

Apparatus and Method For Adapting an Echo Cancellor in a Communication System

5 Background

1. Field of the Invention

The invention relates generally to communication systems, and more particularly to a variable-length echo canceller for a modem.

10 2. Discussion of Related Art

In today's information age, the number of personal computers used in homes, schools, and businesses continues to proliferate with apparently no end in sight. This increasing use of personal computers has prompted the widespread deployment of computer networks such as the Internet which provide an ever-increasing variety of on-line services. While computer and communication technology has advanced greatly over time, it is still very common for users to connect their personal computers to such computer networks over telephone lines using a modem.

20 Historically, the modem has been an add-on feature to the personal computer, either as an external peripheral connected by a cable to the personal computer or as an internal peripheral installed directly into an expansion slot in the personal computer. These modems typically included all of the hardware and software
25 components required to provide modem functionality for the personal computer. Specifically, the modem typically included a microprocessor (and associated memory circuits) for executing general software functions, input/output ports for data and control interfaces, a signal processor for executing computationally
30 intensive signal processing functions, and many special-purpose

hardware components for interfacing with the telephone network and for providing additional signal processing functions. These dedicated hardware components made the modem relatively expensive.

Today, a need has arisen for a less expensive modem that can be built into the personal computer by the manufacturer or packaged together with other external peripherals to provide modem functionality for the personal computer. One solution is to implement predominantly all of the modem processing functions in software (with the remaining functions implemented in specialized hardware) and to run the software as an application on the microprocessor in the personal computer. This type of modem, often referred to as a "software modem," is feasible today due to the amount of processing resources provided by modern microprocessors. By eliminating most of the dedicated hardware components from the modem and utilizing the processing and memory resources of the personal computer, the software modem is relatively inexpensive.

One issue with the software modem is that the modem software must share the processing resources of the personal computer with other application software such as a word processor, spreadsheet program, or Internet browser. This "feature" of the software modem, while making the software modem inexpensive, consumes processing resources that otherwise would be available to the other application software. As a result, the software modem may noticeably affect the performance of the other application software. Thus, one goal of the software modem is to provide full modem functionality using as little processing resources as possible.

In a software modem, the echo cancellers which are used to compensate for signal reflections in the telephone network represent a substantial proportion of the total amount of processing resources

required by the modem. Two independent echo cancellers are typically employed, one for cancelling a near-end echo and one for cancelling a far-end echo. These echo cancellers are typically designed and implemented to compensate for a worst-case (i.e.,
5 maximum span) echo. Because the amount of processing resources needed to implement an echo canceller in software is directly proportional to the number of taps in the echo canceller filter, which in turn is directly proportional to the span over which the echo canceller works, an echo canceller that is designed and implemented
10 to compensate for the worst-case echo necessarily consumes a maximum amount of processing resources. Thus, a need remains for a system, device, and method for reducing the number of taps in the echo canceller filter to reduce the amount of processing resources consumed by the echo canceller.

15

Brief Description of the Drawing

In the Drawing,

FIG. 1 shows an analog-to-analog modem connection as is known in the art;

20 FIG. 2 shows a digital-to-analog modem connection as is known in the art;

FIG. 3 shows an exemplary relationship between the near-end and far-end echoes;

25 FIG. 4 shows an exemplary embodiment of a symbol-driven echo canceller;

FIG. 5 shows the relationship between the bulk-delay line, the near-end echo canceller, and the far-end echo canceller;

FIG. 6A shows a conceptual view of an echo canceller filter having X filter taps;

FIG. 6B shows a preferred embodiment of a near-end or far-end echo canceller implemented using three subfilters;

FIG. 7 shows generally the relevant steps of the prior art V.34 startup sequence;

5 FIG. 8A shows an exemplary positioning of the near-end and far-end echo canceller filters upon initialization;

FIG. 8B shows positioning of an echo canceller filter upon initialization in more detail;

10 FIG. 9 shows an exemplary positioning of the near-end and far-end echo canceller filters in which the far-end filter span overlaps the near-end filter span;

FIG. 10 shows an exemplary positioning of the near-end and far-end echo canceller filters in which the far-end filter span is truncated; and

15 FIG. 11 shows an exemplary embodiment including a number of steps which are taken during the startup sequence in order to adapt the echo canceller filters to the reflections in the network

Detailed Description

20 As discussed above, the need remains for an apparatus and method for reducing the number of taps in the echo canceller filter in order to reduce the amount of processing resources consumed by the echo canceller. Embodiments of the present invention eliminate the near-end or far-end echo canceller completely under certain
25 conditions, and also reduce the number of taps in the near-end and/or far-end echo canceller filters under yet other conditions. As a result, these embodiments satisfy the need for reducing the amount of processing resources consumed by the echo cancellers.

FIG. 1 shows a modem-to-modem connection as is known in the art in which modems 102 and 108 are connected to the telephone network over traditional analog local loops 103 and 107, respectively. Both modems send and receive analog signals. The analog signals sent by modem 102 over analog local loop 103 are digitized at interface 104 for transport across the digital backbone network 110 and converted back to analog signals at interface 106 for transport over analog local loop 107 to modem 108. Likewise, the analog signals sent by modem 108 over analog local loop 107 are digitized at interface 106 for transport across the digital backbone network 110 and converted back to analog signals at interface 104 for transport over analog local loop 103 to modem 102.

FIG. 2 shows a modem-to-modem connection as is known in the art in which modem 208 is connected to the telephone network over a traditional analog local loop 207 and modem 202 is connected to the telephone network by way of a digital connection 203. Modem 202 sends and receives digital signals, while modem 208 sends and receives analog signals. The digital signals sent by modem 202 over digital connection 203 are transported across the digital backbone network 210 and converted to analog signals at interface 206 for transport over analog local loop 207 to modem 208. The analog signals sent by modem 208 over analog local loop 207 are digitized at interface 206 for transport across the digital backbone network 210 and the digital connection 203 to modem 202.

Communications over the telephone network are affected by various types of distortions in both the analog and digital portions of the network. Of particular relevance to the present invention are signal reflections which echo the transmitted signal back to the transmitting modem. These signal reflections are induced at various

points within the network, specifically at the hybrid interfaces within the modems and within the telephone central office equipment, and arrive at the receiver of the transmitting modem at varying times based on the distance between the transmitting modem and each reflection point. Reflections generated within the local loop are received sooner, and with greater signal power, than reflections generated at the far end of the network. For convenience, reflections generated within the local loop are referred to as "near-end reflections" or "near-end echoes," and reflections generated at the far end of the network are referred to as "far-end reflections" or "far-end echoes."

An exemplary relationship between the near-end and far-end echoes is shown in FIG. 3. As discussed previously, the near-end echoes generated within the local loop are received sooner, and with greater signal power, than the far-end echoes generated at the far end of the network. The peak of the near-end echo is substantially synchronized with the receipt of the echo produced from the hybrid in the transmitting modem, and therefore the peak of the near-end echo occurs after a fixed, internal delay (referred to as "sys_delay"). The sys_delay is typically a constant value, and can be determined by placing the transmitting modem in a local analog loopback and measuring the loopback delay. Likewise, the peak of the far-end echo is substantially synchronized with the receipt of the echo produced from the hybrid in the remote modem, and therefore the peak of the far-end echo is based on the round-trip delay between the transmitting modem and the remote modem. The round-trip delay is typically measured during an initial startup sequence, and is described in more detail below.

The span of each echo is relatively short, on the order of 10 to 20 milliseconds, and the far-end echo is delayed by an amount B which is based on the round-trip delay between the transmitting modem and the remote reflection points. If the round-trip delay is relatively long, as shown in FIG. 3, then there will be substantially no reflective energy between the near-end echo and the far-end echo. However, if the round-trip delay is relatively short, then the far-end echo may overlap the near-end echo either partially or completely.

As can be seen in FIGs. 1 and 2, the types of reflections that occur in the network depends on the topology of the network. In an analog-to-analog connection as shown in FIG. 1, each modem receives both near-end and far-end reflections. The near-end reflections are produced at the transmitting modem's hybrid interface and at the central office equipment on the local side of the network, while the far-end reflections are produced at the central office equipment on the remote side of the network and at the remote modem's hybrid interface. In a digital-to-analog connection as shown in FIG. 2, however, the digital side modem receives only far-end reflections, while the analog side modem receives only near-end reflections.

With respect to a particular modem, reflections of the modem's transmitted signals have the effect of distorting the signals received by the modem. Specifically, the signals received by the modem include signals from a remote modem combined with the reflected signals from the modem's own transmissions. Unless the reflected signals are removed from the received signal or are otherwise compensated for by the modem, incorrect information may be extracted from the received signals resulting in corrupted data.

Accordingly, the modem includes logic for cancelling the near-end and far-end reflections, typically in the form of a symbol-driven

echo canceller. The symbol-driven echo canceller emulates the signal reflections and produces an echo cancelling signal which, when subtracted from the signals received by the modem, substantially eliminates the echoes from the received signals. Because the near-end and far-end reflections are delayed in time and are dependent upon the length of the signal paths in the network, the total echo channel response varies from network to network (and even connection to connection) and can be very long (on the order of 600 milliseconds) in certain situations. Emulation of such channels using a single filter is possible but unrealistic due to the computational complexity. Therefore, the symbol-driven echo canceller typically includes two separate filters for constructing the near-end and far-end echo channels, respectively, each capable of emulating the span of its respective echo channel. For convenience, the filter for constructing the near-end echo channel is referred to as the "near-end echo canceller" and the filter for constructing the far-end echo channel is referred to as the "far-end echo canceller."

An exemplary embodiment of a symbol-driven echo canceller is shown in FIG. 4. Specifically, the symbol-driven echo canceller includes a symbol modulator 402, a bulk-delay line 404, and two complex-valued adaptive transversal filters 406 and 408 for constructing the near-end and far-end echo channels, respectively. Each symbol from the symbol modulator 402, representing a transmitted symbol, is placed in the bulk-delay line 404 where it is held until needed by the near-end echo canceller 406 and the far-end echo canceller 408. Each symbol is delayed by a first delay time before being processed by the near-end echo canceller 406 and by a second delay time before being processed by the far-end echo canceller 408. The outputs of the near-end echo canceller 406 and

the far-end echo canceller 408 are combined to form echo cancelling signal 410. The first and second delay times are determined according to measured delays such that the echo cancelling signal 410 is provided to the receiver substantially in synchronization with the corresponding near-end and far-end echoes received by the modem, and are described in more detail below.

The relationship between the bulk-delay line, the near-end echo canceller, and the far-end echo canceller is shown in FIG. 5. Modulated symbols are inserted into the bulk-delay line at the insertion point and remain in the bulk-delay line for a predetermined number of symbol times. The echo canceller filters can be conceptualized as being projected onto the bulk-delay line such that each filter begins at a particular symbol offset from the insertion point and spans a predetermined number of symbols, where the number of symbols spanned by each filter is determined by the number of taps in the filter. A particular symbol is inserted at the insertion point, is operated upon by the near-end echo canceller after the first delay time, and is operated upon by the far-end echo canceller again after the second delay time.

In a preferred embodiment, the symbol-driven echo canceller is implemented in software. The bulk-delay line is implemented as a circular buffer, and three pointers are used to indicate the point where a symbol is inserted (put_ptr) and the two points where symbols are read for processing by the near-end echo canceller (ne_get) and the far-end echo canceller (fe_get). The near-end echo canceller and the far-end echo canceller each process one symbol per symbol time according to the get pointer for the respective echo canceller. Each time a symbol is inserted at the put_ptr, all three

pointers are advanced by one symbol such that the ne_get and fe_get pointers remain at fixed offsets from the put_ptr.

FIGs. 6A and 6B show details of a preferred embodiment of a near-end or far-end echo canceller. In the preferred embodiment, the receiver sampling rate is three times the transmitter symbol rate. Therefore, for each modulated symbol S_n received from the bulk-delay line, the echo canceller produces three echo canceller output signals, one for each of three consecutive receiver samples $R_n(0)$, $R_n(1)$, and $R_n(2)$ corresponding to symbol S_n .

Because only one input symbol is available for generating the three echo canceller output signals, the echo canceller filter interpolates the three echo canceller output signals from the single input symbol. Specifically, the echo canceller filter is updated three times for each input symbol, where each update is based on the input symbol and one of three consecutive feedback samples $e_n(0)$, $e_n(1)$, and $e_n(2)$ which are the three echo cancelled receiver samples immediately preceding receiver samples $R_n(0)$, $R_n(1)$, and $R_n(2)$.

As shown in FIG. 6A, the echo canceller filter can be conceptualized as a shift register having X filter taps which is shifted and updated three times for each input symbol S_n . For the first update, symbol S_n is effectively placed in filter tap 0, and the filter is updated using feedback sample $e_n(0)$. At that time, there is effectively no symbol available for updating filter taps 1 and 2. Therefore, the filter update affects only filter tap 0 and every third filter tap therefrom (corresponding to previously received symbols), while the rest of the filter taps remain unchanged (indicated by the value "0"). For the second update, symbol S_n is effectively shifted to filter tap 1, and the filter is updated using feedback sample $e_n(1)$. At that time, there is effectively no symbol available for updating filter

taps 0 and 2. Therefore, the filter update affects only filter tap 1 and every third filter tap therefrom. For the third update, symbol S_n is effectively shifted to filter tap 2, and the filter is updated using feedback sample $e_n(2)$. At that time, there is effectively no symbol available for updating filter taps 0 and 1. Therefore, the filter update affects only filter tap 2 and every third filter tap therefrom. The filter update cycle then begins again with the next symbol S_{n+1} using feedback samples $e_{n+1}(0)$, $e_{n+1}(1)$, and $e_{n+1}(2)$ which are the echo cancelled receiver samples $R_n(0)$, $R_n(1)$, and $R_n(2)$, respectively.

10 In the preferred embodiment, the echo canceller filter is not implemented as a shift register having X filter taps. Such an implementation would be wasteful of processing resources because only one-third of the filter taps actually change during each filter update. Instead, it is preferable to implement the echo canceller filter using three subfilters, as shown in FIG. 6B. In this embodiment, each subfilter is a complex-valued adaptive transversal filter supporting $X/3$ filter taps which is updated using the Least Mean Squared algorithm using feedback 606 from the receiver. For each symbol 602 received from the bulk-delay line, for example, 15 symbol S_n , each subfilter is updated and produces one echo canceller output signal 604. Specifically, subfilter 608 (corresponding to filter tap 0 and every third filter tap therefrom) is updated using symbol S_n and feedback sample $e_n(0)$ and produces echo canceller output signal $Y_n(0)$, subfilter 610 (corresponding to filter tap 1 and every third filter tap therefrom) is updated using symbol S_n and feedback sample $e_n(1)$ and produces echo canceller output signal $Y_n(1)$, 25 and subfilter 612 (corresponding to filter tap 2 and every third filter tap therefrom) is updated using symbol S_n and feedback sample $e_n(2)$ and produces echo canceller output signal $Y_n(2)$. The echo canceller

output signals $Y_n(0)$, $Y_n(1)$, and $Y_n(2)$ from both echo cancellers are effectively subtracted from the receiver samples $R_n(0)$, $R_n(1)$, and $R_n(2)$, respectively, thereby substantially removing the reflected signals from the receiver samples $R_n(0)$, $R_n(1)$, and $R_n(2)$. The

5 resulting echo cancelled receiver samples are then used as feedback samples $e_{n+1}(0)$, $e_{n+1}(1)$, and $e_{n+1}(2)$, respectively, for the next filter update using the next input symbol S_{n+1} .

In a typical modem, the near-end and far-end echo canceller filters have a fixed number of filter taps sufficient to compensate
10 for reflections having a predetermined maximum span. For convenience, these filters are referred to as "fixed-length" filters. In these modems, the amount of processing resources consumed by the echo canceller filters is not an issue, since the echo canceller filters are typically implemented using hardware which dedicates a
15 sufficient amount of processing resources to the echo canceller filters. Therefore, even if the span of signal reflections encountered in the telephone network are of shorter duration than the predetermined maximum span, there is no compelling reason to eliminate or reduce the number of taps in the echo canceller filters
20 in order to save processing resources.

In an embodiment of the present invention, the echo canceller filters are designed and implemented so that the number of taps is variable and adaptable to the actual span of the reflections. For convenience, these filters are referred to as "variable-length"
25 filters. Each variable-length filter is designed and implemented to support the same number of taps as the corresponding fixed-length filter, and therefore each variable-length filter remains capable of compensating for reflections having a predetermined maximum span. However, the number of taps actually used for echo cancellation is

adapted based on the power level and span of the actual signal reflections as measured by the modem. Because the span of signal reflections encountered in modern telephone networks is typically of shorter duration than the predetermined maximum span, it is
5 generally possible to eliminate at least some of the echo canceller filter taps and to thereby reduce the amount of processing resources consumed by the echo canceller filters.

In another embodiment of the present invention, either one or both of the echo cancellers can be disabled completely such that the
10 eliminated echo canceller(s) consumes no processing resources (other than perhaps some overhead for conditional execution of software). Referring to FIG. 2, modem 202 receives only a far-end echo, while modem 208 receives only a near-end echo. Therefore, in a preferred embodiment, the near-end echo canceller of a modem such as modem
15 202 is disabled completely, while the far-end echo canceller of a modem such as modem 208 is disabled completely. Furthermore, an alternate embodiment completely disables the far-end echo canceller if the delay between the near-end echo and the far-end echo is either zero or is so small that any far-end echo can be satisfactorily
20 handled by the near-end echo canceller alone.

Programming of the echo cancellers is typically performed during a startup sequence in which the modem characterizes the communication link and exchanges operating parameters with the remote modem. In a preferred embodiment, the modem supports a
25 startup sequence such as the ITU V.34 startup sequence which includes, among other things, a Phase 2 for measuring round-trip delay and for exchanging certain operating parameters, and a Phase 3 for training the echo cancellers.

FIG. 7 shows generally the steps of the prior art V.34 startup sequence that are relevant to the present discussion. After beginning in step 702 and going through at least a first startup phase (not shown), the modem performs the Phase 2 sequence which includes measuring a round-trip delay (RTDE) and exchanging certain operating parameters in step 704. Of particular relevance to the present invention is the exchange of information indicating whether or not one of the modems is operating on the digital side of a digital-to-analog connection (for example, modem 202 in FIG. 2). Subsequent to measuring the RTDE, the modem sets up its echo cancellers for a worst-case echo in step 706 and performs the Phase 3 sequence which includes training the echo cancellers in step 710. The modem then continues on with remaining phases of the startup sequence (not shown), ultimately terminating in step 799.

In a preferred embodiment of the present invention, a number of echo canceller adjustments are made during the startup sequence in order to adapt the echo canceller filters to the reflections in the network.

In one embodiment, the modem completely disables its near-end echo canceller if it is operating on the digital side of a digital-to-analog connection. As discussed previously, a modem operating on the digital side of a digital-to-analog connection does not require a near-end echo canceller, since only far-end reflections are present. By completely eliminating the near-end echo canceller, the amount of processing resources required by the symbol-driven echo canceller is substantially reduced.

In another embodiment, the modem completely disables its far-end echo canceller following Phase 2 if the remote modem indicates that the remote modem is operating on the digital side of a digital-

to-analog connection. As discussed previously, a modem operating on the analog side of a digital-to-analog connection does not require a far-end echo canceller, since only near-end reflections are present. By completely eliminating the far-end echo canceller, the amount of processing resources required by the symbol-driven echo canceller is substantially reduced.

In yet another embodiment, the modem completely disables its far-end echo canceller following Phase 2 if the RTDE is zero or is less than a predetermined value such that the span of the far-end echo canceller is entirely within the span of the near-end echo canceller. When the round-trip delay is very small, there is significant overlap between the near-end echo and the far-end echo. In this case, utilization of the near-end echo canceller alone is sufficient to compensate for both near-end and far-end echoes. By completely eliminating the far-end echo canceller, the amount of processing resources required by the symbol-driven echo canceller is substantially reduced.

After any adjustments are made following Phase 2, the remaining near-end and far-end echo canceller filters are initialized (step 706 in FIG. 7). The step of setting up the echo canceller for a worst-case echo involves projecting the remaining near-end and far-end echo canceller filters onto the bulk-delay line at appropriate offsets from the insertion point. Each echo canceller filter is initialized for its maximum span (48 symbols for the near-end filter and 35 symbols for the far-end filter in the preferred embodiment) and is projected onto the bulk-delay line such that the peak of each echo response falls roughly at a predetermined point within the span of the respective filter ($2/5$ of the filter span for the near-end filter and $1/3$ of the filter span for the far end filter in the preferred

embodiment). In a preferred embodiment, positioning the echo canceller filters involves converting the `sys_delay` and RTDE times into an integral number of symbols for determining a first symbol boundary and a second symbol boundary, respectively, and then

5 placing the $2/5$ point of the near-end filter at the first symbol boundary and placing the $1/3$ point of the far-end filter at the second symbol boundary. Since neither the $2/5$ point of the near-end filter nor the $1/3$ point of the far-end filter is an integral number of symbols, the numbers are rounded down to the nearest integral

10 number of symbols. Thus, in a preferred embodiment, the near-end echo canceller filter is positioned onto the bulk-delay line such that the `ne_get` pointer precedes the first symbol boundary by 19 symbols ($\cdot 0.4 \cdot 48^a$), while the far-end echo canceller filter is positioned onto the bulk-delay line such that the `fe_get` pointer precedes the second

15 symbol boundary by 11 symbols ($\cdot 0.333 \cdot 35^a$). This positioning of the filters, shown generally in FIG. 8A and in more detail in FIG. 8B, places the peak of each echo response at an appropriate point within each filter in order to more effectively emulate the echo channels. For convenience, the symbol boundary for a filter is referred to as the

20 "center" of the filter, and the taps on either side of the center are referred to as a "tail" of the filter.

After initializing the remaining filter(s), and before training the filter(s), a determination is made whether or not the far-end echo canceller filter can be eliminated or truncated. This adjustment is

25 only possible if both the near-end and far-end echo canceller filters remain (i.e., neither was previously eliminated) and span of the far-end echo canceller filter overlaps the span of the near-end echo canceller filter.

The far-end echo canceller filter is eliminated if the entire span of the far-end echo canceller filter is within the span of the near-end echo canceller filter. This can happen, for example, when the RTDE is very small. In the preferred embodiment where the near-end echo canceller filter span is 48 symbols and the far-end echo canceller filter span is 35 symbols, this situation occurs if the start of the far-end echo canceller filter (i.e., the fe_get pointer) is less than or equal to 13 symbols behind the start of the near-end echo canceller filter (i.e., the ne_get pointer). By completely eliminating the far-end echo canceller, the amount of processing resources required by the symbol-driven echo canceller is substantially reduced.

Even if the entire span of the far-end echo canceller filter is not within the span of the near-end echo canceller filter, the far-end echo canceller filter may be truncated if its span partially overlaps the span of the near-end echo canceller filter (i.e., at least one symbol is within the span of both filters). In the preferred embodiment where the near-end echo canceller filter span is 48 symbols and the far-end echo canceller filter span is 35 symbols, this situation occurs if the start of the far-end echo canceller filter (i.e., the fe_get pointer) is less than 48 symbols behind the start of the near-end echo canceller filter (i.e., the ne_get pointer). An example of overlapping filter spans is shown in FIG. 9. In this case, the far-end echo canceller filter is shortened by moving the fe_get pointer to the symbol following the end of the near-end echo canceller filter as shown in FIG. 10. Truncating the far-end echo canceller filter reduces the span of the filter and thus reduces the number of filter taps. By truncating the far-end echo canceller filter, the amount of processing resources required by the symbol-

driven echo canceller is reduced by an amount proportional to the number of filter taps eliminated from the far-end echo canceller filter.

5 In the example shown in FIG. 10, the truncated far-end echo canceller filter continues to have two tails. It is important to note, however, that if the far-end echo canceller filter overlaps the near-end echo canceller filter such that the center of the far-end filter is within the span of the near-end filter, then truncating the far-end filter will eliminate the tail shown on the right-hand side of the
10 filter and may eliminate some or all of the taps on the left-hand side of the filter (all of the far-end filter taps may be eliminated, for example, if the RTDE is zero). Thus, in a preferred embodiment, the step of eliminating the far-end echo canceller filter if the RTDE is below a predetermined threshold (described above) is not employed,
15 and instead this step of truncating the far-end echo canceller filter is used and will have the effect of completely eliminating the far-end echo canceller filter if the RTDE is indeed below the predetermined threshold.

20 After making any echo canceller filter adjustments prior to training the echo canceller filters, the remaining echo canceller filter(s) are trained as part of Phase 3 of the startup sequence (step 710 in FIG. 7). This step of training the echo canceller filters allows each filter to adapt to its respective echo channel and to thereby produce an appropriate echo canceller output signal.

25 After training the remaining echo canceller filter(s), a final adjustment is made to eliminate any unessential filter taps from each remaining filter. In a preferred embodiment, the filter taps are examined in groups of three (where each group of three filter taps corresponds to one symbol) beginning at the end of each tail and

working in toward the center. A group of three echo canceller filter taps is deemed to be unessential if the echo power level for the corresponding symbol is below a predetermined threshold, and is otherwise deemed to be essential. In the preferred embodiment, the predetermined threshold for the near-end echo canceller is $1.0e-5$, and the predetermined threshold for the far-end echo canceller is $1.0e-3$. Where the magnitude M of a filter tap n is a complex coefficient of the form:

$$M_n = R_n + jI_n,$$

the echo power level P for three taps is determined according to the formula:

$$P = \sum R_n^2 + I_n^2$$

for $n = 1$ to 3 . Paring of a particular tail is stopped once an essential group of three filter taps is found or if the entire tail is eliminated.

By reducing the number of filter taps in the remaining filter(s), the amount of processing resources required by the symbol-driven echo canceller is reduced by an amount proportional to the number of filter taps eliminated.

It is important to note that this final adjustment may be used as the default mechanism for eliminating an echo canceller under certain circumstances. This is particularly important for a modem operating on the analog side of a digital-to-analog connection because the information exchanged during Phase 2 may be insufficient to identify the connection as a digital-to-analog connection when the connection rate is less than 56 kilobits/second. Therefore, the modem may initially leave its far-end echo canceller enabled, but subsequently eliminate its far-end echo canceller when all of the filter taps are found to be unessential due to the absence of a far-end echo.

Thus, in an exemplary embodiment as shown in FIG. 11, a number of steps are taken during the startup sequence in order to adapt the echo canceller filters to the reflections in the network. After beginning in step 1102 and going through at least a first startup phase (not shown), the modem disables its near-end echo canceller in step 1104 if the modem is operating on the digital side of a digital-to-analog connection. The Phase 2 sequence, which includes measuring a round-trip delay (RTDE) and exchanging certain operating parameters, is performed in step 1106. Subsequent to performing the Phase 2 sequence, the modem disables its far-end echo canceller in step 1108 if the modem is operating on the analog side of a digital-to-analog connection as determined by the information exchanged during Phase 2. Optionally, the modem may also disable its far-end echo canceller in step 1110 if the RTDE is less than a predetermined threshold, although this step is not used in the preferred embodiment. In preparation for Phase 3, the modem then sets up the remaining echo canceller(s) for a worst-case echo in step 1112. After setting up the echo canceller(s), the modem eliminates the far-end echo canceller in step 1114 if the span of the far-end echo canceller is entirely within the span of the near-end echo canceller, and truncates the far-end echo canceller in step 1116 if the span of the far-end echo canceller only partially overlaps the span of the near-end echo canceller. Steps 1114 and 1116 are not applicable if either of the echo canceller filters had been eliminated prior to step 1112. The modem then performs the Phase 3 sequence which includes training the echo canceller(s) in step 1118. After training the echo canceller(s), the modem adapts the echo canceller(s) in step 1120 by eliminating any unessential filter taps. The modem then continues on with remaining phases of the startup sequence (not shown),

ultimately terminating in step 1199. It will be apparent to a skilled artisan that the techniques for eliminating, truncating, and adapting the echo cancellers may be applied individually or in combination with each other.

5 The techniques for eliminating, truncating, and/or adapting the echo cancellers may be embodied in hardware, firmware, or software. An exemplary implementation of a variable-length echo canceller such as a software modem or other modem includes logic or computer readable program code means for disabling an echo canceller,
10 truncating a far-end echo canceller, and/or adapting an echo canceller by eliminating any unessential filter taps. Again, it will be apparent to a skilled artisan that the logic or computer readable program code means for eliminating, truncating, and/or adapting the echo cancellers may be applied individually or in combination with
15 each other.

 While the techniques of the present invention have been described in relation to a software modem for the purpose of reducing the amount of processing resources required by the software modem, it will be apparent to a skilled artisan that the same
20 techniques are applicable to echo cancellers generally. For instance, it is well known that filters are typically less than perfect, and that a filter which is too long can actually amplify noise. Therefore, the techniques for eliminating, truncating, and/or adapting the echo canceller filters are useful for adapting each filter to an appropriate
25 length so that the filter does not generate excessive noise.

 Finally, it will be apparent to a skilled artisan that, by eliminating or reducing the number of echo canceller filter taps when an echo is present, any reduction in the consumption of processing resources may be realized at a cost in echo canceller performance.

As described above, a group of three echo canceller filter taps is deemed unessential and is eliminated if the echo power level at the symbol corresponding to the filter taps is below a predetermined threshold. Therefore, it is possible that, by eliminating certain filter taps, echo signal energy which would have been cancelled by those taps is left in the received signal as noise which must then be overcome by the receiver. Thus, there is generally a tradeoff between the filter length and the filter performance, and the predetermined threshold is selected so that the echo canceller performance remains at all times within acceptable operating limits.

The present invention may be embodied in other specific forms without departing from the essence or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive.

What is claimed is:



1. A device for operation over a communication system, the device comprising:

5 a near-end echo canceller having a first predetermined maximum number of filter taps for emulating a first echo channel received at a first delay time and having a first predetermined maximum span;

10 a variable-length far-end echo canceller having a second predetermined maximum number of filter taps for emulating a second echo channel received at a second delay time and having a second predetermined maximum span;

logic for determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller; and

15 logic for truncating the far-end echo canceller if the span of the far-end echo canceller overlaps the span of the near-end echo canceller.

2. The device of claim 1 wherein the logic for determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller comprises:

logic for determining whether the span of the far-end echo canceller is entirely within the span of the near-end echo canceller.

25 3. The device of claim 2 wherein the logic for truncating the far-end echo canceller comprises:

logic for eliminating the far-end echo canceller if the span of the far-end echo canceller is entirely within the span of the near-end echo canceller.



4. The device of claim 1 wherein the logic for determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller comprises:

5 logic for determining whether at least one symbol is within the span of both the near-end echo canceller and the far-end echo canceller.

5. The device of claim 4 wherein the logic for truncating the far-end echo canceller comprises:

10 logic for shortening the far-end echo canceller such that the span of the far-end echo canceller begins at a symbol immediately following a last symbol within the span of the near-end echo canceller if at least one symbol is within the span of both the near-end echo canceller and the far-end echo canceller.



6. An apparatus comprising a computer usable medium having computer readable program code means embodied therein for adapting a far-end echo canceller, the computer readable program code means comprising:

5 computer readable program code means for a near-end echo canceller having a first predetermined maximum span;

computer readable program code means for a variable-length far-end echo canceller having a second predetermined maximum span;

10 computer readable program code means for determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller; and

computer readable program code means for truncating the far-end echo canceller if the span of the far-end echo canceller overlaps the span of the near-end echo canceller.

15

7. The apparatus of claim 6 wherein the computer readable program code means for determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller comprises:

20 computer readable program code means for determining whether the span of the far-end echo canceller is entirely within the span of the near-end echo canceller.

25 8. The apparatus of claim 7 wherein the computer readable program code means for truncating the far-end echo canceller comprises:

computer readable program code means for eliminating the far-end echo canceller if the span of the far-end echo canceller is entirely within the span of the near-end echo canceller.



9. The apparatus of claim 6 wherein the computer readable program code means for determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller comprises:

5 computer readable program code means for determining whether at least one symbol is within the span of both the near-end echo canceller and the far-end echo canceller.

10. The apparatus of claim 9 wherein the computer readable
10 program code means for truncating the far-end echo canceller comprises:

computer readable program code means for shortening the far-end echo canceller such that the span of the far-end echo canceller begins at a symbol immediately following a last symbol within the
15 span of the near-end echo canceller if at least one symbol is within the span of both the near-end echo canceller and the far-end echo canceller.



11. In a device comprising a near-end echo canceller and a variable-length far-end echo canceller, the near-end echo canceller having a first predetermined maximum number of filter taps for emulating a first echo channel received at a first delay time and having a first
5 predetermined maximum span, the far-end echo canceller having a second predetermined maximum number of filter taps for emulating a second echo channel received at a second delay time and having a second predetermined maximum span, a method for adapting the far-end echo canceller comprising the steps of:

- 10 determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller; and
 truncating the far-end echo canceller if the span of the far-end echo canceller overlaps the span of the near-end echo canceller.

- 15 12. The method of claim 11 wherein the step of determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller comprises:

 determining whether the span of the far-end echo canceller is entirely within the span of the near-end echo canceller.

20

13. The method of claim 12 wherein the step of truncating the far-end echo canceller comprises:

 eliminating the far-end echo canceller if the span of the far-end echo canceller is entirely within the span of the near-end echo
25 canceller.



14. The method of claim 11 wherein the step of determining whether the span of the far-end echo canceller overlaps the span of the near-end echo canceller comprises the step of:

5 determining whether at least one symbol is within the span of both the near-end echo canceller and the far-end echo canceller.

15. The method of claim 14 wherein the step of truncating the far-end echo canceller comprises the step of:

10 shortening the far-end echo canceller such that the span of the far-end echo canceller begins at a symbol immediately following a last symbol within the span of the near-end echo canceller if at least one symbol is within the span of both the near-end echo canceller and the far-end echo canceller.

15 16. A device for operation over a communication system, the device comprising:

a variable-length echo canceller having a predetermined maximum number of filter taps;

20 logic for training the variable-length echo canceller based on echo characteristics of the communication system; and

logic for eliminating unessential filter taps from the trained variable-length echo canceller.



17. The device of claim 16 wherein the logic for eliminating unessential filter taps comprises:

logic for computing an echo power level corresponding to a predetermined symbol time;

5 logic for determining whether said echo power level is less than a predetermined threshold; and

logic for eliminating the filter taps corresponding to the predetermined symbol time if the echo power level is less than the predetermined threshold.

10

18. The device of claim 17 wherein the logic for computing the echo power level comprises:

logic for computing the sum of the squares of the magnitudes of each filter tap corresponding to the predetermined symbol time.

15

19. The device of claim 17 wherein the logic for eliminating the filter taps comprises:

logic for shortening the span of the echo canceller by one symbol

20

20. An apparatus comprising a computer usable medium having computer readable program code means embodied therein for adapting an echo canceller, the computer readable program code means comprising:

25 computer readable program code means for a variable-length echo canceller having a predetermined maximum number of filter taps;



computer readable program code means for training the variable-length echo canceller based on echo characteristics of the communication system; and

5 computer readable program code means for eliminating unessential filter taps from the trained variable-length echo canceller.

21. The apparatus of claim 20 wherein the computer readable program code means for eliminating unessential filter taps
10 comprises:

computer readable program code means for computing an echo power level corresponding to a predetermined symbol time;

computer readable program code means for determining whether said echo power level is less than a predetermined threshold; and

15 computer readable program code means for eliminating the filter taps corresponding to the predetermined symbol time if the echo power level is less than the predetermined threshold.

22. The apparatus of claim 21 wherein the computer readable program code means for computing the echo power level comprises:

20 computer readable program code means for computing the sum of the squares of the magnitudes of each filter tap corresponding to the predetermined symbol time.



23. The apparatus of claim 21 wherein the computer readable program code means for eliminating the filter taps comprises:

computer readable program code means for shortening the span of the echo canceller by one symbol

5

24. In a device comprising a variable-length echo canceller having a predetermined maximum number of filter taps, a method for adapting the echo canceller comprising the steps of:

training the variable-length echo canceller based on echo characteristics of the communication system; and

10

eliminating unessential filter taps from the trained variable-length echo canceller.

15

25. The method of claim 24 wherein the step of eliminating unessential filter taps comprises the steps of:

computing an echo power level corresponding to a predetermined symbol time;

determining whether said echo power level is less than a predetermined threshold; and

20

eliminating the filter taps corresponding to the predetermined symbol time if the echo power level is less than the predetermined threshold.

25

26. The method of claim 25 wherein the step of computing the echo power level comprises the step of:

computing the sum of the squares of the magnitudes of each filter tap corresponding to the predetermined symbol time.



27. The method of claim 25 wherein the step of eliminating the filter taps comprises the step of:

logic for shortening the span of the echo canceller by one symbol

5

28. A device for operation over a communication system, the device comprising:

an echo canceller capable of filtering an echo in the communication system;

10 logic for determining whether the echo canceller is unnecessary; and

logic for disabling the echo canceller if the echo canceller is determined to be unnecessary.

15 29. The device of claim 28 comprising:

logic for disabling the echo canceller if the echo canceller is a far-end echo canceller and the device is determined to be operating on the analog side of a digital-to-analog connection.

20 30. The device of claim 28 comprising:

logic for disabling the echo canceller if the echo canceller is a near-end echo canceller and the device is determined to be operating on the digital side of a digital-to-analog connection.

25



31. The device of claim 28 comprising:

logic for determining a round trip delay over the communication system; and

5 logic for disabling the echo canceller if the echo canceller is a far-end echo canceller and the round-trip delay is determined to be less than a predetermined threshold.

32. An apparatus comprising a computer usable medium having

10 computer readable program code means embodied therein for eliminating echo canceller logic, the computer readable program code means comprising:

computer readable program code means for an echo canceller capable of filtering an echo in the communication system;

15 computer readable program code means for determining whether the echo canceller is unnecessary; and

computer readable program code means for disabling the echo canceller if the echo canceller is determined to be unnecessary.

20 33. The apparatus of claim 32 comprising:

computer readable program code means for disabling the echo canceller if the echo canceller is a far-end echo canceller and the device is operating on the analog side of a digital-to-analog connection.

25

34. The apparatus of claim 32 comprising:

computer readable program code means for disabling the echo canceller if the echo canceller is a near-end echo canceller and the



device is operating on the digital side of a digital-to-analog connection.

35. The apparatus of claim 32 comprising:

5 computer readable program code means for determining a round trip delay over the communication system; and

computer readable program code means for disabling the echo canceller if the echo canceller is a far-end echo canceller and the round-trip delay is determined to be less than a predetermined
10 threshold.

36. In a device having an echo canceller capable of filtering an echo in the communication system, a method for eliminating echo canceller logic comprising the steps of:

15 determining whether the echo canceller is unnecessary; and
disabling the echo canceller if the echo canceller is determined to be unnecessary.

37. The method of claim 36 comprising:

20 disabling the echo canceller if the echo canceller is a far-end echo canceller and the device is determined to be operating on the analog side of a digital-to-analog connection.

38. The method of claim 36 comprising:

25 disabling the echo canceller if the echo canceller is a near-end echo canceller and the device is determined to be operating on the digital side of a digital-to-analog connection.



39. The method of claim 36 comprising:

determining a round trip delay over the communication system;

5 and

disabling the echo canceller if the echo canceller is a far-end echo canceller and the round-trip delay is determined to be less than a predetermined threshold.