



US00RE48832E

(19) **United States**
(12) **Reissued Patent**
Marrow et al.

(10) **Patent Number:** **US RE48,832 E**
(45) **Date of Reissued Patent:** **Nov. 23, 2021**

(54) **MEASURING ANGLE OF INCIDENCE IN AN ULTRAWIDEBAND COMMUNICATION SYSTEM**

(71) Applicant: **DECAWAVE LTD.**, Dublin (IE)

(72) Inventors: **Gavin Marrow**, Dublin (IE); **Michael McLaughlin**, Dublin (IE); **Ciaran McElroy**, Dublin (IE)

(73) Assignee: **Decawave, Ltd.**, Dublin (IE)

(21) Appl. No.: **16/437,695**

(22) Filed: **Jun. 11, 2019**

Related U.S. Patent Documents

Reissue of:

(64) Patent No.: **10,075,210**
Issued: **Sep. 11, 2018**
Appl. No.: **15/310,548**
PCT Filed: **May 23, 2014**
PCT No.: **PCT/EP2014/060722**
§ 371 (c)(1),
(2) Date: **Nov. 11, 2016**
PCT Pub. No.: **WO2015/176776**
PCT Pub. Date: **Nov. 26, 2015**

U.S. Applications:

(63) Continuation-in-part of application No. 13/775,282, filed on Feb. 25, 2013, now Pat. No. 8,760,334, which
(Continued)

(51) **Int. Cl.**
H04B 1/7163 (2011.01)
G01S 3/02 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC **H04B 1/71637** (2013.01); **G01S 3/023** (2013.01); **G01S 3/46** (2013.01); **G01S 3/48** (2013.01)

(58) **Field of Classification Search**
CPC **H04B 1/71637**; **G01S 3/023**; **G01S 3/46**; **G01S 3/48**

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,249,542 B1 6/2001 Kohli et al.
6,693,954 B1 2/2004 King et al.

(Continued)

FOREIGN PATENT DOCUMENTS

WO 2014001651 A1 1/2014
WO WO2014001651 A1 1/2014
WO 2015051821 A1 4/2015

OTHER PUBLICATIONS

Loeffler, Andreas, et al., "Localizing with Passive UHF RFID Tags Using Wideband Signals," Radio Frequency Identification from System to Applications, Chapter 5, Jun. 2013, 26 pages.

(Continued)

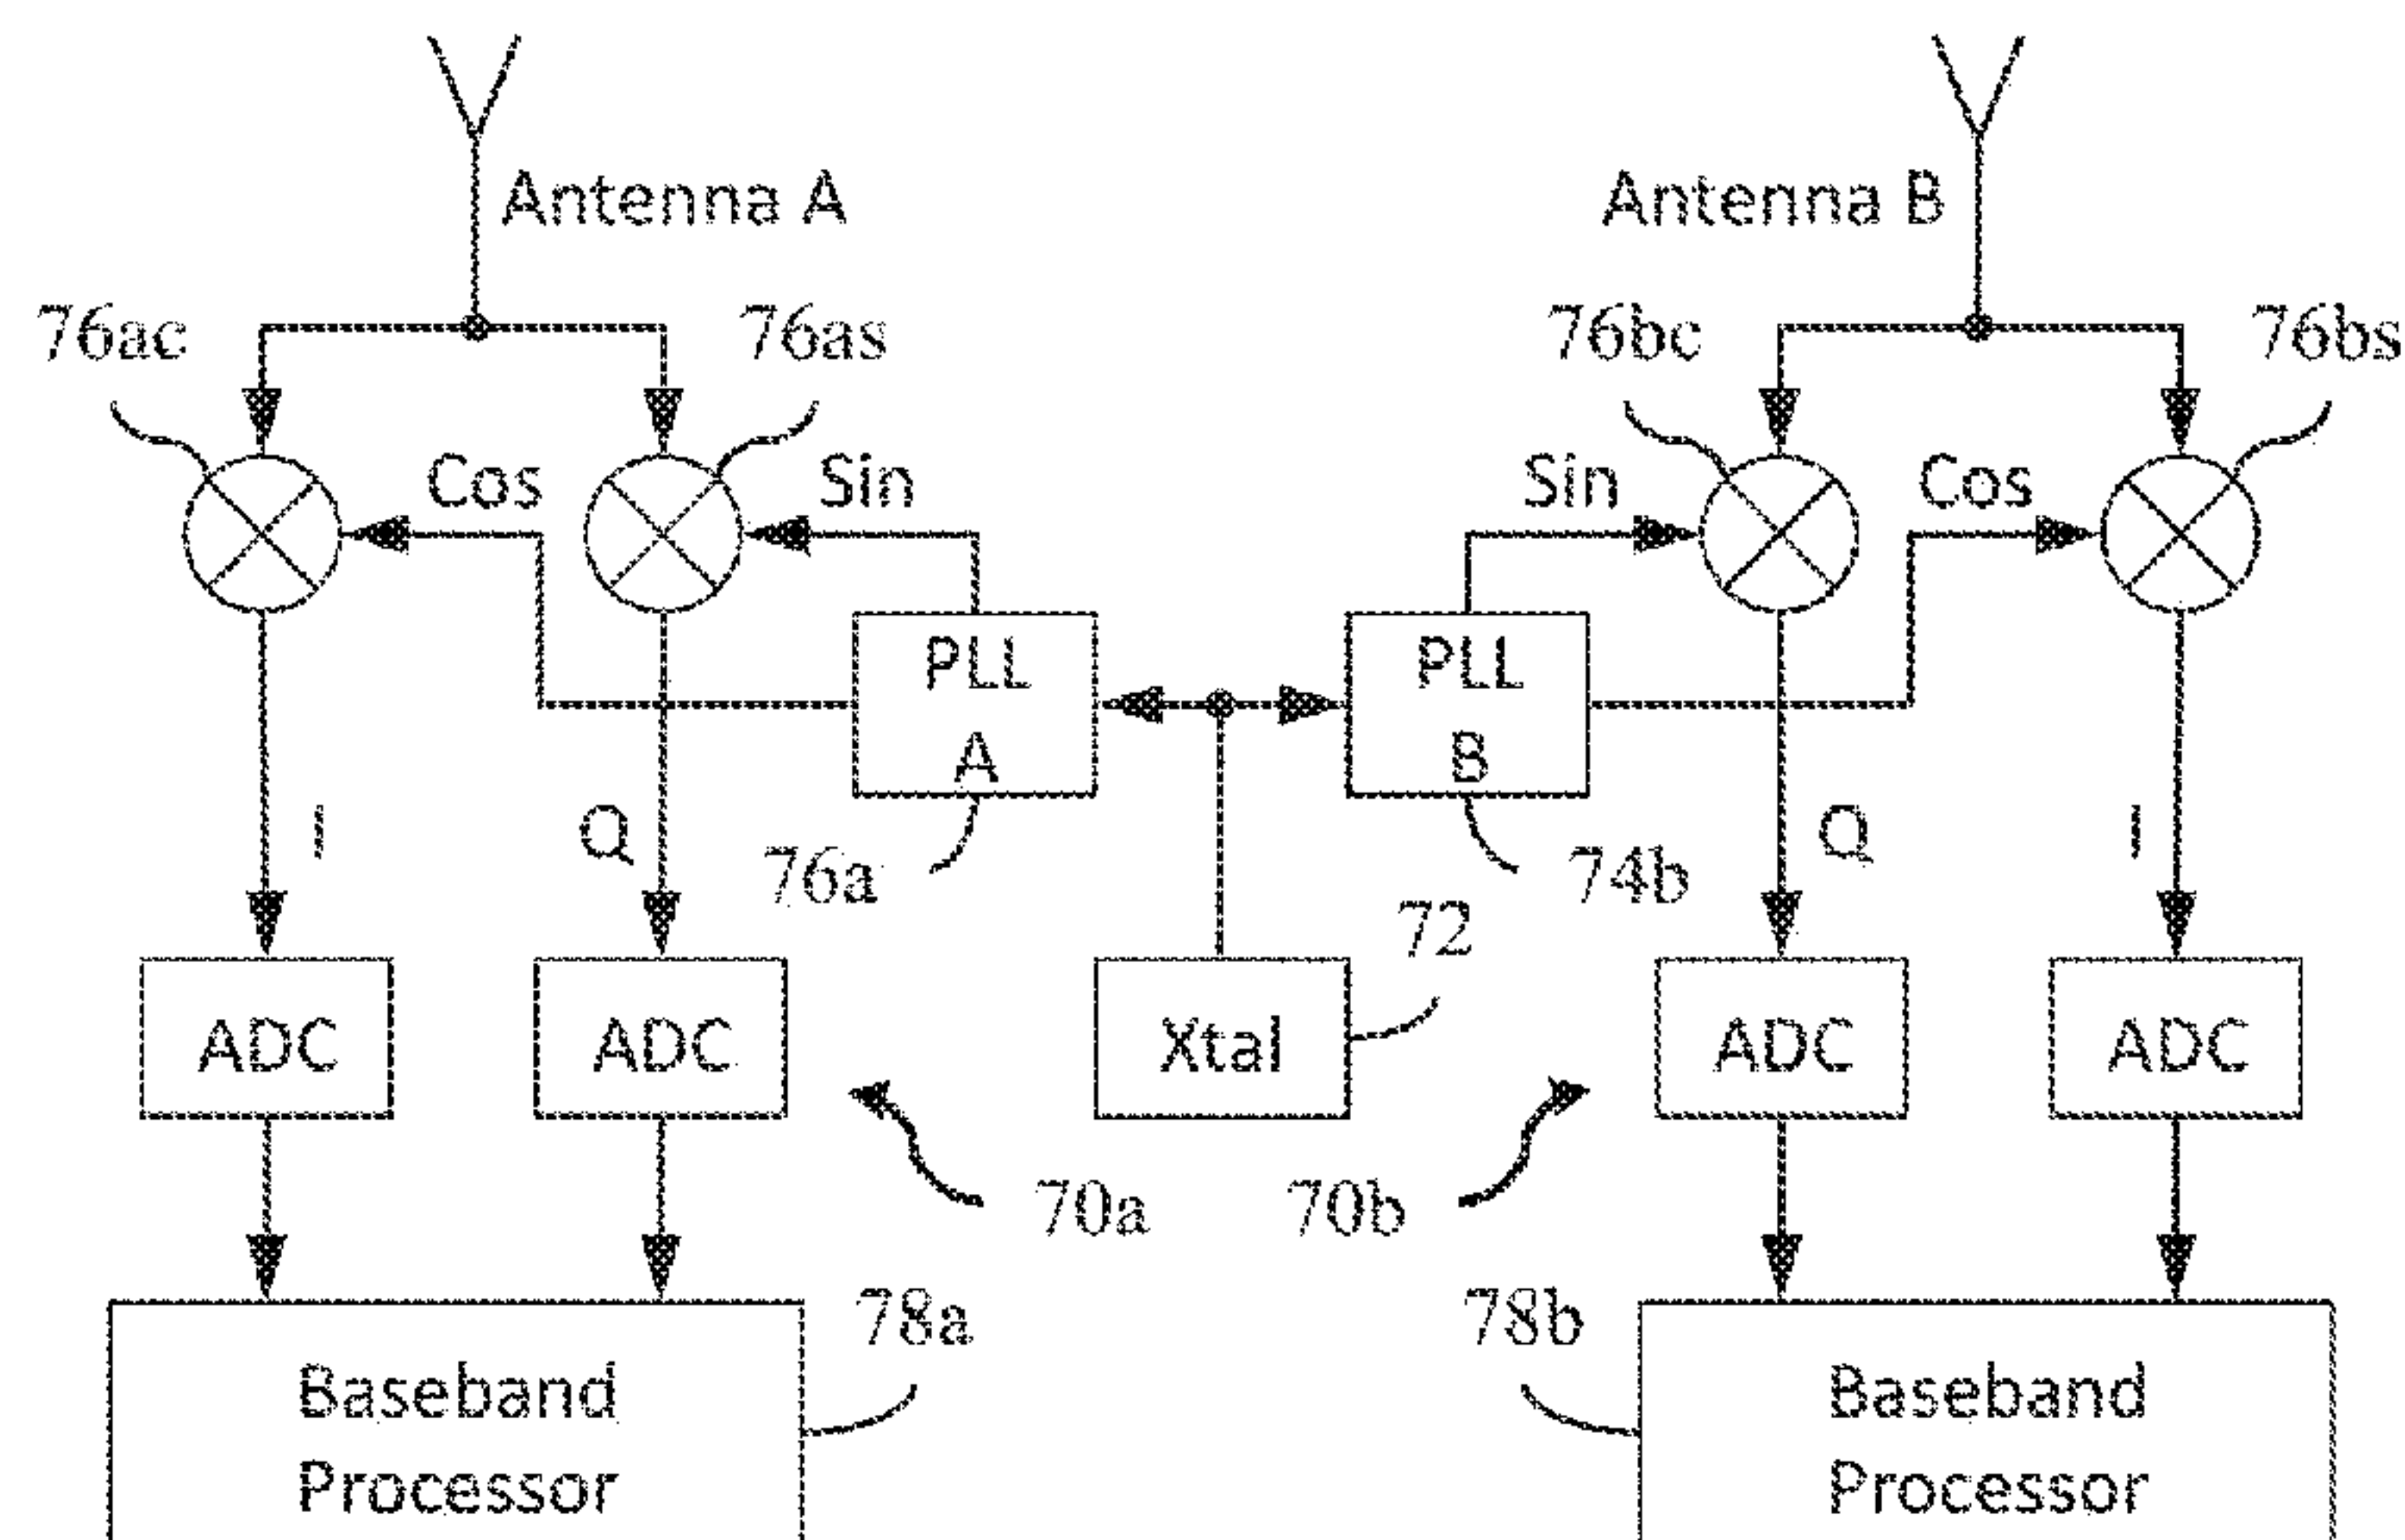
Primary Examiner — Yuzhen Ge

(74) *Attorney, Agent, or Firm* — Withrow & Terranova, PLLC

(57) **ABSTRACT**

In an ultra-wideband ("UWB") receiver, a received UWB signal is periodically digitized as a series of ternary samples. During a carrier acquisition mode of operation, the samples are continuously correlated with a predetermined preamble sequence to develop a correlation value. When the value exceeds a predetermined threshold, indicating that the preamble sequence is being received, estimates of the channel impulse response ("CIR") are developed. When a start-of-frame delimiter ("SFD") is detected, the best CIR estimate is provided to a channel matched filter ("CMF"). During a data recovery mode of operation, the CMF filters channel-injected noise from the sample stream. Both carrier phase errors and data timing errors are continuously detected and corrected during both the carrier acquisition and data recovery modes of operation. The phase of the carrier can be determined by accumulating the correlator output before it is rotated by the carrier correction. By comparing the carrier phases of two receivers separated by a known distance, d , the angle of incidence, θ , of the signal can be determined.

3 Claims, 15 Drawing Sheets



Related U.S. Application Data

is a continuation-in-part of application No. 12/885,517, filed on Sep. 19, 2010, now Pat. No. 8,437,432.

(60) Provisional application No. 61/316,299, filed on Mar. 22, 2010.

(51) **Int. Cl.**

G01S 3/46 (2006.01)

G01S 3/48 (2006.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,735,242	B1	5/2004	Kenney et al.	
7,436,359	B1 *	10/2008	Rose	G01S 5/12 342/424
8,436,758	B2	5/2013	McLaughlin et al.	
8,437,432	B2	5/2013	McLaughlin et al.	
8,760,334	B2	6/2014	McLaughlin et al.	
9,778,340	B2	10/2017	Mutz et al.	
10,075,210	B2	9/2018	Marrow et al.	
2005/0018762	A1 *	1/2005	Aiello	H04B 1/7183 375/219
2005/0152326	A1	7/2005	Vijayan et al.	
2006/0050822	A1 *	3/2006	Panikkar	H04L 7/041 375/354
2007/0197262	A1 *	8/2007	Smith	H04B 3/542 455/562.1
2008/0072025	A1	3/2008	Staszewski et al.	
2011/0142093	A1	6/2011	De Rosa	
2011/0188618	A1 *	8/2011	Feller	H04B 1/10 375/346
2011/0199263	A1 *	8/2011	Kang	G01S 3/04 342/442
2012/0069868	A1	3/2012	McLaughlin et al.	
2013/0163638	A1	6/2013	McLaughlin	
2014/0093023	A1	4/2014	Park	
2014/0112378	A1	4/2014	Ji et al.	
2014/0125518	A1 *	5/2014	Feller	H04B 7/00 342/357.51
2016/0242135	A1	8/2016	McLaughlin et al.	
2017/0085293	A1	3/2017	Marrow et al.	
2018/0138941	A1	5/2018	McLaughlin et al.	
2018/0331718	A1	11/2018	McLaughlin et al.	

OTHER PUBLICATIONS

Manteuffel, et al., "Antenna and propagation impairments of a UWB localization system integrated into an aircraft cabin", Antennas and Propagation Conference, Nov. 2010, pp. 589-592.

Wu, et al., "Carrier-less, Single and Multi-Carrier UWB Radio Technology," International Workshop on Ultra Wideband Systems Joint with Conference on Ultra Wideband Systems and Technologies, 2004, pp. 192-196.

Zhonghai, Wang, et al., "A New Multi-Antenna Based LOS -NLOS Separation Technique", Digital Signal Processing Workshop and 5th IEEE Signal Processing Education Workshop, Jan. 2009, pp. 331-336.

First Office Action for Chinese Patent Application No. 201480079211.7, dated Jun. 4, 2018, 15 pages.

Second Office Action for Chinese Patent Application No. 201480079211.7, dated Apr. 2, 2019, 14 pages.

Third Office Action for Chinese Patent Application No. 201480079211.7, dated Jan. 6, 2020, 8 pages.

International Search Report and Written Opinion for International Patent Application No. PCT/EP2014/060722, dated Jan. 28, 2015, 10 pages.

International Preliminary Report on Patentability for International Patent Application No. PCT/EP20141060722, dated Dec. 8, 2016, 8 pages.

Non-Final Office Action for U.S. Appl. No. 15/310,548, dated Jun. 30, 2017, 11 pages.

Final Office Action for U.S. Appl. No. 15/310,548, dated Feb. 27, 2018, 9 pages.

Non-Final Office Action for U.S. Appl. No. 15/814,711, dated Feb. 8, 2019, 15 pages.

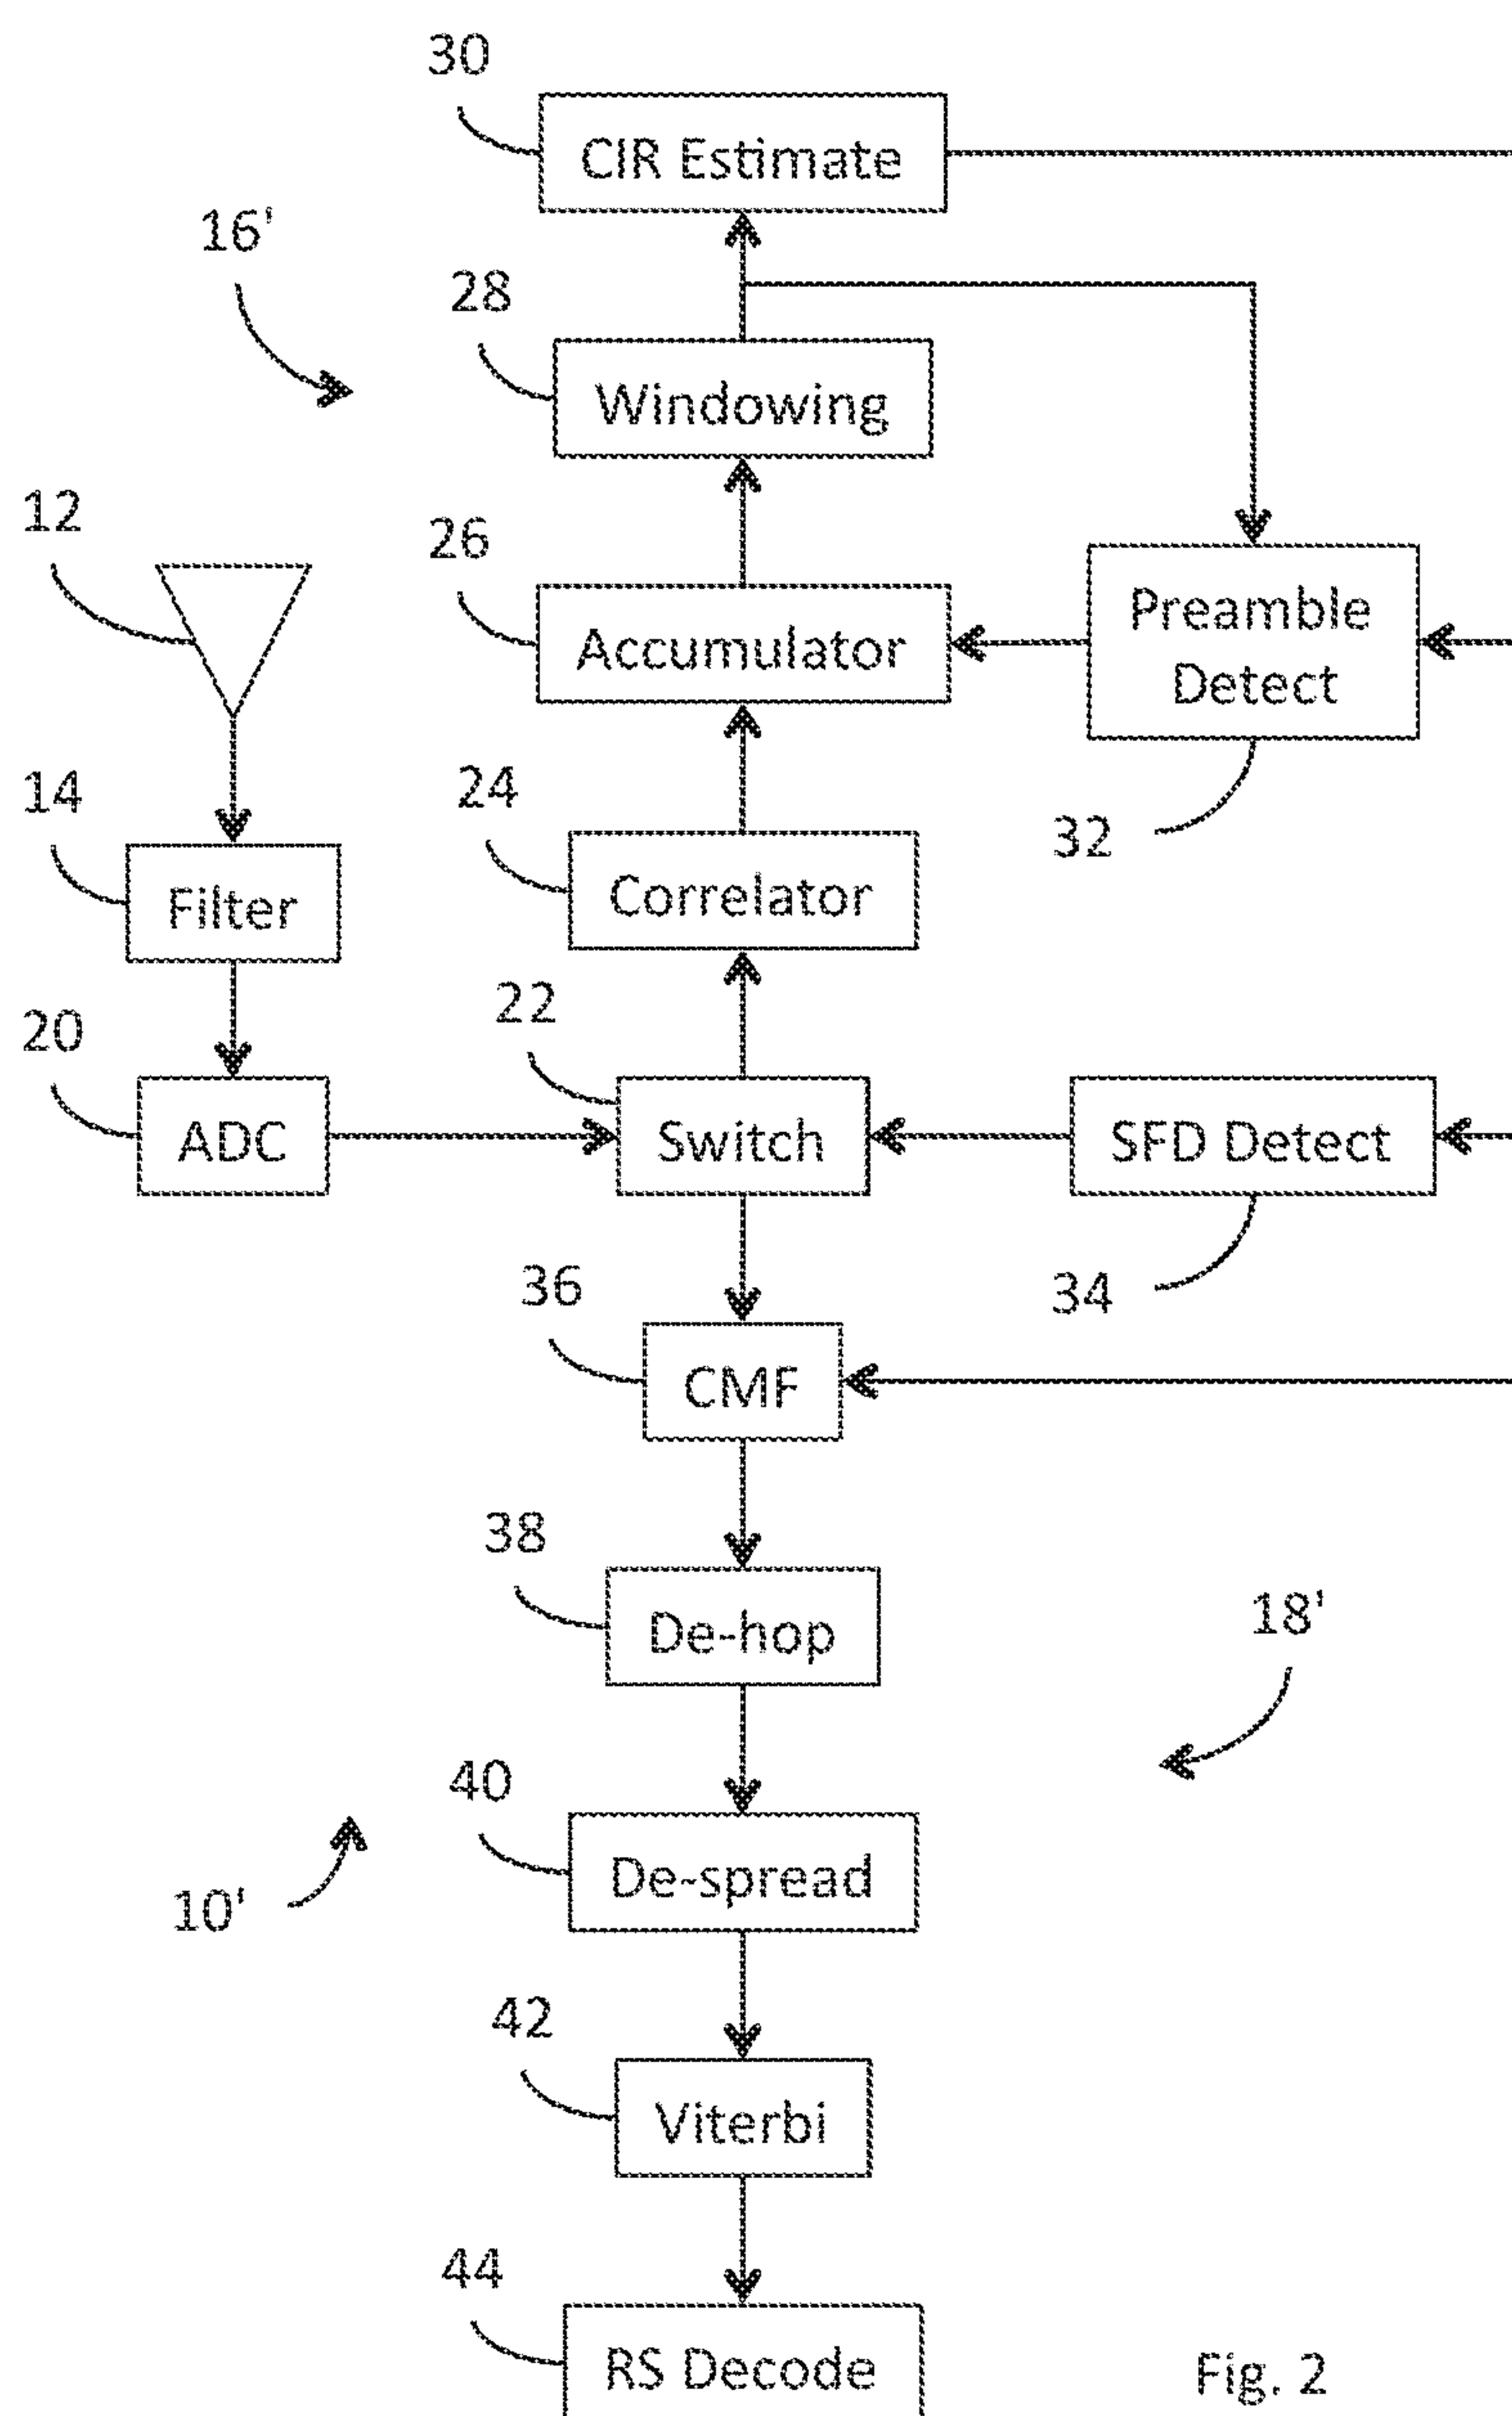
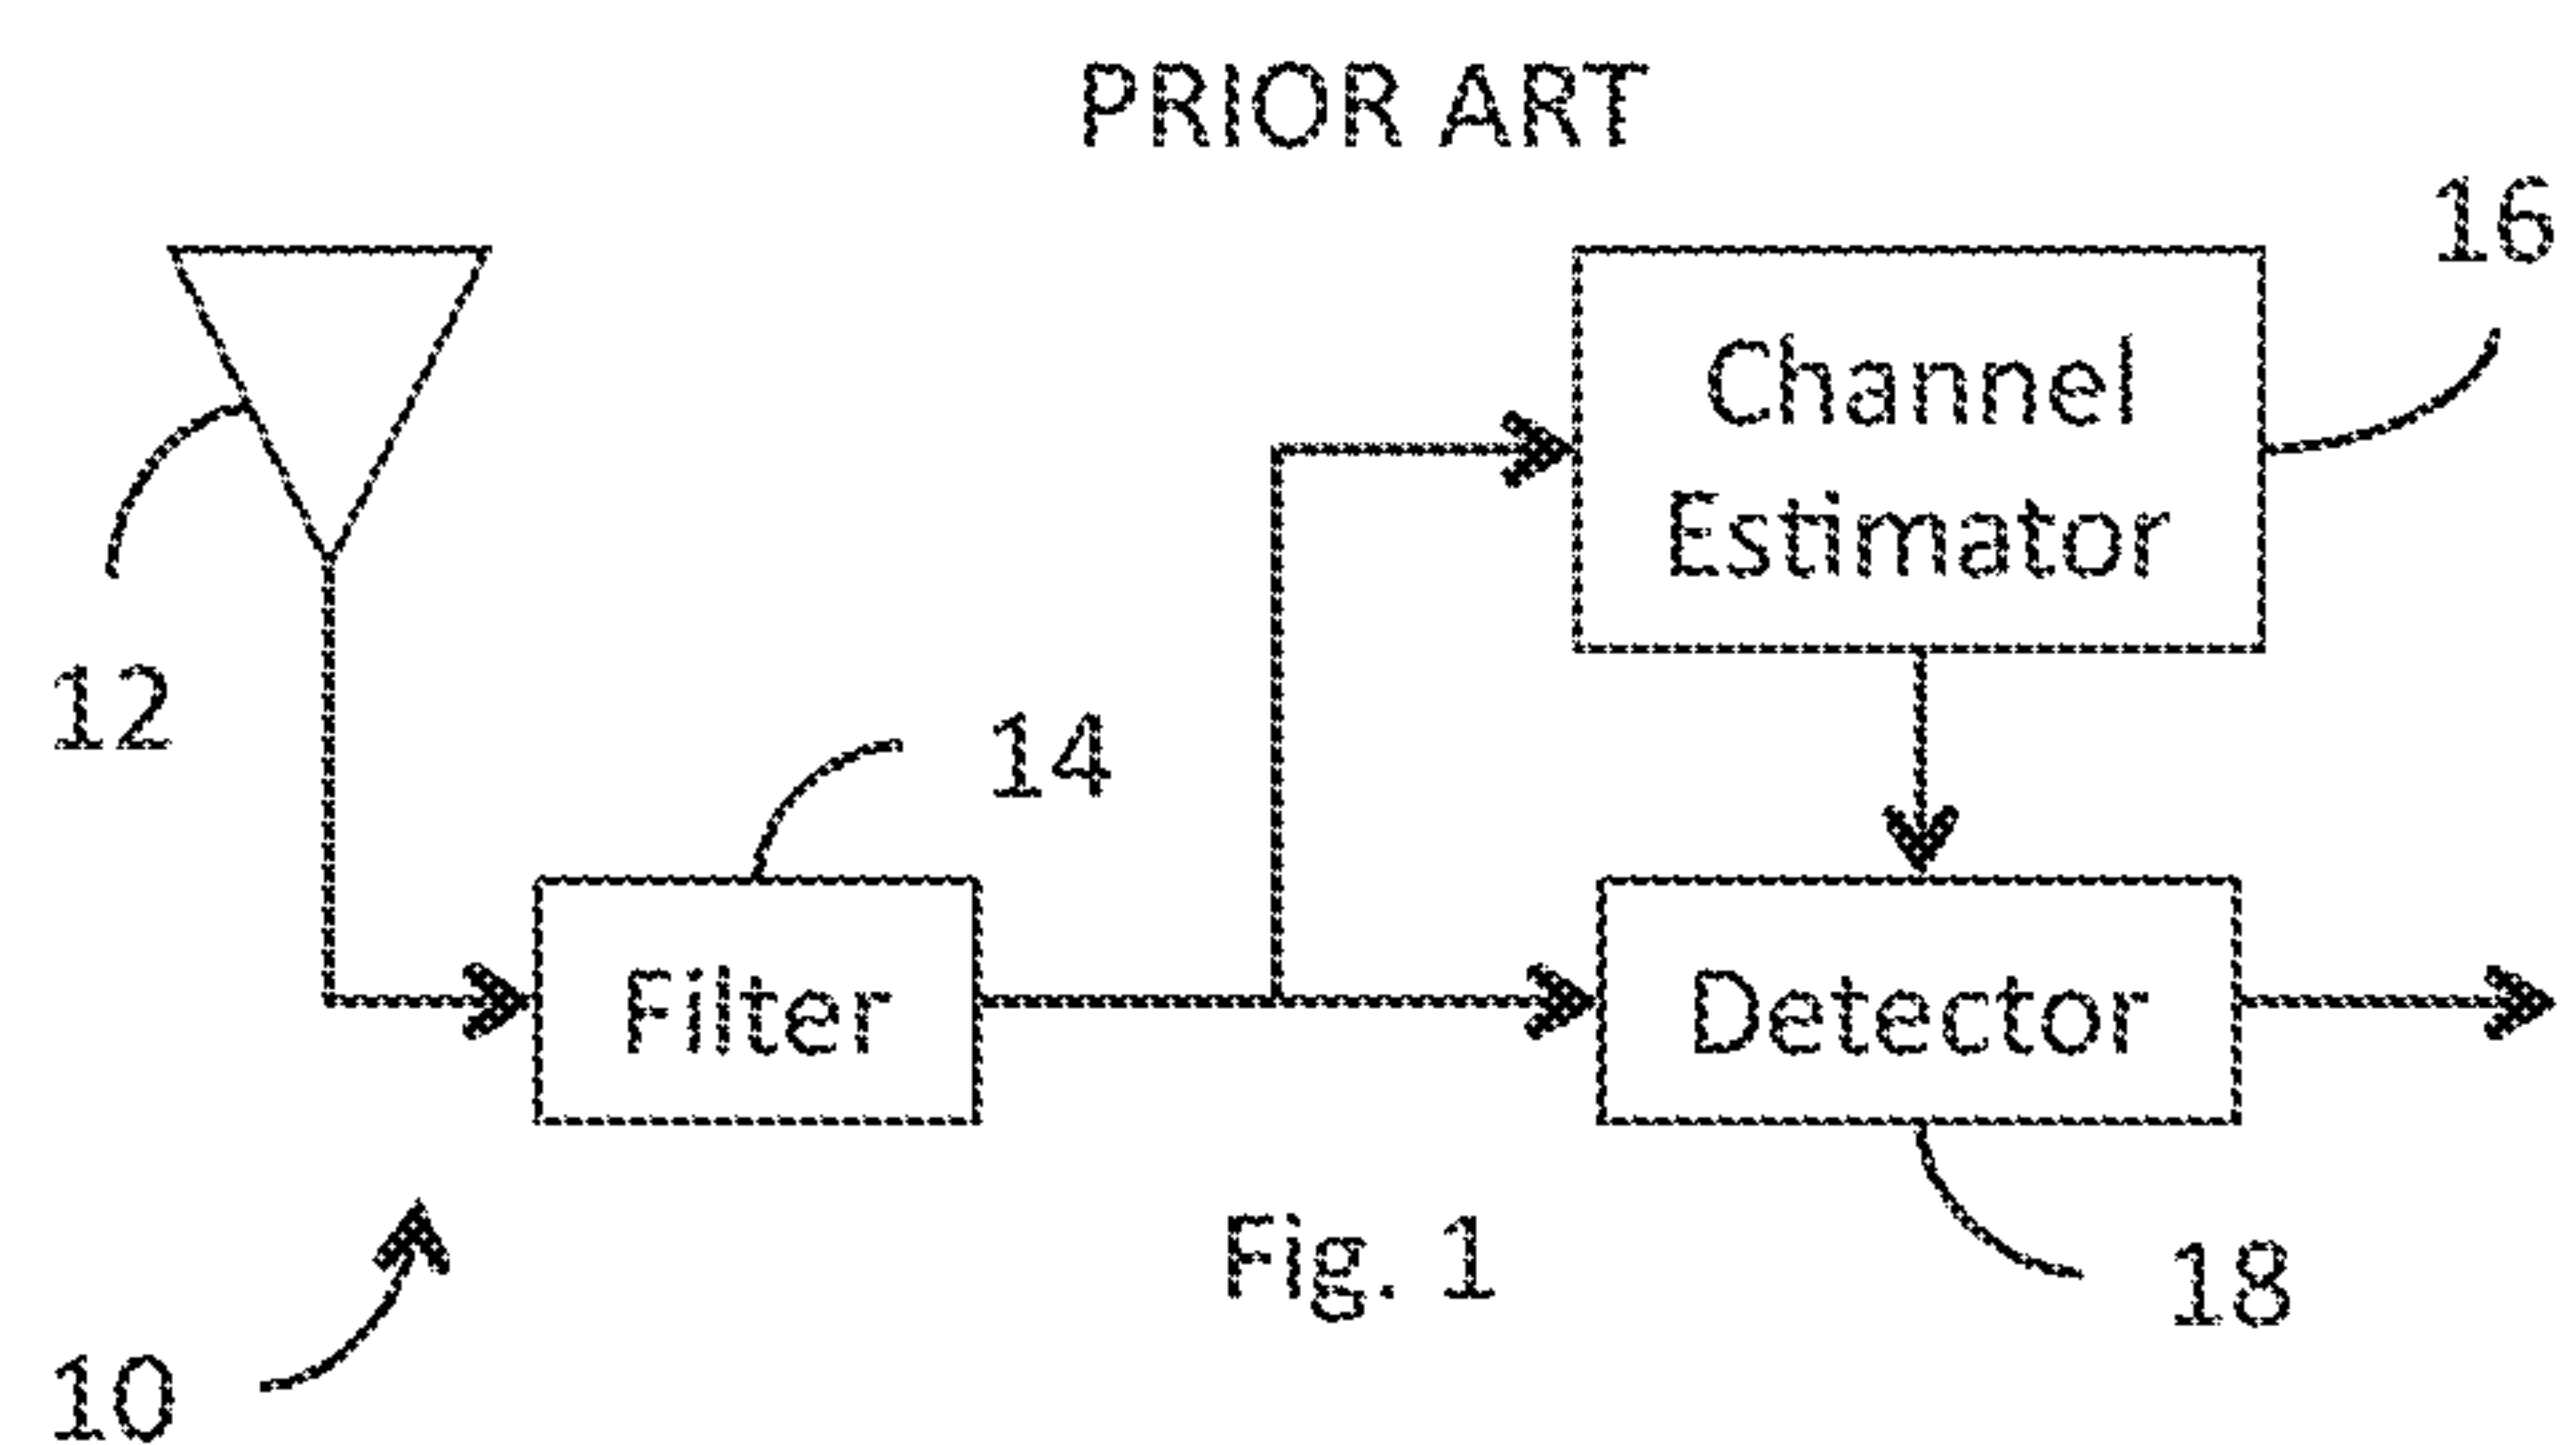
Final Office Action for U.S. Appl. No. 15/814,711, dated Aug. 29, 2019, 10 pages.


Non-Final Office Action for U.S. Appl. No. 15/868,026, dated Mar. 8, 2019, 11 pages.

Notice of Allowance for U.S. Appl. No. 15/868,026, dated Sep. 3, 2019, 8 pages.

First Office Action in Application No. 201480079211.7, dated Jun. 4, 2018, State Intellectual Property Office of the People's Republic of China, Related Office Action in Foreign Counterpart of DCWV019US0.

* cited by examiner










Fig. 4



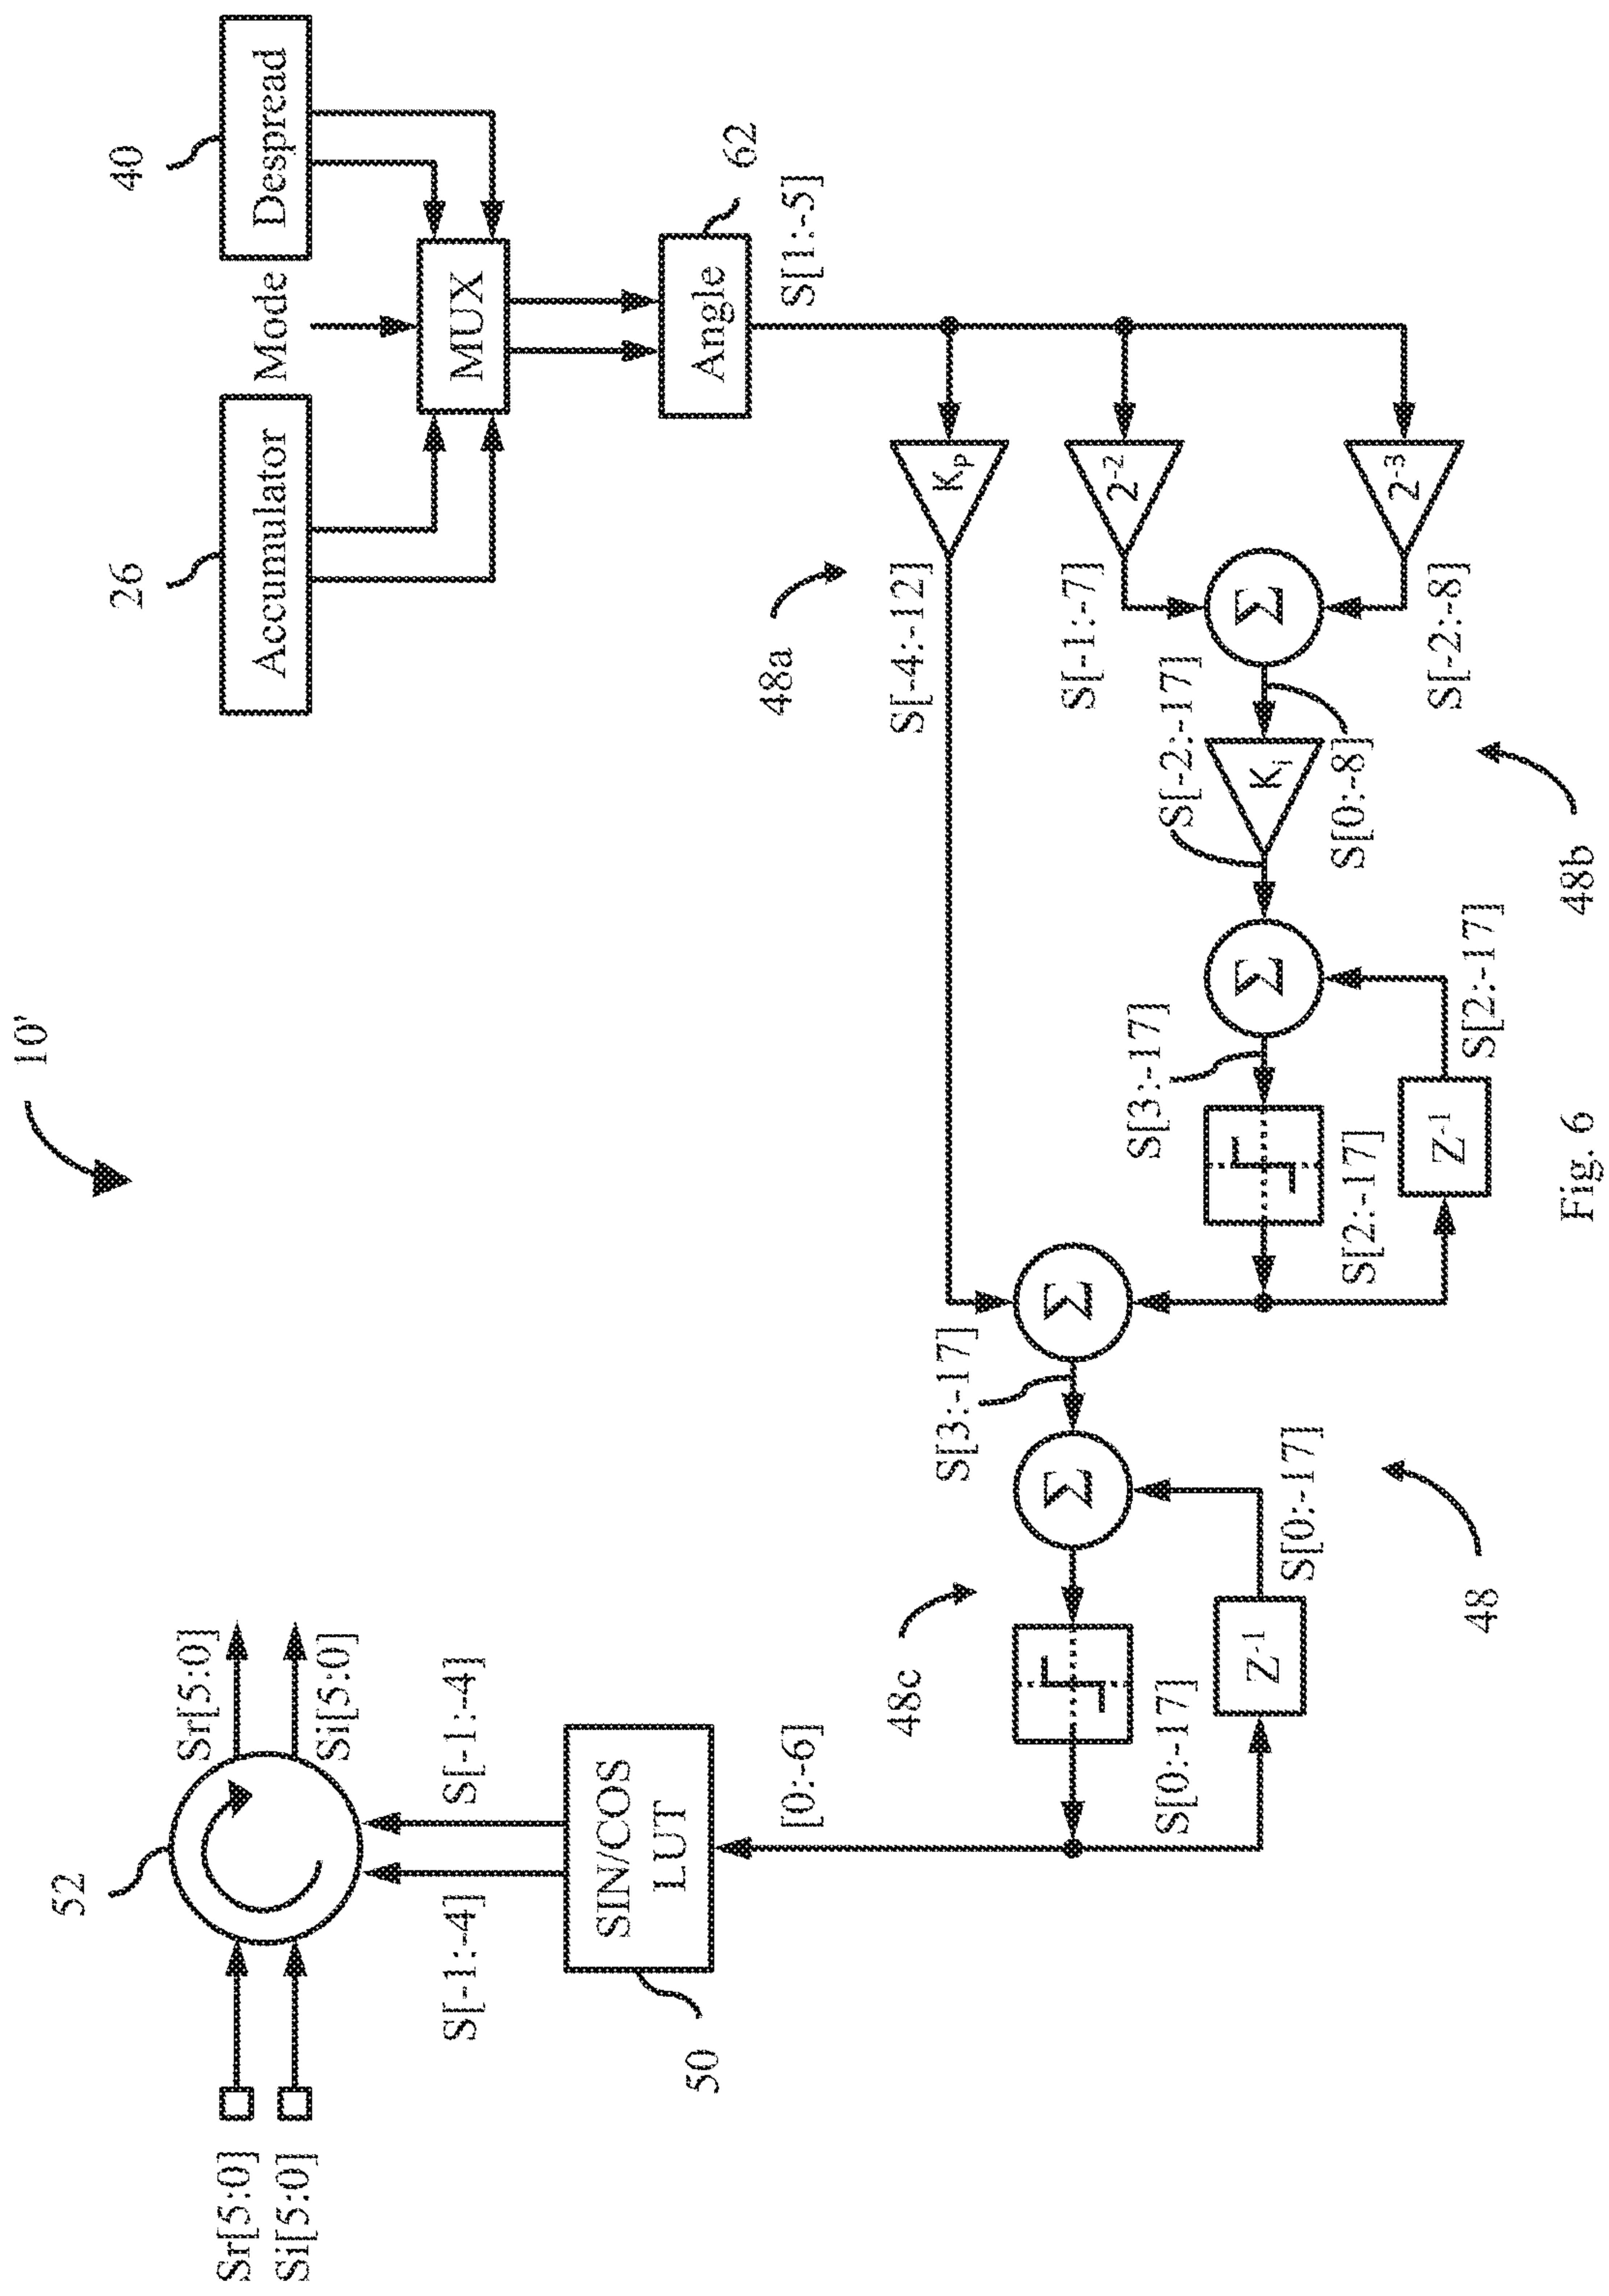


Fig. 6




Fig. 8




Fig. 9




Fig. 10

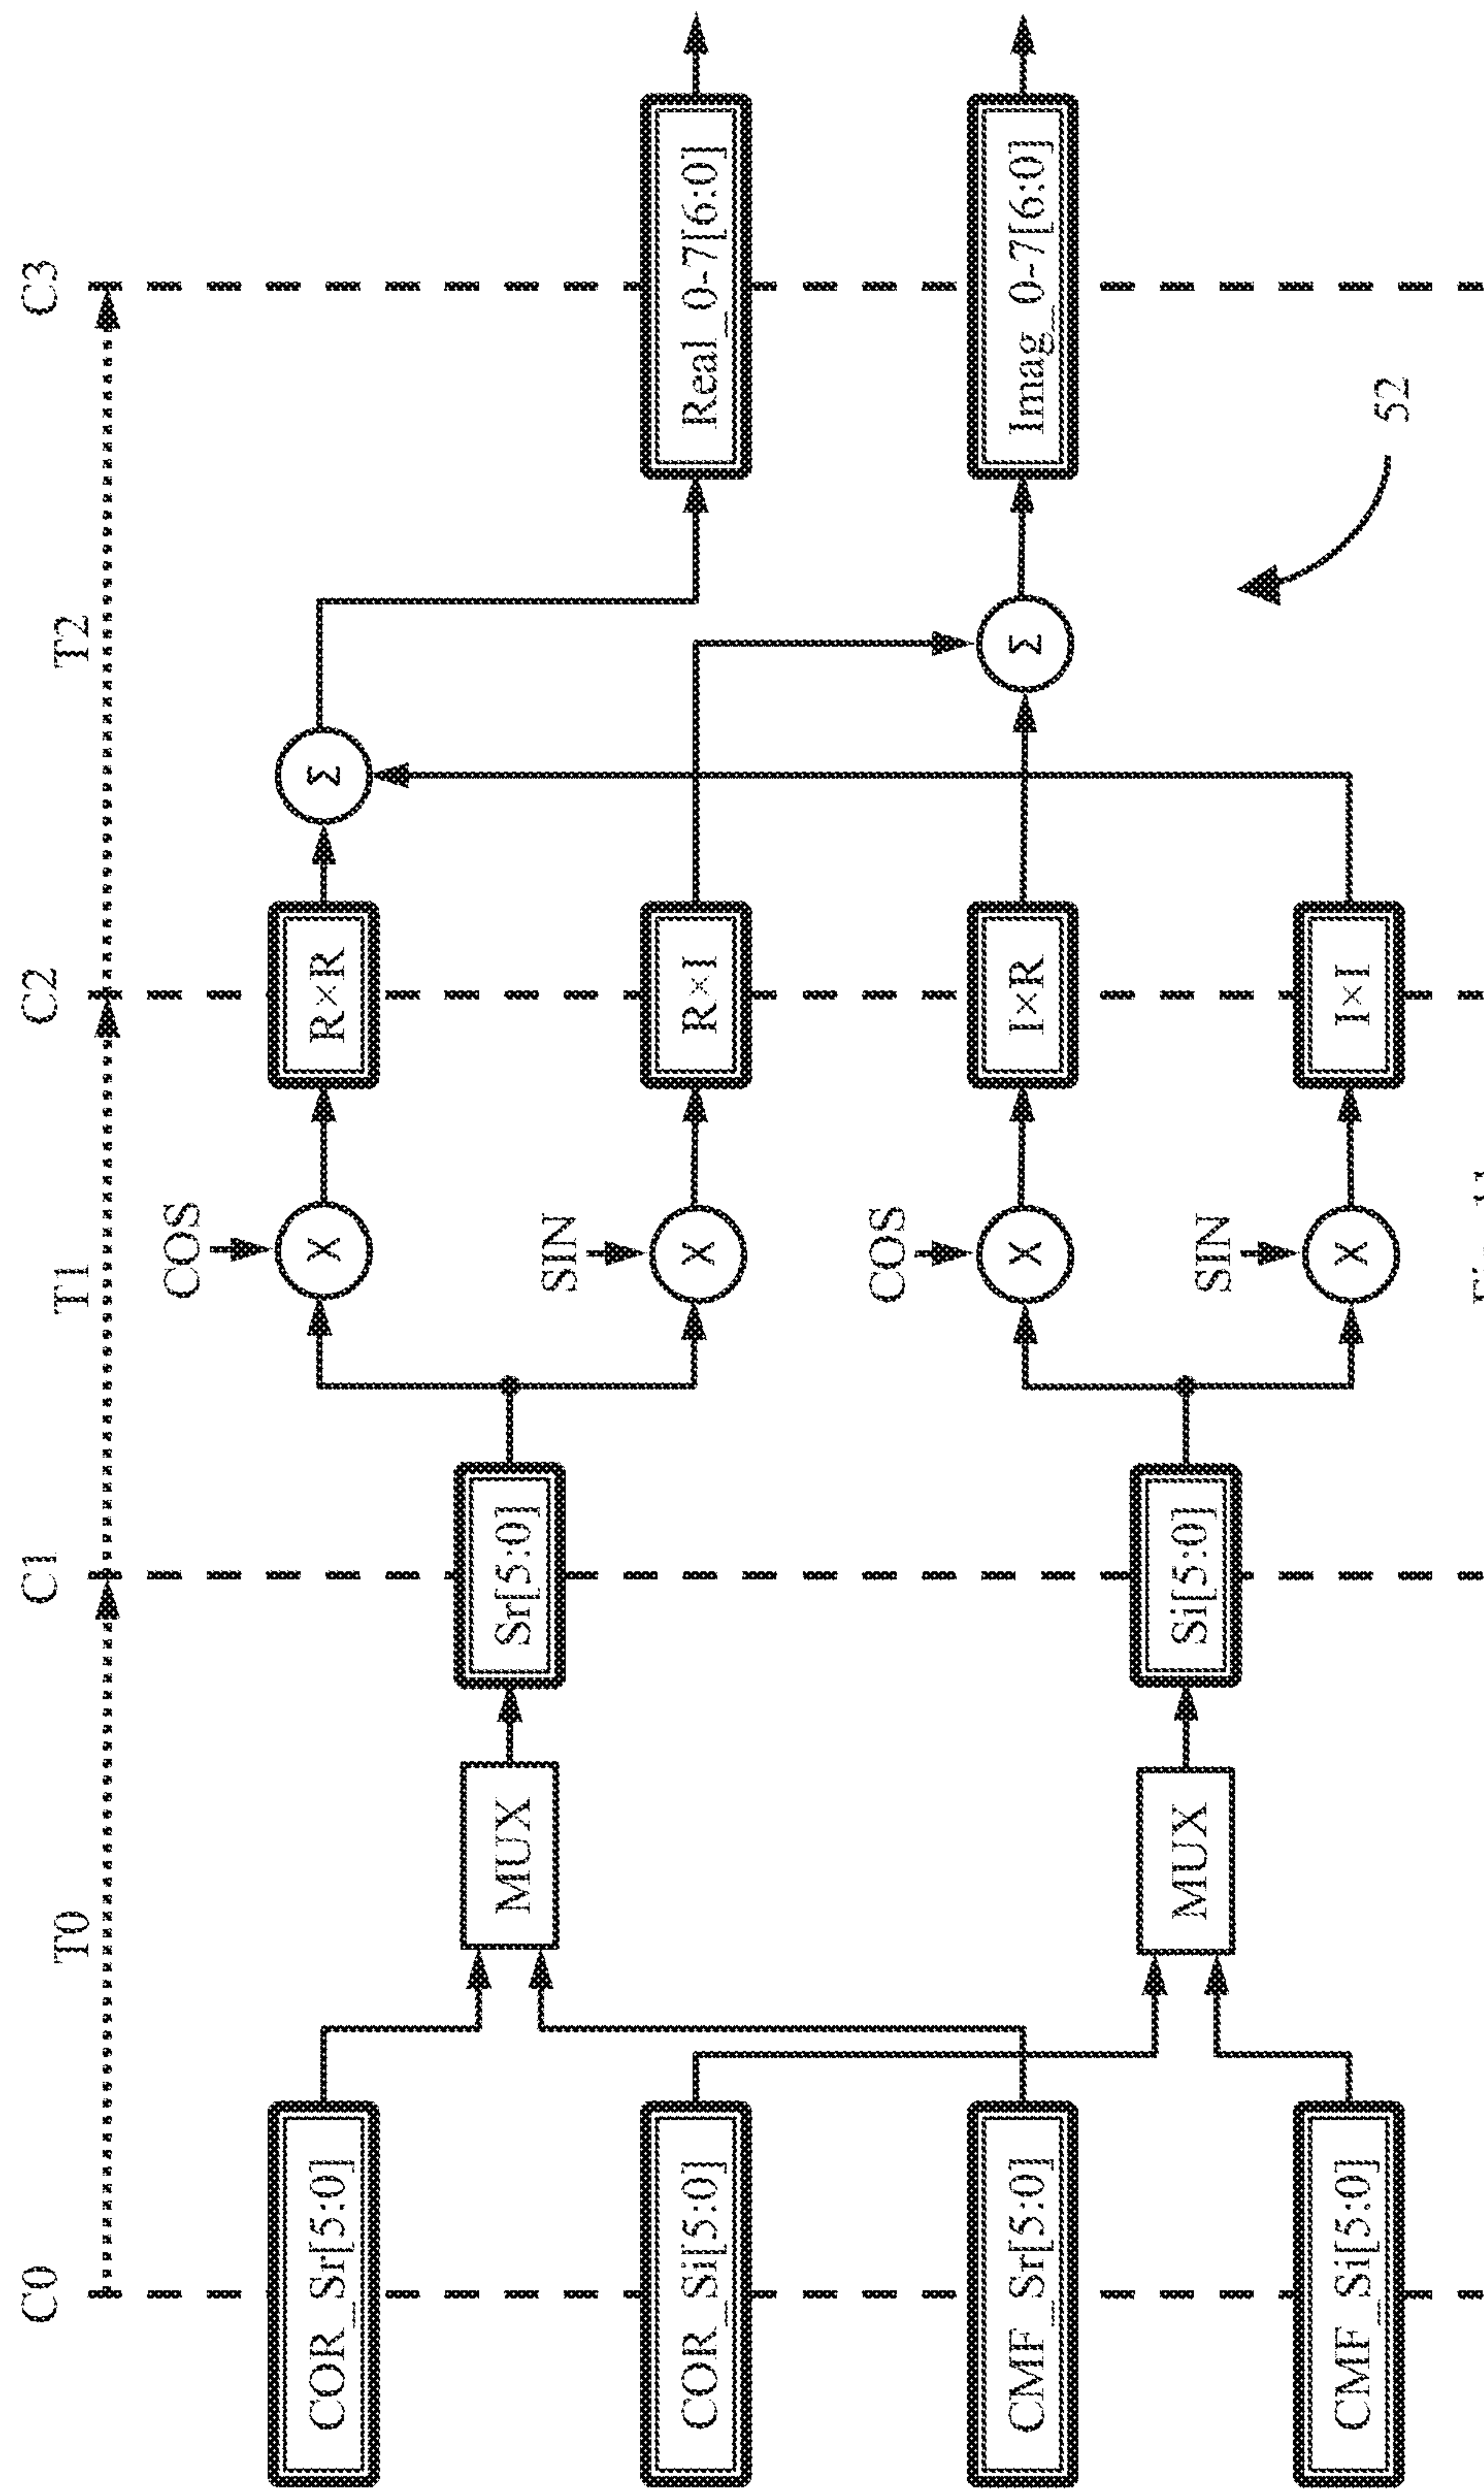


Fig. 11

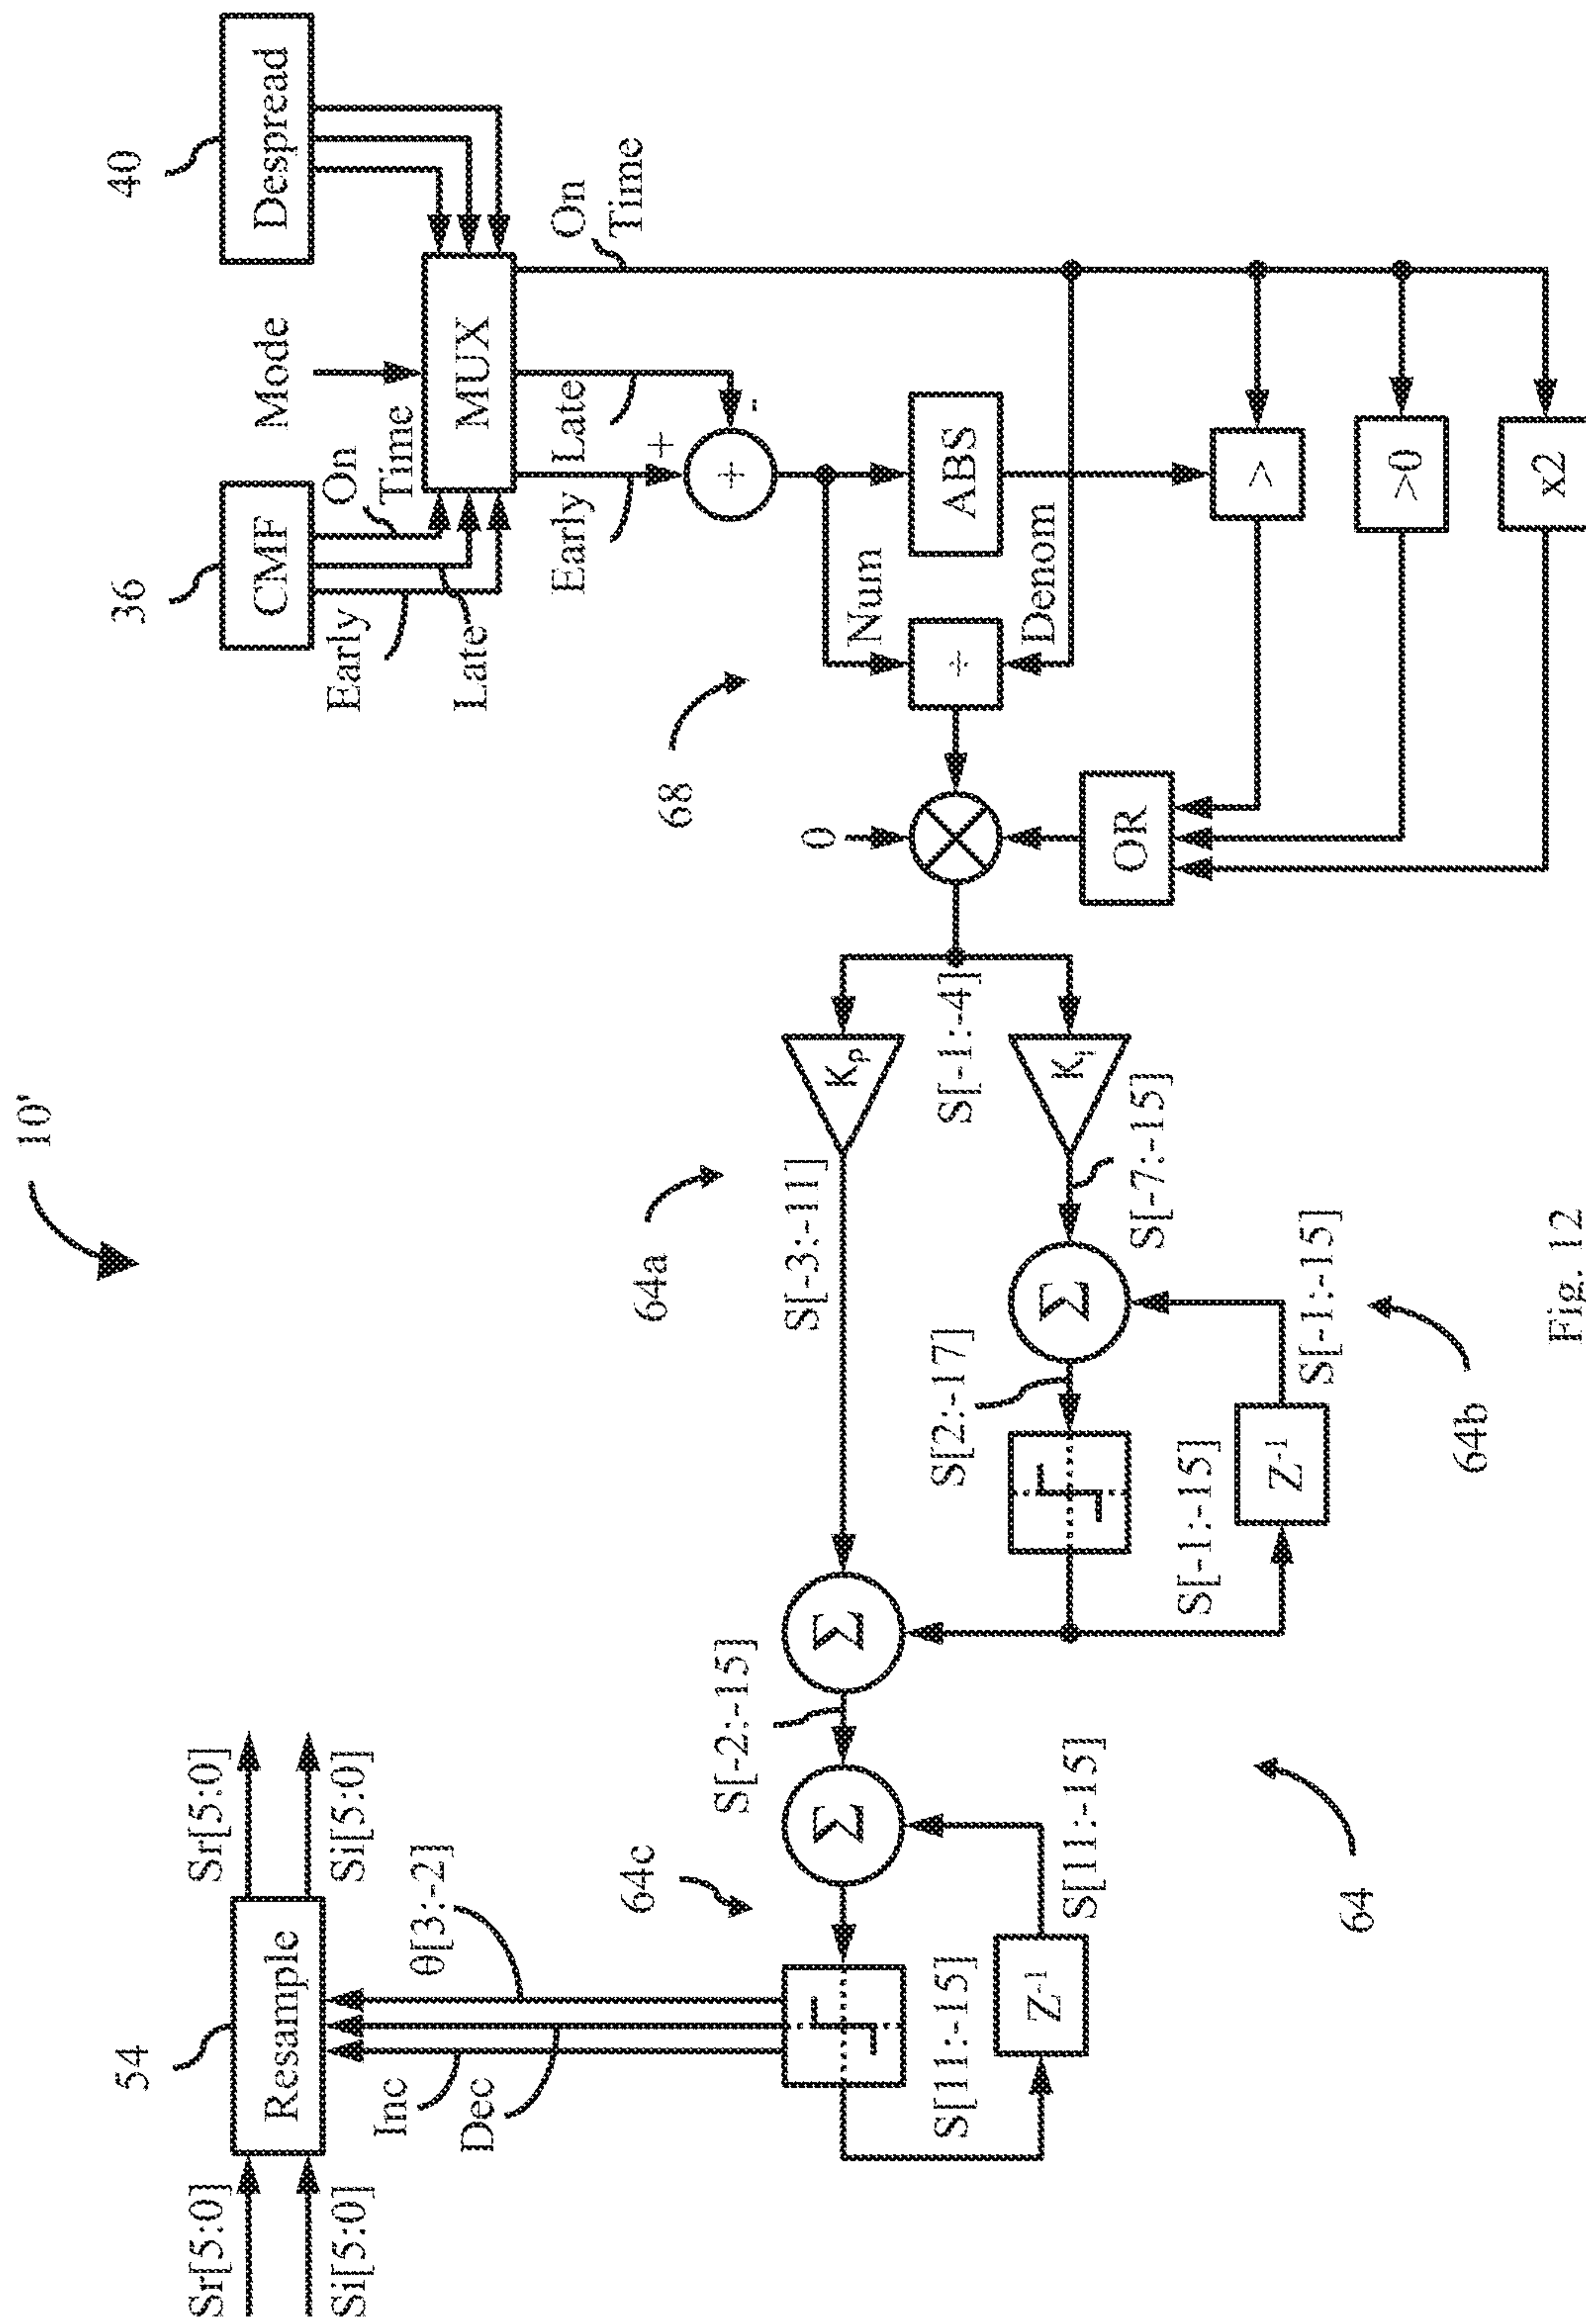


Fig. 12




Fig. 13

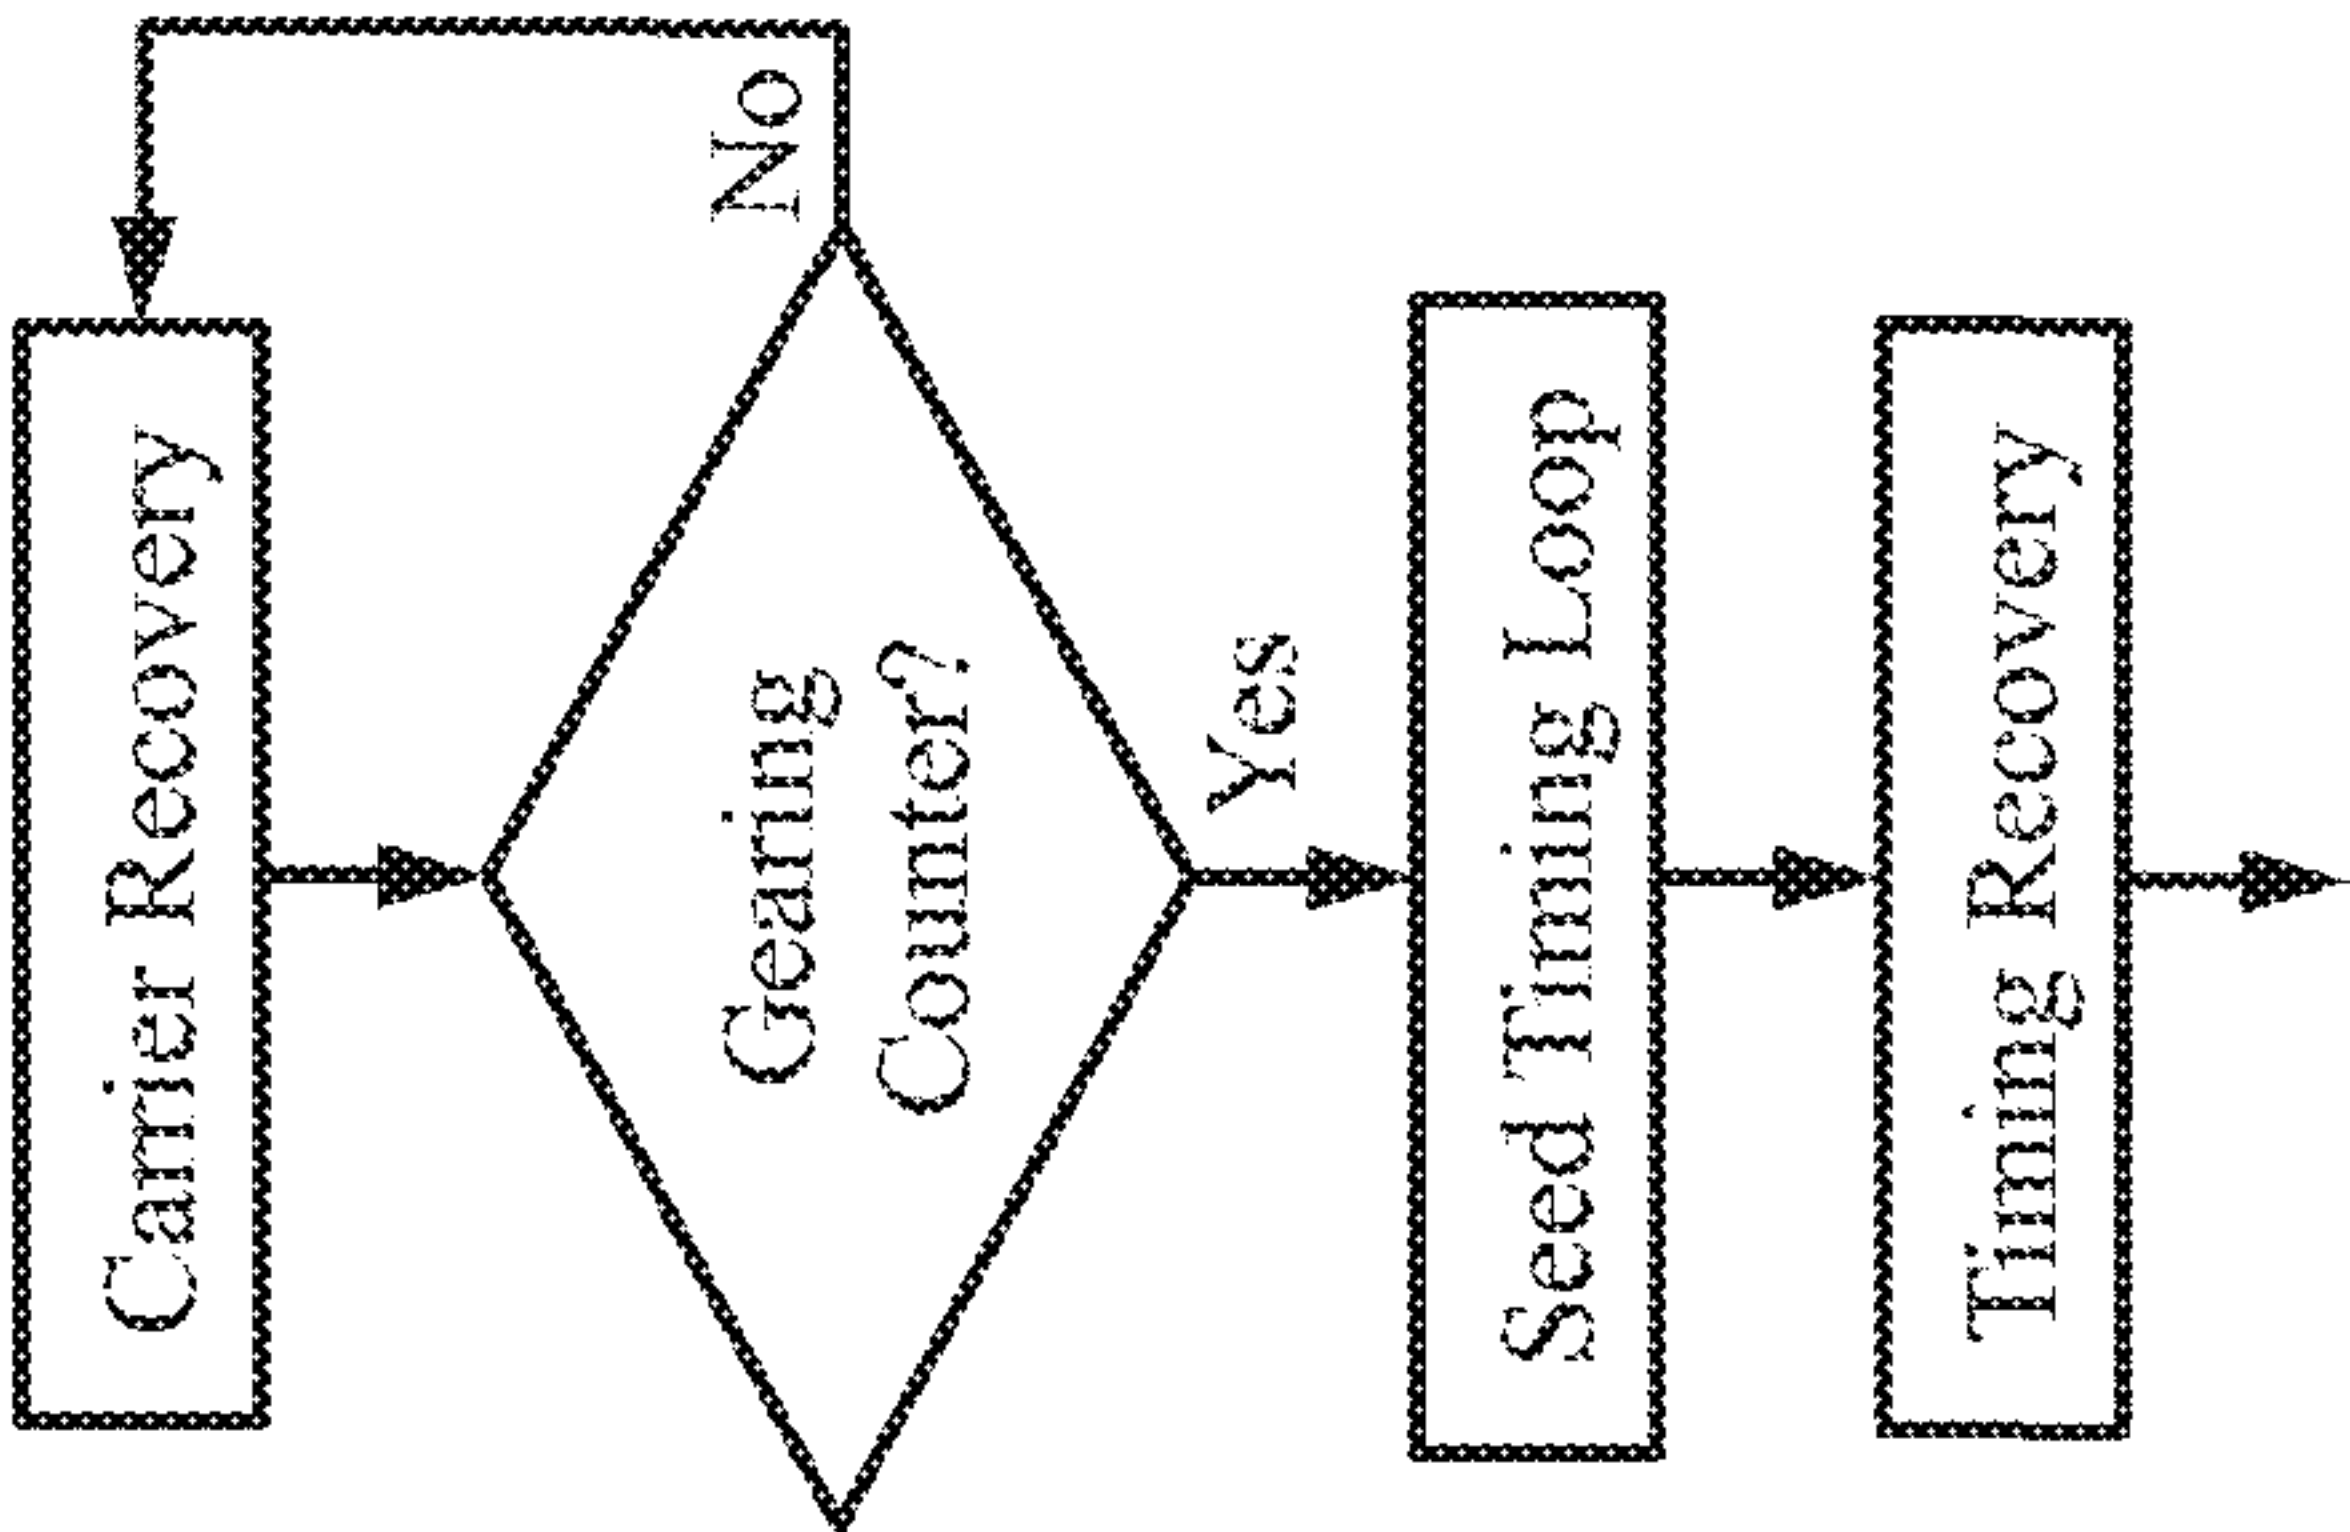


Fig. 14




Fig. 15

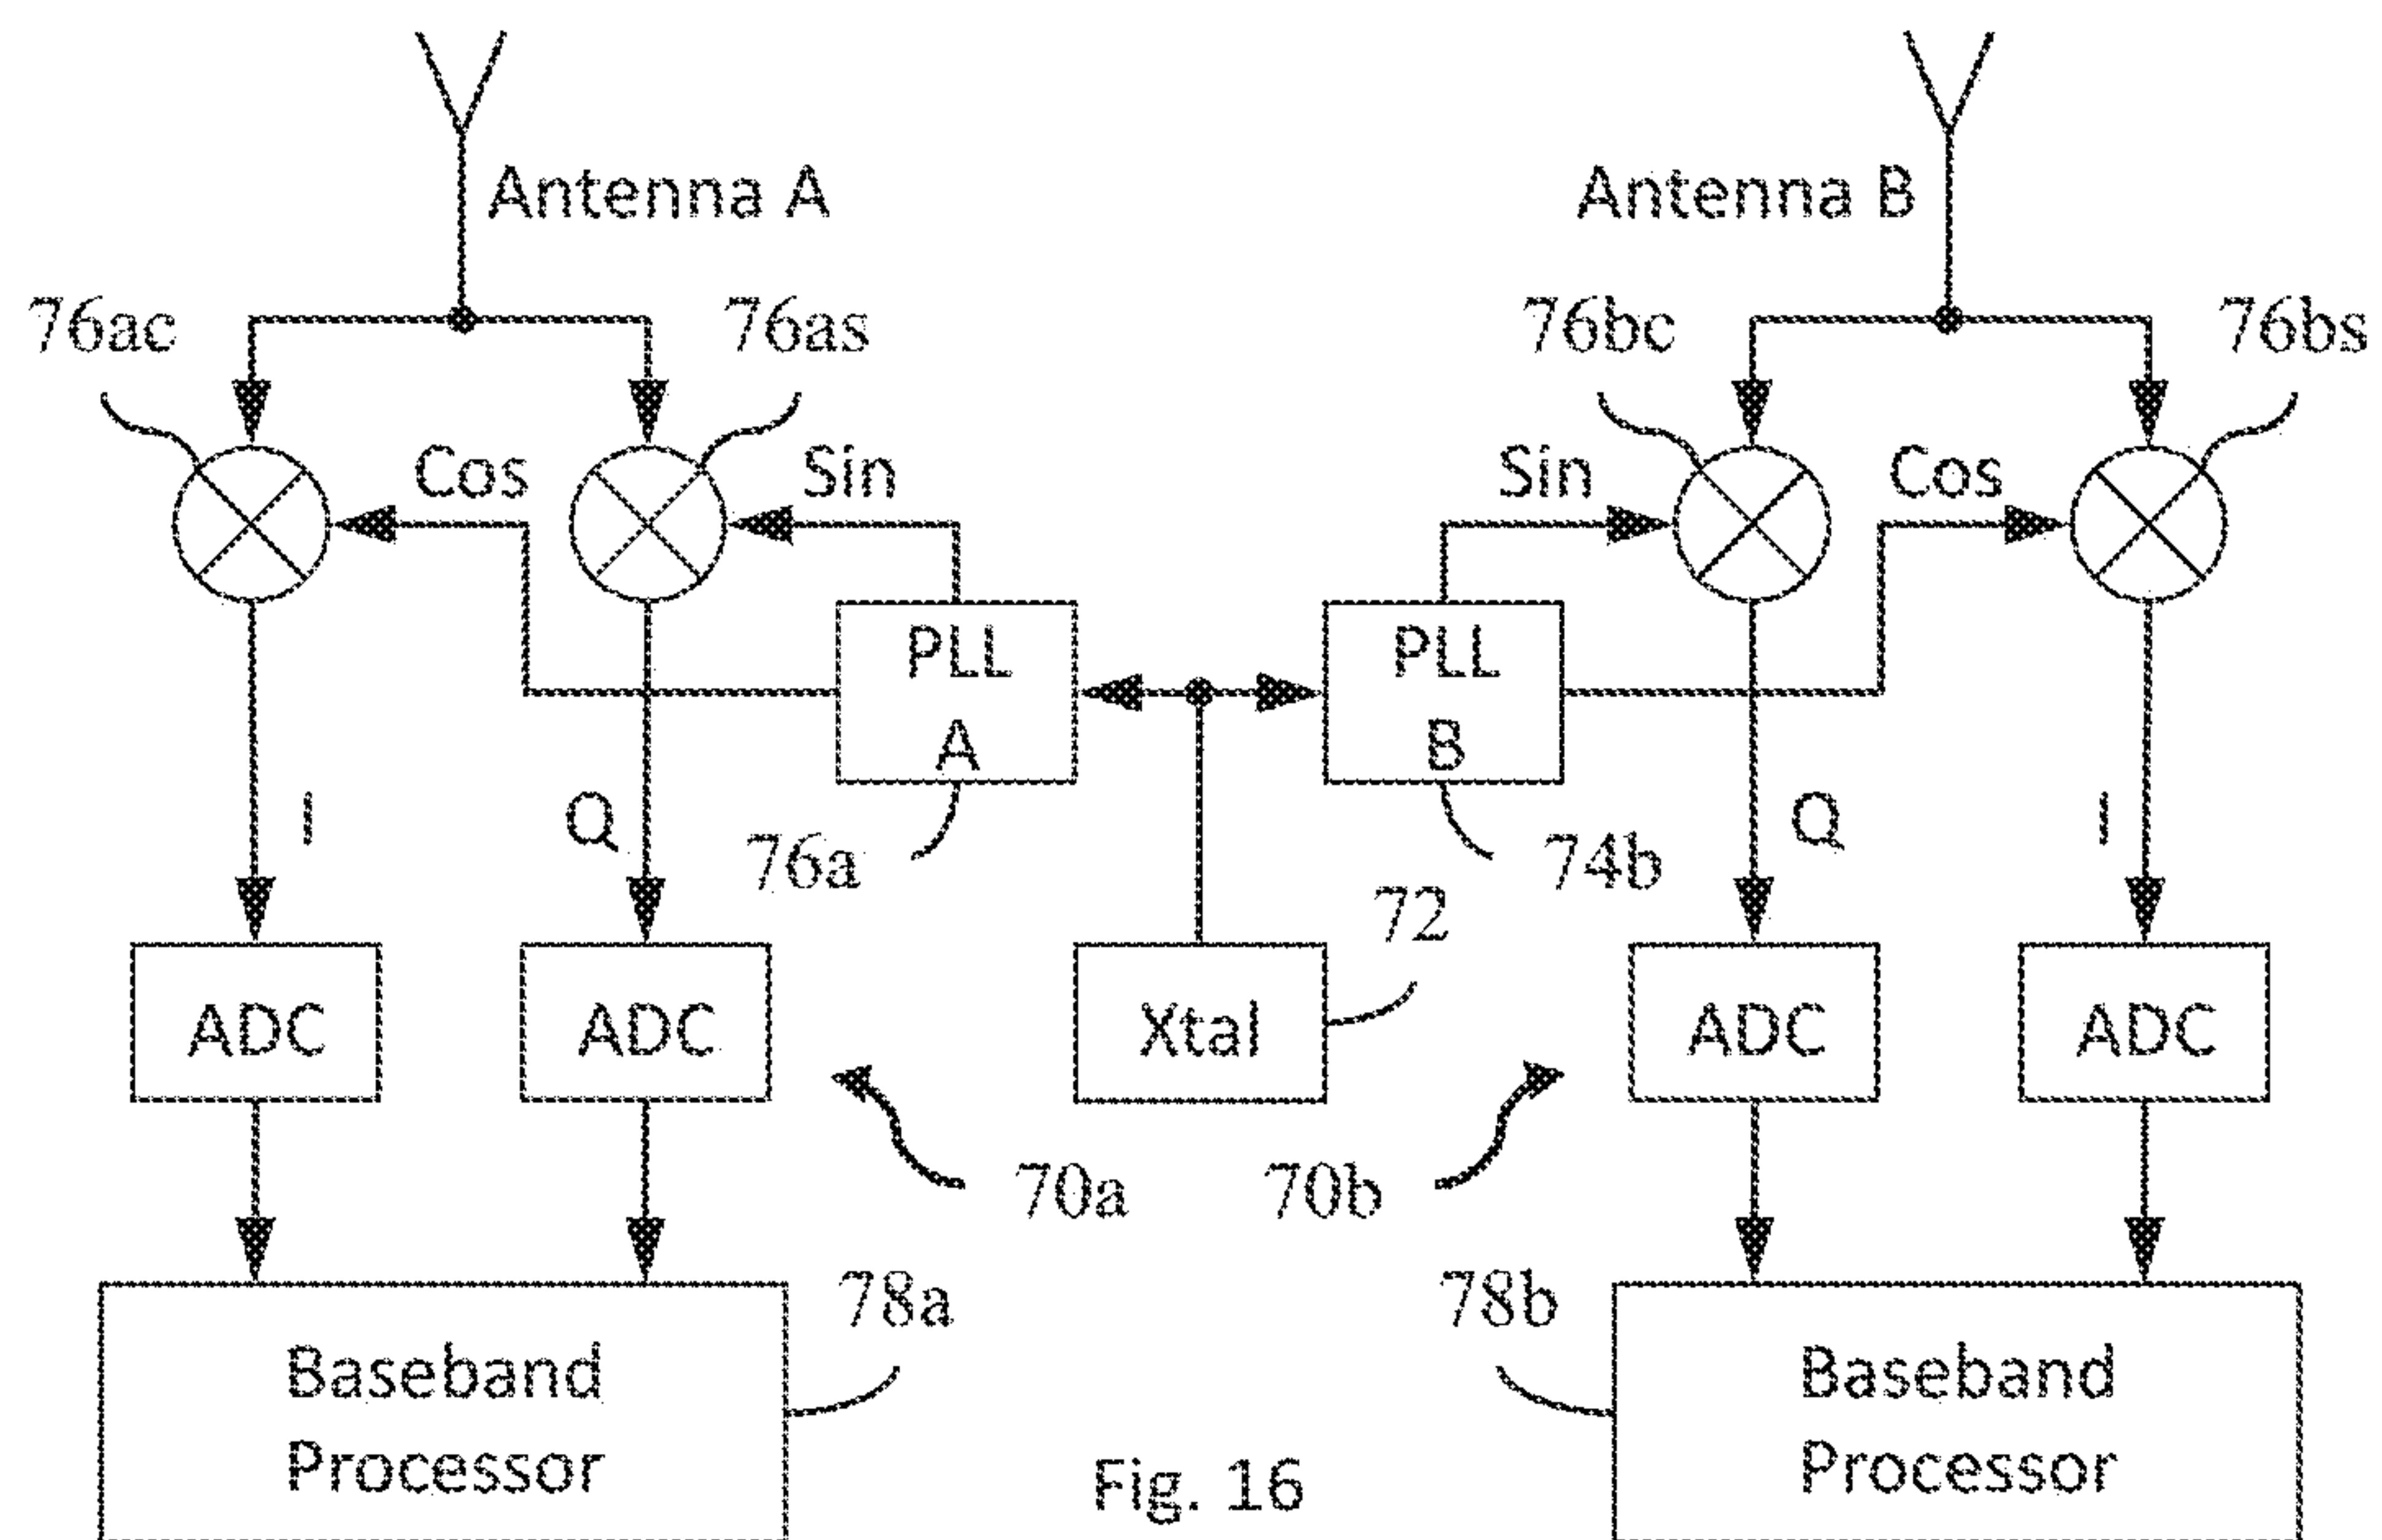


Fig. 16




Fig. 17: Angle of Incidence to two ½ wavelength separated antennas




Fig. 18: Angle of Incidence to two 1 wavelength separated antennas




Fig. 19: Angle of Incidence to two antennas separated by 3 wavelengths




Fig. 20: Angle of Incidence to two antennas separated by 7.5cm for two different center frequencies




Fig. 21

MEASURING ANGLE OF INCIDENCE IN AN ULTRAWIDEBAND COMMUNICATION SYSTEM

Matter enclosed in heavy brackets [] appears in the original patent but forms no part of this reissue specification; matter printed in italics indicates the additions made by reissue; a claim printed with strikethrough indicates that the claim was canceled, disclaimed, or held invalid by a prior post-patent action or proceeding.

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is the U.S. National Stage of PCT Application No. PCT/EP2014/060722, filed 23 May 2014 (“Parent PCT Application”).

The [subject matter of this] *Parent PCT Application is a Continuation-in-Part of and* is related to application Ser. No. 13/775,282, filed 28 Mar. 2013, *now U.S. Pat. No. 8,760,334, issued 24 Jun. 2014* (“First Related Application”).

The subject matter of this Application is related to application Ser. No. 13/033,098, filed 23 Feb. 2011 (“Second Related Application”), which is in turn related to Provisional Application Ser. No. 61/316,299, filed 22 Mar. 2010 (“Related Parent Provisional”).

The [subject matter of this] *First Related Application is a Continuation-in-Part of and* is related to application Ser. No. 12/885,517, filed 19 Sep. 2010, *now U.S. Pat. No. 8,437,432, issued 7 May 2013* (“Related Parent Patent”), which is in turn also related to the Related Parent Provisional.

The subject matter of this Application is also related to the subject matter of PCT Application Serial No. PCT/EP2013/070851, filed 7 Oct. 2013 (“Third Related Application”).

This application, and the subject matter herein, [is related] *claims priority to:*

1. The Parent PCT Application;
2. The First Related Application;
- [3. The Second Related Application;]
- [4] 3. The Related Parent Patent; *and*
- [5] 4. The Related Parent Provisional; [and]
- [6. The Third Related Application;]

collectively, “[Related] *Priority References*”, *and hereby claims the benefit of the filing date thereof pursuant to* 37 C.F.R. § 1.78.(a)(4).

The subject matter of the [Related] *Priority References, the Second Related Application, and the Third Related Application*, each in its entirety, is expressly incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to ultra-wideband communication systems, and, in particular, to a receiver for use in an ultra-wideband communication system adapted to determine the angle of arrival of an RF signal.

2. Description of the Related Art

In general, in the descriptions that follow, we will italicize the first occurrence of each special term of art which should be familiar to those skilled in the art of ultra-wideband (“UWB”) communication systems. In addition, when we first introduce a term that we believe to be new or that we will use in a context that we believe to be new, we will hold the term and provide the definition that we intend to apply to that term. In addition, throughout this description, we will

sometimes use the terms assert and negate when referring to the rendering of a signal, signal flag, status bit, or similar apparatus into its logically true or logically false state, respectively, and the term toggle to indicate the logical inversion of a signal from one logical state to the other. Alternatively, we may refer to the mutually exclusive boolean states as logic_0 and logic_1. Of course, as is well known, consistent system operation can be obtained by reversing the logic sense of all such signals, such that signals described herein as logically true become logically false and vice versa. Furthermore, it is of no relevance in such systems which specific voltage levels are selected to represent each of the logic states.

In general, in an ultra-wideband (“UWB”) communication system, a series of special processing steps are performed by a UWB transmitter to prepare payload data for transmission via a packet-based UWB channel. Upon reception, a corresponding series of reversing steps are performed by a UWB receiver to recover the data payload. Details of both series of processing steps are fully described in IEEE Standards 802.15.4 (“802.15.4”) and 802.15.4a (“802.15.4a”), copies of which are submitted herewith and which are expressly incorporated herein in their entirety by reference. As is known, these Standards describe required functions of both the transmit and receive portions of the system, but specify implementation details only of the transmit portion of the system, leaving to implementers the choice of how to implement the receive portion.

One or more of us have developed certain improvements for use in UWB communication systems, which improvements are fully described in the following pending applications or issued patents, all of which are expressly incorporated herein in their entirety:

“A Method and Apparatus for Generating Codewords”, U.S. Pat. No. 7,787,544, issued 31 Jul. 2010;

“A Method and Apparatus for Generating Codewords”, application Ser. No. 11/309,222, filed 13 Jul. 2006, now abandoned;

“A Method and Apparatus for Transmitting and Receiving Convolutionally Coded Data”, U.S. Pat. No. 7,636,397, issued 22 Dec. 2009;

“A Method and Apparatus for Transmitting and Receiving Convolutionally Coded Data”, U.S. Pat. No. 8,358,709, issued 22 Jan. 2013; and

“Convolution Code for Use in a Communication System”, U.S. Pat. No. 8,677,224, issued 18 Mar. 2014.

One particular problem in multi-path, spread-spectrum systems, including UWB, is channel-induced noise present in the received signal. One common technique for significantly reducing the noise level relative to the receive level is to develop, during reception of a training sequence portion of the preamble of each transmitted packet, an estimate of the channel impulse response (“CIR”). Following detection in the received packet of the start-of-frame delimiter (“SFD”), the best CIR estimate is reversed in time and the complex conjugate is developed. This conjugate CIR estimate is thereafter convolved with the payload portion of the packet using a channel matched filter (“CMF”). Shown in FIG. 1 is a UWB receiver 10 adapted to operate in this manner. As is known, the signal received via an antenna 12 is continuously conditioned by a filter 14. During reception of the training sequence, channel estimator 16 develops from the conditioned signal the conjugate CIR estimate. During reception of the payload data, detector 18 employs a CMF (not shown) to convolve the conditioned signal with the conjugate CIR estimate, thereby significantly improving the signal-to-noise ratio (“SNR”) and facilitating recovery of

the payload data. See, also, "Efficient Back-End Channel Matched Filter (CMF)", U.S. Pat. No. 7,349,461, issued 25 Mar. 2008.

As noted in 802.15.4a, § 5.5.7.1, "UWB devices that have implemented optional ranging support are called ranging-capable devices ("RDEVs")." (Emphasis in original.) For certain applications, such RDEVs are commonly implemented in the form of a relatively compact, autonomous radio-frequency identification ("RFID") tag or the like. Due to the small form factor and limited power supply, it is especially important to select circuit implementations that provide maximum performance at minimum power. Unfortunately, in known implementations of the UWB receiver, improvements in performance usually come at the expense of power. For example, it is known that a rake filter provides good performance in multi-path, spread-spectrum systems such as UWB. See, e.g., slide 21 of "The ParthusCeva Ultra Wideband PHY Proposal", IEEE P802.15 Working Group for Wireless Personal Area Networks, March 2003, a copy of which is submitted herewith and which is expressly incorporated herein in its entirety by reference. However, known rake filter implementations tend to consume significantly more power than other prior art techniques.

In ranging systems, as in other RF systems, the receiver must coordinate its internal operation to the signal being received from the transmitter. In general, the receiver must achieve synchronism with the received carrier signal, a process referred to as carrier recovery. In addition, the receiver must further achieve synchronism with the information signals superimposed on the carrier, a process referred to as timing recovery. We submit that prior art techniques for performing both carrier recovery and timing recovery in the digital domain are less than optimum.

In the RF system topology shown in FIG. 15, it can be seen that, because of the non-zero angle of incidence, θ , the RF signal will arrive at one antenna before the other. In particular, it can be seen that the path to antenna A is greater than to antenna B by $p=d*\sin(\theta)$. In order to calculate θ , the angle of incidence, the time difference of arrival could be found. If d is relatively large then this would provide quite an accurate estimate of θ . On the other hand, if d is small the estimate turns out to be highly error prone.

FIG. 16 shows two receivers, 70a and 70b, which are clocked from the same crystal 72. If the same crystal 72 clocks identical phase locked loops ("PLLs"), 74a and 74b, the generated carriers that are supplied to the respective down converter mixers, 76ac-76as and 76bc-76bs, will have the same phase. The RF signal will arrive at a slightly later time at antenna A than antenna B, so it will encounter a down converter carrier phase that is different in each of the mixers 76. If the baseband processors, 78a and 78b, are capable of calculating the complex impulse response of the channel, that impulse response will have a different in-phase ("I") to quadrature ("Q") ratio I/Q which is equal to the phase delay caused by the signal travelling the extra distance, p , before encountering the mixer 70a and being down-converted by the carrier. If the carrier frequency is high, e.g., 4 GHz or 6.5 GHz, then quite small distances, p , will lead to a relatively large carrier phase difference.

$$\sin\theta = \frac{p}{d} \quad [\text{Eq. 1}]$$

$$\lambda = \frac{c}{f} \quad [\text{Eq. 2}]$$

where:

f is the carrier frequency,
 c is the speed of light, and
 λ is the carrier wavelength.

$$\frac{\alpha}{2\pi} = \frac{\rho}{\lambda} \quad [\text{Eq. 3}]$$

where:

α is the phase difference between the two carriers for the same point on the incident RF signal.

$$p = \frac{\alpha\lambda}{2\pi} \quad (\text{from Eq. 2 and Eq. 3}) \quad [\text{Eq. 4}]$$

$$\sin\theta = \frac{\alpha\lambda}{2\pi d} \quad (\text{from Eq. 1 and Eq. 4}) \quad [\text{Eq. 5}]$$

$$\theta = \sin^{-1} \frac{\alpha\lambda}{2\pi d} \quad (\text{from Eq. 5}) \quad [\text{Eq. 6}]$$

If, in Eq. 6, d is set to be a half wavelength, then FIG. 17 shows the relationship between α , the phase difference of the impulse responses, and θ , the angle of incidence. Note that the slope of the dark grey section is approximately 3, whereas the slope of the lighter grey section is 0.6, i.e., 5 times worse. If, however, d is set to be a one wavelength, then FIG. 18 shows the relationship between α and θ . Note that, at this separation, there is an ambiguity in that each phase relationship has two possible angles of incidence. As can be seen from FIG. 19, as the antennae are moved further apart, say to 3 wavelengths, the ambiguity only increases.

We submit that the larger separation of one wavelength or more is advantageous for two reasons: first, the slope of the angle of incidence curve versus phase change curve is larger and stays larger for longer, thereby allowing more accurate determination of angle of incidence; and second, as the antennas get closer together, their near fields interfere and their performance starts to affect each other. This is particularly the case when the separation is lower than one wavelength.

We submit that what is needed is an improved method and apparatus for use in the receiver of a UWB communication system to determine angle of incidence. In particular, we submit that such a method and apparatus should provide performance generally comparable to the best prior art techniques but more efficiently than known implementations of such prior art techniques.

BRIEF SUMMARY OF THE INVENTION

In accordance with a preferred embodiment of our invention, we provide a method for use in an RF system comprising first and second RF receivers separated by a predetermined distance, d , and an RF transmitter, the method comprising the steps of: [1] in the first and second receivers, synchronizing the first and second receivers to a predetermined time base; [2] in the transmitter, transmitting an RF signal having a predetermined carrier wavelength, λ ; [3] in the first receiver, receiving the transmitted signal and developing a first phase value as a function of the complex baseband impulse response of the received signal; [4] in the second receiver, receiving the transmitted signal and developing a second phase value as a function of the complex baseband impulse response of the received signal; [5] devel-

5

oping a phase difference value, α , as a function of the first and second phase values; and [6] developing an angle of arrival, θ , of the transmitted signal relative to the first receiver according to the following:

$$\theta = \sin^{-1} \frac{\alpha\lambda}{2\pi d}.$$

The methods of our invention may be embodied in computer readable code on a suitable computer readable medium such that when a processor executes the computer readable code, the processor executes the respective method.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

Our invention may be more fully understood by a description of certain preferred embodiments in conjunction with the attached drawings in which:

FIG. 1 illustrates, in block diagram form, a prior art receiver adapted for use in a UWB communication system;

FIG. 2 illustrates, in block diagram form, one embodiment of the receiver shown in FIG. 1, but constructed in accordance with our invention;

FIG. 3 illustrates, in block diagram form, the portion of the receiver shown in FIG. 2 that performs carrier and timing recovery when the receiver is operating in a carrier acquisition mode;

FIG. 4 illustrates, in block diagram form, the portion of the receiver shown in FIG. 2 that performs carrier and timing recovery when the receiver is operating in a data recovery mode;

FIG. 5 illustrates, in block diagram form, the error angle calculation block shown, e.g., in FIG. 3;

FIG. 6 illustrates, in relative isolation, those components of the receiver shown in FIG. 3 adapted to perform carrier recovery;

FIG. 7 illustrates, in block diagram form, a programmable scaler adapted for use with our carrier loop filter shown in FIG. 6;

FIG. 8 illustrates, in signal wave form, our approach to scale the filtered phase error estimate and apply it to the accumulator more often than once per symbol in the case of 850 kb/s;

FIG. 9 illustrates, in signal wave form, our approach to scale the filtered phase error estimate and apply it to the accumulator more often than once per symbol in the case of 6.8 Mb/s;

FIG. 10 illustrates, in signal wave form, our approach to scale the filtered phase error estimate and apply it to the accumulator more often than once per symbol in the case of 110 kb/s;

FIG. 11 illustrates, in block diagram form, a rotator adapted for use with our carrier loop filter shown in FIG. 6;

FIG. 12 illustrates, in relative isolation, those components of the receiver shown in FIG. 4 adapted to perform timing recovery;

FIG. 13 illustrates, in block diagram form, a resampler adapted for use with our timing recovery loop shown in FIG. 12;

FIG. 14 illustrates, in flow diagram form, a process for seeding the timing loop filter shown in FIG. 12.

FIG. 15 illustrates, generally in topographic perspective, an RF communication system, and, in particular, illustrates

6

the different angles of incidence of the transmitted RF signal on two antennas spaced apart by a distance d ;

FIG. 16 illustrates, in block diagram form, the antennas of FIG. 15, together with the respective RF receivers;

FIG. 17 illustrates, in waveform, the ratio of the phase shift of the carrier as a function of angle of incidence with respect to two $\frac{1}{2}$ wavelength separated antennas;

FIG. 18 illustrates, in waveform, the ratio of the phase shift of the carrier as a function of angle of incidence with respect to two 1 wavelength separated antennas;

FIG. 19 illustrates, in waveform, the ratio of the phase shift of the carrier as a function of angle of incidence with respect to two antennas separated by 3 wavelengths;

FIG. 20 illustrates, in waveform, the ratio of the phase shift of the carrier as a function of angle of incidence with respect to two antennas separated by 7.5 cm at two different carrier frequencies; and

FIG. 21 illustrates, in waveform, the resultant calculated angle of incidence for a 100 packet test.

In the drawings, similar elements will be similarly numbered whenever possible. However, this practice is simply for convenience of reference and to avoid unnecessary proliferation of numbers, and is not intended to imply or suggest that our invention requires identity in either function or structure in the several embodiments.

DETAILED DESCRIPTION OF THE INVENTION

Shown in FIG. 2 is a UWB receiver 10' constructed in accordance with our invention. As in the prior art system shown in FIG. 1, the signal received by antenna 12 is continuously conditioned by filter 14. The conditioned signal is then periodically sampled by an analog-to-digital converter ("ADC") 20 and provided as a continuous series of digital samples. In accordance with a preferred embodiment of our invention, ADC 20 is specially adapted to provide each digital sample in ternary form, i.e., $[-1, 0, +1]$. In view of the difficulty of currently available standard digital circuit technology efficiently to represent a 3-value variable in the form of a single ternary trit, we anticipate, at least in the near term, such variables will require representation using 2 conventional, binary bits, wherein a first one of the bits represents the numeric component of the variable, i.e., $[0, 1]$, and the second bit represents the sign of the variable, i.e., $[+, -]$. In this regard, it could be argued that circuit technology has not progressed all that much since Soviet researchers built the first (perhaps only?) documented ternary-based computer systems. See, "A Visit to Computation Centers in the Soviet Union," Comm. of the ACM, 1959, pp. 8-20; and "Soviet Computer Technology—1959", Comm. of the ACM, 1960, pp. 131-166; copies of which are submitted herewith and which are expressly incorporated herein in their entirety by reference.

In the context of our invention, our trit can be distinguished from a conventional sign+magnitude implementation such as that described in Amoroso83, cited above. Consider the strategy for A/D conversion shown in FIG. 5 of Amoroso83; and, note, especially, that there are three separate and distinct switching thresholds: (i) a sign threshold $[T_o]$; (ii) a positive magnitude threshold $[T_o + \Delta]$; and (iii) a negative magnitude threshold $[T_o - \Delta]$. (See, also, Amoroso83, p. 1119, lines 21-24.) We have discovered that adapting the ADC to use ONLY a positive magnitude threshold $[T_o + \Delta]$ and a negative magnitude threshold $[T_o - \Delta]$ results in only a very small loss in resolution, while improving the performance of an impulse radio UWB receiver.

Accordingly, in our preferred embodiment, ADC **20** implements only positive/negative magnitude thresholds $[T_0 \pm \Delta]$, thereby simplifying the circuit while simultaneously improving both the conversion time of the ADC **20** and, in general, the performance of the receiver. Such an implementation lends itself naturally to our trit-based scheme, wherein the three defined states indicate, for example, that:

$[-1] \Rightarrow$ the input is below the negative magnitude threshold $[T_0 - \Delta]$;

$[0] \Rightarrow$ the input is between the negative magnitude threshold $[T_0 - \Delta]$ and the positive magnitude threshold $[T_0 + \Delta]$; and

$[+1] \Rightarrow$ the input is above the positive magnitude threshold $[T_0 + \Delta]$.

In contrast to a conventional sign+magnitude implementation, our trit-based ADC **20** can be readily adapted to operate either at a higher sample rate (improved performance but with more power) or at an equivalent sample rate (substantially equivalent performance but with less complexity, thereby reducing both circuit size and power consumption).

Additional details relating to the construction and operation of our UWB receiver **10'** can be found in the Related References. As explained in the Related References, receiver **10'** initially operates in an acquisition mode, during which the components are configured to detect an incoming UWB signal transmitted by a remote UWB transmitter (not shown), and to achieve synchronism with that transmission, a process referred to as acquisition. Having achieved acquisition, receiver **10'** transitions into a data mode, during which the components are configured to recover data contained within each transmitted packet, a process referred to as demodulation or data recovery.

Shown in FIG. **3** are the components of our receiver **10'** that perform carrier and timing recovery when the receiver **10'** is operating in the acquisition mode. As explained in our Related References, during operation, ADC **20** develops trit-valued samples of both the in-phase, $Sr[5:0]$, and quadrature, $Si[5:0]$, components of the received signals. In the acquisition mode, these samples are passed to the correlator **24**, where they are correlated against the pre-defined preamble code. If a valid preamble is present, the output of the correlator **24** comprises a noisy estimate of the channel impulse response ("CIR"). This noisy CIR estimate is passed, without carrier and timing recovery, to accumulator **26**. If a preamble is present, then, as accumulator **26** adds symbols together, the CIR estimate will grow faster than the noise floor. By comparing successive accumulated groups of CIR estimates, the acquisition control logic (described in our Related References) can determine if there is a valid preamble present.

Once a preamble has been identified, the carrier recovery logic is activated to correct carrier error in the received data. Logic **46** computes an instantaneous phase error estimate of the received UWB signal by performing an inverse tangent operation on the in-phase, i.e., real, and quadrature, i.e., imaginary, components of the signal phasor. A carrier loop filter **48** uses this estimate to compute a correction angle to be applied to the current input to the accumulator **26**. This correction angle is developed as a carrier recovery phase signal (7-bits unsigned) wherein the output range is 0.0 to almost 2.0, and where the value 2.0 is equivalent to one revolution. A look-up-table ("LUT") **50** converts the correction angle to a corresponding pair of cosine (5-bits, including a sign bit) and sine (5-bits, including a sign bit) values. Using these sine and cosine values, a rotator **52** rotates the correlated samples by implementing a complex multiplication for each sample, followed by rounding to

return the output real and imaginary samples to 7-bit signed values. The rotated correlated samples are then resampled by resampler **54** for use by the accumulator **26**.

There is a finite probability that a detected preamble is not valid. Accordingly, our accumulator **26** spends some time assessing the quality of the incoming signal. If the quality is found to be poor, the preamble is rejected and accumulator **26** resumes searching for a preamble. If the quality is found to be sufficiently high, the next task is to search for a start-of-frame-delimiter ("SFD") by comparing the incoming correlator samples against the accumulated CIR estimate. This is performed for each symbol, and the result quantised to a 2-bit signed value. This 2-bit signed value should be +1 throughout the preamble, but then, once the SFD pattern is received, it should follow the pattern of the SFD. For example, for a short SFD, this would be $[0, +1, 0, -1, +1, 0, 0, -1]$. A corresponding search pattern is used to find this sequence on the 2-bit quantised values. This allows the SFD pattern to be identified, and the time to transition to data mode determined.

Shown in FIG. **4** are the components of our receiver **10'** that perform carrier and timing recovery when the receiver **10'** is operating in the data mode. No later than the end of the acquisition mode, the CIR estimate developed by accumulator **26** is loaded into the channel-matched-filter ("CMF") **36**. Once receiver **10'** switches to data mode, the ADC **20** samples are passed to the CMF **36**, which gathers the energy of the received UWB pulses together, thereby compensating for the smearing effect of the multi path channel. The resulting output pulses are passed through the timing and carrier recovery loops to compensate for any offset, and then into despreader **40**. Despreader **40** applies the same pseudo-random noise ("PN") spreading sequence used in the transmitter (not shown) to generate the burst, and integrates over the burst. As despreader **40** does not know if a logic_1 or a logic_0 is being transmitted, it despreads at both possible locations to produce two estimates. These two estimates provide the soft input to the Viterbi decoder **42**, which produces a maximum likelihood estimate of the bit sequence. The first section of the packet to be demodulated is the PHY header ("PHR"), which contains information about the data rate and the number of octets transmitted. As is known, this is protected by a single-error-correct-double-error detect ("SECDED") code. Once the PHR is decoded, the relevant information is extracted to allow the demodulation, i.e., recovery, of the data payload. The remaining payload is then passed through the Viterbi decoder **42** to Reed-Solomon ("RS") decoder **44** to correct any errors or report errors that cannot be corrected.

As shown in FIG. **3**, in the acquisition mode, the instantaneous phase error estimate is derived from the output of correlator **24** as follows:

1. The accumulated samples, $Sr[14:0]$ and $Si[14:0]$, are quantized and the conjugate calculated in logic block **56**;
2. The complex correlated samples from resampler **54** are multiplied by the conjugate by multiplier **58**; and
3. The products are summed by summer **60** to produce the complex instantaneous phase error estimate phasor, comprising $Sr[15:0]$ and $Si[15:0]$.

As shown in FIG. **4**, in the data mode, the complex instantaneous phase error estimate phasor simply comprises the outputs of despreader **40**, comprising, $Sr[15:0]$ and $Si[15:0]$.

The instantaneous phase error estimate phasor is then converted to a corresponding angle by logic block **62**. As shown in FIG. **5**, the complex error estimate phasor consists of two 17-bit signed numbers. The angle of the phasor is

determined by dividing the imaginary component of the phasor, Si[15:0], by the real component, Sr[15:0]. Ideally the resulting quotient would then be passed to an inverse tangent function to compute the exact angle; however, for the purposes of our carrier recovery algorithm, we have determined that the quotient per se is a sufficient approximation to the angle. To simplify the division computation, we first identify the sign and quadrant of the phasor, and then compute the absolute value of the real and imaginary parts. These absolute values are then passed to the division computation, and the resulting angle is post-processed to map it into the appropriate quadrant/sign. The remapped estimate is then subjected to rounding and saturation to produce the signed 7-bit instantaneous phase error estimate, S[0:-5], in radians.

The structure of carrier loop filter **48** can be seen in greater detail in FIG. 6. As shown, the instantaneous phase error estimate is passed through a proportional gain arm **48a** and an integral gain arm **48b**. The gains, K_p and K_i , are controlled by a gear-shifting algorithm to allow the loop to lock quickly, and then to rapidly tighten to a narrow bandwidth to minimise the impact of noise on the carrier recovery algorithm. The gear shifting is controlled by a counter (e.g., see, FIG. 7) that counts the number of phase error estimates supplied to the loop, and selects a respective scaling factor. (The default gear shifting table is described below.) The filtered phase error estimate is accumulated in a carrier phase accumulator **48c** to track the over-all phase error, and to adjust the instantaneous phase error to track frequency errors of up to ± 40 ppm.

When receiver **10'** transitions from acquisition to data mode, the rate of update of the carrier recovery loop **48** will change (from the preamble symbol interval to the data symbol interval); this requires that the integral term in the carrier loop filter **48b** be scaled to compensate for this change as follows:

TABLE 1

Carrier Loop Filter Scaling Table		
Preamble Symbol	Data Symbol	Scale Factor
992	1024	33/32
1016	10124	129/128
992	8192	8 * (33/32)
1016	8192	8 * (129/128)

In data mode, we have determined that the application of the phase rotation in a single "lump" at the end of each symbol has a negative impact on the performance of the receiver. For the 6.8 Mb/s case, the phase is applied at the end of a group of 8 symbols, so the symbols towards the end of this group suffer from increased phase error as compared to those at the start. Similarly for the 850 kb/s case, the symbols representing a logic_1 will have a higher phase error than those representing a logic_0. Worst of all is the 110 kb/s case, which will suffer from a phase error increase throughout a symbol and, depending on the hop position of the symbol, will have an effectively random phase error to start with if the carrier offset is high enough. To compensate for this, our preferred embodiment will smooth the phase rotation during the data demodulation phase.

As shown in FIG. 7, our approach is to scale the filtered phase error estimate and apply it to the accumulator more often than once per symbol. In the case of 850 kb/s, the filtered phase error will be scaled down by a factor of 2 and accumulated twice (i.e., every 64 clock cycles) during the phase update interval. (See, e.g., FIG. 8.) This means that

instead of the phase adjustment being applied in a single lump at the end of the interval, it will be distributed over the course of the symbol, thereby allowing the despread data towards the end of the interval to have a more accurate phase correction. In the case of 6.8 Mb/s, the filtered phase error will be scaled down by a factor of 8 and accumulated 8 times (i.e., every 16 clock cycles) during the phase update interval. (See, e.g., FIG. 9.) In the case of the 110 kb/s, the filtered phase error estimate will be scaled down by a factor of 64 and accumulated 64 times (i.e., every 16 clock cycles) during the phase update interval. (See, e.g., FIG. 10.) The despreader **40** will control when the phase rotation is applied, so that the phase rotation is updated just before it is needed for demodulation of a burst.

In our preferred embodiment, we implement a register-based field programmable gear shifting mechanism. Ten gears may be configured; one is reserved for demodulation mode, allowing nine acquisition gears. Each gear is assigned: a count at which it is activated; a K_p value; and a K_i value. Writing a value of logic_0 as the count for a gear other than the first gear terminates the gear shifting table; whilst still switching to the demodulation gear when the acquisition phase is over. Note that two sets of demodulation coefficients must be specified, one for the 110 Kbps data rate case, and one for the 850K and 6.81 Mbps cases. The default values for each of the available programmable registers are given in the following table:

TABLE 2

Carrier Recovery Loop Gear Shifting Table			
Register	Count	K_i	K_p
CR0	0	0x8: 3 * 2 ⁻⁶	0x7: 2 ⁻³
CR1	12	0x7: 3 * 2 ⁻⁷	0x7: 2 ⁻³
CR2	20	0x6: 3 * 2 ⁻⁸	0x6: 2 ⁻⁴
CR3	32	0x6: 3 * 2 ⁻⁸	0x6: 2 ⁻⁴
CR4	40	0x5: 3 * 2 ⁻⁹	0x5: 2 ⁻⁵
CR5	64	0x4: 3 * 2 ⁻¹⁰	0x5: 2 ⁻⁵
CR6	128	0x3: 3 * 2 ⁻¹¹	0x4: 2 ⁻⁶
CR7	192	0x2: 3 * 2 ⁻¹²	0x4: 2 ⁻⁶
CR8	256	0x1: 3 * 2 ⁻¹³	0x3: 2 ⁻⁷
CR9	1023	0x0: 3 * 2 ⁻¹⁴	0x3: 2 ⁻⁷

The K factors are coded as follows:

	Minimum	Maximum	Decode
Count	1	1023	Sample count on which to apply gearing values
Ki	0x0	0xA	0x0 = 2 ⁻¹¹ * (2 ⁻² + 2 ⁻³) 0x1 = 2 ⁻¹⁰ * (2 ⁻² + 2 ⁻³) 0x2 = 2 ⁻⁹ * (2 ⁻² + 2 ⁻³) 0x3 = 2 ⁻⁸ * (2 ⁻² + 2 ⁻³) 0x4 = 2 ⁻⁷ * (2 ⁻² + 2 ⁻³) 0x5 = 2 ⁻⁶ * (2 ⁻² + 2 ⁻³) 0x6 = 2 ⁻⁵ * (2 ⁻² + 2 ⁻³) 0x7 = 2 ⁻⁴ * (2 ⁻² + 2 ⁻³) 0x8 = 2 ⁻³ * (2 ⁻² + 2 ⁻³) 0x9 = 2 ⁻² * (2 ⁻² + 2 ⁻³) 0xA = 2 ⁻¹ * (2 ⁻² + 2 ⁻³) 0xB-0xF - invalid
Kp	0x0	0x7	0x0 = 00 0x1 = 2 ⁻⁹ 0x2 = 2 ⁻⁸ 0x3 = 2 ⁻⁷ 0x4 = 2 ⁻⁶ 0x5 = 2 ⁻⁵ 0x6 = 2 ⁻⁴ 0x7 = 2 ⁻³

Under very noisy conditions, the carrier recovery loop may fail to lock correctly. This will result in a preamble

11

rejection in the accumulator 26 (if this mode is enabled), effectively giving the carrier recovery loop another shot. The timing recovery loop can still fail to lock, however, and it does not get another chance since by the time this has an effect the preamble will likely be confirmed.

Preferably, LUT 50 updates the SIN (5-bits signed) and COS (5-bits signed) values under the control of accumulator 26 during acquisition mode and by despreaders 40 during data mode. This is in order to prevent phase changes being applied at times when the data is important to the algorithm in question, so it must be applied outside of the impulse response during acquisition, and outside of a burst position during demodulation.

Rotator 52 takes the SIN and COS values from LUT 50 and applies them to the incoming data vectors. This rotation is applied to the correlator 24 outputs during acquisition, and to the CMF 36 outputs during data demodulation. As shown in FIG. 12, the first stage, T0, selects between the registered outputs of correlator 24 and CMF 36 depending on the mode. The second stage, T1, performs the 4 multiplications that are required for a full complex multiplication, these multiplication results are then registered. The third stage, T2, sums the real components and the imaginary components of the complex multiplication, applies rounding to return the rotated samples to a precision of 7-bits (signed) and registers them for the output. Therefore the total latency of the rotator 52 is four clock cycles, C0-C3.

Our timing estimation is based on an early-late gating algorithm. As illustrated in FIG. 3 and FIG. 4, the source of data for this algorithm is dependent upon the phase of reception: during acquisition mode, the necessary timing information is derived from the output of CMF 36; and, during data mode, the despreaders 40 provide the necessary timing information. In both cases the information is used to generate an instantaneous timing estimate. This estimate is conditioned and passed into a timing loop filter 64 typical of a second order loop. The timing loop filter 64 output provides sample-level timing adjustment information to the resampler 54 during acquisition mode, and clock cycle level adjustments to the accumulator 26 and despreaders 40 during data mode.

As shown in FIG. 12, during acquisition mode, the timing estimate is derived from early/on-time/late samples from the output of the CMF 36. Once the preamble has been detected, the accumulator 26 will start writing CIR coefficients to the CMF 36. As noted before, the output of the correlator 24 provides an instantaneous noisy estimate of the CIR. As the timing offset error accumulates in the accumulator 26, the position of this estimate within the symbol will slowly vary. Quantizing the correlator 24 output and feeding it into the CMF 36 result in a correlation peak at the output of the CMF 36 corresponding to the position of the noisy impulse response estimate in the correlator 24 output symbol. In our preferred embodiment, the functions of quantisation of the correlator 24 output and routing to CMF 36 are performed by resampler 54.

As the timing of the incoming signal changes, the position of this correlation peak in the output of the CMF 36 will move. In general, this movement will be gradual relative to one preamble symbol duration and therefore can be tracked. The accumulator 26 provides a timing flag to indicate when the correlation peak is expected at the output of the CMF 36. This flag is based on the latency of the accumulator 26 and CMF 36 relative to the estimated impulse response location. Initially, this will be very accurate as the timing phase error will be negligible, but, as the phase error accumulates, the peak will move, thus providing the desired timing informa-

12

tion. The correlation peak is ideally the on-time sample, with the early and late samples being the one immediately preceding and following the on-time sample respectively. These three samples are passed to a phase detector 66 (see, FIG. 3) that computes an instantaneous timing phase error for the timing loop filter 64.

During demodulation, the despreaders 40 provide dedicated early and late outputs in addition to the normal on-time output used as part of the demodulation process. These outputs are provided for both possible burst positions (depending on if the data is a logic_0 or a logic_1); thus, an early instantaneous decision is needed in order to identify which of the two sets of early/on-time/late samples to use in computing the instantaneous timing phase error estimate. Once this decision is made, the computation is performed and the error passed to the timing loop filter 64.

The phase error estimation is based on the previously described early/on-time/late samples. A difference computation is performed by computational block 68. The difference between the early and late values is computed and checked against the on-time (on-time should be greater). If the on-time value is negative or zero (after conditioning), then the data is unreliable and the phase estimate zeroed. The difference is divided by twice the peak and the result checked to be less than 0.75 (otherwise it is considered unreliable) and passed out as the instantaneous timing error estimate in the format S[-1:-4]. This instantaneous timing error estimate is then passed on to the timing loop filter 64.

As shown in FIG. 12, timing loop filter 64 comprises a proportional gain arm 64a and an integral gain arm 64b. The gains, Kp and Ki are controlled by a gear-shifting algorithm to allow the loop to lock quickly and then rapidly to tighten to a narrow bandwidth to minimise the impact of noise on the timing recovery algorithm. The gear shifting is controlled by a counter (e.g., see, FIG. 7) that counts the number of timing error estimates supplied to the loop and operates, in one example, as follows:

TABLE 3

Timing Estimation Loop Gear Shifting Table		
Number of samples	K _p	K _i
0-5	0	0
5-28	31/(2 ⁻⁷)	0
28-80	31/(2 ⁻⁷)	20/(2 ⁻¹¹)
80-120	15/(2 ⁻⁷)	5/(2 ⁻¹¹)
120 onwards	11/(2 ⁻⁷)	1/(2 ⁻¹¹)

Timing loop filter accumulator 64c accumulates the lower noise estimate of the timing phase error to track the timing phase error and adjust the sampling phase error between 0 and +15.875 samples. Timing loop filter 64 also develops phase increment ("Inc") and phase decrement ("Dec") signals that are used to adjust the timing in units of 16 samples because they either drop or add a clock cycle delay in the accumulator 26 (during acquisition) or despreaders 40 (during demodulation). Therefore, if an adjustment of -2.5 samples is required, then phase_inc is used to introduce an offset of -16 samples, while the phase error driving resampler 40 will apply a correction of +13.5 samples, giving the overall required phase adjustment of -2.5 samples. Similarly, for example, an adjustment of +19.125 samples may be achieved by applying a +16 sample adjustment with the phase_dec signal, followed by an additional +3.125 sample correction in the resampler 54. By way of illustration, a suitable embodiment of resample 54 is illustrated in FIG. 13.

13

The carrier recovery loop needs to lock quickly in order to successfully receive the signal, whereas the timing recovery loop can take longer. As a result, if the carrier recovery loop fails to lock soon after the preamble is found, then the preamble will be rejected, allowing the carrier recovery loop another chance to lock. The timing recovery loop, however, does not get a second chance, so, to improve the chances of lock, the timing recovery loop can be seeded with an estimate based on the carrier loop integrator. Preferably, seeding is enabled via the use of a control signal: if the state of this signal is, e.g., logic_0, then timing seeding does not take place and the gearing table (see, below) must be set up to allow for this (initially wide bandwidth to allow acquisition, then narrowing as the lock improves); but, if the state of this signal is, e.g., logic_1, then timing seeding is enabled and the loop is assumed close to lock from the outset, and a more aggressive gearing table can be used. In our preferred embodiment, seeding is enabled by default.

We have determined that the value of the loop integrator in the carrier recovery loop can be used to seed the loop integrator in the timing recovery loop, thereby giving the timing recovery loop a jump start and enhancing the chances of it locking. The formula we prefer to use to compute the seed value is:

$$I_{tim} = (\text{Scale})(I_{car})$$

where:

I_{tim} = timing loop integrator;

I_{car} = carrier loop integrator;

$$\text{Scale} = \frac{F_s}{2F_c}$$

F_s = sampling rate; and

F_c = carrier frequency.

Since F_s and F_c are related this is simpler in practice:

F _c (MHz)	Channel	Scale Factor
3494.4	1	1/7
3993.6	2, 4	1/8
4492.8	3	1/9
6489.6	5, 7	1/13

As shown in FIG. 14, we prefer to allow the carrier recovery loop to settle on a good quality estimate before seeding the timing recovery loop; we define the delay in terms of the carrier recovery loop gearing counter and prefer to make this threshold value programmable. Once this gearing counter threshold value is reached, the value held in the carrier recover loop integrator is scaled according to the table above (depending on the channel setting) to produce a timing seed value with the precision S[−1:−15]. Rounding does not need to be applied to this computation, just truncation, because the computation is a one-time event in the receiver 10' and does not take part in a recursive loop—therefore the bias introduced by not rounding will not accumulate to cause significant inaccuracy.

In our preferred embodiment, we implement a register-based field programmable gear shifting mechanism. Ten gears may be configured; one is reserved for demodulation mode, allowing nine acquisition gears. Each gear is assigned: a count at which it is activated; a K_p value; and a K_i value. Writing a value of logic_0 as the count for a gear other than the first gear terminates the gear shifting table;

14

whilst still switching to the demodulation gear when the acquisition phase is over. Note that two sets of demodulation coefficients must be specified, one for the 110 Kbps data rate case, and one for the 850K and 6.81 Mbps cases. The default values for each of the available programmable registers are given in the following table:

TABLE 4

Timing Estimator Default Programmable Gear Shifting Register Values				
Register	Value	Default		
		Count [9:0]	K_i [14:10]	K_p [19:15]
TR0	0XF8000	0	0	31
TR1	0x8141E	30	5	16
TR2	0X58428	40	1	11
TR3	0X00000	0	0	0
TR3	0X00000	0	0	0
TR5	0X00000	0	0	0
TR6	0X00000	0	0	0
TR7	0X00000	0	0	0
TR8	0X00000	0	0	0
TR9	0X5A161	N/A	1/8 (110 Kbps)	11

where:

Value comprises a 20-bit variable expressed in hexadecimal format;

Count comprises bits [9:0] of the Value;

K_i comprises bits [14:10] of the Value; and

K_p comprises bits [19:15] of the Value.

The K factors are coded as follows:

TABLE 5

Gear Shifting Register Value Decode			
	Minimum	Maximum	Decode
Count	1	1023	Sample count on which to apply gearing values
K_i	0x00	0x1F	0x00 = 00 0x01 = 1×2^{-7} 0x1F = 31×2^{-7}
K_p	0x00	0x1F	0x00 = 00 0x01 = 1×2^{-7} 0x1F = 31×2^{-7}

Computing Angle of Incidence:

In a practical coherent receiver, it is necessary to track the carrier of the transmitter. For example, in the system of FIG. 16, both receivers 70 will be doing this, but because of different noise input mixed with the received signal, they will not necessarily do it identically. In, for example, the clock tracking loop shown in FIG. 3, the correlator output is accumulated to identify the channel impulse response, but before it is accumulated it is rotated by a carrier correction. Because this rotation is likely to be different in each of the receivers 70, it must be undone in order to calculate the phase difference between the two carriers. For example, with reference to FIG. 3, during an angle of incidence calculation mode, the rotator 52 can be rendered inoperative (or, alternatively, LUT 50 can be configured to output a fixed rotation of 0°); otherwise, the logic 46 works as described above. The phase difference between the two carriers is therefore the difference between the angles of the first paths in the accumulators minus the individual phase corrections that have been applied at the time the accumulation of the channel impulse response stops and is measured. FIG. 21 shows the calculated angle of arrival when repeated, using this method, on 100 separate packets. In the test, the actual angle of arrival was −5° and the antennas were separated by

15

one wavelength. The carrier frequency used in this test was 4 GHz, and the standard deviation of the error was 2.1°.

We propose two ways to solve the ambiguity in solutions that occurs at an antenna separation of more than $\frac{1}{2}$ a wavelength. First, we measure the time of arrival of the packet at each antenna. The angle of incidence that is most consistent with the measured time of arrival differences is the one chosen. Take the example shown in FIG. 18. If α is measured as -125° then there are two possible values for θ either -20° or $+40^\circ$. If the signal arrived at antenna A first, then -20° is the correct but if the signal arrives at antenna B first, then $+40^\circ$ is correct. Second, we resolve the ambiguity by sending two packets but at different carrier frequencies. FIG. 20 shows an example of the relationship between α and θ for two different carrier frequencies. Because the curves are different, only one of the possible solutions occurs at both carrier frequencies. For example, if θ , the angle of incidence was -50° then α at 4 GHz would be measured as $+90^\circ$ which could correspond to an α of either -50° or of $+15^\circ$. At 6.5 GHz would be measured at -90° which could correspond to an α of either -50° , 10° or $+24^\circ$. Since -50° is the only solution in common, it must be the correct one. Of course in practice, noise in the system means that the common solution will not be exactly the same so it will be necessary to choose the solution which has the smallest difference in the two sets of possible solutions.

Even if the two receivers 70a and 70b are fed from the same clock, it may happen that the delay of this clock to one receiver is different to the delay to the other receiver. In this case there will be a fixed phase difference between the carriers. However, this phase difference can be calibrated, e.g., by measuring α at a known angle of arrival and subtracted from α , before applying the formula of Eq. 6.

Rather than supplying the two different PLLs, 76a and 76b, with the common crystal 72, there are other ways to synchronize the receivers 70, e.g., the two receivers 70a and 70b could be synchronized by supplying both with a clock from a single PLL, e.g., the PLL 76a.

Although we have described our invention in the context of particular embodiments, one of ordinary skill in this art will readily realize that many modifications may be made in such embodiments to adapt either to specific implementations. By way of example, it will take but little effort to adapt our invention for use with a different ADC scheme when it can be anticipated that the target application will not be subject to significant levels of in-channel CW interference. Further, the several elements described above may be implemented using any of the various known semiconductor manufacturing methodologies, and, in general, be adapted so as to be operable under either hardware or software control or some combination thereof, as is known in this art. Alternatively, the several methods of our invention as disclosed herein in the context of special purpose receiver apparatus may be embodied in computer readable code on a suitable computer readable medium such that when a general or special purpose computer processor executes the computer readable code, the processor executes the respective method.

Thus it is apparent that we have provided an improved method and apparatus for use in the receiver of a UWB communication system to determine angle of incidence. In particular, we submit that such a method and apparatus should provide performance generally comparable to the best prior art techniques but more efficiently than known implementations of such prior art techniques.

16

The invention claimed is:

1. In a radio frequency ("RF") system comprising an RF transmitter, a first RF receiver having a first antenna, and a second RF receiver having a second antenna, the first and second antennas being separated by a predetermined distance, d , a method comprising the steps of:

- [1] in the first [and second receivers] receiver, using [respective] a first [and second] clock tracking [loops] loop to synchronize the carrier phase of [said] the first receiver to the transmitter and in the second receiver, using a second clock tracking loop to synchronize the carrier phase of the second receiver to the transmitter;
- [2] in the transmitter, transmitting an ultra-wideband ("UWB") signal having a predetermined carrier wavelength, λ ;
- [3] in the first receiver:
 - [3.1] receiving the transmitted UWB signal;
 - [3.2] developing a first phase value as a function of the complex baseband impulse response of said received UWB signal; and
 - [3.3] correcting the first phase value by subtracting the phase of the first clock tracking loop;
- [4] in the second receiver:
 - [4.1] receiving the transmitted UWB signal;
 - [4.2] developing a second phase value as a function of the complex baseband impulse response of said received UWB signal; and
 - [4.3] correcting the second phase value by subtracting the phase of the second clock tracking loop;
- [5] developing a phase difference value, α , as a [function of] difference between the corrected first and second phase values; and
- [6] developing an angle of arrival, θ , of the transmitted UWB signal relative to the first receiver as a function of d , λ and α .

2. The method of claim 1 wherein, if d is at least $\lambda/2$, step [6] is further characterized as:

- [6] developing a plurality of angles of arrival, θ , of the transmitted UWB signal relative to the first receiver according to the following:

$$\theta = \sin^{-1} \frac{\alpha \lambda}{2\pi d};$$

the method further comprising the step of:

- [8] in the first and second receivers, determining respective first and second times of arrival of the transmitted UWB signal; and
- [9] selecting one of the plurality of angles of arrival as a function of the first and second times of arrival.

3. The method of claim 1 wherein, if d is at least $\lambda/2$, step [6] is further characterized as:

- [6] developing a plurality of angles of arrival, θ , of the transmitted signal relative to the first receiver according to the following:

$$\theta = \sin^{-1} \frac{\alpha \lambda}{2\pi d};$$

the method further comprising the step of:

- [8] performing steps [1] through [6] using a first carrier wavelength λ_1 ;
- [9] performing steps [1] through [6] using a second carrier wavelength λ_2 different than the first carrier wavelength λ_1 ; and

17

[10] selecting one of the plurality of angles of arrival as
a function of the first and second carrier wavelengths.

[4. A non-transitory computer readable medium including
executable instructions which, when executed in a process-
ing system, cause the processing system to perform all of the 5
steps of a method according to any one of claims **1, 2** and
3.]

* * * * *

18

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : RE48,832 E
APPLICATION NO. : 16/437695
DATED : November 23, 2021
INVENTOR(S) : Gavin Marrow et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the Specification

Column 1, Line 65, replace “we will hold” with --we will bold--.

Column 3, Line 26, replace “received front the transmitter” with --received from the transmitter--.

Column 4, Line 36, replace “more is advantageous two reasons” with --more is advantageous for two reasons--.

Column 6, Lines 61, 62, and 65, replace “T_o” with --T₀--.

Column 7, Line 2, replace “T_o” with --T₀--.

Column 8, Line 51, replace “derived front the output” with --derived from the output--.

Signed and Sealed this
Eighth Day of February, 2022



Drew Hirshfeld
*Performing the Functions and Duties of the
Under Secretary of Commerce for Intellectual Property and
Director of the United States Patent and Trademark Office*