed. Then, monitoring for a transmit period and responsive to a transmit period, receiving one or more idle or training frames. This method optionally adapts the communication system based on one or more of the one or more idle or training frames. After optional adaptation, the system may re-enter a reduced power consumption state in which idle frames are not processed.

[010] In one embodiment, the predetermined number of idle frames comprises one idle frame. In addition, the step of monitoring for a transmit period comprises receiving a timer output such that the timer output signals the transmit period. It is also contemplated that optionally adapting may comprise analyzing the idle or training system to determine channel conditions and then comparing the channel conditions with prior channel conditions. Based on this, the method then adapts the communication system if the comparing reveals a change in channel conditions

greater than a threshold. This method may receive and detect an activate frame and, responsive thereto, restore data communication.

[011] Also disclosed herein is a method for maintaining a communication link during periods of non-use of the communication system comprising detecting an idle
5 period in a communication system. Then, responsive to the idle period, initiating a counter or timer for a first time period, as part of an idle frame sequence, and sending idle frames during the idle frame sequence. Then, monitoring for an end of the first time period and responsive to the end of the first time period, stopping the sending of
10 idle frames. This method then initiates a counter or timer to start a second time period and monitors for the end of the second time period and responsive to the end of the second time period sending a sync or idle frame to an opposing receiver. Then the method receives and processes the sync or idle frame at the opposing receiver to maintain one or more settings of the opposing receiver.

[012] In one embodiment, the counter or timer times a predetermined period of time.
15 For example the counter or timer may count or track frames. In addition, stopping the sending of idle frames may comprise sending no signal. In addition, maintaining one or more settings of the opposing receiver may comprise performing timing synchronization, adaptation, or both. In addition, the step of stopping the sending of
20 idle frames further may comprise powering down one or more aspects of a transmitter.

[013] Also disclosed herein is a system for saving power in a communication device comprising an input/output port configured to connect to one or more channels and a transmitter configured to transmit data over the channel via the input/output port.

Also part of this embodiment is a controller configured to interface with the transmitter. The controller may be configured to detect idle periods, control the transmitter to initiate a silent period by no longer sending a signal over the channel after a predetermined idle period, and reactivate the transmitter periodically to send
5 sync or idle information, and reactivate the transmitter in response to active data to be transmitted. The sync or idle information to be used by an opposing receiver to maintain the opposing receiver settings.

[014] In one embodiment the system further comprises a detector, the detector configured to detect incoming or outgoing signals. It is further contemplated that the
10 system further comprises a timer configured to time idle periods. The memory may be configured to store information regarding the length of a predetermined idle period. In addition, the system may further comprise a clocking system having a clock signal, and during a silent period the clock signal is not provided to two or more transmitter apparatus. In one embodiment during the silent period, the controller turns
15 off one or more aspect of the transmitter and one or more aspects of the opposing receiver are also turned off.

BRIEF DESCRIPTION OF THE DRAWINGS

[015] The components in the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention. In the figures, like reference numerals designate corresponding parts throughout the different views.

5 [016] Figure 1 illustrates a prior art frame stream.

[017] Figure 2 illustrates a frame stream in accord with one embodiment of the present invention.

[018] Figure 3 illustrates a block diagram of an example embodiment of the present invention.

10 [019] Figure 4 illustrates a detailed block diagram of an example embodiment of a transceiver in an on state.

[020] Figure 5 illustrates a detailed block diagram of an example embodiment of a transceiver in a power save state.

15 [021] Figure 6 illustrates an operation flow diagram of an example method of idle state transition and operation.

DETAILED DESCRIPTION

[022] Figure 1 illustrates a prior art frame stream. In this prior art frame stream, during periods when data is not transmitted, idle frames are transmitted to maintain the communication device synchronization and equalizers. Data is defined herein as user information which is transmitted between communication devices for higher level application levels. Idle frames are digital content sent during periods when data is not sent, such as when the communication device is not in use by a user.

[023] During a frame sequence 104 a user or the communication device is utilizing the communication channel for the exchange of data. As such, data is being transmitted and receive by opposing transceivers and the communication system is functioning as desired to perform data communication. Thereafter, during a frame sequence 108, the communication devices stop transmitting data and become idle. During idle periods, idle frames, containing scrambled 1's and 0's are transmitted between stations and processed. The idle frames do not contain user data and as such, are transmitted to maintain the communication link. Maintaining the communication link may include maintaining timing, maintaining equalizer adaptation, maintaining echo and crosstalk canceller adaptation, as well as physical link security keeping the system informed that the far-end device is present and still connected by physical cabling. As a drawback, transmitting and processing idle frames during inactive periods undesirably maintains power consumption at the transmitter and receiver. As can be appreciated, transmitting idle frames during periods of non-use is wasteful. Some network connected computer may be left on all night and consume power the

entire night. Power consumption also leads to heat generation and additional cooling requirements. Upon a user or other machine restoring active communication, data frames are transmitted again at frame sequence 112.

[024] To overcome the drawbacks of the prior art and to provide additional advantages, a method and apparatus is disclosed herein to reduce power consumption in communication system by entering a reduced power mode (state) during idle periods. Idle periods may be any time when data is not transmitted, received, or both, or when the rate of data transfer is sufficiently low to allow for a reduce power mode.

[025] In one embodiment, during periods of reduced traffic, the transmitter, receiver, and received signal processing elements may be powered down or power down for time frames during the idle periods. These periods may be referred to herein as silent periods. Figure 2 illustrates a frame stream in accord with one embodiment of the present invention. As shown, frame sequence 204 comprises data frames as discussed above. Thereafter, at frame sequence 208, the data is complete and the system enters an idle state. During this idle state, idle frames are transmitted. Any number of idle frames may be transmitted, but after a predetermined number of frames, the system may be configured to transmit a start silent frame during frame sequence 212. The silent start frame comprises a frame containing a code that initiates the silent period. Alternatively, the silent period may start automatically after the predetermined number of idle frames.

[026] During the silent frame sequence, 216 the system may enter a sleep or power down mode. It is contemplated that numerous different levels of operational

reduction may occur. This may include shutting select certain systems down, slowing operation, or initiating periodic operation.

[027] Next, during frame sequence 220, the operation may send a sync frame between the transmitter and the receiver. The sync frame as shown in sequence 220
5 may comprise one frame or any number of frames. It is contemplated that the sync frame may be sent at a predetermined time which is commonly know by the transmitter and the receiver and referred to herein as a sync frame time. It is also contemplated that an idle frame may be sent instead of the sync frame. Upon receipt
10 of a sync frame, the receiver may perform processing to maintain the clock synchronization, timing, and perform adaptation to adjust equalizer and canceller coefficients. The sync frame could also be a wake-up code frame to restore active communication.

[028] Thereafter, in Figure 2, the frame sequence 224 contains further silent frames. During this period, the system is still not in use and as such, it continues in silent
15 mode by not transmitting frames. Power consumption is reduced. At frame sequence 228, which may comprise a single frame or multiple frames, a stop silent frame is transmitted. The stop silent frame comprises a frame containing a code that is transmitted to the receiver and indicates to the receiver that active data communication is requested. It is thus contemplated that the receiver may monitor for
20 an incoming frame and upon detection of an incoming frame, such as the stop silent frame sequence, then the receiver may resume active communication.

[029] Alternatively, the system, namely transmitter / receiver pair, may be configured to await the next sync frame 220 to restore active communication. Thus,

in this configuration, the receiver and transmitter can only resume active communication at one of the predetermined times.

[030] Thereafter, at frame sequence 232 active communication resumes. In one embodiment, it is contemplated that data may be contained in the sync or idle frames.

5 This allows for a slow rate data exchange, but still maintains the power saving benefits. This may be beneficial when only a low data rate is required or only a small amount of data need be periodically sent. Thus, data could be transferred for M 'on frames' for very N 'off frames'. Overhead and housekeeping data, such as timing information or precoder coefficients could also be sent. In either situation, the timing
10 lock and frame sync may be maintained. And a wake cycle to refresh DSP coefficients is not necessary.

[031] In one example embodiment, during idle mode, instead of continually transmitting, receiving, and processing idle frames, which consumes substantial power, the system enters power save mode and only transmits idle or sync frames
15 every M out of X frames and is silent, i.e. powered down to some extent, N out of the X frames. M plus N equal X and the value of N is contemplated as being greater than the value of M. The values M, N, and X may be any positive whole number.

[032] By way of example, during a 10 frame window during idle mode (power save mode), one (N) idle symbol frame may be sent and the system will be silent (powered
20 down) for the remaining 9 (M) frames. As a result, notwithstanding standing leakage current, for the components that are powered down there is a 90% power savings. In addition, the transmit processing, receive processing is also the most power intensive aspects of the communication system. Even though other aspects of the transceiver

may remain active and powered up, by substantially reducing the power consumption of these aspects, overall power consumption is greatly reduced. It is also contemplated that leakage current may be reduced or eliminated by gating off power connections to one or more devices.

5 [033] For communication systems that remain idle for long periods, such as when the user is on the phone, away at lunch, in meetings, and at nights, the power savings can be substantial. The power savings are further increased if the ratio of idle frames to silent frames is increased, such as 100 to 1 or 1000 to 1. Based on testing and estimation, the time between transmission of an idle frame during idle silent periods
10 can be up to minutes. It is contemplated that the period between transmission of an idle or sync frame during a silent periods could be any ratio or time depending on network and communication usage, priority, and desired power saving, keeping as an outer time limit a time which maintains system timing, synchronization, and adaptation.

15 [034] One factor which may determine the ratio of silent periods to idle frame transmit periods, i.e. time periods between transmission of frame(s) containing idle frames is the rate of change of the channel in terms of the channel echo and crosstalk signature. It is contemplated that adaptation of one or more filters within the communication system may be sufficiently maintained by the exchange of the idle
20 symbols during transmit frames even though the transmit frames may be spaced between significant periods of silent frame, i.e. time periods when the system is silent and powered down to some extent. In such an embodiment, seconds or minutes may pass between transmission of the idle symbols and adaptation will occur based on

analysis of the idle symbol during the transmit frame sufficient to maintain the communication link and the optimal (full rate) transmit rate. Due to the short duration in which sufficient synchronization and adaptation may occur, only a few sync or idle frames need be transmitted to maintain the transceiver.

5 [035] During such idle states, it is contemplated that the timing circuit may be maintained in an active state, to maintain synchronization, but the clock to the echo and next cancellers may be disabled. Other components of the system may be enabled and disabled as desired. As can be appreciated in CMOS applications, eliminating the clock eliminates the power draw, except for leakage current.

10 [036] As a result of the sufficient adaptation during the transmit frame during idle state and maintaining the timing circuit active during idle state, the communication system may rapidly reactivate and enter full rate communication with minimal latency. For example, in contrast to other proposed power reduction schemes, which may take in the milliseconds or more to restart to full rate mode, the method and
15 apparatus for power reduction disclosed herein is able wake to full rate communications in within microseconds or less. Recovery time may depend on the number of silent frames in relation to the number of idle symbol transmit frames. For example, a high performance network or user may recover faster as based on administrator settings which set a fewer number of silent frames between transmit
20 frames. A standard or low level network or user may increase the number of silent frames between idle symbols (transmit frames). Power savings vary accordingly. The settings may be established by a network administrator and be based on user, network link, network, time of day, date, or type of frame or activity.

[037] In addition, even during idle periods, when a user is not demanding data exchange, a physical change in the echo and crosstalk signature which is sufficient to require significant adaptation or retraining, can be detected upon reception of a transmit frame, containing an idle symbol(s) and the significant adaptation or retraining may occur before a user even requires full rate data exchange. If a need for retraining is detected, additional symbols may be sent to achieve such retrain. Thereafter, the system may return to idle mode. As a result, even a physical change to channel or surrounding channels may be detected and adapted to before a user requires full rate data exchange. This is true because it is unlikely the physical change (or any significant change requiring a full retrain) to the channel will occur at the generally same exact time that a user initiates a return to active mode.

[038] It is further contemplated that there may be different levels of idle or power saving states. In one such state, the clock may be disabled from select components. In another state, all power is cut to select components, which will eliminate power consumption from leakage current.

[039] Figure 3 illustrates a block diagram of an example embodiment of the present invention. This is but one example embodiment of a power saving transceiver and as such, it is contemplated that other embodiments may be created which do not depart from the claims that follow. In this example embodiment an input/output (I/O) port 304 provides a signal to a transceiver 306. Receiving the signal inside the transceiver 306 is a hybrid 308, which is known by those of skill in the art as a signal splitter. The I/O port 304 may comprise any type input and have any type connector or number of channels. In one embodiment, the input comprises one or more twisted

pair conductor. The hybrid 308 connects to a transmitter 316 and receiver as shown and as is understood in the art. The hybrid 308 has an output that connects to the receiver 312. The receiver 312 includes apparatus that process incoming signal to establish the signal for use by higher level applications. The hybrid 308 receives an
5 input from the transmitter 316 as shown. The transmitter 316 includes apparatus that process incoming signals from higher level application to establish signal for transmission over a channel. These elements are not shown in detail in this example block diagram. In one embodiment the receiver and transmitter operation is according to Ethernet protocol. In one embodiment the receiver and transmitter
10 operate at 10Gbits per second. The transmitter 316 and receiver 312 connect to the upper level application layers via input/output ports 350A, 350B.

[040] A detector 320 connects to and monitors the transmit and receive paths as shown. The detector may comprise any hardware, software, or combination thereof, configured to detect frames or communication on the monitored paths. The detector
15 320 may also connect directly to the receiver and/or transmitter as shown to thereby access the internal components of the receiver and transmitter.

[041] Connecting to the detector 320 is a controller 324. The controller 324 may comprise any controller or processor capable of monitoring for idle symbols, causing the communication system to enter a silent period, and thereafter, restoring active data
20 communication.

[042] The controller 324 also connects to a counter or timer 328, a clock 332, and a memory 340. It is contemplated that the controller 324 may connect to or communication with additional components in the transmitter 316 and receiver 312.

In this example embodiment the counter or timer 328 is configured to determine when to enter silent power save mode, and determine when to synchronize and adapt during silent power save mode. It is also contemplated that instead of a timer or counter, a clock signal may be used. The memory 340 may store machine readable code for use
5 by the processor or controller 324. The memory may also store operation parameters including, but not limited to time before entry into power saving mode, the sync and adaptation schedule, such as how often to sync or adapt during silent periods, and how long to signal before resuming operation or any other operation parameter. The clock 332 may comprise a common clock shared by the transmitter 316 and the receiver
10 312. In one embodiment, the controller 324 may shut down the clock, or switch the clock away from the transceiver components as part of the silent power save mode.

[043] In operation, the transmitter 316 transmits data over the channel. The receiver 312 may be receiving data during transmission depending on various factors. During operation, the system may enter an idle state, wherein a user or machine is not
15 transmitting active data over the I/O port 304. At such time, the transmitter 316 will enter a idle state and begin transmitting idle frames. Concurrently the timer 328 and controller 324 may communicate to initiate a timer running. Alternatively, the controller 324 may count or track the number of idle frames sent over the channel. When a predetermine amount of time has passed while transmitting only idle frames
20 or a predetermined number of idle frames have been sent, then the controller initiate a power down operation to reduce power consumption during periods of non-use. In one embodiment, this comprises powering down the transmitter. In one embodiment, this comprises powering down the transmitter and receiver. In one embodiment, this

comprises powering down the clock or switching the clock away from the transmitter and/or receiver. In one embodiment, this comprises removing bias from the transmitter and/or receiver.

[044] Upon entering power down mode, the transmitter 316 stops transmitting
5 frames for a predetermine time period or frame period. Consequently, it is contemplated that the receiver processing at the opposing transceiver may likewise be synchronized and placed in reduce power consumption mode. This state may be considered a silent state and may last for any predetermined period of time.

[045] At the start of the silent state, the timer 328 or other counting element may
10 monitor the time from the start of the silent period. After a predetermined amount of time, the controller 324 will cause the transmitter to generate and sent a sync frame or an idle frame. Upon receipt, the opposing receiver can process this known sync frame or an idle frame to maintain and update the receiver settings. For example, it is contemplated that timing and clock synchronization may occur as well as adaptation
15 for changing channel conditions.

[046] This pattern may continue wherein the system does not transmit idle frames when in power savings mode but will periodically, at a time known to both the transmitter and receiver, send and process sync or idle frames. This sync or idle frames are processed so that when active data communication is restored, the system
20 may quickly restore communication without having to undergo a full retrain sequence.

[047] Active data communication may be restored in two different ways. In one embodiment, the restoration of active data communication can only occur at the predetermined time for a sync or idle frame when in silent mode. In such an

embodiment, the sync or idle frame may be supplemented or replaced with a restart code frame that signals the opposing receiver to restore active data communication.

[048] Alternatively, the detector 320 may monitor the transmit path or receive path as shown in Figure 3 for incoming idle or other type frames during a silent period. If

5 idle or other type frames are detected then the detector 320 may signal the controller 324 to restore active data communication. This may occur by reversing the steps taken to enter power saving mode. In either embodiment, an acknowledgement may be sent from the opposing transceiver that the request to restore active communication has been received and acknowledged.

10 [049] It is also contemplated that the predetermined time period between syncing or adaptation when in silent mode may be set by the user or the communication system controller 324. In such an embodiment, the channel characteristics and rate of change may be monitored. If the change conditions are subject to change, then the sync frame and adaptation may occur more frequently.

15 [050] In the case of Ethernet, 10GBASE-T type Ethernet utilizes a fixed-length LDPC PHY frame, which may be referred to a set framing. This set framing interval may be utilized to signal reduced frame speed. As such, the receiver can alternate between 'off' or silent frames without burning signal processing decoding power. IN addition, near and far end frames can be staggered for additional savings when frames
20 are received. Stated another way, opposing transceivers may alternate the time at which they send sync frames. In this manner, it is contemplated that returning to high rate active data communication may occur within a single frame. Likewise, other PHYs can be structured similarly using period sync or idle symbols and their PCS's.

[051] Figure 4 illustrates an example communication system receiver, such as may be implemented in an Ethernet communication system in an on or active state. In this embodiment, communication channel (not shown) connects to a connector 404. The signal passes through the connector to a data path 408 where processing of the signal may occur to recover the data. Echo cancellation and crosstalk cancellation may occur within the data path 408. Each element is descriptively labeled within Figure 4 and as such, each element is not discussed in detail herein beyond that described above and below. A control section 112 generates a clock signal and maintains synchronization of the receiver based on a reference frequency. The data path 108 receives the clock signal from the control section 112. During active data exchange operation occurs in the system of Figure 4 as is understood in the art. Also included in this embodiment is one or more control elements configured to monitor for idle periods and enter/exit idle state as described herein. This may be enabled in hardware, software, or a combination of both.

[052] In this example embodiment, the transceiver is shown in an on state and it is contemplated that in this on state, all of the elements shown in Figure 4 are on, powered up, and active. This state would be considered a prior art power consumption state since all elements are active.

[053] Figure 5 illustrates a detailed block diagram of an example embodiment of a transceiver in a power save state such that the elements of the data path of powered down. Powering down these elements may occur in any manner known now or in the future. In one embodiment, the elements are powered down by switching out the clock signal to these elements which substantially reduces power consumption. In a second

embodiment, the elements are powered down or have the clocks or supply voltages switched or gates. This may eliminate leakage power consumption as well, at the cost of slightly longer recovery time. In one embodiment of the power save state, the block elements shown in dashed lines may be powered down during power save state. This save substantial amount of power and reduces heat generation. In one embodiment, the PCS and PMA synchronization is maintained. This enables blinder return to active data communication at a high rate. This method and apparatus also overcomes the drawback with simply reducing data transmit rate. In particular, because different data rates utilize different PCS elements restoring status to a higher rate requires a time consuming re-activation of the higher rate PCS. Thus, when transitioning up from one data rate to a higher rate, the high rate PCS would have to be activated, which takes longer than the method and apparatus proposed herein.

[054] Figure 6 illustrates an operation flow diagram of an example method of transition into and out of power save state. At a step 600 the system monitors the activity status of the communication link. The monitoring occurs to determine if the link is active, such as exchanging data, or idle and simply consuming valuable power transmitting, receiving and processing idle symbols. After a set time period of exchanging idle symbols the system will enter idle state, which may also be referred to herein as power save state. This occurs at step 604. The amount of time may be established by the manufacture, network installer, or network administrator.

[055] At a step 608, the system determines the idle state operational parameters. The idle state parameters may include the degree to which systems are modified or shut down to reduce power consumption. For example, there may be varying levels

of power consumption which correspond to different changes to the communication system. In one embodiment, the clock to the data processing path is switched off. In addition, the idle state parameters may include the time between transmission of idle frames and silent frames during idle state. This may be referred to as the silent frame
5 ratio.

[056] At a step 612, the system turns off the clock to the data path and stops transmitting and idle symbols. The transmitter may go silent for any number of silent frames depending on the silent frame ratio. This reduces power consumption. The timing functions of the receiver/transmitter may remain on to maintain
10 synchronization. At a step 616, the system initiates a timer or other monitoring system to determine when to receive and/or transmit one or more idle symbol frames(s). In one embodiment, the system performs time synchronized monitoring to determine when to send a silent frame. In one embodiment, a predetermined number of silent frames may be sent before a transmit frame occurs.

15 [057] Then, at a step 620, the system detects or determines the assigned time period, (or frame number) to transmission/reception of an idle symbol and turns on the clock to the data path. In other embodiments, other events may occur to activate the receiver/transmitter system. Then, at step 624, the communication system receives a frame. It is contemplated that at this stage, the frame may be analyzed to determine if
20 the frame is data or idle symbols. This occurs at decision step 628. If at decision step 628 a data frame is detected, then the operation advances to step 632, wherein the communication system is restarted or is maintained in active mode and continues

communication at full rate. Thereafter, the operation returns to step 600. After active communication, the system may return to power save mode as described above.

[058] Alternatively, if at decision step 628 the system detects an idle symbol, then the system may stay in power save mode. The operation advances to step 636 to
5 process the idle frame to maintain system synchronization and adapt to any changes in channel condition. By processing the idle symbols, which are periodically received, the communication system is capable of returning to full rate communication at any time with minimal delay.

[059] The method and apparatus disclosed herein overcomes the challenges of how
10 to transition rapidly from power save state to active data communication. By sending a sync or idle symbol for synchronization and adaptation periodically during silent mode, the settings of the transceivers may be maintained current. Restoring active data communication may occur quickly. It is contemplated that with this method and apparatus autonegotiation can be bypassed. In one embodiment, transition time of 1
15 microsecond may be possible.

[060] The following definitions relate to the abbreviations utilized in the slides, which are part of this provisional application.

PHYs – Physical layer transceiver devices

10GBASE-T – 10 gigabit Ethernet transceiver designed to run on twisted-pair copper
20 media

PCS – Physical Coding Sublayer – a sublayer of a PHY that defines how bits are encoded and grouped for transmission, and decoded on reception

PMA – Physical Media Adaptation (sublayer) – a sublayer of a PHY that defines how encoded and grouped bits are mapped to electrical signals for transmission and received/demapped on reception. This sublayer is also responsible for any signal processing and conditioning necessary for reception.

- 5 LDPC – Low Density Parity Check code. LDPC codes are powerful forward error correction codes used in the 10GBASE-T PCS.

DSP – Digital Signal Processor (a way of implementing the PMA)

Acronyms in the block diagram:

DAC – Digital to analog converter

- 10 TxTH Precoder – Transmitter Tomlinson-Harashima Precoder – a signal conditioning device used to provide high-performance equalization for coded systems.

XGMII Interface – The 10(X) Gigabit Media Independent Interface. A PHY-neutral interface defined by IEEE Std. 802.3-2005 between the MAC and PHY layers. The XGMII interface is highly parallel (32 bits wide), hence the use of the XAUI interface
15 externally.

XAUI – The 10(X) Gigabit Attachment Unit Interface SERIALizer/DESerializer. A 4-lane serial interface (and it's associated serializer/deserializer hardware) defined to efficiently transmit 10 Gigabit Ethernet interfaces across printed circuit board traces.

- LNA – Low Noise Amplifier – used in the receiver front end to amplify the weak
20 received signal without adding excess noise.

PGA – Programmable Gain Amplifier – used to adjust the level of the receive signal.

RxADC – Receiver Analog to Digital Converter – converts the received signal to digital words for digital signal processing.

PLL – Phase Locked Loop – a device used for timing recovery and clock generation.

MDIO – Management Data Input/Output – an interface defined by 802.3-2005 used to access management registers in an Ethernet device.

[061] The principles discussed herein may apply to any other type communication
5 system, including but not limited to 10GBASE-T, 1000BASE-T, and 100BASE-T.
Other systems, methods, features and advantages of the invention will be or will become apparent to one with skill in the art upon examination of the following figures and detailed description. It is intended that all such additional systems, methods, features and advantages be included within this description, be within the scope of the
10 invention, and be protected by the accompanying claims.

[062] While various embodiments of the invention have been described, it will be apparent to those of ordinary skill in the art that many more embodiments and implementations are possible that are within the scope of this invention. In addition, the various features, elements, and embodiments described herein may be claimed or
15 combined in any combination or arrangement.

CLAIMS

What is claimed is:

- 5 1. A method for reducing power consumption in an Ethernet communication device comprising:
- monitoring for an idle frame;
 - upon detecting a predetermined number of idle frames, entering a reduce power consumption state in which idle frames are not processed;
 - 10 monitoring for a transmit period;
 - responsive to a transmit period, receiving one or more idle or training frames;
 - optionally adapting the communication system based on one or more of the one or more idle or training frames; and
 - re-entering a reduced power consumption state in which idle frames are not
 - 15 processed.
2. The method of Claim 1, wherein the predetermined number of idle frames comprises one idle frame.
- 20 3. The method of Claim 1, wherein monitoring for a transmit period comprises receiving a timer output, wherein the timer output signals the transmit period.
4. The method of Claim 1, wherein optionally adapting comprises:
- analyzing the idle or training system to determine channel conditions;
 - 25 comparing the channel conditions with prior channel conditions;
 - adapting the communication system if the comparing reveals a change in channel conditions greater than a threshold.
5. The method of Claim 1, further comprising receiving an activate frame and,
- 30 responsive thereto, restoring data communication.

6. A method for maintaining a communication link during periods of non-use of the communication system comprising:

detecting an idle period in a communication system;

responsive to the idle period, initiating a counter or timer for a first time period, as part of an idle frame sequence;

sending idle frames during the idle frame sequence;

monitoring for an end of the first time period and responsive to the end of the first time period;

stopping the sending of idle frames;

initiating a counter or timer to start a second time period;

monitoring for the end of the second time period and responsive to the end of the second time period sending a sync or idle frame to an opposing receiver;

receiving and processing the sync or idle frame at the opposing receiver to maintain one or more settings of the opposing receiver.

15

7. The method of Claim 6, wherein counter or timer times a predetermined period of time.

8. The method of Claim 6, wherein counter or timer counts or track frames.

20

9. The method of Claim 6, wherein stopping the sending of idle frames comprises sending no signal.

10. The method of Claim 6, further comprising resetting the second time period.

25

11. The method of Claim 6, wherein maintaining one or more settings of the opposing receiver comprises performing timing synchronization, adaptation, or both.

12. The method of Claim 6, wherein stopping the sending of idle frames further comprise powering down one or more aspects of a transmitter.

30

13. A system for saving power in a communication device comprising
an input/output port configured to connect to one or more channels;
a transmitter configured to transmit data over the channel via the input/output
port;

5 a controller configured to interface with the transmitter, the controller
configured to:

detect idle periods;

control the transmitter to initiate a silent period by no longer sending a
signal over the channel after a predetermined idle period;

10 reactivate the transmitter periodically to send sync or idle information,
the sync or idle information to be used by an opposing receiver to maintain the
opposing receiver settings;

reactivate the transmitter in response to active data to be transmitted.

15 14. The system of Claim 13, further comprising a detector, the detector configured
to detect incoming or outgoing signals.

15. The system of Claim 13, further comprising a timer configured to time idle
periods.

20

16. The system of Claim 13, further comprising a memory configured to store
information regarding the length of a predetermined idle period.

17. The system of Claim 13, wherein the system further comprises a clocking
25 system having a clock signal, and during a silent period the clock signal is not
provided to two or more transmitter apparatus.

18. The system of Claim 13, wherein during the silent period, the controller turns
off one or more aspect of the transmitter.

30

19. The system of Claim 18, wherein during the silent period, one or more aspects of the opposing receiver are also turned off.

5




Fig. 1
Prior Art




Fig. 2




Fig. 3

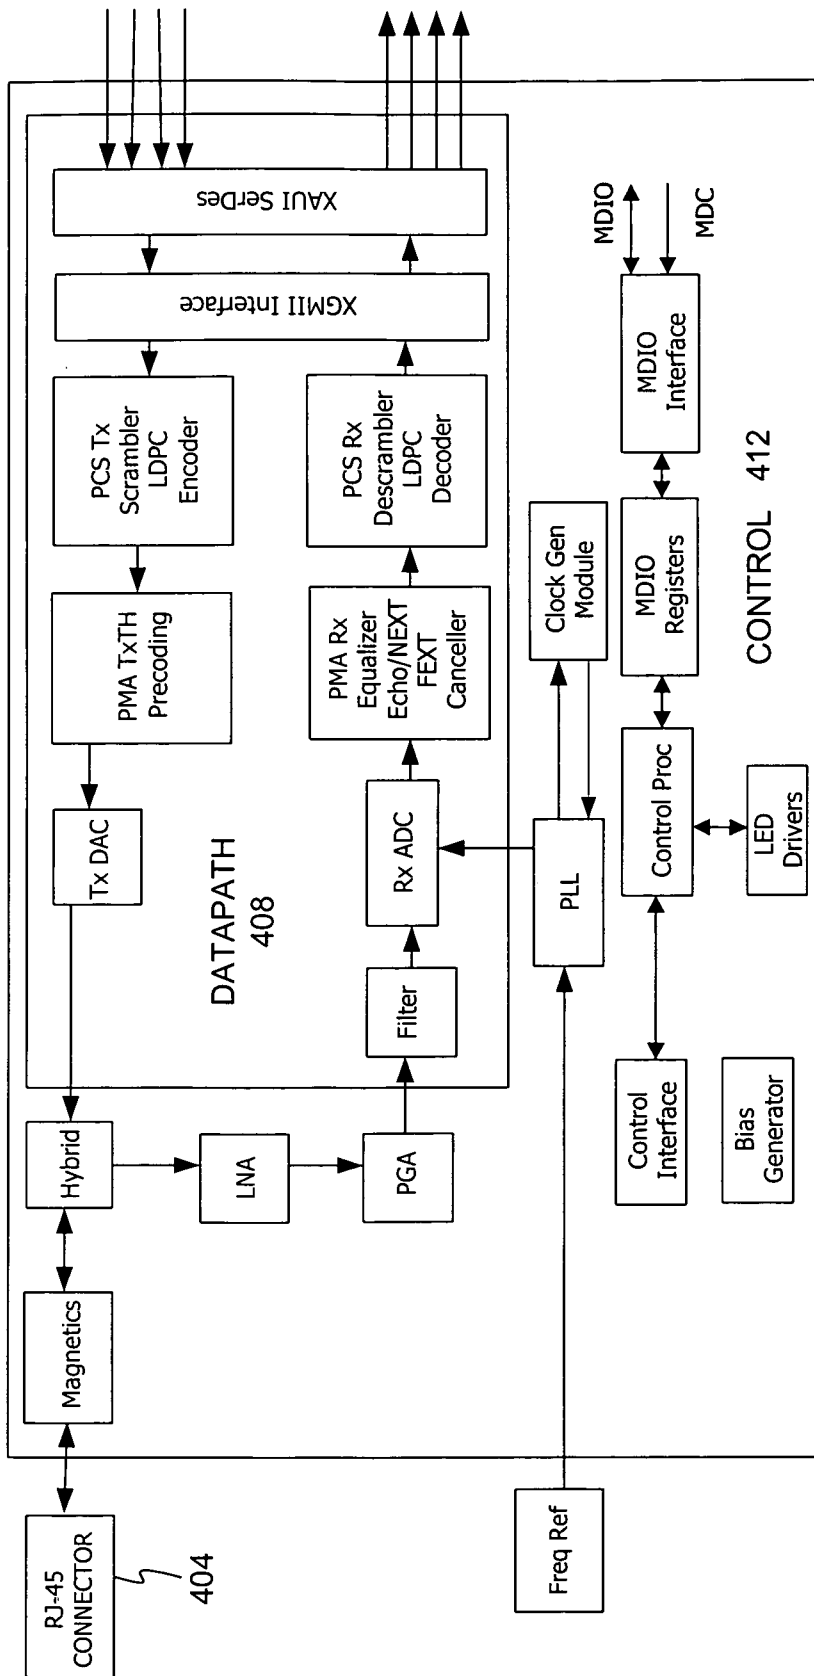


Fig. 4

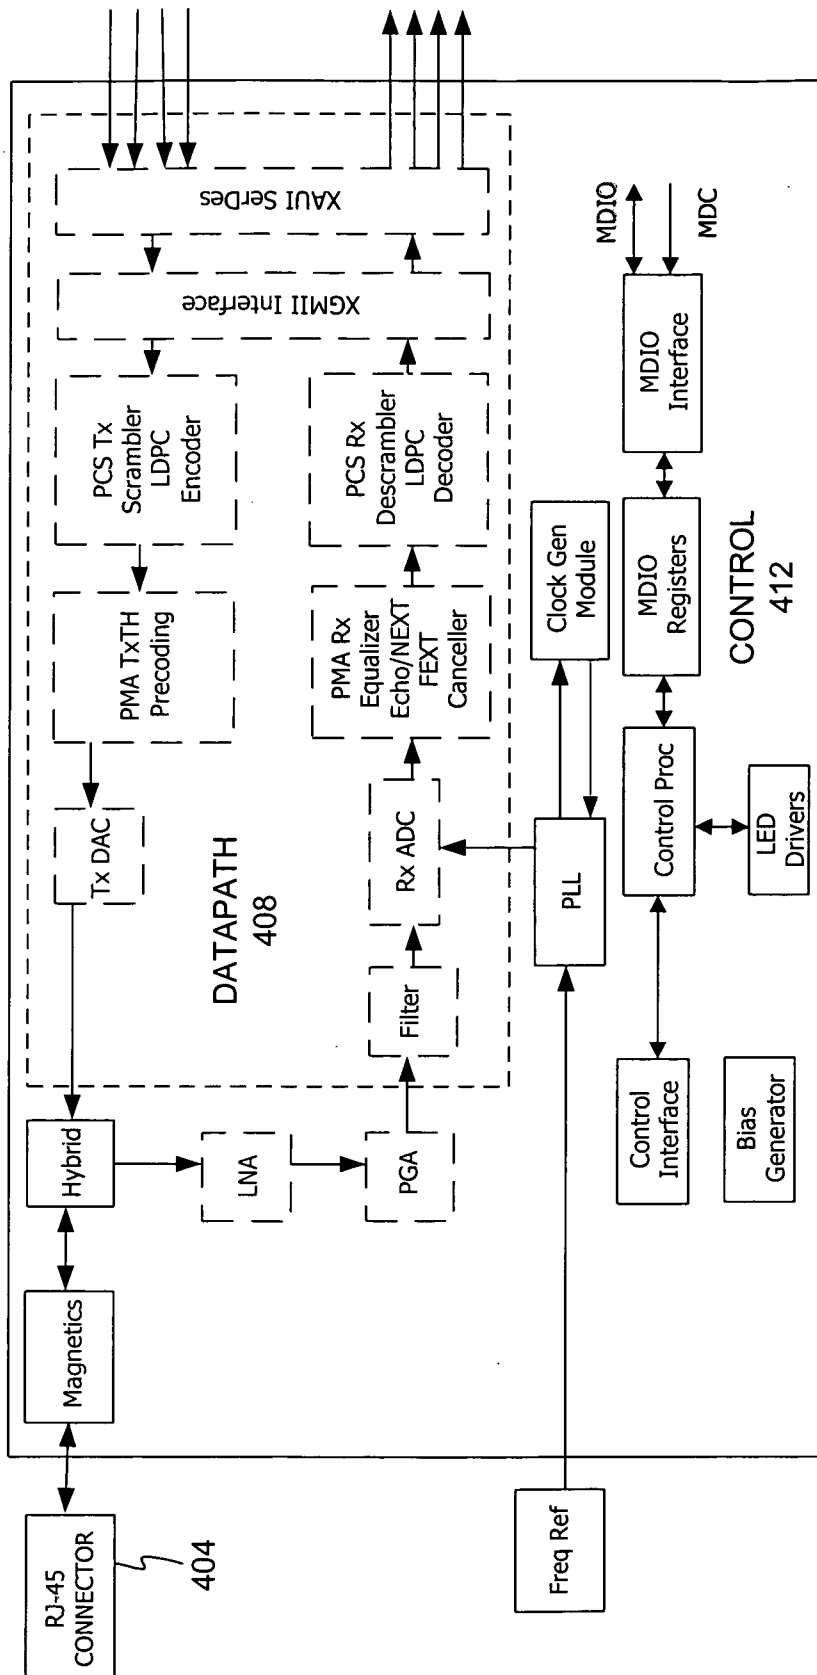


Fig. 5

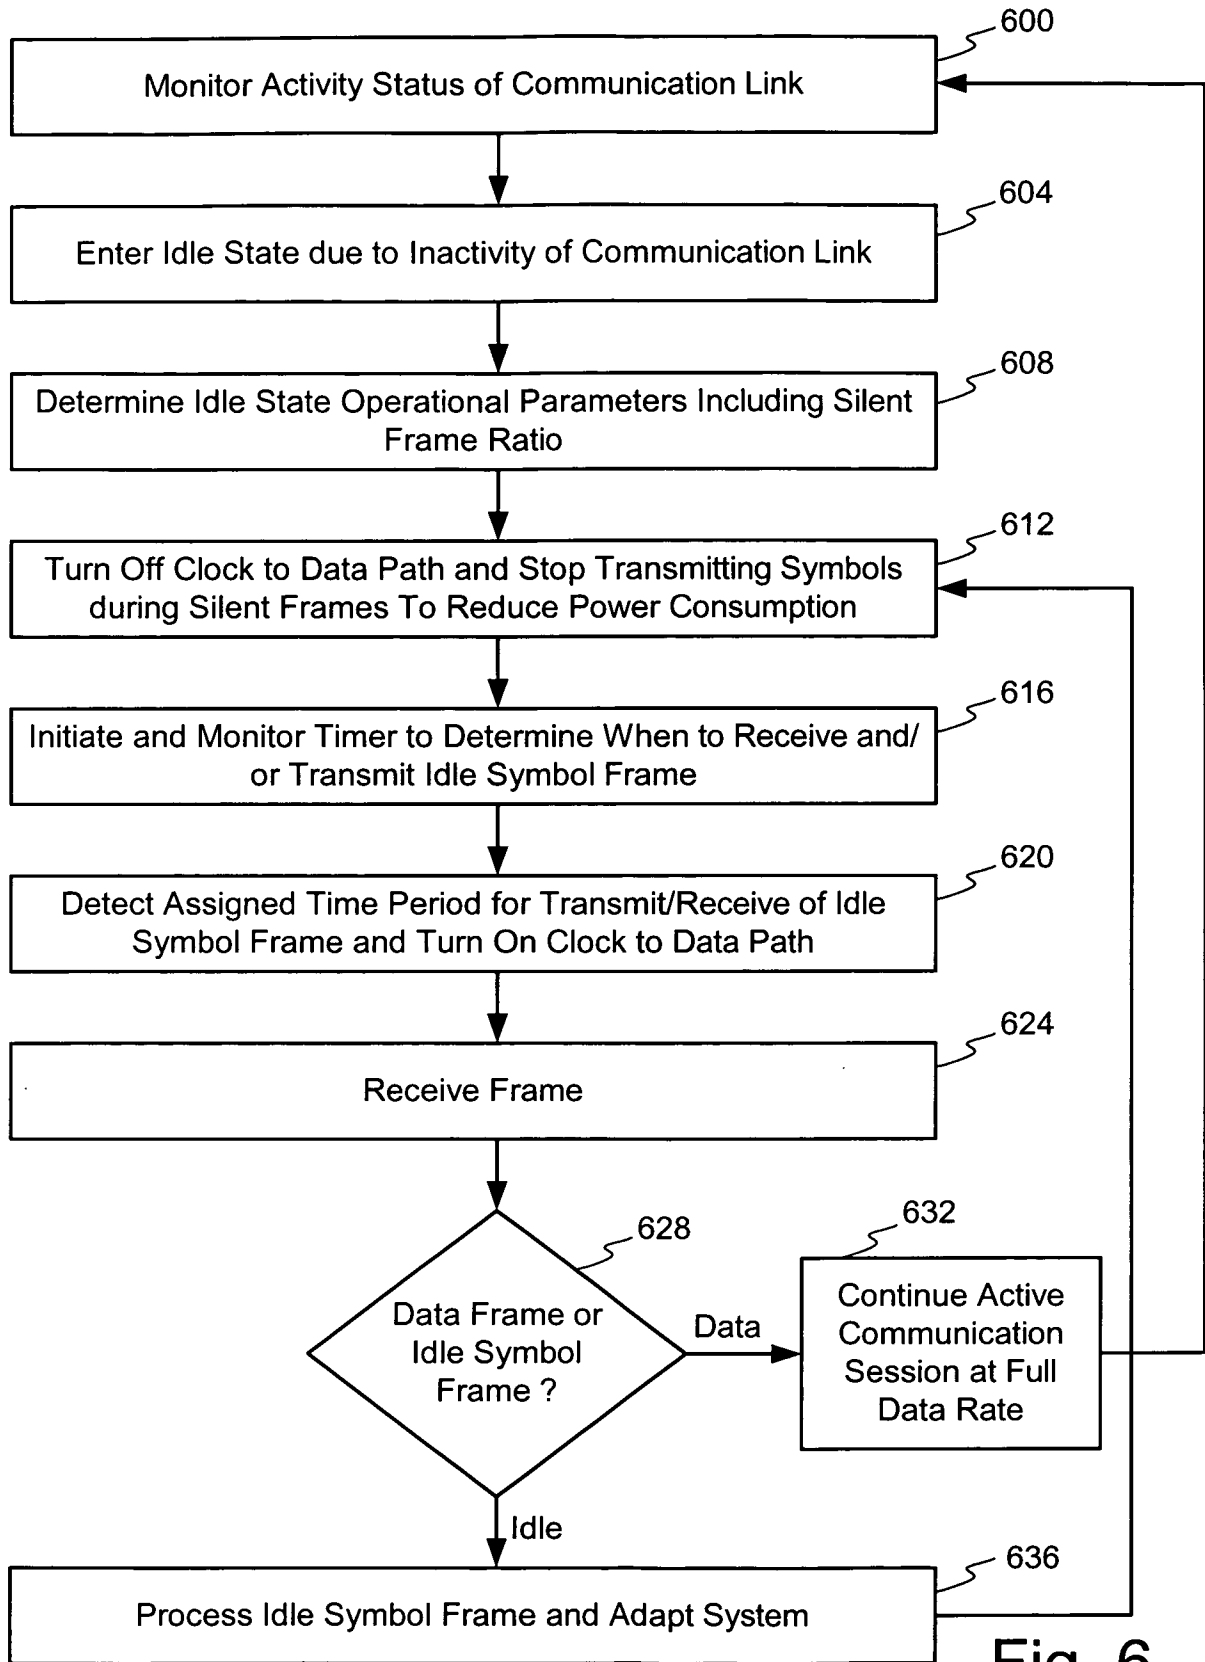


Fig. 6