



(22) Date de dépôt/Filing Date: 1999/03/08

(41) Mise à la disp. pub./Open to Public Insp.: 1999/09/16

(62) Demande originale/Original Application: 2 320 701

(30) Priorités/Priorities: 1998/03/09 (09/037,328) US;
1998/05/14 (09/078,466) US; 1998/05/14 (09/078,933) US;
1998/08/28 (09/143,476) US

(51) Cl.Int.⁷/Int.Cl.⁷ H04L 1/20, H04L 29/02


(71) Demandeur/Applicant:
BROADCOM CORPORATION, US

(72) Inventeurs/Inventors:
AGAZZI, OSCAR E., US;
CREIGH, JOHN L., US;
HATAMIAN, MEHDI, US

(74) Agent: FETHERSTONHAUGH & CO.

(54) Titre : APPAREIL ET METHODE POUR LA REDUCTION DU GASPILLAGE DE POUVOIR DANS UN SYSTEME DE COMMUNICATION

(54) Title: APPARATUS AND METHOD FOR POWER DISSIPATION REDUCTION IN A COMMUNICATION SYSTEM



(57) Abrégé/Abstract:

A method for reducing power dissipation within a communications system having a plurality of adaptive filters with a plurality of taps, each tap switchable between an active and an inactive state and each having a coefficient. The method involves specifying an acceptable error for the system, and setting a tap threshold for each active tap. For each active tap, those taps having a coefficient with an absolute value less than the tap threshold set for the active tap are deactivated. A system error is computed and compared to the acceptable system error. If the computed system error is less than the acceptable system error, the tap threshold for each active tap is increased. The above process is repeated until the computed system error approaches the acceptable system error without exceeding the acceptable system error.



ABSTRACT

A method for reducing power dissipation within a communications system
5 having a plurality of adaptive filters with a plurality of taps, each tap switchable
between an active and an inactive state and each having a coefficient. The method
involves specifying an acceptable error for the system, and setting a tap threshold
for each active tap. For each active tap, those taps having a coefficient with an
absolute value less than the tap threshold set for the active tap are deactivated. A
10 system error is computed and compared to the acceptable system error. If the
computed system error is less than the acceptable system error, the tap threshold for
each active tap is increased. The above process is repeated until the computed
system error approaches the acceptable system error without exceeding the
acceptable system error.

15

20

APPARATUS AND METHOD FOR POWER DISSIPATION IN A COMMUNICATION SYSTEM

This application is divided from Canadian Patent Application Serial Number
5 2,320,701, filed March 8, 1999.

BACKGROUND OF THE INVENTION

The invention relates to systems for, and methods of, reducing the power
dissipation in devices within a communications system and to systems for, and
10 methods of, reducing such power dissipation in communications systems having high
throughputs. A "high throughput" as used within the context of this disclosure may
include, but is not limited to, one gigabit (GB) per second.

A basic communications system is illustrated in FIG. 1. The system includes a
hub and a plurality of computers serviced by the hub in a local area network (LAN).
15 Four computers are shown by way of illustration but a different number of computers
may be contained within the system. Each of the computers is usually displaced from
the hub by a distance which may be as great as approximately one hundred meters
(100 m.). The computers are also displaced from each other. The hub is connected to
each of the computers by a communications line. Each communication line includes
20 unshielded twisted pairs of wires or cables. Generally, the wires or cables are formed
from copper. Four unshielded twisted pairs of wires are provided in each
communication line between each computer and the hub. The system shown in FIG. 1
is operative with several categories of unshielded twisted pairs of cables designated as
categories 3, 4, 6 and 7 in the telecommunications industry. Category 3 cables are the
25 poorest quality (and lowest cost) and category 6 and 7 cables are the best quality (and
highest cost).

Associated with each communications system is a "throughput". The
throughput of a system is the rate at which the system processes data and is usually
expressed in bits/second.

30

1 Most communications systems have throughputs of 10 megabits (Mb)/second or 100 Mb/second. A rapidly evolving area of communications system technology enables 1 Gb/second full-duplex communication over existing category-5 unshielded twisted pair cables. Such a system is commonly referred to as "Gigabit Ethernet."

5 A portion of a typical Gigabit Ethernet is shown in FIG. 2. The Gigabit Ethernet provides for transmission of digital signals between one of the computers and the hub and the reception of such signals at the other of the computer and the hub. A similar system can be provided for each of the computers. The system includes a gigabit medium independent interface (GMII) block which receives data in byte-wide format at a specified rate, for example 125 MHz, and
10 passes the data onto the physical coding sublayer (PCS) which performs scrambling, coding, and a variety of control functions. The PCS encodes bits from the GMII into 5-level pulse amplitude modulation (PAM) signals. The five symbol levels are -2, -1, 0, +1, and +2. Communication between the computer and hub is achieved using four unshielded twisted pairs of wires or cables, each operating at 250 Mb/second, and eight transceivers, one positioned at each end of a
15 unshielded twisted pair. The full-duplex bidirectional operation provides for the use of hybrid circuits at the two ends of each unshielded twisted pair. The hybrid controls access to the communication line, thereby allowing for full-duplex bidirectional operation between the transceivers at each end of the communications line.

A common problem associated with communications systems employing multiple
20 unshielded twisted pairs and multiple transceivers is the introduction of crosstalk and echo noise or impairment signals into the transmission signals. Noise is inherent in all such communications systems regardless of the system throughput. However, the effects of these impairment signals are magnified in Gigabit Ethernet. Impairment signals include echo, near-end crosstalk (NEXT), and far-end crosstalk (FEXT) signals. As a result of these impairment
25 signals the performance of the transceivers, particularly the receiver portion, is degraded.

NEXT is an impairment signal that results from capacitive and inductive coupling of the signals from the near-end transmitters to the input of the receivers. The NEXT impairment signals encountered by the receiver in transceiver A are shown in FIG. 3. The crosstalk signals from transmitters B, C, and D appear as noise to receiver A, which is attempting to detect the
30 direct signal from transmitter E. Each of the receivers in the system encounters the same effect and accordingly the signals passing through the receivers experience signal degradation due to NEXT impairment signals. For clarity of FIG. 3, only the NEXT impairment experienced by receiver A is illustrated.

Similarly, because of the bidirectional nature of the communications systems, an echo
35 impairment signal is produced by each transmitter on the receiver contained within the same transceiver as the transmitter. The echo impairment signal encountered by the receiver in each transceiver is shown in FIG. 4. The crosstalk signals from transmitters appear as noise to the receivers, which are attempting to detect the signal from the transmitter at the opposite end of

1 the communications line. Each of the receivers in the system encounters the same effect and accordingly the signals passing through the receivers experience signal distortion due to the echo impairment signal.

5 Far-end crosstalk (FEXT) is an impairment that results from capacitive coupling of the signal from the far-end transmitters to the input of the receivers. The FEXT impairment signals encountered by the receiver in transceiver A are shown in FIG. 5. The crosstalk signals from transmitters F, G, and H appears as noise to receiver A, which is attempting to detect the direct signal from transmitter E. Each of the receivers in the system encounters the same effect and accordingly the signals passing through the receivers experience signal distortion due to the FEXT impairment signal. For clarity of FIG. 5 only the FEXT impairment experienced by receiver A is illustrated.

10 As a result of these noise impairment signals the performance of the communication system is degraded. The signals carried by the system are distorted and the system experiences a higher signal error rate. Thus there exists a need in the art to provide a method of, and an apparatus for, compensating for the degradation of communication system performance caused by noise impairment signals and to provide a method of, and apparatus for, reducing such noise in a high throughput system such a Gigabit Ethernet. Aspects of the present invention fulfil these needs.

20 Four transceivers at one end of a communications line are illustrated in detail in FIG. 6. The components of the transceivers are shown as overlapping blocks, with each layer corresponding to one of the transceivers. The GMII, PCS, and hybrid of FIG. 6 correspond to the GMII, PCS, and hybrid of FIG. 2 and are considered to be separate from the transceiver. The combination of the transceiver and hybrid forms one "channel" of the communications system. Accordingly, FIG. 6 illustrates four channels, each of which operates in a similar manner. The transmitter portion of each transceiver includes a pulse-shaping filter and a digital-to-analog (D/A) converter. The receiver portion of each transceiver includes an analog-to-digital (A/D) converter, a first-in first-out (FIFO) buffer, a digital adaptive equalizer system including a feed-forward equalizer (FFE) and a detector. The receiver portion also includes a timing recovery system and a near-end noise reduction system including a NEXT cancellation system and an echo canceller. The NEXT cancellation system and the echo canceller typically include numerous adaptive filters.

35 Characteristics of the communication line, *e. g.*, length, may impact the ability of the NEXT cancellation system and echo cancellers to effectively cancel NEXT and echo noise. Measurements of typical cable responses, as well as simulation, show that in order to provide an adequate level of cancellation of these sources of interference, "long" echo and NEXT cancellers are required. The term "long" is used to describe a canceller having a large number of taps as necessitated by the characteristics of the cable. For example, FIG. 7 shows the echo impulse response for a 100m cable with a characteristic impedance of 85 ohm and 100 ohm terminations.

1 Although the nominal characteristic impedance is 100 ohm, manufacturing standards allow for
a 15% tolerance. The mismatch in impedance may result in a reflection at the far-end of the
cable, which causes a secondary pulse with a delay of about one microsecond. Because of the
long delay, cancelling this pulse requires an echo canceller with about 140 taps (125 taps to
5 cover the one microsecond delay, plus approximately 15 additional taps to cancel the secondary
pulse).

Often the echo impulse response has additional reflections at intermediate values of delay.
In addition, structural return loss of the cable may cause continuous variations of the
characteristic impedance along the cable, which results in a large number of smaller reflections
10 at intermediate points. These intermediate reflections mean that the echo canceller should not
be configured to cancel only the initial impulse and the end reflection but instead should be
configured to cover the full span of the impulse response. Varying cable characteristics result
in a wide variability of cable impulse responses. In addition, the response of a particular cable
may change as a result of its operating environment. For example, a change in operating
15 temperature may change the impulse response of the cable. Accordingly, it is difficult to
precompute the locations at which taps are required and to build these locations into the design
of the echo and NEXT cancellers. FIG. 8 shows the NEXT impulse response for a 100m cable.
As indicated, the NEXT response can also be long, requiring a large number of taps in the NEXT
cancellers which make up the NEXT cancellation system. The combination of the NEXT and
20 echo cancellers consumes the majority of the DSP operations in a Gigabit Ethernet.

The large number of taps required to achieve a satisfactory performance in these systems
results in high power dissipation. This high power dissipation is undesirable in that it may make
high throughput communication systems, particularly Gigabit Ethernet, inoperable and
unmarketable. Thus there exists a need in the art to provide a method of, and an apparatus for,
25 reducing the power dissipation of communication systems employing a large number of taps and
to provide a method of, and apparatus for, reducing such power dissipation in a high throughput
system such a Gigabit Ethernet. Aspects of the present invention fulfil these needs.

One of the most critical phases of the operation of a Gigabit Ethernet transceiver is the
startup. During this phase adaptive filters contained within the transceiver converge, the timing
30 recovery subsystem acquires frequency and phase synchronization, the differences in delay
among the four wire pairs are compensated, and pair identity and polarity is acquired. Successful
completion of the startup allows normal operation of the transceiver to begin.

In one startup protocol, known as "blind start", the transceivers converge their adaptive
filters and timing recovery systems simultaneously while also acquiring timing synchronization.
35 A disadvantage of such a startup is that there is a high level of interaction among the various
adaptation and acquisition algorithms within the transceiver. This high level of interaction
reduces the reliability of the convergence and synchronization operations which occur during
startup.

5

Thus there exists a need in the art to provide a startup protocol for use in a
10 high throughput communications system, such as a Gigabit Ethernet, that uses the
optimal sequence of operations and minimizes the interaction among the various
adaptation and acquisition algorithms. Aspects of the present invention fulfil these
needs.

15 SUMMARY OF THE INVENTION

Briefly, and in general terms, the invention relates to systems for, and methods
of, reducing the power dissipation within communications systems.

20

25

30

5

In a first aspect, the invention involves a method for reducing power dissipation within a communications system having a plurality of adaptive filters with a plurality of taps, each tap switchable between an active and an inactive state, each
10 having a coefficient. The method includes the steps of a) specifying an acceptable error for the system; b) for each active tap, setting a tap threshold; c) for each active tap, deactivating those taps having a coefficient with an absolute value less than the tap threshold set for the active tap; d) computing a system error; e) comparing the
15 computed system error to the acceptable system error; f) if the computed system error is less than the acceptable system error, increasing the tap threshold for each active tap; and g) repeating steps c) through f) until the computed system error approaches the acceptable system error without exceeding the acceptable system error.

20

25

30

5

10

In a second facet, the invention involves a method for reducing power dissipation within a communications system having at least one adaptive filter with a plurality of taps , each tap switchable between an active and an inactive state, each tap having a coefficient. The method includes the steps of a) computing an initial system error; b) for each active tap, setting a tap en-or threshold; c) for each active tap, deactivating those taps having a coefficient with an absolute value less than the tap error threshold set for the active tap; d) computing a subsequent system error, e) if the difference between the subsequent system error and the initial system error is less than a prespecified value, increasing the tap error threshold for each active tap; and f) repeating steps c) through e) until the difference between the subsequent system error and the initial system error exceeds the prespecified value.

25

30

5

BRIEF DESCRIPTION OF THE DRAWINGS

10

FIGURE I is a schematic block diagram of a communications system providing a plurality of computers connected to a hub by communications lines to form a local area network (LAN);

FIG. 2 is a schematic block diagram of a communications system providing a
15 gigabit medium independent interface (GMII), a physical coding sublayer (PCS) and a plurality of unshielded twisted pairs of wires, each with a transceiver at each end;

20

25

30

5

FIG. 3 is a schematic block diagram of a portion of the communications system of FIG. 2 depicting the NEXT impairment signals received by receiver A from adjacent transmitters B, C, and D;

10

FIG. 4 is a schematic block diagram of a portion of the communications system of FIG. 2 depicting the echo impairment signal received by receiver A from transmitter A;

15

FIG. 5 is a schematic block diagram of a portion of the communications system of FIG. 2 depicting the FEXT impairment signals received by receiver A from opposite transmitters F, G and H;

20

25

30

5

10 FIG. 6 is a schematic block diagram of a communications system including a plurality of transceivers, each having a NEXT cancellation system, an echo canceller, a feed forward equalizer, digital adaptive filter system including one detector, and a timing recovery circuit;

15 FIG. 7 depicts an impulse response for an echo signal passing through a lOOm communications line;

 FIG. 8 depicts an impulse response for an NEXT signal passing through a lOOm communications line;

20 FIG. 9 is a schematic block diagram of a communications system including a plurality of transceivers each having a NEXT cancellation system, an echo canceller, and a FEXT cancellation system, a digital adaptive filter system including a plurality of detectors and a skew adjuster, and a timing recovery circuit;

25

30

WO 99/4680 /

1 FIG. 10 is a schematic block diagram of a symbol-by-symbol detector of FIG. 9, each including a plurality of slicers, feedback filters and adders and receiving as input a soft decision;

FIG. 11 is a schematic block diagram of the NEXT cancellation systems of FIG. 9, each including a plurality of adaptive transversal filters (ATF) and adders and receiving as input
5 transmitted signals from adjacent transmitters;

FIG. 12 is a schematic block diagram of the echo cancellers of FIG. 9, each including an ATF and receiving as input transmitted signals from same transmitters;

FIG. 13 is a schematic block diagram of the FEXT cancellation systems of FIG. 9, each including a plurality of ATFs and an adder and receiving as input transmitted signals from
10 opposite transmitters;

FIG. 14 depicts a direct impulse response arriving at the receiver after a FEXT impulse response;

FIG. 15 depicts a direct impulse response and FEXT impulse response arriving at the receiver at substantially the same time;

15 FIG. 16 depicts a direct impulse response arriving at the receiver before a FEXT impulse response;

FIG. 17 is a schematic block diagram of a communications system in including a plurality of transceivers, each having a NEXT cancellation system, an echo canceller, and a FEXT cancellation system, digital adaptive filter system including one detector, and a timing recovery
20 circuit;

FIG. 18 is a schematic diagram of an ATF, including a cascade of taps, present in the NEXT cancellation system of FIG. 11, the echo cancellers of FIG. 12, and the FEXT cancellation system of FIG. 13;

25 FIG. 19 is a flow chart illustrating one embodiment of a power dissipation reduction method;

FIG. 20 depicts the mean squared error (MSE) to signal ratio as a function of time during the initial convergence of a communications system;

FIG. 21 depicts the taps of an echo canceller which remain active after convergence of a communications system having an error threshold of 24 dB;

30 FIG. 22 depicts the taps of an echo canceller which remain active after convergence of a communications system having an error threshold of 26 dB;

FIG. 23 is a flow chart illustrating another embodiment of a power dissipation reduction method;

35 FIG. 24 is a schematic block diagram depicting the master-slave relationship between the transceivers of each of the transceiver channels of FIG. 2; and

FIG. 25 is a timing diagram depicting the stages of one embodiment of a startup protocol; and

1 FIG. 26 is a timing diagram depicting the stages of another embodiment of a startup protocol.

5 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

5 The discussion in this specification may be considered to relate specifically to a Gigabit Ethernet for the purposes of explanation and understanding of the invention. However, it will be understood that the concepts of this invention and the scope of the claims apply to other types of communications systems than a Gigabit Ethernet.

10 Communications System Overview:

10 A communications system incorporating the features of this invention is generally indicated at 10 in FIG. 1. The system 10 includes a hub 12 and a plurality of computers serviced by the hub in a local area network (LAN). Four computers 14 are shown by way of illustration but a different number of computers may be used without departing from the scope of the invention. Each of the computers 14 may be displaced from the hub 12 by a distance as great as approximately one hundred meters (100 m.). The computers 14 are also displaced from each other.

15 The hub 12 is connected to each of the computers 14 by a communications line 16. The communication line 16 comprises a plurality of unshielded twisted pairs of wires or cables. Generally, the wires or cables are formed from copper. Four unshielded twisted pairs of wires are provided in the system 10 between each computer and the hub 12. The system shown in FIG. 1 is operative with several categories of twisted pairs of cables designated as categories 3, 4, 5, 6 and 7 in the telecommunications industry. Category 3 cables are the poorest quality (and lowest cost) and category 6 and 7 cables are the best quality (and highest cost). Gigabit Ethernet uses category 5 cables.

20 FIG. 2 illustrates, in detail, a portion of the communications system of FIG. 1 including one communications line 16 and portions of one of the computers 14 and the hub 12. The communications line 16 includes four unshielded twisted pairs of wires 18 operating at 250 Mb/second per pair. A transceiver 20, including a transmitter (TX) 22 and receiver (RX) 24, is positioned at each unshielded end of each twisted pair 18. Between each transceiver 20 and its associated unshielded twisted pair 18 is a hybrid 26. The hybrid 26 is the interface to the communication line 16 which allows for full-duplex bidirectional operation between the transceivers 20 at each end of the communications line. The hybrid also functions to isolate the transmitter and receiver associated with the transceiver, from each other.

25 The communications system includes a standard connector designated as a gigabit media independent interface (GMII) 28. The GMII 28 may be an eight bit wide data path in both the transmit and receive directions. Clocked at a suitable frequency, such as 125 MHz, the GMII results in a net throughput in both directions of data at a suitable rate such as 250 Mb/second per

1 pair. The GMII provides a symmetrical interface in both the transmit and receive directions. A
 physical coding sublayer (PCS) 30 receives and transmits data between the GMII 28 and the
 transceivers 20. The PCS 30 performs such functions as scrambling and encoding/decoding data
 before forwarding the data to either the transceiver or the GMII. The PCS encodes bits from the
 5 GMII into 5-level pulse amplitude modulation (PAM) signals. The five symbol levels are -2, -1,
 0, +1, and +2. The PCS also controls several functions of the transceivers, such as skew control
 as explained below.

Transceiver Circuitry:

10 Four of the transceivers 20 are illustrated in detail in FIG. 9. The components of the
 transceivers 20 are shown as overlapping blocks, with each layer corresponding to one of the
 transceivers. The GMII 28, PCS 30, and hybrid 26 of FIG. 9 correspond to the GMII, PCS, and
 hybrid of FIG. 2 and are considered to be separate from the transceiver. The combination of the
 transceiver 20 and hybrid 26 forms one "channel" of the communications system. Accordingly,
 15 FIG. 9 illustrates four channels, each of which operate in a similar manner.

The transmitter portion of each transceiver 20 includes a pulse shaping filter 32 and a
 digital to analog (D/A) converter 34. In a preferred embodiment of the invention the D/A
 converter 34 operates at 125 MHz. The pulse shaping filter 32 receives one one-dimensional (1-
 D) symbol from the PCS. This symbol is referred to as a TXData_x symbol 36, where x is 1
 20 through 4 corresponding to each of the four channels. The TXData_x symbol 36 represents 2 bits
 of data. The PCS generates one 1-D symbol for each of the channels. The symbol for each
 channel goes through a spectrum shaping filter of the form $0.75 + 0.25z^{-1}$ at the pulse shaping
 filter 32 to limit emissions within FCC requirements. This simple filter shapes the spectrum at
 the output of the transmitter so that its power spectral density falls under that of communications
 25 systems operating at 100 Mb/second on two pairs of category -5 twisted pair wires. The symbol
 is then converted into an analog signal by the D/A converter 34 which also acts as a lowpass
 filter. The analog signal gains access to the unshielded twisted pair wire 18 through the hybrid
 circuitry 26.

The receiver portion of each transceiver includes a signal detector 41, an A/D converter
 30 42, a FIFO 44, a digital adaptive equalizer system, a timing recovery circuit and noise reduction
 circuitry. The digital adaptive equalizer system includes a feed-forward equalizer (FFE) 46, two
 devices 50, 56, a skew adjuster 54 and two detectors 58, 60. The functions of these components,
 as related to the present invention, are explained below. The noise reduction circuitry includes
 a NEXT cancellation system 38, an echo canceller 40, and a FEXT cancellation system 70.

35 The A/D converter 42 provides digital conversions of the signals received from the hybrid
 26 at a suitable frequency, such as 125 MHz, which is equal to the baud rate of the signals. The
 A/D converter 42 samples the analog signals in accordance with an analog sample clock signal
 78 provided by the decision-directed timing recovery circuit 64. The FIFO 44 receives the

1 digital conversion signals from the A/D converter 42 and stores them on a first-in-first-out basis.
 The FIFO 44 forwards individual signals to the FFE 46 in accordance with a digital sample clock
 signal 80 provided by the timing recovery circuit 64. The FFE 46 receives digital signals from
 the FIFO 44 and filters these signals. The FFE 46 is a least mean squares (LMS) type adaptive
 5 filter which performs channel equalization and precursor inter symbol interference (ISI)
 cancellation to correct for distortions in the signal.

It is noted that the signal introduced into the A/D converter 42 and subsequently into the
 FIFO 44 and FFE 46 has several components. These components include the direct signal
 received directly from the transmitter 22 at the opposite end of the unshielded twisted pair wire
 10 18 with which the receiver 24 is associated. Also included are one or more of the NEXT, echo,
 and FEXT impairment signals from other transmitters 22 as previously described. The signal
 including the direct signal and one or more of the impairment signals is referred to as a
 "combination signal."

The FFE 46 forwards the combination signal 48 to a second device 50, typically a
 15 summing device. At the second device 50 the combination signal 48 is combined with the
 outputs of the NEXT cancellation system 38 and echo canceller 40 to produce a signal which is
 substantially devoid of NEXT and echo impairment signals. This signal is referred to as a "first
 soft decision" 52. The signal detector 41 detects the signals from the second device 50 and
 forwards the signals to the skew adjuster 54. Upon signal detection, the signal detector 41
 20 initiates various system operations, one of which -- as described below -- includes transitioning
 between phases of a startup protocol. The skew adjuster 54 receives the first soft decision 52
 from the second device 50 and outputs a signal referred to as a "second soft decision" 66. The
 skew adjuster 54 performs two functions. First, it compensates for the difference in length of the
 unshielded twisted pairs 18 by delaying the first soft decision 52 so that the second soft decisions
 25 66 from all of the receivers in the system are in sync. Second, it adjusts the delay of the first soft
 decision 52 so that the second soft decision 66 arrives at the first device 56 at substantially the
 same time as the output of the FEXT cancellation system 70. The skew adjuster 54 receives
 skew control signals 82 from the PCS 30.

The skew adjuster 54 forwards the second soft decision 66 to a first device 56, typically
 30 a summing device. At the first device 56 the second soft decision 66 is combined with the output
 of the FEXT cancellation system 70 to produce a signal which is substantially devoid of FEXT
 impairment signals. This signal is referred to as a "third soft decision" 68. The first detector
 58 receives the third soft decision 68 from the first device 56. The first detector 58 provides an
 output signal, *i. e.*, a "final decision" 72. The first detector 58 may be a slicer which produces
 35 a final decision 72 corresponding to the analog signal level closest in magnitude to the level of
 the third soft decision 68. The first detector 58 may also be either a symbol-by-symbol detector
 or a sequential detector which operates on sequences of signals across all four channels
 simultaneously, such as a Viterbi decoder.

1 In one configuration of the transceiver the first detector 58 is a symbol-by-symbol
 detector. A group of symbol-by-symbol detectors 58, one for each channel, is shown in FIG. 10.
 Each first detector 58 includes a slicer 98, adaptive feedback filter 100 and an adder 102. The
 adder 102 combines the third soft decision 68 with the output of the adaptive feedback filter 100
 5 to provide an output which is introduced to the slicer 98. The output of the slicer 98 is
 introduced to the adaptive feedback filter 100. The first detector 58 provides an output signal
 72 which corresponds to the discrete level from the set $[-2, -1, 0, 1, 2]$ which is closest to the
 difference between the third soft decision 68 and the output of the feedback filter 100. The
 adaptive feedback filter 100 corrects for distortion in the third soft decision 68. This filter 100
 10 uses past slicer 98 decisions to estimate postcursor ISI caused by the channel. This ISI is
 canceled from the third soft decision 68 to form the final decision signal 72.

In another configuration of the transceiver the first detector 58 is a combination of a
 sequential decoder with a decision feedback equalizer (DFE) using the architecture usually
 known as multiple DFE architecture (MDFE) sequential detector. The sequential decoder 58
 15 looks at all signals from all four channels at the same time and at successive samples from each
 channel over several periods of unit time. A sequential decoder receives as input at least one
 signal from each of the first devices 56. The sequential decoder 58, in general, is responsive to
 the sequences of the output signals from the first devices 56 for (1) passing acceptable sequences
 of such signals and (2) discarding unacceptable sequences of such signals in accordance with the
 20 constraints established by the code standard associated with the system. Acceptable sequences
 are those which obey the code constraints and unacceptable sequences are those which violate
 the code constraints.

The second detector 60 (FIG. 9) receives the first soft decision 52 from the second device
 50. The second detector 60 is a symbol-by-symbol detector similar to the first detector 58 (FIG.
 25 10). It provides an output signal 74 which corresponds to the discrete level from the set $[-2, -1, 0, 1, 2]$
 which is closest to the difference between the first soft decision 52 and the output of the
 feedback filter 100. The second detector 60 produces output signals 74 without the benefit of
 FEXT cancellation, as a result, these decisions have a higher error rate than those made by the
 first detector 58, which enjoys the benefits of FEXT cancellation. Because of this fact, these
 30 decisions are called "tentative decisions". It is important to note that the postcursor ISI present
 in the input to the second detector 60 is canceled using the adaptive feedback filter 100, (FIG.
 10) contained within the second detector, whose inputs are the tentative decisions 74. The
 coefficients of this adaptive feedback filter 100 are the same as those of the adaptive feedback
 filter associated with the first detector 58 (FIG. 9).

35 A third device 62, typically a summing device, receives the first soft decision signal 52
 from the second device 50 and the tentative decision signals 74 from the second detector 60.
 At the third device 62 the first soft decision 52 is combined with the tentative decision signal 74
 to produce an error signal 76 which is introduced into the timing recovery circuit 64. The timing

1 recovery circuit 64 receives the tentative decision 74 from the second detector 60 and the error
 signals 76 from the third device 62. Using these signals as inputs the timing recovery circuit 64
 outputs an analog clock sync signal 78 which is introduced to the A/D converter 42 and a digital
 clock sync signal 80 which is introduced into the FIFO 44. As previously mentioned, these
 5 signals control the rate at which the A/D converter 42 samples the analog input it receives from
 the hybrid 26 and the rate at which the FIFO forwards digital signals to the FFE 46.

As mentioned before, the symbols sent by the transmitters 22 (FIG. 2) in the
 communications system cause NEXT, echo and FEXT impairments in the received signal for
 each channel. Since each receiver 24 has access to the data for the other three channels that
 10 cause this interference, it is possible to nearly cancel each of these effects. NEXT cancellation
 is accomplished using three adaptive NEXT cancelling filters 84 as shown in the block diagram
 of FIG. 11. Each NEXT cancellation system 38 receives three TXData symbols 36 from each
 of the transmitters at the same end of the communications line 18 as the receiver with which the
 NEXT cancellation system is associated. Each NEXT cancellation system 38 includes three
 15 filters 84, one for each of the TXData symbols 36. These filters 84 model the impulse
 responses of the NEXT noise from the transmitters and may be implemented as adaptive
 transversal filters (ATF) employing, for example, the LMS algorithm. The filters 84 produce a
 replica of the NEXT impairment signal for each TXData symbol 36. A summing device 86
 combines the three individual replica NEXT impairment signals 92 to produce a replica of the
 20 NEXT impairment signal contained within the combination signal received by the receiver with
 which the NEXT cancellation system 38 is associated. The replica NEXT impairment signal 88
 is introduced into the second device 50 (FIG. 9) where it is combined with the combination
 signal 48 to produce a first soft decision signal 52 which is substantially devoid of NEXT
 impairment signals.

25 Echo cancellation is accomplished with an adaptive echo cancelling filter 85 as shown in
 the block diagram of FIG. 12. Each echo canceller 40 receives the TXData symbols 36 from
 the transmitter at the same end of the twisted wire pair 18 as that of the receiver with which the
 echo canceller is associated. As shown in FIG. 12, each echo canceller 40 includes one filter
 85. These filters 85 model the impulse responses of the echo noise from the transmitter and may
 30 be implemented as ATFs employing, for example, the LMS algorithm. The filter produces a
 replica of the echo impairment signal contained within the combination signal received by the
 receiver with which the echo canceller 40 is associated. The replica echo impairment signal 90
 is introduced into the second device 50 (FIG. 9) where it is combined with the combination
 signal 48 to produce a first soft decision signal 52 which is substantially devoid of echo
 35 impairment signals.

FEXT cancellation is accomplished with three adaptive FEXT cancelling filters 87 as
 shown in the block diagram of FIG. 13. Each FEXT cancellation system 70 receives three
 tentative decision symbols 74, one from each of the receivers at the same end of the

communications line as the receiver with which the FEXT cancellation system is associated. Each FEXT cancellation system 70 includes three filters 87, one for each of the tentative decision symbols 74. These filters 87 model the impulse responses of the FEXT noise from transmitters and may be implemented as ATFs employing, for example, the LMS algorithm. The filters 87 produce a replica of the FEXT impairment signal 96 for each individual tentative decision symbol 74. A summing device 108 combines the three individual replica FEXT impairment signals 96 to produce a replica of the FEXT impairment signal contained within the combination signal 48 received by the receiver with which the FEXT cancellation system is associated. The replica FEXT impairment signal 94 is introduced into the first device 56 (FIG. 9) where it is combined with the second combination signal 66 to produce the third soft decision signal 68 which is substantially devoid of FEXT impairment signals. It is important to note that the higher error rate of the tentative decisions 74 does not degrade the performance of the FEXT cancellation system 70, because the decisions used to cancel FEXT are statistically independent from the final decisions 72 made by the receiver whose FEXT is being canceled.

The symbols provided by the first detector 58 are decoded and descrambled by the receive section of the PCS 30 before being introduced to the GMII. Variations in the way the wire pairs are twisted may cause delays through the four channels by up to 50 nanoseconds. As a result, the symbols across the four channels may be out of sync. As previously mentioned, in the case where the first detector is a sequential detector, the PCS also determines the relative skew of the four streams of 1-D symbols and adjusts the symbol delay, through the skew adjuster 54, prior to their arrival at the first detector 58 so that sequential decoder can operate on properly composed four-dimensional (4-D) symbols. Additionally, since the cabling plant may introduce wire swaps within a pair and pair swaps among the four unshielded twisted pairs, the PCS 30 also determines and corrects for these conditions.

Noise Reduction:

As previously mentioned, FEXT is an impairment that results from capacitive coupling of the signal from the far-end transmitters to the input of the receivers, as shown in FIG 5. The crosstalk signals from transmitters F, G, and H appear as noise to receiver A, which is attempting to detect the signal from transmitter E. A similar situation applies to all other receivers regarding the signals from the appropriate transmitters located at the opposite end of the line.

The FEXT noise experienced by receiver A and originating from transmitter F can be modeled as the convolution of the data symbols transmitted by F with a certain impulse response that depends on the properties of the cable and models the coupling characteristics of the unshielded twisted pairs used by transmitter F and receiver A. A typical measured FEXT impulse response 104 is shown in FIGS 14-16. A similar description can be given for all the other possible receiver-transmitter combinations. Therefore, there are a total of twelve FEXT impulse

1 responses 104 describing the FEXT noise signals from transmitters E, F, G, and H to receivers
A, B, C, and D. These twelve impulse responses are not identical, although each has a general
shape similar to that shown in FIGS 14-16.

5 Although FEXT is an impairment for many communications systems other than Gigabit
Ethernet, in these systems a given receiver usually does not have access to the symbols detected
by the other receivers, because these receivers may not be physically located in the same place,
and/or because they operate at rates that are not synchronized to the data rate of the receiver
suffering from FEXT. Aspects of the present invention takes advantage of the fact that in
10 Gigabit Ethernet transceivers the decisions that correspond to all four channels are available to
the four receivers and the decisions may be made synchronous.

In operation there may be delays associated with the transmission of signals across the
communications line. The synchronization of the signals within the system is crucial to effective
cancellation of the noise. It is important that the replica noise impairment signals arrive at the
summing devices at substantially the same time as the combination signal and/or soft decision
15 signals. With regard to the FEXT impairment signal, because the impairment is caused by the
transmitters at the opposite end of the receiver there is likely to be a delay between the time that
the second soft decision signal 66 arrives at the first device 56 and the time at which the replica
FEXT impairment signal 94 arrives. In some channels, as illustrated in FIG. 14, the group delay
of the FEXT signal 104 could be smaller than the group delay of the desired signal 106. In this
20 case the tentative decisions 74 provided by receivers B, C, and D of FIG. 5 arrive at the FEXT
cancellation system 70 of receiver A too late to be effective in canceling the FEXT impairment.

To compensate for this delay, the invention employs a skew adjuster 54 which, as
previously stated, delays the first soft decision signal 52 by a time substantially equal to or
greater than the time delay between the arrival at the receiver of the direct signal and the FEXT
25 impairment signals associated with such receiver. If the output is delayed by an amount greater
than the time delay, which would result in the situation illustrated in FIG. 16, the adaptive
feedback filter 87 (FIG. 13) within the FEXT cancellation system 70 compensates for the over
delay by delaying the replica FEXT impairment signal 94 so that it arrives at the first device 56
(FIG. 9) at substantially the same time as the second soft decision signal 66.

30 The third soft decision 68 resulting from FEXT cancellation allows the first detector 58
to make more reliable final decisions 72, with a greatly reduced error rate. Computer simulations
show that a typical improvement achievable with the invention described herein is approximately
2 to 3dB when the signal to noise ratio at the input of the first detector 58 before FEXT
cancellation is approximately 25dB. This corresponds to a reduction of the symbol error rate of
35 a factor 1000 or larger.

If there is no delay associated with the transmission of signals across the communications
line both the FEXT impairment signal 104 and the direct signal 106 arrive at the receiver at the

1 substantially the same time, as shown in FIG. 15. In this situation, the delay of the skew adjuster
 54 is set to zero. In the alternative, an embodiment of the invention, as shown in FIG. 17, with
 only one detector 110 and one summing device 112 may be used. In this configuration, the
 summing device 112 receives the replica NEXT, echo, and FEXT impairment signals 88, 90, 94
 5 and the combination signal 48 and produces a first soft decision 52 substantially devoid of
 impairment signals. This first soft decision 52 is introduced into the detector 110 and the third
 device 62. The detector 110 may include either a single symbol-by-symbol detector or both a
 symbol-by-symbol detector and a sequential detector. In the case of a symbol-by-symbol
 detector the final decision 72 and second output 114 of the detector 110 are identical. In the
 10 case of both a symbol-by-symbol detector and a sequential detector, the final output is provided
 by the sequential detector and is introduced to the PCS 30. The second output 114 is provided
 by the symbol-by-symbol detector and is introduced to the timing recovery circuitry 64 and the
 third device 62 for use in determining the error signal 76.

15 Power Dissipation Reduction:

As previously mentioned, the NEXT cancellation system, echo canceller and FEXT
 cancellation system use ATFs to effectively cancel the noise from the combination signal. An
 example of an ATF which may be employed is shown in FIG. 18. The ATF 120 includes a
 plurality of taps 122 each including a multiplier 124 and an adder 126. Associated with each tap
 20 122 is a coefficient C_n , where n is 0 through $x-1$ where x is the number of taps in the ATF. The
 circuitry associated with each tap 122 includes a 1-bit storage (not shown) that allows for
 activation and deactivation of the tap. The values of the coefficients C_n are adjusted in
 accordance with an LMS algorithm as mentioned before. Interposed between the taps 122 are
 registers 128. These registers 128 provide data to the taps 122 at timed intervals in accordance
 25 with a clock signal.

The impulse responses of an echo and NEXT, as shown in FIGS. 7 and 8, indicates that
 not all taps 122 in the NEXT and echo cancellers 38, 40 are contributing significantly to the
 performance of the communications system. Aspects of the present invention determines what
 taps 122 are not contributing significantly to the reduction of the mean squared error (MSE) of
 30 the system and deactivates these taps, thereby eliminating them from the filtering computation
 and thus reducing considerably the power dissipation of the system. Furthermore, as shown by
 the impulse response of FIGS. 7 and 8, the need to build NEXT and echo cancellers 38, 40 with
 a long span is difficult to avoid. Specific cable responses may differ from the one depicted in
 FIGS. 7 and 8, and accordingly require more or fewer taps 122 than that required for the cable
 35 of FIGS. 7 and 8. As mentioned before it is difficult to determine a priori what taps 122 are
 needed with a particular cable.

1 In accordance with aspects of the present invention, the NEXT , echo, and FEXT
 cancellers 38, 40, 70 are configured with ATFs 120 which employ a sufficient number of taps
 122 to provide adequate cancellation with the worst-case expected impulse responses. This may
 5 require 140 taps 122 as in the example of FIG. 7, or even more for longer cables. To reduce the
 power dissipation, the taps 122 are examined after convergence, and those taps that are found
 not to contribute significantly to the performance of the system are deactivated. When the tap
 122 is deactivated, it is removed from the NEXT, echo and FEXT replica computation and from
 the adaptation and its contribution to the overall power dissipation of the system is substantially
 eliminated.

10 When the system is initially converged, all taps 122 are active, so the NEXT, echo and
 FEXT cancellers 38, 40, 70 are converged along their entire length. After convergence, the taps
 122 are examined to determine which ones can be deactivated using the tap scanning algorithm
 depicted in FIG. 19. At step S1, an acceptable level of error for the system S_{ea} is specified. At
 step S2, a tap coefficient threshold T_{th} is set for each active tap. While each individual tap may
 15 have a unique tap threshold T_{th} , in a preferred embodiment of the invention the tap thresholds
 for all taps are substantially the same. The initial value of the tap coefficient threshold T_{th}
 is sufficiently low such that only a few taps 122 are deactivated and the performance of the system
 is not significantly affected. In a preferred embodiment of the invention, the tap coefficient
 threshold T_{th} is initially set equal to the tap coefficient C_n having the minimum absolute value.
 20 Alternatively, a reasonable value can be determined by simulation. This initial value is not
 critical, as long as it is sufficiently low to avoid a large degradation of the performance of the
 system the first time the tap scanning procedure is applied.

At step S3, the absolute value of the tap coefficient C_n for each active tap is compared
 to the tap coefficient threshold T_{th} . If the tap coefficient C_n is less than the tap coefficient
 25 threshold T_{th} the tap 122 is deactivated at step S4. This process is repeated for each tap 122
 in the filter 120. Preferably, the determination of whether to deactivate a tap 122 is done in a
 sequential manner starting at the input end of the filter 120. At step S5, the error for the system
 S_{ec} is computed. This error is computed by first computing the MSE for each active tap 122
 by multiplying the absolute value of the tap coefficient C_n by the average energy signal. The
 30 error of the filter 120 associated with the tap 122 is determined by summing the individual tap
 errors. The error of the system is then determined by summing the individual filter errors.

In step S6, the computed system error S_{ec} is compared to the specified acceptable system
 error S_{ea} . If the computed system error S_{ec} is less than the acceptable system error S_{ea} , the tap
 threshold T_{th} for each active tap is increased by a small amount at step S7 and steps S3 through
 35 S6 are repeated. As a result, some additional taps 122 are deactivated, but the number of taps
 deactivated is usually not very large because of the small increase in the tap threshold T_{th} .
 Steps S3 through S6 are repeated until the computed system error S_{ec} approaches the

1 acceptable system error S_{ea} without exceeding the acceptable system error. If the computed system error S_{ec} is greater than the acceptable system error S_{ea} , the tap scanning algorithm stops.

5 As an alternative to determining whether to deactivate a tap based on the MSE of the communications system, a determination may be made based on the MSE of an individual filter. In this embodiment of the invention an acceptable level of error for a filter F_{ea} is specified. Individual taps are deactivated, as previously described, and a computed filter error F_{ec} is calculated. If the deactivation of a tap does not cause the computed filter error F_{ec} to exceed the acceptable filter error F_{ea} the tap remains inactive.

10 In yet another embodiment of the invention, the contribution of each deactivated tap to the MSE of the filter and, in turn, the system is calculated. If the MSE contribution of the tap is determined to be an acceptable amount the tap remains deactivated. This method is generally preferred during the initial startup of the system when the overall MSE of the system is large due to the nonconverged state of the filter coefficients within the system. Once the filter coefficients of the system have initially converged, the determination of whether to deactivate a tap is generally made based on the taps contribution to the MSE of the system as previously described with reference to FIG. 19.

20 The final result of the tap scanning algorithm is that, in typical channels of the communication system, a large number of taps 122 is deactivated, and power dissipation is reduced by a large factor. As an example, computer simulations of the tap scanning algorithm when operating on the channel whose echo response is shown in FIG. 7 are presented in FIG. 20. This figure shows both the master and slave MSE to signal ratios as a function of time during the initial convergence of the system. At time $t=360,000$ bauds, the tap scanning algorithm begins, and as a result the MSE to signal ratio increases to the prespecified target of 24dB. FIG. 21 shows the taps of the echo canceller after convergence with a threshold of 24 dB and FIG. 22 shows the taps with a threshold of 26 dB. The deactivated taps are shown as zeros. In FIG. 21, the total number of active taps for the echo canceller is twenty-two. Similarly, the number of active taps 122 for the three NEXT cancellers (not shown) forming the NEXT cancellation system is six, two, and zero, respectively. In FIG. 22, the total number of active taps for the echo canceller is forty-seven. Similarly, the number of active taps 122 for the three NEXT cancellers (not shown) is six, two, and zero, respectively.

30 For the case of the 24 dB threshold, out of 440 initially active taps, only 30 remain active after the application of the tap scanning algorithm, while maintaining a 5dB margin for required bit error rate. Notice from FIGS. 21 and 22 that those taps 122 which remain active occur at sparse locations, and it would have been difficult to statically allocate these taps during the design of the NEXT and echo cancellers, because the location of taps is highly dependent on the specific cable response.

1 In another embodiment of the invention, the tap scanning algorithm monitors the change
in MSE as the tap scanning algorithm progresses. The algorithm is applied until the change in
MSE -- rather than the MSE itself -- exceeds a prespecified value, for example 1dB. If the
MSE to signal ratio before the scanning is applied is 25dB, the final MSE to signal ratio is
5 24dB, and a large number of taps is deactivated. This embodiment of the tap scanning
algorithm is depicted in FIG. 23. In step S10, the initial system error is S_{ei} computed. In step
S11, an acceptable system error differential D_e is specified. In step S12, a tap coefficient
threshold T_{th} is set for each active tap. As with the other tap scanning algorithm, the initial
value of the tap threshold T_{th} is sufficiently low such that only a few taps 122 are deactivated
10 and the performance of the system is not significantly affected.

At step S13, the absolute value of the tap coefficient C_n , for each active tap is
compared to the tap coefficient threshold T_{th} . If the tap coefficient C_n is less than the tap
coefficient threshold T_{th} the tap 122 is deactivated at step S14. This process is repeated for
each tap 122 in the filter 120. Preferably, the determination of whether to deactivate a tap 122
15 is done in a sequential manner starting at the input end of the filter 120. At step S15, the
subsequent error for the system S_{es} is computed.

In step S16, the subsequent error for the system S_{es} is compared to the to the initial
system error is S_{ei} . If the difference between the subsequent system error S_{es} and the initial
system error S_{ei} is less than a prespecified value, the tap threshold T_{th} is increased by a small
20 amount at step S17 and steps S13 through S16 are repeated. As a result, some additional taps
122 are deactivated. Steps S13 through S16 are repeated until the difference between the
subsequent system error S_{es} and the initial system error S_{ei} exceeds the prespecified value.

In some cases the impulse responses of NEXT, echo and FEXT may change during
normal operation, for example as a result of temperature changes. It is therefore desirable to
25 periodically activate previously deactivated taps 122, preferably in a sequential manner, and
recheck if the absolute value of the tap coefficient C_n is below the tap threshold T_{th} . If a tap
coefficient C_n has grown to a value above the tap threshold T_{th} , the tap 122 remains active,
otherwise it is deactivated. Similarly, those taps 122 that were active may fall below the tap
threshold T_{th} in which case they are deactivated. All this can be accomplished with a periodic
30 reapplication of the sequential tap scanning algorithm during normal operation.

In an alternate embodiment of the invention a select number of taps 122, for example
ten, positioned at the input end of the filter 120 are not subject to deactivation. Usually there
is a large slew rate in these first few taps 122, which means that their numerical value could
change significantly if the sampling phase changes. This sampling phase could change
35 dynamically as a result of jitter, causing some previously deactivated taps 122 to become
significant. By maintaining a number of taps 122 at the input end of the filter in an active state,
potential degradations in the presence of jitter are avoided.

1 When the total number of taps is very large, for example 440, the power
dissipation could be large during the initial convergence transient before the tap
scanning algorithm has had a chance to deactivate the taps. Although the average
power dissipation of the system is still greatly reduced, the peak power is not. A
5 preferred embodiment of the invention compensates for this by converging the NEXT,
echo and FEXT cancellers in stages. For example, a block of 20 taps is converged at
a time, and the tap scanning algorithm is then applied to these taps on a per-block
basis. In cases where the MSE during the initial convergence is large, for example as
10 a result of the fact that the initial block of 20 taps may not be large enough to provide
a lower MSE, it may be better to monitor the sum of the squared values of the
coefficients of deactivated taps 122 as a measure of whether the algorithm can be
terminated.

Start-Up Protocols:

15 One of the most critical phases of the operation of the communications system
is the transceiver startup. During this phase the adaptive filters contained within the
FFE 46 (FIG. 9), echo canceller 40, NEXT cancellation system 38, FEXT cancellation
system 70, timing recovery system 64 and detector 58 of the receiver portion of each
transceiver converge. During convergence the actual output of the adaptive filters are
20 compared to expected output of the filters to determine the error. The error is
reduced to substantially zero by adjusting the coefficients of the algorithm which
defines the transfer function of the filter. Similarly, the timing recovery system is
converged by adjusting the frequency and phase of the phase lock loop and the local
oscillator contained within the timing recovery system so that the signal-to-noise ratio
25 of the channel is optimized. In addition, the differences in delay among the four wire
pairs are compensated, and pair identity and polarity, are acquired. Successful
completion of the startup ensures that the transceiver can begin normal operation.

 In accordance with aspects of the present invention each of the transceiver
channels operate in a loop-timed fashion, as shown in FIG. 24. The transceivers 20
30 at the two ends of the each twisted wire pair 18 assume two different roles as far as
synchronization is concerned. One of the transceivers, called the master 130, transmits
data using an independent clock GTX_CLK provided through the GMII interface 28
(FIG. 9). This clock signal is fixed in both frequency and phase and is provided to
the master transceiver 130 of each the four transceiver channels in the communications
35 system. In actuality the transmit clock used by the master 130 may be a filtered
version of GTX_CLK, obtained using a phase locked loop with a very narrow
bandwidth, to reduce jitter. The transceiver 20 at the other end of the twisted wire

1 pair 18, called the slave 132, synchronizes both the frequency and phase of its receive
 and transmit clocks to the signal received from the master 130, using the timing
 recovery system 64 (FIG. 9) located in the receiver 24. The slave 132 transmit clock
 maintains a fixed phase relationship with the slave receive clock at all times. The
 5 receive clock at the master 130 synchronizes, in phase but not in frequency, with the
 signal received from the slave transmitter 22. Thus, after an initial acquisition period,
 the master 130 receive clock follows the master transmit clock with a phase difference
 determined by the round trip delay of the loop. This phase relationship may vary
 dynamically as a result of the need of the master 130 receive clock to track jitter
 10 present in the signal received from the slave 132.

Slave Transmit Clock Fixed:

The sequence of events during the startup protocol of the present invention is
 shown in FIG. 25. The protocol consists of three phases 134, 136, 138 during which
 15 the receivers are trained, *e. g.*, adaptive filters are converged, timing synchronization
 is acquired, etc., followed by normal operation which begins during phase four 140.
 During the first phase 134, the master begins transmitting to the slave using a transmit
 clock signal that is fixed in both frequency and phase. The master trains its near-end
 noise reduction system by converging the adaptive filters contained within its echo
 20 canceller and NEXT cancellation system (E). At the same time, the slave trains its
 equalizers and far-end noise reduction system by converging the adaptive filters
 contained within its DFE, FFE and FEXT cancellation system (D). While training its
 equalizer and far-end noise reduction system the slave simultaneously acquires timing
 synchronization in both frequency and phase (T). It may also at this time compensate
 25 for the differential delay among the four twisted wire pairs, identify the four pairs, and
 correct the polarity of the pairs.

In one embodiment of the protocol, the transition from the first phase 134 to
 the second phase 136, at both master and slave, occurs after a fixed and prespecified
 period of time. In a preferred embodiment, however, the slave transitions from the
 30 first phase 134 to second phase 136 when it detects that its receiver has converged the
 adaptive filters contained within its DFE, FFE and FEXT cancellation system (D) and
 has acquired timing synchronization (T). As previously mentioned, the master
 receiver includes a signal detector 41 (FIG. 9) which detects energy in the line coming
 from the slave. The master transitions from the first phase 134 to the second phase
 35 136 when it detects this energy from the slave. Therefore, the slave takes the initiative
 in transitioning from the first phase 134 to the second phase 136, and the master
 follows when it detects the signal from the slave.

1 The convergence of the echo canceller and NEXT cancellation system during
the first phase 134 at the master is done with the objective of allowing the signal
detector at the master to detect the signal from the slave. Without proper echo and
NEXT cancellation, the signal detector would be triggered by the echo and NEXT
5 noise present in the receiver. After the transition has occurred, the master discards the
echo canceller and NEXT cancellation system coefficients which result from the
converging in the first phase 134. This may be done by resetting the adaptive filters
in the echo canceller and NEXT cancellation system. It is important to note that the
correct sampling phases for the four receivers at the master is obtained during the third
10 phase 138, therefore the echo canceller and NEXT cancellation system coefficients
obtained during the first phase 134 may differ from the final values to be reacquired
in the third phase 138.

During the second phase 136, the slave trains its near-end noise reduction system
by converging the adaptive filters contained within its echo canceller and NEXT
15 cancellation system (E). At the same time, the master trains its equalizers and far-end
noise reduction system by converging the adaptive filters contained within its DFE,
FFE and FEXT cancellation system (D). While training its equalizers and far-end
noise reduction system, the master simultaneously acquires timing synchronization in
phase only (P). The master may also at this time compensate for the differential delay
20 among the four twisted wire pairs, identify the four pairs, and correct the polarity of
the pairs. During the second phase, the slave saves the timing recovery state variables
that had been acquired during the first phase 134, and freezes its frequency and phase.
By doing this, the slave is guaranteed to sample with the correct phase, the signal
coming to it from the master when the master resumes transmission at the beginning
25 of the third phase 138. The slave also freezes the coefficients of the DFE, FFE and
FEXT cancellation system acquired during the first phase 134.

Similar to the transition from the first phase 134 to the second phase 136, the
transition from the second phase 136 to the third phase 138 may occur after a fixed
and prespecified period of time. While the duration of the first, second, and third
30 phases 134, 136, 138 is fixed, the duration is not necessarily equal for all phases. In
a preferred embodiment, however, the master transitions from the second phase 136
to third phase 138 when it detects that its receiver has converged the adaptive filters
contained within its DFE, FFE and FEXT cancellation system (D) and has acquired
timing synchronization (P). Like the master, the slave receiver includes a signal
35 detector 41 (FIG. 9) which detects energy in the line coming from the master. The
slave transitions from the second phase 136 to the third phase 138 when it detects this
energy from the master. Therefore the master takes the initiative in transitioning from

1 the second phase 136 to the third phase 138, and the slave follows when it detects the signal from the master.

During the third phase 138 the slave freezes the coefficients of the echo cancellers and NEXT cancellation system and maintains a steady state condition during which the operating characteristics of the slave are not adjusted. Similarly, the master freezes the coefficients of the DFE, FFE and FEXT cancellation system and the phase of its clock signal. The master also retrain its near-end noise reduction system by reconverging its echo canceller and NEXT cancellation system (E) during the third phase 138. It is important to note that in the third phase 138 the slave resumes transmission using the clock recovered from the signal transmitted by the master, and therefore the master already knows the correct frequency with which to operate its receiver. The "relative sampling phases" of the four receivers, *i. e.*, the differences in sampling phases of three of the receivers versus one of them arbitrarily used as reference, are also known, because they were acquired during the second phase 136. However, the "overall sampling phase" of the receivers, *i. e.*, the sampling phase of the receiver arbitrarily chosen as reference, is not yet known and has to be acquired during the third phase 138. When both master and slave have completed their training operations, they exchange messages indicating that they are ready to transmit valid data. During phase four 140, all coefficients of the adaptive filters previously frozen are unfrozen and the transmission of data is ready to take place.

Slave Transmit Clock Free-Running:

The sequence of events during the startup protocol of the present invention is shown in FIG. 26. The protocol consists of three phases 144, 146, 148 during which the receivers are trained, *e. g.*, adaptive filters are converged, timing synchronization is acquired, etc., followed by normal operation which begins during phase four 150. The startup protocol is initiated by the master when it starts transmitting a signal to the slave using a transmit clock signal which is fixed in both frequency and phase. During the first phase 144, the master trains its near-end noise reduction system by converging the adaptive filters contained within its echo canceller and NEXT cancellation system (E). As previously mentioned, the slave receiver includes a signal detector 41 (FIG. 9) which detects energy in the line coming from the master. Upon detection of the signal from the master, the slave trains its equalizers and far-end noise reduction system by converging the adaptive filters contained within its DFE, FFE and FEXT cancellation system (D). While training its equalizers and far-end noise reduction system the slave simultaneously acquires timing synchronization in both frequency and phase (T). It may also at this time compensate for the differential

1 delay among the four twisted wire pairs, identify the four pairs, and correct the polarity of the pairs.

5 In one embodiment of the protocol, the transition from the first phase 144 to the second phase 146, at both master and slave, occurs after a fixed and prespecified period of time. In a preferred embodiment, however, the slave transitions from the first phase 144 to second phase 146 when it detects that its receiver has converged the adaptive filters contained within its DFE, FFE and FEXT cancellation system (D) and has acquired timing synchronization (T). The master receiver also includes a signal detector 41 (FIG. 9) which detects energy in the line coming from the slave. The master transitions from the first phase 144 to the second phase 146 when it detects this energy from the slave. Therefore, the slave takes the initiative in transitioning from the first phase 144 to the second phase 146, and the master follows when it detects the signal from the slave.

15 The convergence of the echo canceller and NEXT cancellation system during the first phase 144 at the master is done with the objective of allowing the signal detector at the master to detect the signal from the slave. Without proper echo and NEXT cancellation, the signal detector would be triggered by the echo and NEXT noise present in the receiver. After the transition has occurred, the master discards the echo canceller and NEXT cancellation system coefficients which result from the converging in the first phase 144. This may be done by resetting the adaptive filters in the echo canceller and NEXT cancellation system. It is important to note that the correct sampling phases for the four receivers at the master is obtained again in the third phase 148, therefore the echo canceller and NEXT cancellation system coefficients obtained during the first phase 144 may differ from the final values to be reacquired in the third phase 148.

25 During the second phase 146, the slave trains its near-end noise reduction system by converging the adaptive filters contained within its echo canceller and NEXT cancellation system (E). The slave also freezes the coefficients of the DFE, FFE and FEXT cancellation system acquired during the first phase 144. At the same time, the master trains its equalizers and far-end noise reduction system by converging the adaptive filters contained within its DFE, FFE and FEXT cancellation system (D). While training its equalizers and far-end noise reduction system, the master simultaneously acquires timing synchronization in both frequency and phase (T). The master may also, at this time, compensate for the differential delay among the four twisted wire pairs, identify the four pairs, and correct the polarity of the pairs. The slave may run off of any stable clock that has a frequency offset of less than a prespecified limit, e. g., for example 200ppm, when compared to the master transmit

1 clock. During the second phase 146, the slave transceivers transmit using a free-
 running clock. As a result, the master acquires both frequency and phase
 synchronization during the second phase 146. Phase synchronization is a normal
 5 function that the master performs in any form of startup protocol, however, frequency
 synchronization is not a usual master function, and it is only performed during the
 second phase 146 of the startup with the objective of interoperating properly with
 slave transceivers that do not save the timing recovery state variables at the transition
 from the first phase 144 to the second phase 146.

10 Similar to the transition from the first phase 144 to the second phase 146, the
 transition from the second phase 146 to the third phase 148 may occur after a fixed
 and prespecified period of time. While the duration of the first, second, and third
 phases 144, 146, 148 is fixed, the duration is not necessarily equal for all phases. In
 a preferred embodiment, however, the master transitions from the second phase 146
 15 to third phase 148 when it detects that its receiver has converged the adaptive filters
 contained within its DFE, FFE and FEXT cancellation system (D) and has acquired
 timing synchronization (T). The master begins transmitting a signal to the slave. The
 slave transitions from the second phase 146 to the third phase 148 when it detects this
 signal from the master. Therefore the master takes the initiative in transitioning from
 the second phase 144 to the third phase 148, and the slave follows when it detects the
 20 signal from the master.

During the third phase 148, the slave freezes the coefficients of the echo
 canceller and the NEXT cancellation system and maintains its near-end noise reduction
 system in a steady state condition. Because the timing recovery state variables had
 not been saved at the transition to the second phase 146, the slave also reacquires
 25 timing synchronization in both frequency and phase (T) during the third phase 148.

The master, during the third phase 148, freezes the coefficients of its DFE and FEXT
 cancellation system and the frequency of its clock signal. The master also retrains its
 near-end noise reduction system by reconverging its echo canceller and NEXT
 cancellation system (E). The master also reacquires timing synchronization in phase
 30 only (P). It is important to note that in the third phase 148 the slave resumes
 transmission using the clock recovered from the signal transmitted by the master, and
 therefore the master already knows the correct frequency with which to operate its
 receiver. The "relative sampling phases" of the four receivers, *i. e.*, the differences in
 sampling phases of three of the receivers versus one of them arbitrarily used as
 35 reference, are also known, because they were acquired during the second phase 146.
 However, the "overall sampling phase" of the receivers, *i. e.*, the sampling phase of the
 receiver arbitrarily chosen as reference, is not yet known and has to be acquired

1 during the third phase 148. When both master and slave have completed their training
operations, they exchange messages indicating that they are ready to transmit valid
data. During phase four 150, all coefficients of the adaptive filters previously frozen
are unfrozen and the transmission of data is ready to take place.

5 Although this invention has been disclosed and illustrated with reference to
particular embodiments, the principles involved are susceptible for use in numerous
other embodiments which will be apparent to persons of ordinary skill in the art.
The invention is, therefore, to be limited only as indicated by the scope of the
appended claims.

10

15

20

25

30

35

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

- 5 1. A method for reducing power dissipation within a communications system having a plurality of adaptive filters with a plurality of taps, each tap switchable between an active and an inactive state, each having a coefficient, said method comprising the steps of:
- 10 a) specifying an acceptable error for the system;
- b) for each active tap, setting a tap threshold;
- c) for each active tap, deactivating those taps having a coefficient with an absolute value less than the tap threshold set for the active tap;
- d) computing a system error;
- e) comparing the computed system error to the acceptable system error;
- 15 f) if the computed system error is less than the acceptable system error, increasing the tap threshold for each active tap; and
- g) repeating steps c) through 1) until the computed system error approaches the acceptable system error without exceeding the acceptable system error.

20

2. The method of claim 1 wherein the determination of whether to deactivate a tap is done in a sequential manner starting at the input end of the filter.

- 25 3. The method of claim 1 further comprising the steps of:
periodically activating previously deactivated taps;
repeating steps b) through g).

4. The method of claim 3 wherein the previously deactivated taps are activated in a sequential manner starting at the input end of the filter,

30

5. The method of claim 1 wherein the step of computing the system error comprises the steps of:

- 5 determining the error of each individual filter; and
summing the individual filter errors.

6. The method of claim 5 wherein the step of computing the filter error comprises the steps of:

- 10 determining the error of each individual tap; and
summing the individual tap errors.

7. The method of claim 6 wherein the error of each individual tap is the mean square error for the tap and is determined by multiplying the absolute value of
15 the tap coefficient by the average energy signal.

8. The method of claim 1 further comprising the steps of:
specifying an acceptable level of error for the taps;
for each active tap, calculating the error of deactivating the tap; and if the
20 calculated error
is less than the acceptable error, deactivating the tap.

9. The method of claim 8 wherein said steps are performed prior to steps a)
through g).
25

10. A method for reducing power dissipation within a communications system having at least one adaptive filter with a plurality of taps , each tap switchable between an active and an inactive state, each tap having a coefficient, said method comprising the steps of:

- 30 a) computing an initial system error

- b) for each active tap, setting a tap error threshold;
- 5 c) for each active tap, deactivating those taps having a coefficient with an absolute value less than the tap error threshold set for the active tap;
- d) computing a subsequent system error;
- e) if the difference between the subsequent system error and the initial system error is less than a prespecified value, increasing the tap error threshold for
- 10 each active tap; and
- f) repeating steps c) through e) until the difference between the subsequent system error and the initial system error exceeds the prespecified value.

11. The method of claim 10 wherein the determination of whether to
 15 deactivate a tap is done in a sequential manner starting at the input end of the filter.


12. The method of claim 10 wherein each tap threshold is initially set equal to the tap coefficient having the minimum absolute value.

20 13. The method of claim 10 wherein the tap error threshold is the substantially the same for each tap in the filter.


25

30

FIG. 1
PRIOR ART



2/23






FIG. 3

4/23




FIG. 4

5/23







FIG. 5


6/23

FIG. 6
PRIOR ART

WO 99/46867


PCT/US99/05145





9/23

FIG. 10



10/23




FIG. 11

11/23






FIG. 12

12/23



13/23

FIG. 14




WO 99/46867

PCT/US99/05145

14/23

FIG. 15



15/23

FIG. 16






FIG. 17




17/23



18/23

FIG. 19



19/23




FIG. 20

WO 99/46867

PCT/US99/05145

20/23

FIG. 21






FIG. 22



21/23

FIG. 23








FIG. 25



E = NEAR-END NOISE REDUCTION SYSTEM CONVERGENCE
D = EQUALIZER/FAR-END NOISE REDUCTION SYSTEM CONVERGENCE
P = ACQUIRE TIMING (PHASE ONLY)
T = ACQUIRE TIMING (FREQUENCY AND PHASE)

FIG. 26



E = NEAR-END NOISE REDUCTION SYSTEM CONVERGENCE
D = EQUALIZER/FAR-END NOISE REDUCTION SYSTEM CONVERGENCE
P = ACQUIRE TIMING (PHASE ONLY)
T = ACQUIRE TIMING (FREQUENCY AND PHASE)

