

DOWNLINK COMMUNICATION

CROSS-REFERENCE TO RELATED PATENT APPLICATIONS

[0001] This application claims priority to United States Patent Application No. 12/408,538, filed March 20, 2009, granted on December 29, 2009 as U.S. Patent No. 7,639,726 and to United States Application No. 12/647,756, filed December 28, 2009, the entire disclosure of which is incorporated herein by reference.

FIELD

[0002] Embodiments of the present application relate to the field of communications. More specifically, exemplary embodiments relate to downlink communication in a slot-based communication system.

BACKGROUND

[0003] A number of modulation techniques have been developed for facilitating communications in a network where multiple users are present. Such techniques include code division multiple access (CDMA), time division multiple access (TDMA), and frequency division multiple access (FDMA). CDMA is a spread-spectrum technology that uses pseudo-random number sequences to modulate incoming data, multiple transmitters transmitting on the same signal, and orthogonal codes (Walsh codes) to correlate different communication channels. TDMA uses time slots to coordinate multiple uplink transmitters that are transmitting in the same sub-slots. Users transmit in rapid succession, one after the other, each using his/her own time slot, allowing multiple stations to share the same transmission medium (e.g., radio frequency channel) while using only a portion of the total available bandwidth. FDMA allocates different users with different carrier frequencies of the radio spectrum.

[0004] In addition to modulation techniques, protocols exist for determining how network devices respond when two devices attempt to use a data channel simultaneously (called a collision). CSMA/CD (Carrier Sense Multiple Access / Collision Detection) is used by Ethernet networks to physically monitor the traffic on the line at participating stations. If

no transmission is taking place at the time, the particular station can transmit. If two stations attempt to transmit simultaneously, this causes a collision, which is detected by all participating stations. After a random time interval, the stations that collided attempt to transmit again. If another collision occurs, the time intervals from which the random waiting time is selected are increased step by step. This is known as exponential back off.

SUMMARY

[0005] An exemplary embodiment uses a random phase multiple access communication interface. The interface can communicatively connect to systems and devices using spread spectrum modulation methods without the use of orthogonal codes.

[0006] An exemplary random phase multiple access communication interface communicatively connects systems and devices using spread spectrum modulation methods. The random selection of chip (or timing) offsets as a multiple access scheme allows for non-coordinated data transmission without needing to be assigned a unique “code.” All users transmit using the same PN (pseudo noise) code such that a PN array despreader at the access point can be used. If two signals are received at the access point at the same PN offset (or the sum of the PN offset with the transmission delay in number of chips yields the same value for 2 or more transmissions), then a “collision” has occurred and it may not be possible to demodulate these 2 or more signals. The randomization of timing offsets each time means that any “collisions” that occur only occur during that frame. A retransmission scheme and a new randomized offset is used to get through in a subsequent attempt.

[0007] An exemplary embodiment includes a transmitter at the tag (uplink) and a method of transmitting signals from the tag to an access point. Each tag includes its own transmitter which transmits information in the form of frames. A frame can be formed from information provided on a channel having a fixed data rate. The data can be spread using the same pseudo-noise (PN) code, and can have a randomly selected chip offset. The transmitter also applies frequency rotation and sample clock correction to match the reference oscillator of the access point. A plurality of tags is associated with a single access point to form the network. Each of the plurality of tags transmits information using the same PN code along

with a randomly selected chip offset. The phase is randomly selected each frame over a large number of chips (i.e., 8192).

[0008] Another exemplary embodiment includes a transmitter at an access point (downlink) and a method for transmitting signals from the access point to the tags. The access point transmitter can be similar to that of the tags. However, the access point transmitter uses a unique PN code for each of the tags with which it communicates. The use of distinct PN codes for each tag provides security and allows each tag to ignore signals which are directed toward other tags. The frames transmitted by the access point also include a preamble of approximately 9 symbols to allow for rapid acquisition at the tags.

[0009] Another exemplary embodiment includes a demodulator at the tag and a method for demodulating signals received by the tag. An automatic frequency control (AFC) derotator multiplication is applied to signals received at the tag. The AFC derotator multiplication is a 1 bit complex operation with a 1 bit complex output such that gate count is improved. The tag uses a PN array despreader that takes advantage of the huge computational savings in the 1 bit data path.

[0010] Another exemplary embodiment includes a demodulator at the access point and a method for demodulating signals received at the access point. The access point demodulator has capacity to simultaneously demodulate several thousand or more links received from tags. To demodulate such a large number of links, the access point demodulator includes a PN array despreader.

[0011] Another exemplary embodiment includes synchronization of the tag with a master timing of the access point. The access point can periodically transmit a broadcast frame. During a 'cold' timing acquisition, the tag uses its PN despreader to analyze the broadcast frames and identify the master timing of the access point. Cold timing acquisition is expected to occur one time when the tag is first introduced into the system. After the initial cold acquisition, the tag can perform a 'warm' timing acquisition each time the tag wakes up to transmit or receive a signal. The warm timing acquisition utilizes less power than the cold timing acquisition.

[0012] In at least one exemplary embodiment, each tag separately generates a PN code. A gold code is an example of a PN code that is parameterizable such that each user has its own. As such, only data destined for a particular user is visible to it. Using unique PN codes, a tag does not process data that is not its own.

[0013] An exemplary method for communicating through a multiple access communication interface includes receiving a first signal from a first tag, where the first signal is spread using a predetermined pseudo-noise (PN) code, and further where the first signal includes first payload data. A second signal is received from a second tag. The second signal is spread using the predetermined PN code, and the second signal includes second payload data. The first payload data from the first signal is identified at least in part with a PN array despreader. The second payload data from the second signal is also identified at least in part with the PN array despreader.

[0014] An exemplary system for communicating through a multiple access communication interface includes a first tag, a second tag, and an access point. The first tag has a first transmitter configured to transmit first payload data in a first signal, wherein the first signal is spread using a predetermined pseudo-noise (PN) code. The second tag has a second transmitter configured to transmit second payload data in a second signal, wherein the second signal is spread using the predetermined PN code. The access point is in communication with the first tag and the second tag and includes a receiver and a despread array. The receiver is configured to receive the first signal and the second signal. The despread array is configured to despread the first signal and the second signal.

[0015] An exemplary access point for use in a multiple access communication system includes a processor, a receiver in communication with the processor, and a transmitter in communication with the processor. The receiver is configured to receive a first signal from a first tag, wherein the first signal includes first payload data, and further wherein the first signal is spread using a predetermined pseudo-noise (PN) code. The receiver is also configured to receive a second signal from a second tag, wherein the second signal includes second payload data, and further wherein the second signal is spread using the predetermined PN code. The transmitter is configured to transmit a third signal to the first tag, wherein the

third signal is spread with a second PN code, and further wherein the second PN code is specific to the first tag.

[0016] An exemplary method for conserving bandwidth in a communication system is provided. The method includes spreading a data frame using a first pseudo-noise (PN) spreader. A broadcast frame is spread using a second PN spreader. A complex data stream having a first component and a second component is generated. The data frame is assigned to the first component and the broadcast frame is assigned to the second component. The complex data stream is transmitted to a tag.

[0017] An exemplary access point is also provided. The access point includes a process and a transmitter operatively coupled to the processor. The processor is configured to spread a data frame using a first pseudo-noise (PN) spreader. The processor is also configured to spread a broadcast frame using a second PN spreader. The processor is further configured to generate a complex data stream having a first component and a second component, where the data frame is assigned to the first component and the broadcast frame is assigned to the second component. The transmitter is configured to transmit the complex data stream to a tag.

[0018] An exemplary computer-readable medium is also provided. The computer-readable medium has computer-readable instructions stored thereon that, upon execution by a processor, cause an access point to spread a data frame using a first pseudo-noise (PN) spreader. The computer-readable instructions also cause the access point to spread a broadcast frame using a second PN spreader. The computer-readable instructions also cause the access point to generate a complex data stream having a first component and a second component, where the data frame is assigned to the first component and the broadcast frame is assigned to the second component. The computer-readable instructions further cause the access point to transmit the complex data stream to a tag.

[0019] These and other features, aspects and advantages will become apparent from the following description, appended claims, and the accompanying exemplary embodiments shown in the drawings, which are briefly described below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] Fig. 1 is a diagram depicting an uplink transmitter according to an exemplary embodiment.

[0021] Fig. 2 is a diagram depicting a downlink transmitter according to an exemplary embodiment.

[0022] Fig. 3 is a diagram depicting slot structures and assignments in an exemplary embodiment.

[0023] Fig. 4 is a diagram depicting a PN (pseudo noise) despread array in an exemplary embodiment.

[0024] Fig. 5 is a flow diagram depicting operations performed in the tag processing of a broadcast channel from a cold start in an exemplary embodiment.

[0025] Fig. 6 is a flow diagram depicting operations performed in the tag processing of a dedicated channel from a warm start in an exemplary embodiment.

[0026] Fig. 7 is a diagram depicting a tag receive data path in an exemplary embodiment.

[0027] Fig. 8 is a diagram depicting time tracking in an exemplary embodiment.

[0028] Fig. 9 is a diagram depicting an AFC (automatic frequency control) rotation in an exemplary embodiment.

[0029] Fig. 10 is a diagram depicting a dedicated communication finger in an exemplary embodiment.

[0030] Fig. 11 is a flow diagram depicting operations performed during access point receive processing in an exemplary embodiment.

[0031] Fig. 12 is a diagram depicting an access point receive data path in an exemplary embodiment.

[0032] Fig. 13 is a diagram depicting asynchronous initial tag transmit operations in an exemplary embodiment.

[0033] Fig. 14 is a diagram depicting interactions between an access point and a tag in a slotted mode according to an exemplary embodiment.

[0034] Fig. 15 is a diagram depicting data transfer between an access point and a tag according to an exemplary embodiment.

[0035] Fig. 16 is a diagram illustrating contents of a complete slot in accordance with a first exemplary embodiment.

[0036] Fig. 17 is a diagram illustrating contents of a complete slot in accordance with a second exemplary embodiment.

[0037] Fig. 18 illustrates a fundamental downlink slot in accordance with an exemplary embodiment.

[0038] Fig. 19 is a diagram illustrating preamble frame processing in accordance with an exemplary embodiment.

[0039] Fig. 20 is a diagram illustrating a data sub-slot hierarchy in accordance with an exemplary embodiment.

[0040] Fig. 21A illustrates a fundamental downlink slot with a plurality of sub-slots in accordance with an exemplary embodiment.

[0041] Fig. 21B illustrates a fundamental downlink slot with a single slot in accordance with an exemplary embodiment.

[0042] Fig. 22 is a flow diagram illustrating operations performed to construct a frame in accordance with an exemplary embodiment.

[0043] Fig. 23 is a diagram illustrating a downlink transmission model in accordance with an exemplary embodiment.

DETAILED DESCRIPTION

[0044] Exemplary embodiments are described below with reference to the accompanying drawings. It should be understood that the following description is intended to describe exemplary embodiments, and not to limit the invention defined in the appended claims.

[0045] Fig. 1 illustrates an uplink transmitter 10 which includes structures such as a convolution encoder, an interleave module, a modulator, a pseudo-noise spreader, a filter, a bank of taps, an automatic frequency control (AFC) rotator, and other such structures. These structures perform operations depicted in blocks 12, 14, 16, 18, 20, and 22. The transmit path of uplink transmitter 10 is a coded and spread spectrum waveform. In an exemplary embodiment, the uplink transmitter 10 can be included in a tag that communicates with an access point along with other tags using demodulated communication channels. Additional, fewer, or different operations may be performed by the uplink transmitter 10 depending on the particular embodiment. The operations may also be performed in a different order than that shown and described. As used herein, a tag can refer to any communications device configured to receive signals from and/or send signals to an access point. The access point can refer to any communications device configured to simultaneously communicate with a plurality of tags. In an exemplary embodiment, the tags can be mobile, low power devices which run off a battery or other stored power, and the access point can be located in a central location and receive power from a power source such as a wall outlet or generator. Alternatively, the tags may plug into an outlet and/or the access point may run off of a battery or other stored power source.

[0046] In block 12, a data stream is received by a convolution encoder and interleave module. In one embodiment, the data stream is 128 Bits including the preamble. Alternatively, data streams of other sizes may be used. Once received, the data stream is encoded using the convolution encoder. In an exemplary embodiment, the data stream may be encoded at a rate of $\frac{1}{2}$. Alternatively, other rates may be used. The data stream can also be interleaved using the interleave module. An encoded symbols stream is output to a block 14 in which a differential binary phase shift keying (D-BPSK) modulator is used to modulate the encoded symbols stream. In alternative embodiments, other modulation schemes may be

used. At block 16, the modulated stream is applied to a PN spreader. In an exemplary embodiment, the PN spreader can use a common network gold code channel using a selected spreading factor. The spreading factor can be a member of the set {64, 128, 256, ..., 8192}. Alternatively, any other code and/or spreading factor may be used. Each of the tags at a given spreading factor is spread by the same PN code with a randomly selected chip offset. The large range of possible randomly selected chip offsets increases the probability that a particular frame will not collide (or, in other words, have the same chip timing at the access point) with another frame from another transmitter. The probability of collision in the limit of approaching capacity may become non-negligible (~10 % or less) and can be solved via retransmission of the same frame at a differently drawn random offset. The PN spreader is described in more detail below with reference to Fig. 4. In an exemplary embodiment, an output of block 18 can have a rate of 1 bit at 1 mega-chip per second (Mcps). Alternatively, other rates may be used.

[0047] At block 18, the data stream is upsampled by a 4 x oversample filter and time tracking logic is used to ensure that all of the frames land at the same sample rate consistent with the frequency reference of the AP. Block 18 receives a sample slip/repeat indicator as an input. In one embodiment, an output of block 18 may have a real frequency of approximately 4 megahertz (MHz). At block 20, an automatic frequency control (AFC) rotation is done including a frequency offset to match the access point's timing offset, ensuring that all of the frames from all of the users lands near the same frequency hypothesis. In one embodiment, an output of block 20 may have a complex frequency of approximately 4 MHz. At block 22, a delay is imposed from the start slot until the correct access slot occurs. In addition, a random chip delay is imposed on the signal. In an exemplary embodiment, the random chip delay can be from 0 to the spreading factor minus 1. Alternatively, a different random chip delay may be used. The slot access can be described by $A(i,j)$ where i is related to the spreading factor as $2^{(13-i)}$ and j is the sub-slot number corresponding to non-overlapping slots. Depending upon the selected spreading factor, there are generally multiple transmit opportunities in a given slot. For the uplink, the access slot can be randomly selected along with a chip offset from 0 to spreading factor minus 1. As such, the probability of collision between uplink users is minimized, while allowing for re-selection for cases

where there are collisions. After the signal has been delayed, the signal can be transmitted to an access point.

[0048] Fig. 2 illustrates a downlink transmitter 30 including structures such as a convolution encoder, an interleave module, a modulator, a pseudo-noise spreader, a filter, a bank of taps, and other such structures. Using transmitter 30, the access point (AP) transmits multiple channels each destined for a particular tag or user. These structures perform operations depicted in blocks 32 through 54. Blocks 32 to 40 and blocks 42 to 50 represent distinct data paths that can be replicated for additional data flows. In an exemplary embodiment, blocks 32-38 can perform operations similar to the operations described with reference to Fig. 1 on a first data stream. Similarly, blocks 42-48 can perform operations similar to the operations described with reference to Fig. 1 on an nth data stream, where n can be any value. The input to block 36 can be a gold code specific to the tag which is to receive the first data stream, and the input to block 46 can be a gold code specific to the tag which is receive the nth data stream. Alternatively, other codes such as a broadcast gold code, a non-gold code, or other may be used to spread the first data stream and/or the nth data stream. The output of block 38 and/or block 48 can be weighted in blocks 40 and 50 in case the data links corresponding to the first data stream and the nth data stream are of unequal power. Once weighted, the paths are summed in a block 52. A hard decision is also made in block 52 where all positive numbers are mapped to 0 and all negative numbers are mapped to 1. Alternatively, a different hard decision may be made. In one embodiment, an output of block 52 may have a rate of 1 bit at 10 Mcps. Alternatively, other rates may be used. The sum output from block 52 is upsampled using a 4 x chip filter in block 54. In one embodiment, an output of block 54 can have a real frequency of 40 MHz. Alternatively, other frequencies may be used. Not shown is a transmission on an adjacent frequency that is a single set of broadcast frames at a maximum downlink spreading factor of 2048. Alternatively, a different maximum downlink spreading factor may be used.

[0049] Fig. 3 illustrates slot structures and assignments. In at least one embodiment, data stream 70 includes slot 72, slot 74, and slot 76. Slot 72 is an AP-to-tags communication, slot 74 is a tags-to-AP communication, and slot 76 is an AP-to-tags communication. In an exemplary embodiment, each of the slots can have a duration of 2.1 seconds. Alternatively, any other duration may be used and/or different slots may have

different durations. The data stream 70 can be implemented in a half-duplex communication scheme such that at any given time, either the AP is transmitting and the tags are receiving, or the tags are transmitting and the AP is receiving. In alternative embodiments, other communication schemes may be used. As shown in Fig. 3, data channel 80 depicts processing gain options for data in slot 72. If a data link closes at a particular gain, the tag only needs to be ready to receive (in AP to tags mode) during the duration of the slot with the corresponding gain. In transmit mode, the slot selection governs the transmission from the tag to the access point such that the tag can minimize its on time in the power consuming transmit mode. For example, a gain of 18 dB only needs a 1.6 ms slot ($A_{7,0}$). Data channel 82 depicts processing gain options for data in slot 74. As can be seen, the power used by a tag can be selected such that each data link arrives at the AP at the same power.

[0050] There is a symmetry between processing a large number of simultaneous waveforms on the AP side, and the processing of the relative few waveforms on the tag side. Automatic frequency control (AFC), time-tracking drift, and frame timing are known on the AP side due to the fact that the AP is the master of these parameters. However, AFC, time-tracking drift, and frame timing may be determined at acquisition on the tag side. The PN array despreader performs the brute force operation associated with both, which is an efficient implementation for exploring acquisition hypothesis/demodulating. Another aspect of this is that this large power-consuming circuit (when active), though running continuously on the AP (which shouldn't matter because it can be plugged into the wall), is only running during a "cold" acquisition on the tag which should happen rarely. Cold acquisition and warm acquisition are described in more detail with reference to Figs. 5 and 6, respectively.

[0051] Fig. 4 illustrates a PN (pseudo noise) despread array, which facilitates both the acquisition of a single waveform on the tag, and brute-force demodulation of multiple waveforms on the AP. In an exemplary embodiment, the PN despread array can perform a 1 bit dot product of many chip-spaced timing hypotheses simultaneously.

[0052] A PN despread core element can be a simple counter that is incremented or not incremented each clock depending on whether the input is a 0 or a 1. Since it is a complex data path, there are two counters: one for I (in-phase) and one for Q (quadrature-phase). Multiplication by a complex exponential is generally a set of 4 rather large scalar

multipliers (4 x 1000 gates is typical) coupled to a complex exponential table. In contrast, a one bit complex multiplier is basically a simple truth table, such as the example table shown below, where the negative denotes the inverse ($0 \rightarrow 1$ and $1 \rightarrow 0$). This truth table can be implemented using just a few gates.

Phase	0	1	2	3
I'	I	-Q	-I	Q
Q'	Q	I	-Q	-I

[0053] Fig. 4 depicts a PN despread array 100. There can be many instantiations (e.g., 256 or more in one embodiment) of pairs of counters for the complex despread operation. The PN despread array 100 can be fed at chip rate with adjacent instantiations of PN despread elements 102, 104, and 106 working on timing hypotheses that are a chip apart. The 1 bit complex data is sent from a block 114 to elements 102, 104, and 106 where it is combined with a PN signal from PN generator 110. PN signal generator 110 can be hardware that outputs the same sequence of 0s and 1s with which the AP is spreading the data. In the case of element 102, the derotated data is combined (more specifically, 1 bit complex multiplied) with the PN signal at a combiner 122a. Real and imaginary parts of this combination are separately input into counters 118a and 120a. The counters 118a and 120a shift the bit stream out upon receipt of a reset signal 112. More specifically, the data in the counters is valid just prior to the reset signal. The reset signal forces zeros into both counters. The multiplexer 108 allows for output of the currently valid counters for that finger that has uniquely finished its desreading operation at that particular clock. Other elements in the PN despread array 100 operate similarly. Element 104 receives derotated data from block 114 and combines it with a PN signal after a delay is imposed by delay block 116a in element 102. The combination is entered into counters 118b and 120b, which gets shifted out of the counters upon a signal from the reset signal 112 with an imposed delay from a delay block 124a. Likewise, element 106 receives derotated data from block 114 and combines it with a PN signal after a delay is imposed by delay block 116b in element 104. The combination is entered into counters 118c and 120c, which gets shifted out of the counters upon a signal from the reset signal 112 with an imposed delay from a delay block 124b.

[0054] After a number of clocks corresponding to the spreading factor, the PN despread element 102 has valid data which is selected for output by a multiplexer 108. Every clock thereafter, the adjacent despread element 104 or 106 is available until all data has been output which can occur during the number of clocks corresponding to the spreading factor plus a number of PN despread instantiations. The PN code that governs the operation of this mechanism can be a gold code parameterized by a value. In alternative embodiments, other PN codes may be used.

[0055] Fig. 5 illustrates operations performed in the tag modem processing of a broadcast channel to demodulate the access point's transmit waveform. Additional, fewer, or different operations may be performed depending on the particular embodiment. The operations may also be performed in a different sequence than that shown and described.

[0056] Upon the initial power-up of the tag, no parameters are known regarding the waveform except for the broadcast channel PN sequence (e.g., the particular gold code or other code parameter). Additionally, the tag may not know with sufficient precision what the relative frequency offset is between the AP and the tag due to oscillator variance between the AP and the tag. Fig. 5 depicts a scanning mode where the range of uncertainty of parts-per-million (ppm) drift between the AP and the tag are explored. In an operation 150, an iteration is made over two slots to enable the tag to tune to a broadcast channel. For example, processing can begin asynchronous to slot timing. During exploration of one half of the hypotheses, the broadcast channel can be active, and during exploration of the other half of the hypothesis the broadcast channel can be inactive. In a first iteration, all hypotheses can be explored using a first slot timing with an asynchronous starting point. If no energy is found in the first iteration, a second iteration is performed. In the second iteration, the asynchronous starting point can have a one slot offset from the asynchronous starting point used in the first iteration. As such, hypotheses that were explored while the broadcast channel was active can be explored while the broadcast channel is active. Once the energy is found, the tag can tune to the broadcast channel. In an exemplary embodiment, operation 150 can represent a starting point for 'cold acquisition.' In an operation 152, a coarse automatic frequency control (AFC) is initialized. In one embodiment, this initial value is set to a most negative value such as -10 ppm offset. Using a known gold code generated PN sequence for the broadcast channel, in an operation 154, non-coherent metrics for all Cx4 spaced

hypotheses for a given coarse AFC hypothesis are calculated. For example, if the spreading factor has a length of 2048, the non-coherent metric for 8192 hypotheses can be calculated.

[0057] In operations 156 and 158, the coarse AFC hypothesis is incremented until the end of the ppm range. For each coarse AFC hypothesis, the hardware depicted in Fig. 7 is used to undo the frequency offset represented by the current hypothesis. The PN despread array is used to generate the despread output of 8 successive symbols. Alternatively, other numbers of symbols may be used. A non-coherent sum of these 8 symbols is then calculated. A set of N (8 in the one embodiment) top metrics along with their associated parameters are maintained in a data structure. As the flowchart of Fig. 5 indicates, the entire range of oscillator ppm uncertainty along all the timing hypotheses at chip x 4 resolution are explored with the expectation that the winning (i.e., valid) one will be represented in the data structure. Along with the most valid hypothesis there generally tends to be lesser multi-path reflections, adjacent AFC coarse frequency hypotheses where appreciable energy accumulation is still present, as well as entirely invalid hypotheses that have generated anomalously large metrics due to noise variance.

[0058] The non-coherent metrics for all chip x 4 timing hypotheses for each coarse AFC can be communicated to a data structure. In an operation 160, the data structure keeps track of the greatest non-coherent metrics (e.g., coarse AFC value, chip x 4 timing hypothesis, non-coherent metric value). The “finalists” are assigned to the N dedicated fingers in an operation 162. Each finger may be uniquely parameterized by a chip x 4 timing value and a coarse AFC hypothesis which is independent of the current coarse AFC hypothesis governing the PN despread array. Since frame timing is initially unknown, each despread symbol that is output by the dedicated finger is hypothesized to be the last in the frame. Thus, the buffered 256 symbols undergo differential demodulation and an additional set of iterations based on multiplying by a constant complex value to perform fine AFC correction, as shown in operations 164 and 166. An output of operation 164 can be a complex cross product from each dedicated finger. In operation 166, a symbol-by-symbol multiplication by a constant complex rotation (as determined by the fine AFC hypothesis) can be iteratively applied to a postulated frame of information to determine which (if any) of the selection of complex rotation constant values uncovers a frame which passes a cyclic redundancy check (CRC). This can be a brute-force operation where a cyclic redundancy

check (CRC) may be performed for each hypothesis. For any valid CRC, a payload from the signal can be sent to MAC, and network parameters can be considered to be known.

[0059] In an operation 168, other slot timing hypothesis are tried. In an exemplary embodiment, the coarse AFC hypotheses associated with the most successful CRCs can be nominal starting coarse AFC hypotheses. Once the entire range of coarse AFC hypothesis are explored, the tag notes a variable called Nominal_Coarse_AFC which is the relevant state information used in future transactions which greatly narrows the range of coarse AFC hypothesis searches because the part-to-part variation of oscillator ppm deviation is much larger than the oscillator drift over the course of a minute or so.

[0060] Fig. 6 illustrates operations performed in the tag processing of a dedicated channel from a warm start which is to say where relevant state information is known. For example, frame timing can be known and a much tighter range of coarse AFC hypothesis may be explored. The modem begins its processing sufficiently early so that valid finger assignments are made prior to the end of the 9 symbol preamble. Alternatively, any other number of symbols may be used.

[0061] In an operation 200, there is no need to iterate over a two slot timing hypothesis because the frame timing is known. Instead of using a broadcast channel, a dedicated channel is used. In an operation 202, a coarse AFC hypothesis is scanned. In an exemplary embodiment, the coarse AFC can be scanned over a small range to account for small frequency drift since the last time accessed. Using a known gold code generated PN sequence unique to the tag, in an operation 204, a non-coherent metric for all chip x 4 spaced hypotheses is calculated. In operations 206 and 208, the coarse AFC hypothesis is incremented until the end of the small ppm range. In an operation 210, a data structure keeps track of the greatest non-coherent metrics (e.g., coarse AFC value, chip x 4 timing hypothesis, non-coherent metric value, etc.) In an operation 212, dedicated fingers are assigned based on the data structure. In an operation 214, symbol cross products are created using current DBPSK and previous DBPSK. An output of operation 214 can be a complex cross product from each dedicated finger. In an operation 216, frames are interleaved and decoded. For any valid CRC, the payload can be sent to a medium access control (MAC) layer. In an operation 218, other slot timing hypothesis are tried. In an exemplary

embodiment, coarse AFC hypotheses associated with the most successful CRCs can be nominal starting coarse AFC hypotheses.

[0062] Fig. 7 illustrates a tag receive data path depicting the tag's demodulation processing in accordance with an exemplary embodiment. As shown, the one-bit complex samples are buffered in a sample buffer 220 such that enough data is present to make reliable detection of valid energy. Exemplary values are provided in the sample buffer block 220. For example, one embodiment buffers 9 symbols. In alternative embodiments, other values may be used. The samples may be input from the I channel and Q channel into this ping-pong buffer scheme at the synchronous sample rate of $\text{chip} \times 2$ or 2 MHz. Alternatively, other rates may be used. At the fast asynchronous clock, these samples are used to explore the various coarse AFC hypothesis. Based on the current coarse AFC hypothesis, time-tracking is performed at $\text{chip} \times 4$ resolution. Since the same timing reference is used to drive both the carrier frequency and the sample clocks on both the AP and the tag, a coarse AFC hypothesis with a known carrier frequency can uniquely map to a known rate of time tracking.

[0063] The sample buffer 220 receives communication signals over the I channel and the Q channel. These signals are sent to time tracking logic 222 and dedicated fingers 234. The time tracking logic 222 also receives a coarse AFC hypothesis and the logic 222 may reset to zero at $\text{chip} \times 4$ parity. The time tracking logic 222 can have two blocks, one with counters initialized to zero for even $\text{chip} \times 4$ parity, and one with counters initialized to midrange (i.e., 2^{25}) for odd $\text{chip} \times 4$ parity. The output of time tracking logic 222 is provided to a block 224 in which virtual $\text{chip} \times 4$ phases are applied. Block 224 also can receive parity from an acquisition state machine. Automatic frequency control (AFC) rotation logic 226 is applied to an output of block 224.

[0064] Fig. 8 illustrates an exemplary embodiment of the two blocks of time tracking logic 222 described with reference to Fig. 7. Stream 250 is a communication stream with an even $\text{chip} \times 4$ parity. Stream 252 is a communication stream with an odd $\text{chip} \times 4$ parity. Fig. 8 depicts the time-tracking operation where each different shading represents a different $\text{chip} \times 4$ spaced sequence. Samples are either inserted or repeated at a rate directly depending on which current AFC hypothesis is being explored, multiplied by a known ratio

between the sample rate and the carrier frequency. This can be used as a locked clock assumption to collapse a 2-dimensional space down to a single dimension. The value N depicted has a fractional component which is book-kept to allow for sufficient time-tracking precision. A particular parity of the 4 possible chip x 4 phases is selected at a given time. The resultant chip rate sequence is then derotated in a 1-bit data path as shown in Fig. 9.

[0065] Fig. 9 depicts the functionality of the AFC (automatic frequency control) rotation logic 226 of Fig. 7 which operates on one of the 4 virtual chip x 4 phases 224 at a given time. Fig. 9 depicts a one-bit derotation mechanism. This derotation mechanism is designed to undo the AFC rotation due to the relative carrier drift between the receiver and transmitter for the postulated coarse AFC hypothesis. Since it's a one-bit transform (represented by the truth table illustrated above), the 90 degree resolution of the process is +/- 45 degrees relative to the continuum of values of the phase due to the AFC drift from the relative oscillator offset.

[0066] The AFC rotation logic 226 can also receive coarse AFC hypotheses as an input. The PN despreading array 228 (Fig. 7) performs its despread operation for chip spaced hypothesis. The PN despreading array 228 may receive current coarse AFC hypotheses, timing parity, timing phase, spreading factor, and/or gold code selection as inputs. As the values are output for a given symbol, the sum is non-coherently accumulated for better metric reliability with the running sum stored in the non-coherent accumulation buffer 230. The size of the buffer is based on the number of despread elements. In an exemplary embodiment, the PN despreading array 228 may have 256 despread elements such that a pass through the sample buffer completes the non-coherent metric for 256 hypotheses. Alternatively, other numbers of despread elements may be used, and the metric may be completed for other numbers of hypotheses. A signal-to-noise ratio (SNR) metric may be used in transmission power control of the tag and for power control feedback to the AP. The hypotheses with the largest metrics are stored in a top N path data structure 232 which is used to control the assignment of the dedicated fingers 234. The top N paths can be N records including timing hypotheses, timing parity, coarse AFC hypotheses, etc.

[0067] Fig. 10 illustrates a dedicated communication finger. Each dedicated finger has access to each of the 4 phases of chip x 4 samples with a chip x 4 selector 260 set as part

of the parameters of the finger assignment. Each finger has its own dedicated PN generator 262 and AFC generator 264 which is used to despread. The dedicated finger accumulates into the symbol accumulator 266 based on the coarse AFC hypothesis, its chip x 4 timing phase, the dependent variable of time-tracking rate, and then outputs a complex variable every spreading factor number of clocks. The dedicated fingers 234 illustrated with reference to Fig. 7 can also receive inputs from the sample buffer 220, and a PN code selection.

[0068] Referring again to Fig. 7, the output from the dedicated fingers 234 goes through a bit-width squeezer 236 that reduces the bit-widths for efficient storage in the frame buffer 238 without sacrificing performance. The output from the bit-width squeezer 236 is provided to the frame buffer 238, which may be a circular buffer mechanism which allows for the general case of processing a 256 symbol frame as if the current symbol is the last symbol of the frame. When frame timing is known, this memory structure can support the specific processing of a frame with the known last symbol.

[0069] Frame buffer 238 outputs the hypothesized frames to the rest of the receive chain. A cross product multiplication block 240 performs the multiplication of the current symbol with the complex conjugate of the previous symbol which is the conventional metric for D-BPSK demodulation. A residual frequency drift may cause the D-BPSK constellation to be rotated by a fixed phase. The role of the fine AFC multiply block 242 is to take a brute-force approach and try different possible phase rotations such that at least one fine AFC hypothesis yields a valid CRC as it passes through a de-interleaver and viterbi decoder 244. The fine AFC multiply block 242 can also receive fine AFC hypotheses as inputs. The output from the de-interleaver and Viterbi decoder 244 is provided to a CRC checker 246. If the CRC is valid, the payload is sent up to the MAC layer.

[0070] Fig. 11 depicts exemplary operations performed during access point receive processing. Additional, fewer, or different operations may be performed depending on the embodiment. Further, the operations can be performed in a different order than that which is described here. The AP performs a brute-force operation checking all possible chip x 2 timing hypothesis, spreading factors, and access slots within spreading factors. This allows for uncoordinated access by the tag. Fortunately, since the AP is the master of frame-timing and AFC carrier reference (all tags can compensate both their carrier drift and sample clock

to meet the AP's timing), the processing burden on the AP is drastically reduced since the AP need not explore the dimensionality of coarse AFC hypothesis or unknown frame timing.

[0071] The flowchart of Fig. 11 shows an example of the ordering of iterating upon all possible chip x 2 timing offset, spreading factors from the set [8192, 4096, ..., 64], and access slot numbers for spreading factors less than the maximum. The AP then performs the similar fine AFC search that the tag performs to allow for a small amount of frequency drift between the timing sources of the tag and the AP to occur since the last transaction. All valid CRCs are passed up to the MAC layer. The flowchart of Fig. 11 illustrates the searching of a multi-dimensional space. In an outermost loop, all possible spreading factors are searched. In an exemplary embodiment, there may be 8 spreading factors [64, 128, 256, 512, 1024, 2048, 4096, 8192]. Alternatively, other spreading factors and/or numbers of spreading factors may be used. In a second loop, all possible sub-slots for a given spreading factor are searched. For example, there may be 128 possible sub-slots for a 64 chip spreading factor and a single degenerate sub-slot for a 8192 chip spreading factor. In a third loop, all possible chip x 2 timing phases within a given sub-slot are searched. As described in more detail below, the various loops are illustrated by the arrows in Fig. 11.

[0072] In an operation 270, one coarse AFC value is used. In an exemplary embodiment, the one coarse AFC value can be 0 since compensation is performed by the tags. In an operation 272, a largest spreading factor (e.g., 8192) is used as a starting point. In alternative embodiments, the largest spreading factor may be larger or smaller than 8192. In an operation 274, access slots are processed within a spreading factor. This process may be degenerate in the case in which there are 8192 spreading factors. In an operation 276, despreading is performed for all chip x 2 spaced hypotheses at the current spreading factor. For example, 16,384 despread operations may be performed if the spreading factor has a length of 8192. Despread is performed for all elements unless the spreading factor is less than the frame buffer number (e.g., 256). In an operation 278, the spreading factor is reduced in half and processing continues. In an operation 280, a determination is made regarding whether the spread factor has been reduced to 64. In alternative embodiments, other predetermined values may be used. If the spread factor has not been reduced to 64 (or other predetermined value), processing continues at operation 276. If the spread factor has been reduced to 64, the system waits for a next sample buffer to fill in operation 282. Once the

next sample buffer is filled in operation 282, control returns to operation 272. In an operation 284, a frame buffer of despread elements is obtained. In an exemplary embodiment, the frame buffer may be complete after 256 symbols are output from a single pass by the PN despread array. In one embodiment, for a 256 stage PN despread array, a pass through may produce 256 timing hypotheses each having 256 symbols. In alternative embodiments, the PN despread array may have more or fewer stages. A cross product of the current despread DBPSK symbol with the previous symbol is calculated in an operation 286. In one embodiment, the cross product may involve 256 symbols for up to 256 frames. Alternatively, other numbers of symbols and/or frames may be used. In an operation 288, the current frame is decoded and phase multiplied based on the AFC hypothesis. In an operation 290, CRCs are checked and for any valid CRC, the payload is sent out of the physical layer (PHY) and up to the medium access control (MAC). As an example, the CRCs may be checked for 256 times the number of fine AFC hypothesis for each pass of a 256 despread array. Upon completion of the process for a given slot, the process is performed for a subsequent slot as illustrated by the arrow from block 282 to block 272.

[0073] Fig. 12 depicts an access point (AP) receive data path. Unlike the tag, an entire frame at the largest spreading factor may be stored in a ping-pong buffer scheme in a sample buffer 300. This buffer scheme can be a substantial amount of memory (e.g., 16.8 Mbits) and in at least one embodiment, it may be stored in a dedicated off-chip memory device. The sample buffer block 300 includes exemplary values. In alternative embodiments, other values may be used. Unlike the tag, the time tracking logic and the AFC rotation logic may not be used since the AP is the master time reference. The sample buffer 300 passes frames to a PN despread array 302, which can perform brute force testing as described previously herein. The PN despread array 302 may include 256 despread elements. Alternatively, any other number of despread elements may be used. The PN despread array 302 may also receive current timing parity (which may be chip x 2 resolution only), hypothesis phase, and/or spreading factor as inputs. An output from the PN despread array 302 is provided to a bit width squeezer 304. The bit width squeezer 304 reduces the size of the frames, which are then sent to a frame buffer 306. The frame buffer block 306 includes exemplary values. In alternative embodiments, other values may be used. Depending on the embodiment, the frame buffer 306 may also be stored in a dedicated off-

chip memory device. The rest of the system is similar to the tag's receive processing where fine AFC hypothesis are iterated upon (operations 310 and 312) with all payloads with valid CRCs being passed up to the AP's MAC (operations 314 and 316). A non-coherent accumulation 308 is used to determine an SNR metric such as signal strength for use in transmission power-control feedback to the tag.

[0074] Fig. 13 illustrates asynchronous initial tag transmit operations, including two types of interactions which result in data transfers from the tag to the AP. For purposes of illustration and discussion, slots 320 represent tag slots and slots 322 represent access point slots. "Cold Start" is where the tag is coming into the system without any relevant state information and "warm start" is where the tag is aware of the system information such as slot timing and a reduced range of coarse AFC hypothesis to explore.

[0075] In the "Cold Start" scenario, the tag begins seeking access at a slot-asynchronous point in time. Fig. 13 depicts a time where the tag begins attempting to acquire the broadcast channel when the AP isn't even transmitting it (slot 1). Eventually, the tag's processing explores the valid coarse AFC hypothesis during a period of time that the AP is transmitting the broadcast frame. Fig. 13 depicts this occurring during slot 2. At this point, the non-coherent energy metric causes a dedicated finger to explore the correct chip x 4 timing and coarse AFC hypothesis. The finger with the correct hypothesis continually treats each new symbol as the last symbol of the frame and pushes these hypothesized frames through the receive chain where the CRC check indicates failure. At the end of slot 4, the valid frame timing is achieved as the CRC check indicates success. At this point, the tag has the same relevant state information that a tag entering at a "warm-start" would have and continues to complete the same processing that a "warm-start" tag would undergo.

[0076] A tag enters the interaction depicted in slot 6 ("Warm Start") either by a transition through a "Cold Start" procedure or directly upon tag wake-up if relevant state information is appropriately maintained. At this point, the tag makes a measurement of the received strength of the broadcast frame and uses this information to determine the transmit power and spreading factor that the tag subsequently transmits at in slot 7. The tag transmits its message based on: 1) using the measured received broadcast channel signal strength and selecting the minimum spreading factor that can be used to close the link, which minimizes

the tag's on time and is best for minimizing power consumption; 2) using the measured received broadcast channel signal strength and the formerly selected spreading factor, the tag transmits at the optimality condition of reception at the AP which is that all user's are received by the AP at very similar values of energy per bit to spectral noise density ratio (E_b/N_0); 3) for all but the maximum spreading factor, randomly selecting the slot access parameter j ; and 4) randomly selecting the chip offset value from 0 to spreading factor -1 such that "collisions" at the AP are minimized and random selection at each transmission allows "collisions" to be resolved in subsequent transmission opportunities.

[0077] During slots 8 and 9, the AP processes all the signals received during slot 7 and sends a positive acknowledgement back during slot 10. The AP either aggregates several ACKs into a single channel characterized by a gold code, or sends a dedicated message to the tag using its dedicated gold code channel. Note that the former method requires some registration procedure (not shown) to assign the channel. In either case, the tag updates its chip $\times 4$ timing using the preamble of the message.

[0078] Fig. 14 illustrates a simple interaction between an access point and a tag in a slotted mode. In an exemplary embodiment, the simple interaction involves no data for the tag and a relatively static channel. For purposes of illustration and discussion, timeline 330 represents tag processing during the slots and timeline 332 represents access point processing during slots. The nature of the system is that the tag spends a maximum possible time in a low-power state – a state where system timing is maintained via a low-power, low-frequency crystal oscillator which is typically 32 kHz. To support this, a maximum tolerable latency upon AP initiated interaction is identified (i.e., this is the rate cycling in and out of the low power state for the tag to check if any AP action is pending). Fig. 14 shows the relatively simple interaction of a tag coming out of its low power state to check if the AP is wanting to initiate a transaction. This occurs at a slot phase and rate agreed upon between the AP and the tag during registration.

[0079] The tag would typically enter a "warm start" where the frame timing and coarse AFC hypothesis are known to within a tight range. The tag makes a measurement of the received broadcast channel power. Fig. 14 shows the scenario where that power has not changed considerably since the last interaction with the AP. This means that the last transmit

power/spreading factor that the AP transmitted at is sufficient to close the link. In slot 3, the tag attempts to acquire on the preamble and then demodulate the frame using its dedicated gold code. A typical scenario is the AP not having sent information and the tag immediately goes back to sleep.

[0080] Fig. 15 depicts a more complicated interaction which involves data transfer and dynamically changing propagation between an access point and a tag according to an exemplary embodiment. For purposes of illustration and discussion, timeline 340 represents tag processing during the slots and timeline 342 represents access point (AP) processing during the slots. Here, the AP has information to send and the propagation of the channel has changed considerably since the last AP transaction. The current broadcast channel power measurement has changed such that the tag knows that the subsequent transmission would not be appropriate if it transmits at the same transmit power/spreading factor as last time. Thus, the tag will send a re-registration message using the protocol explained in Fig. 13 to alert the AP to use a new transmit power/spreading factor appropriate to the current channel conditions. The new information governs the transmission and reception of the frame occurring in slot N+5. The tag generates an acknowledgement (ACK) message governed by the protocol of Fig. 13 to indicate a successful transmission. If the ACK is successfully received, the transaction is considered complete. Otherwise, the tag attempts a retransmission.

[0081] In one embodiment, the air interface utilized by the systems described herein can be a half duplex time division multiplexed frame format. The access point can transmit for a portion of the time in the downlink direction to the tag, and the tag can transmit for a portion of the time in the uplink direction to the access point. The time allocation between the uplink slot and the downlink slot may be equal (i.e., 50% of the time is allocated to the uplink slot and 50% of the time is allocated to the downlink slot). The frame structure can be centered about a slot structure whose format numerics can be based on a maximum supported uplink spreading factor. In an exemplary embodiment, the maximum spreading factor at the uplink can be that which allows the tag to successfully transmit to the access point when the tag is under the most challenging transmit conditions based on weather, location, etc.

[0082] In general, use of a large spreading factor can allow a given transmitter such as a tag to transmit with less power while still being able to be received by a given receiver such as an access point. However, use of a large spreading factor can also increase the time that it takes to transmit a signal. In an exemplary embodiment, the tag may broadcast at a lower power than that used by the access point. As such, the spreading factor of the uplink signal can be selected as large enough such that the signal transmitted by the tag can be received by the access point even when the tag is physically located in a challenging location and/or under challenging transmission conditions. The access point may transmit with more power than the tag. As a result, if the uplink (i.e., tag to access point) transmissions and downlink (i.e., access point to tag) transmissions are given equal amounts of time on the band in which to transmit, the access point can use a smaller spreading factor than the tag. Since the access point signals are not as widely spread, the access point can transmit in a plurality of fundamental downlink slots in the same amount of time as the tag transmits in a single slot. In one embodiment, the access point can transmit at a constant power at or near the maximum RF transmit power available. If the link between the access point and a given tag is robust, a reduced spreading factor can be used for that tag. Robustness of the link can be determined based on a comparison of the link to a predetermined quality threshold. Because the lower spreading factor takes less total time to transmit, the tag can open its receive window for a relatively short period of time, thereby minimizing power consumption of the tag.

[0083] Fig. 16 is a diagram illustrating contents of a complete slot 400 in accordance with a first exemplary embodiment. In the embodiment of Fig. 16, the access point can transmit with more power and use a smaller spreading factor as compared to the tag. For example, the access point may use a spreading factor of 2048 and the tag may use a spreading factor of 8192. Alternatively, other values may be used. Complete slot 400 includes a downlink slot 402, a downlink to uplink gap 404, an uplink slot 406, and an uplink to downlink gap 408. In an exemplary embodiment, downlink to uplink gap 404 may be 15 symbols at the maximum uplink spreading factor (which may be 8192 in one embodiment). Alternatively, any other length downlink to uplink gap may be used. Downlink to uplink gap 404 can be used to ensure a downlink to uplink ratio of 50%. Downlink to uplink gap 404 can also be used to provide tag receiver to transmitter turnaround processing time. In another exemplary embodiment, uplink to downlink gap 408 may be 1 symbol at the maximum

uplink spreading factor. Alternatively, any other length uplink to downlink gap may be used. Uplink to downlink gap 408 can be used to support random phase multiple access (RPMA) for a given spreading factor. As such, a smaller uplink to downlink gap can be used with smaller spreading factors.

[0084] Because the access point may use a smaller spreading factor than the tag, the downlink slot can include a plurality of fundamental downlink slots. Downlink slot 402 includes a fundamental downlink slot 410, a fundamental downlink slot 412, a fundamental downlink slot 414, and a fundamental downlink slot 416. Each of the fundamental downlink slots includes a broadcast preamble 418, a data slot or subslot(s) 420, and a broadcast slot 422. In an exemplary embodiment, broadcast preamble 418 can be 16 symbols. Alternatively, any other length may be used.

[0085] Fig. 17 is a diagram illustrating contents of a complete slot 430 in accordance with a second exemplary embodiment. In the embodiment of Fig. 17, the access point can transmit with the same power as used by the tag. As such, the same spreading factor may also be used by the access point and the tag. For example, the access point and the tag may both use a maximum spreading factor of 8192. Alternatively, other values may be used. Complete slot 430 includes a downlink slot 432, a downlink to uplink gap 434, an uplink slot 436, and an uplink to downlink gap 438. In an exemplary embodiment, downlink to uplink gap 434 may be 15 symbols x 8192 chips. Alternatively, any other length downlink to uplink gap may be used. In another exemplary embodiment, uplink to downlink gap 438 may be 1 symbol x 8192 chips. Alternatively, any other length uplink to downlink gap may be used. Because the access point uses the same spreading factor as the tag, downlink slot 432 includes a single fundamental downlink slot 440. Fundamental downlink slot 440 includes a broadcast preamble 442, a data slot or subslot(s) 444, and a broadcast slot 446.

[0086] Fig. 18 illustrates a fundamental downlink slot 450 in accordance with an exemplary embodiment. Fundamental downlink slot 450 includes a broadcast preamble 452, a broadcast channel slot 454, and a data channel slot 456. Broadcast preamble 452 can be 16 symbols long, or any other length depending on the embodiment. In an exemplary embodiment, broadcast channel slot 454 can include a single broadcast channel frame. In one embodiment, the broadcast channel frame can be identical in creation to a data channel frame

with the exception that the broadcast channel gold code generator may reset every symbol whereas the data channel gold code generator may run until the end of the data channel frame before resetting.

[0087] In an exemplary embodiment, broadcast preamble 452 can be boosted relative to other transmissions made using broadcast channel slot 454 or data slot 456. As an example, broadcast preamble 452 can be transmitted at a maximum power ($P_{\max}$), and other transmissions can be made at one half of the maximum power ($1/2 P_{\max}$). In one embodiment, broadcast preamble 452 can be boosted by 3 decibels (dB) relative to other transmissions via broadcast channel slot 454 and/or data slot 456. Alternatively, broadcast preamble 452 may be boosted by any other amount. The boosted preamble allows receivers at the tags to robustly estimate chip timing and AFC/time tracking with reference to the access point. The payload of broadcast preamble 452 can be programmable. In one embodiment, no channel coding, interleaving, or cyclic redundancy check (CRC) may be applied to the payload of broadcast preamble 452.

[0088] Fig. 19 is a diagram illustrating preamble frame processing in accordance with an exemplary embodiment. A 16-bit register 460 can provide symbols to a modulator 462 for modulation. Alternatively, any other length register may be used. Modulator 462 can output modulated symbols at a symbol rate. Modulator 462 can be a differential binary phase shift keying (DBPSK) modulator, or any other type of modulator known to those of skill in the art. As a result, each symbol of broadcast preamble 452 can be modulated, spread with the maximum downlink spreading factor, and boosted for transmission. In an exemplary embodiment, broadcast preamble 452 may be boosted by ensuring that no other data is transmitted during transmission of broadcast preamble 452. For example, broadcast preamble 452 may be broadcast at $P_{\max}$ through broadcast channel slot 454 and data channel slot 456 may be turned off while broadcast preamble 452 is being broadcast. In one embodiment, broadcast preamble 452 can be transmitted on one of an I channel or a Q channel. Broadcast preamble 452 can be multiplied by a scale factor of 1 such that broadcast preamble 452 is transmitted at full power. When broadcast preamble 452 is not being broadcast, an attenuating scale factor can be used such that data is transmitted at less than full power. In one embodiment, the attenuating scale factor can be $1/\sqrt{2}$, resulting in a 3 dB attenuation. In an alternative embodiment, the scale factor may not be altered. In such an

embodiment, broadcast preamble 452 can be transmitted on both the I channel and the Q channel such that broadcast preamble 452 is transmitted at full power.

[0089] Referring again to Fig. 18, data channel slot 456 may contain a single data channel frame. Alternatively, data channel slot 456 may contain a plurality of data channel frames in a single fundamental downlink slot such as fundamental downlink slot 450. As a result, data channel slot 456 of fundamental downlink slot 450 can include a plurality of sub-slots corresponding to the plurality of data channel frames (i.e., one sub-slot for each frame). In an exemplary embodiment, the spreading factor of data channel slot 456 can be the same as the spreading factor of broadcast channel slot 454. In another exemplary embodiment, the spreading factor of data channel sub-slots can be less than the spreading factor of broadcast channel slot 454.

[0090] In one embodiment, multiple data channel sub-slots can be created by using smaller spreading factors than those used with either a full size (i.e., single) data channel slot or with the broadcast channel slot. Fig. 20 is a diagram illustrating a data sub-slot hierarchy in accordance with an exemplary embodiment. As illustrated in Fig. 20, if a spreading factor of 8192 and a 39 dB gain are used, the data channel includes a single data channel slot $A_{0,0}$. If a spreading factor of 4096 and a 36 dB gain are used, the data channel slot includes two sub-slots $A_{1,0}$ and $A_{1,1}$. Similarly, if a spreading factor of 16 and a 12 dB gain are used, the data channel slot includes 512 sub-slots $A_{9,0} \dots A_{9,511}$, and so on.

[0091] Fig. 21A illustrates a fundamental downlink slot 470 with a plurality of sub-slots in accordance with an exemplary embodiment. As illustrated in Fig. 21A, sub-slots of different sizes can be combined to form the data channel slot. The plurality of sub-slots includes a sub-slot $A_{4,0}$, a sub-slot $A_{4,1}$, a sub-slot $A_{5,4}$, and a sub-slot $A_{6,10}$. In alternative embodiments, other combinations of sub-slots may be used. A broadcast channel slot 472 of fundamental downlink slot 470 can have a spreading factor of 2048. Alternatively, other values may be used. A tag or other receiving device can turn on its receiver to listen to one or more of these sub-slots as appropriate. Fig. 21B illustrates a fundamental downlink slot 480 with a single slot $A_{0,0}$ in accordance with an exemplary embodiment. A broadcast channel slot 482 of fundamental downlink slot 480 can have a spreading factor of 8192. Alternatively, other values may be used.

[0092] Fig. 22 is a flow diagram illustrating operations performed to construct a frame in accordance with an exemplary embodiment. In alternative embodiments, additional, fewer, or different operations may be performed. Further, the use of flow diagrams herein is not meant to be limiting with respect to the order of operations performed. The frame can be for use in a data channel slot and/or a broadcast channel slot. In an operation 500, a frame base is generated. In an exemplary embodiment, the frame base can be an 88-bit payload. Alternatively, other numbers of bits may be used. In an operation 502, a cyclic redundancy check (CRC) is appended to the frame base. In an exemplary embodiment, the CRC can be 32 bits, resulting in a frame of 120 bits. Alternatively, other values may be used. In an operation 504, tail bits are added to the frame. In an exemplary embodiment, 8 tail bits can be added, resulting in a raw frame of 128 bits. Alternatively, a different number of tail bits may be used. In one embodiment, each of the tail bits can have a value of zero. Alternatively, any or all of the tail bits may have a non-zero value. In an operation 506, the raw frame is convolution encoded. The convolution encoding can be performed at a rate of $1/2$, or any other value depending on the embodiment. In one embodiment, an output of the convolution encoder used to perform the convolution encoding can be 256 bits.

[0093] In an operation 508, the bits of the frame are symbol interleaved. In one embodiment, the bits can be interleaved with a strict bit reversed interleaver which utilizes bit reversed addressing. As an example, in a frame buffer containing 256 symbols, the interleaver input addressing can be linear, from 0 thru 255. Each address can be 8 bits. The interleaver can take data at a particular address, and put it in a new position in an output frame buffer. The output address can be a bit reversed ordering of the input address. For example, the symbol at address 15 (00001111b) can be placed in address 240 (1111000b). In an illustrative embodiment, each interleaver input address can be bit reversed to form an output address. In an operation 510, the bits of the frame are modulated. The modulation can be DBPSK modulation. Alternatively, any other type of modulation may be used. The bits can also be spread with a spreading factor based at least in part on the slot size.

[0094] In one embodiment, a complex data stream can be created to minimize bandwidth usage. The data channel can exist on a real component of the complex data stream and the broadcast channel can exist on an imaginary component of the complex data stream, or vice versa. Fig. 23 is a diagram illustrating a downlink transmission model in accordance

with an exemplary embodiment. A data frame can be constructed by a data frame processor 520, a broadcast frame can be constructed by a broadcast frame processor 522, and a broadcast preamble frame can be constructed by a broadcast preamble frame processor 524. The data frame, broadcast frame, and/or broadcast preamble frame can be constructed in accordance with the operations described with reference to Fig. 22. Alternatively, different operations may be performed to construct the frames.

[0095] The data frame is provided to a PN spreader 526 for spreading. PN spreader 526 can receive a frame boundary input to reset PN spreader 526, a user key input to initialize a state of PN spreader 526, and an enable data input to enable PN spreader 526. For example, the frame boundary input can be an indication that a frame is beginning or ending such that the PN/gold code used by PN spreader 526 is reset for each data frame. The user key input can be tied to a tag identification of a tag which is located in a network of the access point. The user key input (or tag identification) can directly affect the PN/gold code generated, and can allow the tag to decode messages that are targeted for the tag in the downlink. In one embodiment, each frame generated by the access point can be based on a particular user key (or tag identification). The enable data input can window the data channel frame. The enable data input can stay high for the duration of the frame, and may span multiple frames during a downlink slot. In one embodiment, PN spreader 526 can run as long as the enable data input is high. An output of PN spreader 526 can be used as the real component of the complex data stream. Alternatively, the output of PN spreader 526 may be used as the imaginary component of the complex data stream.

[0096] The broadcast frame and the broadcast preamble frame are provided to a selector 528 for provision of one of the broadcast frame or the broadcast preamble frame to a PN spreader 530. Selector 528 can receive a preamble enable input to control whether PN spreader 530 receives data from broadcast frame processor 522 or preamble frame processor 524. Pseudo-noise spreader 530 can receive a symbol boundary input to reset PN spreader 530, a broadcast key input to initialize a state of PN spreader 530, and an enable broadcast input to enable PN spreader 530. For example, the symbol boundary input can be an indication that a symbol is beginning or ending such that the gold/PN code used by PN spreader 530 is reset for each symbol. Resetting the gold/PN code after each symbol can make it easier for the tag to acquire the signal broadcast from the access point. Also, by

resetting the gold/PN code on every symbol of the broadcast frame, the code space that the tag has to search is reduced. The broadcast key input can be common for a given network, and can directly affect the gold/PN code sequence that is generated. As an example, different access point networks may have different broadcast channel keys which are used as network identifications. The enable broadcast input can stay high for the duration of the symbol, and PN spreader 530 can run as long as the enable broadcast input remains high. An output of PN spreader 530 can be used as the imaginary component of the complex data stream. Alternatively, the output of PN spreader 530 may be used as the real component of the complex data stream.

[0097] An output of PN spreader 526 and an output of PN spreader 530 can be provided to an up-sampler 532. In one embodiment, up-sampler 532 can up-sample the received signals to 26 MHz. Alternatively, 40 MHz or any other amount of up-sampling may be performed. Up-sampler 532 has a preamble enable input and a data enable input. The preamble enable input can be activated when selector 528 provides the broadcast preamble frame to PN spreader 530 and the data enable input can be activated when selector 528 provides the broadcast frame to PN spreader 530. In an exemplary embodiment, activation of the preamble enable input can cause the broadcast preamble to be boosted on the broadcast channel (which can be the imaginary component of the complex data stream). In one embodiment, polyphase filter taps can incorporate a $1/\sqrt{2}$ gain on the broadcast preamble (or a $1/\sqrt{2}$ attenuation on transmissions other than the broadcast preamble). Activation of the preamble enable input can also turn off the data channel such that the real component of the complex data stream does not broadcast simultaneously with the broadcast preamble. Activation of the data enable input can cause the broadcast frame to be transmitted on the imaginary component of the complex data stream simultaneously with the data frame on the real component of the complex data stream. As such, when the data enable input is activated, up-sampler 532 can receive the data frame and the broadcast frame. Alternatively, the broadcast preamble may be transmitted on the real component of the complex data stream. In another alternative embodiment, the broadcast preamble may be boosted by simultaneously transmitting the broadcast preamble on both components of the complex data stream.

[0098] If the broadcast preamble is being transmitted, up-sampler 532 can provide the up-sampled broadcast preamble to a converter 534. Converter 534 can convert the

broadcast preamble from digital to analog for transmission on one or both components of the complex data stream. If the broadcast frame and data frame are being transmitted, up-sampler 532 can provide the up-sampled data frame to converter 534 on the real component (i.e., I channel) of the complex data stream and the up-sampled broadcast frame to converter 534 on the imaginary component (i.e., Q channel) of the complex data stream, or vice versa. Converter 534 can convert the data frame and the broadcast frame from digital to analog for transmission. Converter 534 can also provide the data frame and the broadcast frame to an RF up-converter 536 for combination into the single complex data stream for bandwidth savings during transmission. Radio frequency up-converter 536 can be part of the RF chip. In one embodiment, the I data stream (real component) and the Q data stream (imaginary component) can be independently differentially binary phase shift keyed. As a result, bandwidth can be conserved as the broadcast channel does not have to occupy a side channel.

[0099] In one embodiment, to preserve downlink bandwidth, the broadcast channel may be partially or fully utilized as an acknowledgment (ACK) channel. In such an embodiment, each bit of the broadcast frame payload (i.e., up to 88 in one embodiment) can represent an ACK bit or a non-acknowledgement (NACK) bit. After network registration, the tag can have knowledge of which bit of the broadcast frame payload corresponds to an ACK channel for the tag. As a result, the broadcast channel can be used as a multi-tag ACK channel, and the access point can transmit multiple acknowledgements in a single fundamental downlink slot.

[0100] In an alternative embodiment, the access point may send out acknowledgements serially to a plurality of tags, and the tags may not know in which sub-slot its ACK is located. In such an embodiment, the ACKs may be transmitted on the data channel. As the tag may not know the sub-slot of its ACK, the tag can utilize non-deterministic frame processing to identify the ACK. In one embodiment, the tag can keep its reception window open for a plurality of sub-slots until the ACK is received. As an example, the access point may decode transmissions from a first tag and a second tag in the same sub-slot with the same spreading factor. If the downlink is TDMA, the access point can serially send out a first ACK to the first tag and a second ACK to the second tag using the data channel. The first tag and the second tag can each keep their respective receive windows

open until the respective ACK is received or until the end of the downlink slot of the complete frame.

[0101] In one embodiment, the access point may transmit acknowledgements and/or other information to tags serially, with only one acknowledgement or other information transmitted during a given transmission slot of the access point. In such an embodiment, a tag may have to wait more than one slot to receive the acknowledgement or other information. As an example with acknowledgements, three tags may transmit to an access point during a first transmission slot. The access point may transmit a first acknowledgement to the first tag during a second transmission slot, a second acknowledgement to the second tag during a third transmission slot, and a third acknowledgement to the third tag during a fourth transmission slot. As such, the third tag has to wait several transmission slots prior to receiving the third acknowledgement. If the third tag is not still waiting during the fourth transmission slot, the third tag may not receive the third acknowledgement, and may have to retransmit the original message to the access point. In one embodiment, the tags may not know which transmission slot, if any, will include their acknowledgement or other information. As such, tags can be configured to wait a predetermined number of transmission slots based on a probability that multiple tags will transmit to the access point during a given slot. As an example, in a given system, it may be highly improbable that four tags will transmit to the access point during a single transmission slot. As such, the tags may be configured to wait no more than three transmission slots of the access point for an acknowledgement or other information before retransmitting the original message to the access point. In alternative embodiments, the tags may be configured to wait any other number of slots for an acknowledgement or other information.

[0102] In an exemplary embodiment, individual tags can be communicated to in the downlink direction in a time division multiple access (TDMA) scheme. As such, each tag may receive one frame worth of data channel data in a given complete slot duration. The data channel data can be decoded by the tag with a unique gold code key specific to the tag. For increased bandwidth or for specific quality of service (QoS) scenarios, the tag may receive multiple sub-slot allocations in a complete slot duration. The tag can use back-to-back decoding to decode the data from the multiple sub-slots. Details of the bandwidth management can be handled by the medium access control (MAC) layer.

[0103] The uplink frame structure can be an aggregation of multiple tag data channel transmissions. Depending on link conditions, tags may transmit with a plurality of different spreading factors. In one embodiment, the spreading factor may be as low as 16 and as high as 8192. Alternatively, other spreading factor limits may be used. The uplink slot can be partitioned into sub-slots in the same way as the downlink slot as described with reference to Figs. 20 and 21. A difference in the uplink can be the way in which multiple access is achieved. In the downlink, multiple access can be achieved as described above through time division multiple access. In the uplink, multiple tag transmissions may be able to occupy the same sub-slot. In such an embodiment, the ability of the access point to discriminate among a plurality of tags can be realized through a technique called random phase multiple access (RPMA).

[0104] Random phase multiple access can be based on a common gold code sequence mapped to a given spreading factor. For example, all tags transmitting with a spreading factor of 256 can use the same gold code regardless of sub-slot location within the uplink slot. Because of the correlation properties of the gold code, the receiver at the access point can discriminate among a plurality of different tags as long as each tag uses a different chip offset. When two or more tags compete for a transmission sub-slot, frame collisions can occur if the randomly chosen sub-slot and the randomly chosen chip offset are the same. Such an occurrence can be identified by the lack of an acknowledgement response from the access point at the tag. If the tag does not receive the acknowledgement response, the tag can randomly select a new sub-slot and a new gold code chip offset. The tag can retransmit using the new sub-slot and new chip offset, thereby minimizing the likelihood of a subsequent conflict.

[0105] In one embodiment, the systems described herein may also be used where multiple access points are present and/or where multiple micro-repeaters are broadcasting on the same channel. In such a system, a random timing offset may be introduced to the downlink spreading system. Each device that transmits on the same channel with the same pseudo-noise (PN) code can use a different random or pseudo-random timing offset into the PN code. As such, receiving devices such as tags can distinguish the multiple transmitters by de-spreading using the appropriate offset into the PN code.

[0106] It is important to understand that any of the embodiments described herein may be implemented as computer-readable instructions stored on a computer-readable medium. Upon execution by a processor, the computer-readable instructions can cause a computing device to perform operations to implement any of the embodiments described herein.

[0107] The foregoing description of exemplary embodiments has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the present invention to the precise form disclosed, and modifications and variations are possible in light of the above teachings or may be acquired from practice of the present invention. The embodiments were chosen and described in order to explain the principles of the present invention and its practical application to enable one skilled in the art to utilize the present invention in various embodiments and with various modifications as are suited to the particular use contemplated. In addition, one or more flow diagrams were used herein. The use of flow diagrams is not intended to be limiting with respect to the order in which operations are performed.

WHAT IS CLAIMED IS:

1. A method for conserving bandwidth in a communication system, the method comprising:
 - spreading a data frame using a first pseudo-noise (PN) spreader;
 - spreading a broadcast frame using a second PN spreader;
 - generating a complex data stream having a first component and a second component, wherein the data frame is assigned to the first component and the broadcast frame is assigned to the second component, and wherein the first component comprises a real component of the complex data stream and the second component comprises an imaginary component of the complex data stream; and
 - transmitting the complex data stream to a tag.
2. The method of claim 1, wherein the broadcast frame includes a payload with a plurality of bits, wherein one or more of the plurality of bits is utilized to convey an acknowledgement to the tag.
3. The method of claim 2, wherein at least one of the plurality of bits is dedicated to the tag for providing the acknowledgement.
4. The method of claim 1, further comprising:
 - spreading a preamble with one of the first PN spreader or the second PN spreader;
 - boosting the preamble such that the preamble is broadcast at a higher power than the broadcast frame or the data frame; and
 - transmitting the preamble to the tag.
5. The method of claim 4, further comprising turning off a data channel during transmission of the preamble, wherein the preamble is transmitted on a broadcast channel.

6. The method of claim 1, wherein the data frame and the broadcast frame are spread with a spreading factor, and further wherein the spreading factor is based at least in part on a quality of a link with the tag.

7. The method of claim 6, further comprising using a reduced spreading factor if the quality of the link satisfies a quality threshold such that a transmission time is minimized.

8. The method of claim 1, further comprising:
receiving a first transmission from the tag and a second transmission from a second tag; and
serially transmitting, in response to the first transmission and the second transmission, first information to the tag and second information to the second tag, wherein the tag and the second tag are unaware of an order in which the first information and the second information are transmitted.

9. An access point comprising:
a processor configured to
spread a data frame using a first pseudo-noise (PN) spreader;
spread a broadcast frame using a second PN spreader; and
generate a complex data stream having a first component and a second component, wherein the data frame is assigned to the first component and the broadcast frame is assigned to the second component, and wherein the first component comprises a real component of the complex data stream and the second component comprises an imaginary component of the complex data stream; and
a transmitter operatively coupled to the processor and configured to transmit the complex data stream to a tag.

10. The access point of claim 9, further comprising a radio frequency (RF) up-converter, wherein processor uses the RF up-converter to generate the complex data stream.

11. The access point of claim 9, wherein the broadcast frame includes a payload with a plurality of bits, wherein one or more of the plurality of bits is utilized to convey an acknowledgement to the tag.

12. The access point of claim 11, wherein at least one of the plurality of bits is dedicated to the tag for providing the acknowledgement.

13. The access point of claim 9, wherein the processor is further configured to
spread a preamble with one of the first PN spreader or the second PN spreader;
and
boost the preamble such that the preamble is broadcast at a higher power than the broadcast frame or the data frame.

14. The access point of claim 9, wherein the data frame is spread with a spreading factor, and further wherein the spreading factor is based at least in part on a quality of a link with the tag.

15. A computer-readable medium having computer-readable instructions stored thereon that, upon execution by a processor, cause an access point to:
spread a data frame using a first pseudo-noise (PN) spreader;
spread a broadcast frame using a second PN spreader; and
generate a complex data stream having a first component and a second component, wherein the data frame is assigned to the first component and the broadcast frame is assigned to the second component, and wherein the first component comprises a real component of the complex data stream and the second component comprises an imaginary component of the complex data stream; and
transmit the complex data stream to a tag.

16. The computer-readable medium of claim 15, wherein the computer-readable instructions further cause the access point to transmit a preamble on the first component and the second component of the complex data stream such that the preamble is transmitted with higher power than the data frame.

17. The computer-readable medium of claim 15, wherein the computer-readable instructions further cause the access point to transmit a preamble on a broadcast channel of the complex data stream, wherein the broadcast channel corresponds to the second component of the complex data stream.

18. The computer-readable medium of claim 15, wherein the computer-readable instructions further cause the access point to apply a scaling factor of $1/\sqrt{2}$ to the broadcast frame such that the broadcast frame is transmitted with less power than the preamble.

19. The computer-readable medium of claim 15, wherein the broadcast frame includes a payload with a plurality of bits, and further wherein one or more of the plurality of bits is utilized to convey an acknowledgement to the tag.