



- (51) International Patent Classification:
G08B 9/00 (2006.01)
- (21) International Application Number:
PCT/US2014/067609
- (22) International Filing Date:
26 November 2014 (26.11.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
61/910,952 2 December 2013 (02.12.2013) US
- (71) Applicant: BOARD OF TRUSTEES OF MICHIGAN STATE UNIVERSITY [US/US]; 450 Administration Building, East Lansing, Michigan 48824 (US).
- (72) Inventors: BISWAS, Subir; 4505 Mistywood Drive, Okemos, Michigan 48864 (US). LORENZ, Stephan; 2545 Candlewick Drive, Lake Orion, Michigan 48359 (US). DONG, Bo; 471 Acalanes Dr., #46, Sunnyvale, California 94086 (US). HUO, Qiong; 1063 Morse Ave., Apt. 12-105, Sunnyvale, California 94089 (US).
- (74) Agents: FALCOFF, Monte L. et al.; Harness, Dickey & Pierce, P.L.C., P.O. Box 828, Bloomfield Hills, Michigan 48303 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,

[Continued on next page]

(54) Title: ULTRASONIC COMMUNICATION SYSTEM




FIG. 19

(57) Abstract: A sensor network includes a sink and multiple sensor nodes. The sink is coupled to a substrate and configured to transmit a periodic ultrasonic pulse on the substrate. A first one of the sensor nodes is coupled to the substrate. The first sensor node is configured to (i) receive the periodic ultrasonic pulse from the substrate, (ii) synchronize an internal clock of the first sensor node to the sink based on the periodic ultrasonic pulse, (iii) selectively detect an event in a region surrounding the first sensor node, and (iv) in response to detecting the event, transmit a first ultrasonic pulse toward the sink on the substrate.



TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))
- of inventorship (Rule 4.17(iv))

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

ULTRASONIC COMMUNICATION SYSTEM

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 61/910,952, filed on December 2, 2013. The entire disclosure of the application
5 referenced above is incorporated herein by reference.

GOVERNMENT RIGHTS

[0002] This invention was made with government support under Contract No. CNS-0915851 awarded by the National Science Foundation. The government has certain rights in the invention.

10 FIELD

[0003] The present disclosure relates to multi-hop networks and more particularly to multi-hop networks communicating using ultrasonic pulses to encode event information.

BACKGROUND

15 **[0004]** The background description provided here is for the purpose of generally presenting the context of the disclosure. Work of the presently named inventors, to the extent it is described in this background section, as well as aspects of the description that may not otherwise qualify as prior art at the time of filing, are neither expressly nor impliedly admitted as prior art against the present
20 disclosure.

[0005] In wireless sensor networks, sensors are responsible for detecting, monitoring, and tracking events with as low an energy consumption as possible. A specific application of sensor networks is binary event sensing. An example application for such sensing is structural health monitoring. For example, a
25 sensor network may monitor the structural integrity of a bridge. In various implementations, each sensor may simply determine whether or not the presence of an abnormality, such as a crack, has been sensed. The sensor may

be configured to send a fatigue event notice in response to sensing the abnormality and remain silent otherwise.

[0006] In another example, a sensor network may perform intrusion detection. Each sensor may transmit an intrusion event notice in response to sensing intrusion in its vicinity. A sink receives transmitted event notices. The sink, or a further processing system, can analyze the timing and approximate location of multiple binary failure events or intrusion events to study the dynamics of structural failure or the path of an intruding entity.

[0007] Simply transmitting an indication of the occurrence of an event requires a single bit of information (i.e., binary information). Traditional packet communication is highly energy inefficient when transmitting a single bit of actual data. Such inefficiency stems from communication, processing, and buffering overhead for each packet (including header, checksum, and synchronization preamble). Instead of incurring the overhead of a packet, sensors may transmit a single pulse to signify the occurrence of an event. In a multi-hop network (i.e., where one or more sensors cannot transmit directly to the sink), sensors may relay pulses from other sensors toward the sink.

Hop-Angular

[0008] A pulse generated by a sensor represents the occurrence of the event as well as its location of origin. A concept of hop-angular event area can be used for event localization. The network is logically divided into a fixed number of angular sectors. In FIG. 1, for example, there are 16 22.5°-wide sectors. Given a pre-defined sector width (22.5° in FIG. 1), the location of a sensor can be represented by the tuple [sector_id, hop_distance], where hop_distance is the number of transmissions required for a pulse to arrive at the sink. Meanwhile, if the sectors are numbered counter clockwise, with the first sector above 0° being sector 1, the location of the hashed Event Area in FIG. 1 can be represented as [15, 3]. This means the node is located in the 15th sector and has a hop_distance of 3 from the sink. This tuple for an event's origin is carried to the sink by the corresponding pulse.

[0009] The concept of event area does not assume any specific shape (i.e., circular or otherwise) of a node's transmission coverage area. It could be of any arbitrary shape as shown in FIG. 1. While the angle for a node is pre-programmed at deployment time, its hop distance may be dynamically
 5 discovered. The worst-case resolution for event localization increases with increasing sector width and increasing transmission range.

Frame Structure

[0010] Nodes are time-synchronized by the sink at the beginning of each frame, which defines the times during which the sink and the sensors can transmit. See
 10 FIG. 2 for the frame definition. Each slot can be used for sending a single pulse. The slot includes a guard time so that even with the maximum amount of clock drift, a pulse transmission will still occur within the desired slot.

[0011] As shown in FIG. 2, the frame contains an uplink (toward the sink) portion and a downlink (away from the sink) portion. The uplink portion includes
 15 a control sub-frame and an event sub-frame. The downlink portion of the frame includes a synchronization slot in which the sink transmits a full-power pulse to all the nodes for synchronizing the clocks. The two following downlink slots of the downlink portion and the reconfiguration portion of the uplink control sub-frame are used for hop-distance discovery. The reconfiguration area in the control sub-
 20 frame has $(H+1)$ slots, where H is the maximum hop-distance. The forwarding flag area is designed for routing pulses toward the sink. The H -slot-wide routing area of the control sub-frame is used for energy management.

[0012] The event sub-frame contains H slot clusters, with each cluster containing $360/\alpha$ slots, where α is the sector width. Each of the slots therefore
 25 corresponds to a different specific [sector-id, hop-distance] tuple. An event-originating node transmits a pulse during the dedicated event sub-frame slot that corresponds to the [sector-id, hop-distance] of the node's event area.

Pulse Forwarding

[0013] While routing the pulse toward the sink, at each intermediate node, the
 30 pulse is transmitted in the same event sub-frame slot that corresponds to the [sector-id, hop-distance] of its event area of origin. In other words, while being

forwarded, the transmission slot for the pulse at all intermediate nodes does not change with respect to the frame. This is how information about the location of origin of an event is preserved during routing. Upon reception, the sink can infer the event area of origin from the [sector-id, hop-distance] value corresponding to the slot during which the pulse is received.

[0014] A hop-distance discovery process may be periodically executed by the network for each node to discover its own hop-distance from the sink node.

When a pulse is transmitted by a node at hop-distance h , only neighboring nodes at hop area $(h-1)$ forward the pulse toward the sink. In other words, the nodes at hop area h and $(h+1)$ should ignore the pulse. This logic ensures that a pulse is eventually delivered to the sink.

[0015] While transmitting a pulse in the event sub-frame, the sensor also transmits a pulse in the corresponding slot of the forwarding flag area of the control sub-frame. That is, while forwarding a pulse from a hop area h node, a pulse is also transmitted in the h^{th} time slot of the forwarding flag area. By looking at the received pulse in the forwarding area, all the receivers of the pulse can decide if it should be discarded or forwarded toward the sink. This can protect against detection of a false pulse and also ensure that a pulse from hop-area h is forwarded only by nodes in hop-area $(h-1)$.

[0016] For additional information regarding hop-angular routing, see Q. Huo and S. Biswas, "A Novel Concept of UWB Pulse Switching in Sensor Networks," ICWMC 2012 and Q. Huo, A. Plummer, and S. Biswas, "Pulse Switching for Static Event Sensing in Sensors Networks," IEEE Globecom, December 2011, the entire disclosures of which are hereby incorporated by reference into this application.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] The present disclosure will become more fully understood from the detailed description and the accompanying drawings, wherein:

[0018] FIG. 1 is a graphical example of hop-angular network layout according to the prior art;

- [0019]** FIG. 2 is a graphical depiction of a frame definition for use with hop-angular routing according to the prior art;
- [0020]** FIG. 3 is a cross-sectional view of a mechanism for coupling an ultrasonic transducer to a substrate;
- 5 **[0021]** FIG. 4 is a plot of gain versus ultrasonic pulse frequency;
- [0022]** FIG. 5A is a plot of received amplitude versus ultrasonic pulse width;
- [0023]** FIG. 5B is a plot of communication range versus ultrasonic pulse width;
- [0024]** FIG. 6 is a graphical representation of using different transmission mode velocities to estimate distance;
- 10 **[0025]** FIG. 7 is an example trace of transmitted signal strength and received signal strength on a common timescale;
- [0026]** FIG. 8 is an example trace of multipath interference;
- [0027]** FIG. 9 is an example trace, on a common axis, of transmitted signal strength and received signal strength;
- 15 **[0028]** FIG. 10 is an example trace, on a common axis, of received signal strength and a processed signal resulting from application of an envelope function;
- [0029]** FIG. 11 is an example pair of traces comparing a received signal strength from a single transmission to a received signal strength from two
- 20 transmissions;
- [0030]** FIG. 12A is a block diagram of an example transceiver circuit;
- [0031]** FIG. 12B is a block diagram of another example transceiver circuit adapted for sink or node use;
- [0032]** FIG. 12C is a block diagram of an example transceiver circuit adapted
- 25 for node use;
- [0033]** FIGs. 13A-13C are probability density functions for distance estimation performed according to the principles of the present disclosure;

- [0034]** FIG. 14 is a graphical depiction of network cells and approximate transmission regions;
- [0035]** FIGs. 15A-15D are a graphical depiction of a frame definition for cell-based routing;
- 5 **[0036]** FIGs. 16A-16D illustrate route discovery using the frame definition of FIGs. 15A-15D;
- [0037]** FIGs. 17A-17B illustrate event forwarding using the frame definition of FIGs. 15A-15D;
- [0038]** FIG. 18 is a block diagram of an example control circuit;
- 10 **[0039]** FIG. 19 is a graphical example of a cellular sensor network layout according to the principles of the present disclosure;
- [0040]** FIG. 20A is an example method of route discovery using the frame definition of FIGs. 15A-15D;
- [0041]** FIG. 20B is an example method of pulse generation and forwarding
15 using the frame definition of FIGs. 15A-15D;
- [0042]** FIGs. 21A-21C are a graphical depiction of a frame definition for cell-based routing including false pulse and lost pulse protection features;
- [0043]** FIG. 22 is a graphical illustration of a routing loop problem and its resolution;
- 20 **[0044]** FIG. 23 is an example method of route discovery using the frame definition of FIGs. 21A-21C;
- [0045]** FIGs. 24A-24B are an example method of pulse generation and forwarding using the frame definition of FIGs. 21A-21C;
- [0046]** FIG. 25 is a graphical example of a cellular sensor network layout
25 including both sensors and actuators;
- [0047]** FIGs. 26A-26D are a graphical depiction of definitions for a pair of frames that allow direct point-to-point addressing;
- [0048]** FIGs. 27A-27B are graphical examples of transmission collisions;

[0049] FIG. 28A is an example method of route discovery using the frame definition of FIGs. 26A-26D;

[0050] FIG. 28B is an example method of pulse generation and forwarding using the frame definition of FIGs. 26A-26D;

- 5 **[0051]** In the drawings, reference numbers may be reused to identify similar and/or identical elements.

DETAILED DESCRIPTION

Introduction

[0052] This disclosure introduces energy efficient systems and methods for
10 communication in wireless sensor networks. Contrary to standard packet-based data exchange, this disclosure uses pulse switching, in which information is conveyed by either transmission or absence of transmission of a single pulse. This type of communication is directly applicable to event detection and localization, where a binary signal indicating whether an event has occurred is
15 sufficient.

[0053] Traditional event detection sensor networks use multi-bit packets in order to exchange data with each other. By using a single pulse, the overhead per data exchange can be significantly reduced. Multiple nodes, also referred to as sensors, can be distributed about an area to be monitored and each is
20 configured to transmit a pulse when occurrence of an event is sensed. The sensors may communicate with each other and/or may relay information to a controller, also called a sink. The sink may process data and make decisions, such as whether to send commands to an actuator and/or whether to issue an alert based on detected events. In addition, the sink may provide data to a
25 separate processing system. Although the present disclosure discusses a single sink, multiple sinks can be supported using, for example, multi-band operation and syntax extensions.

[0054] For larger areas to be monitored, some of the sensors may not be able to directly communicate with the sink, and therefore multi-hop capability is
30 implemented so that sensors can relay pulses from other sensors toward the

sink. As pulses are forwarded toward the sink, information about their origin is maintained so that the sink can localize the event (i.e., determine where the event occurred).

[0055] Instead of traditional wireless sensor networks, which use radio
5 frequency wireless transmission, sensors mounted to a structure (such as the wing of an aircraft) can use the structure itself to transfer signals. For example, ultrasonic waves can be injected onto the structure by one sensor and detected by the sink or intermediary sensors. In traditional wireless sensor networks, the radio interface often requires the most energy at any given node. This means
10 that reductions in energy consumption for communication can have a significant effect on total energy consumption in the network.

Applications

[0056] The present disclosure is well suited for small networks (e.g., with less than a few dozen event areas) distributed within a restricted geographical area.
15 The energy benefits of the architecture may depend on a relatively low event frequency, in which event occurrence is the exception rather than the rule. These qualities are compatible with non-invasive, persistent structural health monitoring, such as for detecting stresses and cracks in a structure, such as a bridge, aircraft wing, or nuclear power plant. Wireless sensor networks are also
20 used in, for example, the medical, environmental, agricultural, civil engineering, aviation, and industrial fields.

[0057] In addition to detecting events to allow for preventative or mitigating measures to be taken, data may be gathered that would be of interest to the builders and manufacturers of structures being monitored. Since the network
25 performs continuous event monitoring, the sink can maintain a database that allows monitoring the performance of certain structures (e.g., aluminum, composite, or concrete) under real conditions to determine failure modes and estimate usable lifetime.

[0058] In addition to sensors, actuators may be present in the wireless sensor
30 network. Actuators may receive instructions from the sink and/or make decisions based on information received from sensors. For example, sensors in an office

may monitor the temperature, and transmit event pulses in response to the temperature being above a predefined comfortable upper temperature. Actuators may increase an air-conditioning output to address the over-temperature condition. Another example wireless sensor network is a fire extinguishing

5 system, in which sensors detect the presence of a fire based on smoke and/or abnormally high temperature. Actuators may activate water sprinklers in response to the detection of fire. See also J. Rao and S. Biswas, "Analyzing Multi-hop Routing Feasibility for Sensor Data Harvesting Using Mobile Sinks," JPDC Volume 72 Issue 6, June, 2012, J. Rao and S. Biswas, "Network Assisted

10 Sink Navigation for Distributed Data Gathering: Stability and Delay-Energy Tradeoffs," Journal Computer Communications Volume 33 Issue 2, February 2010, and J. Rao and S. Biswas, "Stability Analysis of Multi-hop Routing in Sensor Networks with Mobile Sinks" ICDCN 2010, the entire disclosures of which are hereby incorporated by reference into this application.

15

PHYSICAL LAYER

[0059] Pulses can be transmitted from one sensor to another using various technologies, including ultra-wideband impulse radio and/or ultrasonic transducers. For more information on ultra-wideband impulse radio, see Q. Huo, J. Rao, and S. Biswas, "Pulse Switching: Toward a Packet-Less Protocol

20 Paradigm for Event Sensing," IEEE Transactions on Mobile Computing, vol. 12, no. 1, pp. 35-50, January 2013, and Q. Huo, A. Plummer, and S. Biswas, "Ultra Wide Band Impulse Switching Protocols for Event and Target Tracking Applications," IEEE SECON 2011, Q. Huo, B. Dong, and S. Biswas, "Cellular Pulse Switching: An Architecture for Event Sensing and Localization in Sensor

25 Networks," ICDCN 2013, and Q. Huo, B. Dong, and S. Biswas, "A Cellular Pulse Switching Architecture for Binary Event Sensing," IEEE Globecom 2012 the entire disclosures of which are hereby incorporated by reference into this application.

[0060] An ultrasonic transducer, such as a piezoelectric wafer active sensor (PWAS), can be used to create an ultrasonic wave on a medium. For example

30 only, the PWAS may be obtained from APC International, such as a piezoelectric

disk having an outer diameter of 9.55 mm, a thickness of 1.00 mm, and a wrap-around (also known as with feedback, WFB) using APC material #850, which has a relative dielectric constant of 1900. For testing and evaluation purposes, the PWAS may be controlled by a wireless sensor platform, such as the Mica2
5 mesh network radio module from MEMSIC Inc.

[0061] When performing testing and prototyping, moving the piezoelectric element may be necessary. A mechanism for attaching the piezoelectric element to a substrate so that the piezoelectric element is strongly coupled to the substrate and yet can be relocated easily is desirable. A mechanism that, in
10 some implementations, has both of these characteristics is shown in FIG. 3. A piezoelectric element 100 is coupled to a substrate 104 by the force of opposing magnets 108 and 112. An epoxy 116 secures the piezoelectric element 100 to the magnet 108.

Frequency Response

[0062] To maximize efficiency, a frequency (the resonance frequency) at which
15 a particular PWAS transfers a maximum amount of energy to the substrate may be determined. To determine the resonance frequency, a constant input voltage can be provided to the PWAS while varying the frequency. A receiver, such as a second PWAS, measures the received signal strength, and the resonance
20 frequency can be chosen based on the frequency at which the received signal strength is the greatest.

[0063] For example only, this approach was used with a 2024 alloy aluminum sheet 3.6 m long, 1.2 m wide, and 1 mm thick. In FIG. 4, the measured gain is plotted against frequency, where gain is the ratio of the received power to the
25 transmitted power. From FIG. 4, it can be seen that on this substrate, the PWAS was most sensitive at roughly 245 kHz. Therefore, for the following examples, 245 kHz was used to excite the PWAS.

Pulse Width

[0064] Another metric to determine for transmitting ultrasonic pulses is the
30 pulse width for a transmission. If the chosen pulse width is too small, the communication range may be insufficient, which may require more nodes to

span an object of interest. If the chosen pulse width is too large, the extra energy may not lead to any greater communication capability, and instead is simply wasted energy. One approach to identify the desired pulse width is to vary the pulse width of a transmitted signal, and monitor the received signal at a receiver,
 5 such as by using an oscilloscope.

[0065] FIG. 5A shows traces of received signals at a distance of 1 m for varying pulse widths: specifically, 35 μ s, 70 μ s, 140 μ s, and 210 μ s. From FIG. 5A, it can be seen that the amplitude of the received signal is greater for the 70 μ s pulse width than for the 35 μ s pulse width. However, the amplitude of the received
 10 signals corresponding to 140 μ s pulse width and the 210 μ s pulse width are approximately the same as the 70 μ s pulse width. This implies that the receiver PWAS is maximally excited with a pulse width of approximately 70 μ s. Larger values of the pulse width do not result in larger signal amplitude at the receiver, since maximum excitation has already been achieved. For the further examples
 15 below, therefore, a pulse width of approximately 70 μ s will be used.

[0066] Note that longer distances may require a wider pulse width, and shorter distances may permit narrower pulse widths. See FIG. 5B, which plots maximum communication range versus pulse width based on experimental data. In various implementations, the pulse width may be variable so that sensors can be
 20 installed in various applications with different inter-sensor distances. The sensors can then be instructed to use the correct pulse width according to the given installation.

Propagation Delay

[0067] The propagation of mechanical waves in thin metallic substrates can be
 25 described by Lamb waves. Lamb waves propagate in two different forms: symmetric (longitudinal) waves and asymmetric (transverse) waves. The symmetric waves propagate faster and show lower dispersion. The asymmetric waves, meanwhile, propagate slower and have more dispersion. The different symmetric (S) modes are labeled S_0 , S_1 , ..., and S_N . The asymmetric (A) modes
 30 are labeled A_0 , A_1 , ..., and A_N .

[0068] The number of possible modes depends on the thickness of the medium as well as the signal's frequency. In 2024 alloy aluminum, the theoretical velocity of the symmetric wave is roughly 5400 m/s, while the asymmetric wave propagates at approximately 3100 m/s. In the example system described here, where a 1 mm thick 2024 aluminum alloy has been used and the signal frequency is set to 245 kHz, two wave modes (S_0 and A_0) dominate.

[0069] For various purposes, including synchronizing timing, determining the propagation delay of a wave from transmitter to receiver is valuable. Because the symmetric wave has a higher velocity than the asymmetric wave, the symmetric wave arrives earlier. Therefore, separation in time between the received symmetric and asymmetric waves is an indication of distance between the transmitter and receiver, and therefore propagation delay.

[0070] FIG. 6 shows a graphical depiction of propagation of the S_0 and A_0 waves. Because the velocity of the S_0 wave (v_S) is greater than the velocity of the A_0 wave (v_A), the A_0 wave arrives at the transmitter after a delay of Δt compared with the S_0 wave. The velocities v_S and v_A may be predetermined according to the transmission frequency and the material characteristics.

[0071] If t_S is the arrival time at the receiver of the S_0 wave and t_A is the arrival time of the A_0 wave, the time difference can be defined as $\Delta t = t_A - t_S$. The distance d between the transmitter and receiver can be used to express the time difference as $\Delta t = \frac{d}{v_A} - \frac{d}{v_S}$. The distance d can be solved for, resulting in:

$d = \frac{\Delta t}{v_A^{-1} - v_S^{-1}}$. Then, if the propagation time is defined as the time from the beginning of a transmission until a wave (the S_0 wave) first reaches the receiver, the propagation delay is $t_{prop} = \frac{d}{v_S}$.

[0072] FIG. 7 shows a voltage trace of example S_0 and A_0 waves being received, with an indication of how Δt can be defined. If the distance between transmitter and receiver is too small, the S_0 and A_0 waves may overlap to the extent that it is not possible to differentiate between the waves. Therefore a propagation delay may not be accurately calculated for distances below a certain threshold. For example only, using the parameters above, the threshold is

approximately 0.5 m. To allow for accurate propagation delay estimation, the minimum distance between sensors may be set at or above this threshold.

Delay Spread

[0073] The necessary size of a slot, during which a pulse can be transmitted, is
5 determined by the accuracy of the time at which the pulse will be transmitted,
and the amount of delay in receiving the pulse. In FIG. 8, a plot of received
signal strength over time is shown for a received pulse. After performing a
number of measurements using the above example substrate and pulse
characteristics, it has been determined that the delay spread was below 10 ms.
10 Therefore, the slot size should be at least 10 ms, and in the examples below it is
15 ms.

Tx/Rx and envelope

[0074] In FIG. 9, a transmitted signal is shown on a single time axis along with
the received signal at a distance of 1 m. As seen in FIG. 9, the received signal
15 has one section with a large amplitude, and also has quite a few sections where
the signal amplitude is higher than the noise floor. To convert this analog-domain
information into a form that can be processed, an envelope function can be
applied to the received signal.

[0075] In FIG. 10, operation of an envelope function on an example received
20 signal is shown. The envelope function produces an irregular square wave from
the received signal. The threshold used for envelope detection may be adjusted
to achieve a balance between low false pulse rate and low pulse loss rate.

[0076] In various implementations, only the timing of the first rising edge of the
envelope output is used to determine timing of the received pulse. For example,
25 the rising edge may be used to trigger an interrupt in processing circuitry. Similar
processing may be performed on each of the S_0 and A_0 pulses.

Pulse stacking

[0077] As discussed in more detail below, two different sensors may both
transmit a pulse during a slot corresponding to a certain event area. This may
30 occur, for example, when two sensors in the event area are both indicating
detection of an event, or when two sensors outside of the event area are both

retransmitting a pulse for the same event area. In traditional wireless networking, simultaneous transmission is called a collision, and generally results in a failed transmission.

5 [0078] According to the present disclosure, however, the information conveyed in a transmission is simply a single pulse, and therefore transmitting multiple copies of the same pulse simultaneously results in a single received pulse, often one whose amplitude is magnified. For example only, see FIG. 11. The first plot shows transmitted signal strength of a pulse, which at a distance of 0.6 m results in the received signal strength shown in the second plot. By contrast, when two 10 transmitters both transmitted the signal shown in the first plot, the result was as shown in the third plot. Note that the signal in the third plot has a greater amplitude, which is why transmitting multiple simultaneous pulses may be referred to as pulse stacking.

Transceiver

15 [0079] FIG. 12A is a functional block diagram of an example transceiver circuit 200. The transceiver circuit 200 receives signals coupled to a piezoelectric wafer active sensor (PWAS) 204 from a substrate to which the PWAS 204 is attached. The transceiver circuit 200 also transmits signals to the PWAS 204 for coupling onto the substrate.

20 [0080] The transceiver circuit 200 includes a transmit/receive control circuit 208 that interfaces with the PWAS 204. A mode of the transmit/receive control circuit 208 is controlled by a microcontroller circuit 212. In a receive mode, the transmit/receive control circuit 208 provides signals to an amplifier circuit 216, while in a transmit mode, the transmit/receive control circuit 208 sends signals 25 from a voltage switch circuit 220 to the PWAS 204.

[0081] The amplifier circuit 216 amplifies a received signal and provides an amplified signal to a pulse detector circuit 224. The pulse detector circuit 224 determines start times of pulses, and may apply an envelope filter to the amplified signal. Information related to existence and timing of pulses is provided 30 to the microcontroller circuit 212. In various implementations, the pulse detector

circuit 224 may simply drive a single wire to an active state to trigger an interrupt in the microcontroller circuit 212.

[0082] The microcontroller circuit 212 instructs the voltage switch circuit 220 as to which voltage to use. For example, a high voltage may be used for global
5 synchronization pulses that should be received throughout the entire substrate, while a lower voltage is used for local communication pulses. A pulse generator circuit 228 provides pulse waveforms to the voltage switch circuit 220. Operation of the pulse generator circuit 228 and the voltage switch circuit 220 are controlled by the microcontroller circuit 212. For example, the microcontroller
10 circuit 212 may instruct the pulse generator circuit 228 when to generate a pulse, and may instruct the pulse generator circuit 228 on how wide to make the pulse. The microcontroller circuit 212 may command the voltage switch circuit 220 based on whether the pulse will be for local or global transmission.

[0083] A supervisor circuit 232 may control the sensor node at a high level,
15 receiving sensor inputs from a sensor 236 and communicating to and from the sink using the transceiver circuit 200. For example only, the sensor 236 may sense temperature, pressure, acceleration, humidity, and/or force. A power supply 240 provides power to the transceiver circuit 200, the supervisor circuit 232, and the sensor 236. In various implementations, the responsibilities of the
20 supervisor circuit 232 may be assumed by the microcontroller circuit 212, similar to what is shown in FIG. 12B.

[0084] In FIG. 12B, another example transceiver circuit 300 is shown. A transducer 304, such as a piezoelectric wafer active sensor, is interfaced with the transceiver circuit 300. A transmit/receive (Tx/Rx) multiplexer 308 determines
25 whether a signal from a voltage multiplexer 312 is output to the transducer 304 or whether a signal from the transducer 304 is provided to an input multiplexer 316.

[0085] The voltage multiplexer 312 receives a pulse from a pulse multiplexer 320. The voltage multiplexer 312 also receives a higher voltage version of the
30 pulse from a voltage pump circuit 324. The voltage multiplexer 312 outputs either the pulse or the higher voltage pulse depending on a signal from a control circuit 328.

[0086] A clock circuit 332 generates a sinusoidal clock signal at a high frequency. For example only, the clock circuit 332 may include a crystal oscillator, which may or may not be temperature-compensated. The clock from the clock circuit 332 is provided to a divider circuit 336, which reduces a
5 frequency of the clock signal to a frequency that efficiently creates an ultrasonic wave in the substrate to which the transducer 304 is attached.

[0087] The divided clock signal from the divider circuit 336 is a continuously repeating periodic square wave. In order to turn that square wave into discrete pulses, the pulse multiplexer 320 alternately outputs a burst of the divided clock
10 signal in between outputting a zero voltage. Operation of the pulse multiplexer 320 is controlled by a timer circuit 340. The timer circuit 340 causes the pulse multiplexer 320 to switch from outputting zero to outputting the divided clock signal for a predetermined period of time. That predetermined period of time establishes the pulse width. The beginning of that predetermined period of time
15 may be initiated by a signal from the control circuit 328.

[0088] When the control circuit 328 instructs the Tx/Rx multiplexer 308 to provide signals from the transducer 304 to the input multiplexer 316, the control circuit 328 also instructs the input multiplexer 316 to connect the Tx/Rx
20 multiplexer 308 to an amplifier circuit 344. Otherwise, the input multiplexer 316 connects the Tx/Rx multiplexer 308 to ground. This reduces crosstalk between the inputs to the Tx/Rx multiplexer 308.

[0089] In various implementations, signal conditioning may be performed on the received signal prior to amplification by the amplifier circuit 344. For example only, low-pass and/or high-pass filtering may be performed. The amplifier circuit
25 344 generates an amplified signal, which is further amplified by an amplifier circuit 348. This twice-amplified signal is processed by an envelope circuit 352 and the processed result is fed to a comparator circuit 356. The comparator circuit 356 outputs a compare result to the control circuit 328. For example, the comparator circuit 356 may output an active signal to the control circuit 328
30 when the output of the envelope circuit 352 exceeds a threshold.

[0090] In various implementations, the envelope circuit 352 may be implemented as a germanium diode. Although two amplifier circuits 344 and 348

are shown, more or fewer amplifiers may be used. Multiple amplifiers may be used to improve various parameters, such as frequency response, stability, output swing, etc. when compared to a single amplifier. In various implementations, the amplifier circuits 344 and 348 may each be non-inverting amplifiers, inverting amplifiers, or a combination of the two.

[0091] A sensor 360 provides data to the control circuit 328. The sensor 360 may detect an event or provide information to the control circuit 328 from which the control circuit 328 can determine an event has occurred. A power supply 364 powers the transceiver circuit 300 and the sensor 360.

[0092] In one example implementation, the clock circuit 332 is a 7.86432 MHz crystal oscillator with a sinusoidal output. The divider circuit 336 reduces the frequency to 245 kHz and creates a square wave with a 50% duty cycle at that frequency. The divider circuit 336 may include a phase-locked loop (PLL). In order to create a 70 μ s pulse, the timer circuit 340 outputs an active signal to the pulse multiplexer 320 for a duration of 70 μ s. The timer circuit 340 may output this active signal once every 30 ms or on command from the control circuit 328.

[0093] In various implementations, the voltage pump circuit 324 may double voltage of the pulse. To allow for standard AA batteries to be used, the base voltage may be a multiple of 1.5 V, such as 3 V. The voltage pump circuit 324 may boost that to 6 V. In the example used in this disclosure, a 6 V pulse can be detected at one corner when transmitted from an opposite corner. A 6 V pulse is therefore suitable as the global synchronization pulse.

[0094] The timer circuit 340 may be implemented as 555 monostable integrated circuit. The control circuit 328 may be implemented as a Mica2 mesh network radio module. The control circuit 328 may be programmed with software that controls the select inputs of the Tx/Rx multiplexer 308, the voltage multiplexer 312, the input multiplexer 316, and the timer circuit 340. The output of the comparator circuit 356 may be attached to an interrupt input of the control circuit 328. The outputs from the control circuit 328 may be general purpose input/output (GPIO) pins programmable in software.

[0095] In FIG. 12C, a transceiver circuit 380 adapted for use by a node, but not by a sink, may omit the voltage pump circuit 324 and the voltage multiplexer 312 of FIG. 12B. This is because a node is not responsible for transmitting a global synchronization pulse. To save space and cost, these components may
5 therefore be omitted.

Experimental Results

[0096] To test the above transmission parameters and hardware, test pulses were sent at varying distances: specifically, 0.5 m, 0.75 m, 1.0 m, and 3.8 m. The maximum distance of 3.8 m was chosen based on the size of the piece of
10 aluminum used for testing. For each distance, a number of cycles were performed – in each cycle, a pulse was transmitted. A pulse loss rate is determined by dividing the number of pulses not received by the total number of pulses transmitted. Further, a false pulse is defined to be a pulse being detected even though no pulse was being transmitted at the time. A false pulse rate is
15 determined by dividing the number of false pulses received by the total number of actual pulses transmitted.

Distance	Pulse Loss Rate	False Pulse Rate	Pulse Voltage	# of cycles
0.5 m	< 2.16e-6	4.0e-6	3.0V	925,936
0.75 m	< 1.05e-5	3.37e-4	3.0V	189,880
1.0 m	3.32e-7	1e-6	3.0V	6,026,342
3.8 m	< 1.62e-6	9.4e-5	6.0V	617,190

[0097] FIG. 8 shows the results of distance estimation for measured distances between the transmitter and receiver of 0.5 m, 0.75 m, and 1.0 m. When the physical distance was set to be 0.5 m, the distance was estimated 462,967
20 times, and the resulting estimation was $0.5002 \text{ m} \pm 0.0015 \text{ m}$. This result shows rather high accuracy as well as a low variation in the determination of the distance. See FIG. 13A.

[0098] For a measured distance of 0.75 m, 94,940 pulses were used to estimate distance, with the resulting estimation being $0.7495 \text{ m} \pm 0.002 \text{ m}$. See
25 FIG. 13B. For a set distance of 1.0 m, 3,013,171 pulses were analyzed to estimate distance, with the resulting estimation being $1.09 \text{ m} \pm 0.007 \text{ m}$. Given

the precision of each of these estimations, the consistent offset of 9 cm for the 1.0 m distance may be attributable to setup inaccuracy, where the transmitter and receiver were placed slightly further than 1.0 m apart. See FIG. 13C.

[0099] Additional information regarding the ultrasonic physical layer can be found in S. Lorenz, B. Dong, Q. Huo, W.J. Tomlinson, and S. Biswas, "Pulse Based Sensor Networking Using Mechanical Waves Through Metal Substrates," Proceedings of SPIE, Volume 8753, id. 875306 (May 28, 2013), the entire disclosure of which is hereby incorporated by reference into this application.

CELL-BASED ROUTING

[0100] In contrast to the hop-angular routing shown in FIG. 1, cell-based routing is graphically depicted in FIG. 14. In FIG. 14, a sink and five example cells are shown. Although, as discussed in more detail below, the cells can be of any shape, hexagons are shown for simplicity of illustration. See also Q. Huo, B. Dong, and S. Biswas, "Pulse Switching: A Packet-less Networking Paradigm for Energy-constrained Monitoring Applications," in *Wireless Sensor Networks: Theory and Application*, CRC Press, Sep. 2013, the entire disclosure of which is hereby incorporated by reference into this application.

[0101] The sink and each of the cells is shown with a corresponding approximate circular transmission range. For example, the dashed circle surrounding cell 1 overlaps with cell 2, but does not reach cell 3 or the sink. As a result, cell 1 and cell 2 are considered neighbors, but cell 1 and cell 3 are not considered neighbors.

[0102] Each of these cells may include one or more sensors. In various implementations, the sensors are not individually addressable, and are identified only by cell number. The cell number may need to be preprogrammed into the node during installation and setup. In various implementations, nodes in each cell may also be preprogrammed with a list of IDs indicating neighbor cells. Because multiple nodes in a cell may each detect the same event, multiple nodes in a cell may transmit pulses at the same time. While this may be energy inefficient, the resulting pulse should still be intelligible to the receiver, and may in fact be a stronger signal.

[0103] Each cell may include one or more sensor nodes and/or one or more actuator nodes. However, because the nodes are not individually addressable, and are programmed to operate solely based on the cell in which they are located, this disclosure uses the terms node and cell interchangeably. In other words, when stating that a cell transmits a pulse, this means that one or more of the nodes within the cell transmit that pulse. Similarly, saying that a cell detects an event means that one or more of the nodes within the cell detected that event. Further, stating that a cell operates as a next hop in a forwarding route means that the nodes within that cell will forward on an event.

[0104] Returning to FIG. 14, multi-hop transmission is used to relay information about an event from a cell that is not a neighbor of the sink to the sink. For example, an event detected by cell 1 will first be transmitted to cell 2, which will retransmit the pulse to the sink. As described below, the sink and cells may engage in a self-discovery process on an intermittent basis to determine routes to the sink.

[0105] Cell-based routing does not require predefined shapes of cells as the hop-angular approach of FIG. 1 dictates. Further, the cells can be of any arbitrary size, where with hop-angular routing event areas farther away from the sink are larger in size. Further hop-angular routing assumes that routes to the sink will essentially follow a radial path toward the sink. This may not be possible due to communication issues and/or node failure. By contrast, cell-based routing allows a node to transmit a pulse toward the sink along any route no matter how circuitous as long as that route is shorter than a predetermined number of hops.

[0106] The predetermined number of hops may be imposed as a limitation to ensure that route discovery is completed within a deterministic period of time. Further, in implementations in which the entire route discovery process is completed within a single frame, the number of slots dedicated to the route discovery process is a fixed number and therefore limits the potential number of hops in a route.

Frame Definition

[0107] Communication in the pulse-based system of the present disclosure is synchronized to a repeating frame. All nodes in the network are synchronized to know when each frame begins. An example frame structure shown in FIGs. 15A-15D includes three areas: synchronization, route discovery, and pulse forwarding.

[0108] In the synchronization area, the sink injects a high-power synchronization pulse into the medium, which can be heard by all the nodes in the network. A node may use the arrival time of the synchronization pulse as the start of the frame, which is less accurate. Or a node may correct the arrival time based on propagation delay to achieve a more accurate time reference for the start of the frame. As discussed above, propagation delay can be estimated based on the characteristics of asymmetric versus symmetric modes of ultrasonic waves. If propagation delay is not taken into account, there will be inaccuracies in timing among the nodes and the sink, and as a result, larger slots for pulse transmission will be needed.

[0109] After the synchronization area is a delay of length β followed by the route discovery area. Route discovery allows each node to determine the best node to forward a pulse to so that the pulse will arrive at the sink. Following the route discovery area is another delay period of length β , and then the pulse forwarding area. During the pulse forwarding area, a node can transmit a pulse indicating that an event has been detected. In addition, the pulse forwarding area is used to forward pulses toward the sink from other nodes. A further delay frame of length β ends the frame and, in various implementations, the same frame pattern then immediately repeats. The total length of the frame is λ .

[0110] In FIG. 15B, an example of the synchronization area is shown. A synchronization pulse of length α is transmitted, and a time frame of length τ is allotted to allow the synchronization pulse to travel the length of the medium. In other words, τ may be based on the maximum delay from one end of the medium to the other. Note that FIGs. 15A-15D are not to scale. In fact, for example, in actual implementations, α may be significantly shorter than β .

[0111] In FIG. 15C, an example of the route discovery area is shown. A route discovery pulse is transmitted by the sink, and forwarded by each receiving node. As the pulse propagates outward from the sink, each node can determine a path back to the sink according to where the route discovery pulse was received from. In general, the first time a node receives a discovery pulse, whichever node had sent that discovery pulse represents the best path to get it back to the sink.

[0112] The network may be characterized as having up to a number H hops from the sink to the outermost node. The route discovery area therefore has H sections, one for each possible hop. In each section there are $N+1$ slots corresponding to the N nodes and 1 sink. In the first of the H sections, the sink sends the route discovery pulse. This first section corresponds to the first hop away from the sink. The route discovery pulse is placed in the slot corresponding to the transmitter, which in this case is the sink and, in the convention of this implementation, slot 0.

[0113] The nodes that receive this first hop discovery pulse from the sink will broadcast their own broadcast discovery pulse in the second section of the H sections. Each of these receiving nodes will transmit the route discovery pulse in the slot of the second section that corresponds to that node. Because the number of hops between the sink and any given node is not predefined, the route discovery area includes a slot for each node in every one of the H sections. One possible optimization, however, would be to make the first of the H sections a single slot long, where that single slot corresponds to the sink. Because the first route discovery pulse originates at the sink, the slots for nodes 1 – N in the first section may be eliminated.

[0114] In FIG. 15D, an example of the pulse forwarding area is shown. The pulse forwarding area includes N clusters of slots. Within each of the N clusters are $N+1$ slots, which correspond to the N nodes plus the 1 sink. A pulse located in cluster 1 indicates that an event has been detected by a node within cell 1. The slots within cluster 1 are used as a receiver (or, next hop) address. A pulse is transmitted within the slot corresponding to the node to which the pulse is

being immediately transmitted. On the pulse's final hop to the sink, the pulse will be located in slot 0 because the sink is the receiver.

[0115] All the way from the originating node to the sink, the pulse will be transmitted within the same cluster in each successive frame. As a result, when the sink receives a pulse within, for example, cluster N, the sink knows that an event was detected by node N. In other words, if node N detects an event, node N generates a pulse within cluster N of the immediately subsequent frame. Which slot in cluster N will be used is based on which node is the best next hop toward the sink from node N. When that next hop node receives the pulse in cluster N, the next hop node will generate a pulse in cluster N in the following frame. The slot is chosen to transmit to a node that is one step closer to the sink or is the sink itself.

[0116] Total frame size λ is $(\alpha + \tau) + \beta * H * (N + 1) + \beta * N * (N + 1) + 3 * \beta$. As an example only, pulse width α is 80 μ s, slot size β is 15 ms, and the maximum propagation delay τ is 1 ms. If the number of nodes N in the network is 9, and the maximum number of hops is 3, the resulting frame size λ is $.08 + 15 * (3 * 10 + 9 * 10 + 3) = 1845.08$ ms.

Route Discovery

[0117] An example of route discovery is shown in FIGs. 16A-16D. For simplicity, only 5 nodes are shown. Note that in this discussion, each cell may include multiple nodes. However, for architectural purposes, including pulse forwarding and route discovery, the nodes within a cell are interchangeable and therefore each cell is treated as a single node for ease of explanation. In practice, different nodes within a single cell may build up slightly different routing tables. This may be an advantage by leading to an increase in route diversity. Nevertheless, each of the nodes is programmed to perform the same operations.

[0118] In FIG. 16A, at time T_0 (i.e., a first section of the route discovery area of a first frame) the sink transmits a route discovery pulse. Because the route discovery pulse is transmitted by the sink, the pulse is located in the slot corresponding to the sink (slot 0). Nodes 2 and 3 are within the transmission range of the discovery pulse, and therefore each of nodes 2 and 3 update their

respective routing tables based on the reception of the route discovery pulse. The nodes recognize that the pulse came from the sink, and that the pulse was received at time T_0 .

[0119] In FIG. 16B, nodes 2 and 3, which have each just received a route discovery pulse, transmit their own route discovery pulses in the second section (time T_1) of the route discovery area of the first frame. Node 2 transmits the route discovery pulse in the slot corresponding to node 2, while node 3 transmits a route discovery pulse in the slot corresponding to node 3. The sink may note that nodes 2 and 3 are both neighbors of the sink although in various implementations this information is not used. Node 1 recognizes that node 2 is a neighbor, and that the discovery pulse was received from node 2 at time T_1 . Nodes 2 and 3 recognize that each other are neighbors with corresponding reception times of T_1 . Further, node 4 recognizes that node 3 is a neighbor, and that the route discovery pulse from node 3 was received at time T_1 .

[0120] In FIG. 16C, nodes 2 and 3 have already transmitted route discovery pulses and therefore do not transmit further route discovery pulses in the current frame. Node 1, meanwhile, transmits a route discovery pulse, which is received by node 2. Node 2 updates its routing table to indicate that node 1 is a neighbor, and that the route discovery pulse from node 1 was received at time T_2 . Similarly, nodes 3 and 5 note that node 4 is a neighbor with a corresponding time of T_2 .

[0121] In FIG. 16D, node 5 transmits a route discovery pulse, and the only node to receive the pulse, node 4, notes that node 5 is a neighbor with a reception time of T_3 . Each node now has one or more paths to follow toward the sink. The quickest, and therefore generally best, path to the sink is indicated by the neighbor having the lowest reception time. For example, node 2 has 3 neighbors, with the lowest reception time being the route discovery pulse from the sink. As would be expected, therefore, node 2 will transmit pulses to the sink directly. However, if communication from node 2 to the sink is severed and node 2 is able to detect that failure based on response mechanisms discussed below, node 2 can look at the routing table. Node 2 can then determine that node 3 is

the next best way to reach the sink because it has the second-lowest corresponding reception time.

Pulse Forwarding

- [0122]** In FIGs. 17A-17B, an example of pulse forwarding is presented. Assume
5 that an event is detected at node 1. Based on the routing table, node 1 recognizes that the pulse should be forwarded to node 2. Therefore, in the first cluster, which corresponds to the first node, of the pulse forwarding area, node 1 transmits a pulse. Specifically, node 1 transmits a pulse in slot 2 to indicate that the pulse is being transmitted to node 2.
- 10 **[0123]** In FIG. 17B, node 2 has received the pulse from node 1, and based on node 2's routing table, node 2 decides to forward the pulse directly to the sink. Because the pulse originated at node 1, node 2 will forward the pulse still within cluster 1. Note that this is cluster 1 of the following frame. Because node 2 is transmitting the pulse to the sink, the pulse is placed in slot 0 of the first cluster.
- 15 Because transmission is omnidirectional, node 1, node 3, and the sink all receive the pulse. However, because the pulse is in slot 0, nodes 1 and 3 recognize they are not the intended recipients and can ignore the pulse. Given that the pulse is located in cluster 1, the sink recognizes that the event was detected in node 1.

Block Diagram

- 20 **[0124]** FIG. 18 is an example functional block diagram of a control circuit 400. The control circuit 400 may be used for a sensor node, an actuator node, or a sink. An external communication circuit 404 allows the control circuit 400 to communicate with an external processing system. For example, when the control circuit 400 is used in a sink, the external communication circuit 404 may use a
25 wired or wireless local area network to send data to a monitoring computer. The monitoring computer may be connected to a wide area network to allow for remote monitoring.

- [0125]** In other implementations, the external communication circuit 404 itself includes a wide area network transceiver, such as a cellular radio modem. When
30 the control circuit 400 is used for sensor or actuation nodes, the external communication circuit 404 may be omitted to reduce costs and power

consumption. A coordination circuit 408 controls operation of the external communication circuit 404.

[0126] The coordination circuit 408 also receives signals related to received pulses. In various implementations, signals related to received pulses may be provided to the coordination circuit 408 as interrupts. The coordination circuit 408 may also output control signals to control various aspects of a transceiver, such as multiplexers and switches.

[0127] A frame synchronization circuit 412 receives data regarding the synchronization pulse or pulses, and synchronizes timing of the control circuit 400 with the beginning of the frame. The frame synchronization circuit 412 may also determine which frame definition is currently in use in implementations where different frames are used at different points in time. A pulse transmission control circuit 416 determines when within the frame various pulses should be transmitted for purposes such as route discovery and pulse forwarding.

[0128] A route discovery circuit 420 maintains a routing table 424 indicating neighbor nodes and next hops toward the sink and/or other nodes. The route discovery circuit 420 instructs the pulse transmission control circuit 416 to transmit discovery pulses as dictated by the route discovery process. A pulse forwarding circuit 428 transmits pulses toward specified destinations, either a sink or another node. The pulse forwarding circuit 428 determines next hop information from the routing table 424.

[0129] A queue circuit 432 stores event pulses protected by the node in which the control circuit 400 is located, as well as pulses that have been received from other nodes. The queue circuit 432 instructs the pulse forwarding circuit 428 to transmit the pulses to the next hop or hops. In various implementations, the queue circuit 432 operates in a first-in, first-out manner. In various implementations, the queue circuit 432 may operate in a first-in, first-out manner, but prioritize events for transmission that have occurred in the node that includes the control circuit 400.

[0130] A protection and response circuit 436 signals successful or unsuccessful reception of data, and may further indicate whether a collision has been

detected. The protection and response circuit 436 may also execute processes that protect against false positives, such as by calculating and verifying a false pulse protection code. The protection and response circuit 436 instructs the pulse transmission control circuit 416 regarding when to send pulses indicating successful or unsuccessful reception.

[0131] An energy management circuit 440 transitions some or all components the control circuit 400 and some or all components of a connected transceiver. The energy management circuit 440 may consider the current location within the frame, and whether any transmissions are scheduled or any reception is expected.

Irregular Cells

[0132] FIG. 19 is an example of a network layout in which the sensors are not uniformly distributed and the cells are of differing sizes and shapes. This emphasizes that there is no requirement that the cells be of equal size or shape or that they be oriented in a grid-like pattern, such as a honeycomb. Note that the sink is labeled as cell 0, while the other cells are labeled as 1 to 34.

[0133] For example only, a target path is shown along which a target travels past various ones of the sensors. For example only, the target may be a person and the sensors may include motion detectors. As the target passes sensors within a cell, those sensors will transmit pulses toward the sink indicating that an event has been sensed within that cell. FIG. 19 also demonstrates that each cell may include multiple sensors. The sensors operate independently but are programmed such that each of the sensors within a single cell operates on behalf of the cell. Therefore, the sensors may not be individually addressable and may not include any unique identifier other than which number cell they are installed into.

Methods

[0134] FIG. 20A shows an example method of performing route discovery using the frame definition of FIGs. 15A-15D. Route discovery begins during the route discovery portion of each frame at 504. At 504, control determines whether a pulse has been received in a cluster of slots of the route discovery area of the

frame. If so, control transfers to 508; otherwise, control transfers to 512. At 512, if the final cluster of slots of the route discovery portion has finished, control ends; otherwise, control returns to 504.

- [0135]** At 508, control determines whether the transmitter of the pulse is a predefined neighbor of the current node. If so, control continues at 516; otherwise, control ignores the pulse and transfers to 512. In various implementations, the nodes may not be programmed with neighbor information, and therefore 508 may be omitted, with control proceeding directly from 504 to 516.
- [0136]** At 516, control adds an entry to the routing table based on the node that transmitted the received pulse. The transmitter node is therefore a next hop toward the sink. The time at which the pulse was received is also added to the routing table. An earlier time for a routing table entry indicates that the corresponding next hop is a quicker path to the sink. Control then continues at 520.

- [0137]** At 520, if the node has transmitted a discovery pulse already, control transfers to 512; otherwise, control transfers to 524, where control determines whether the final cluster of the discovery phase has finished. If so, control ends. Otherwise, control transfers to 528 to schedule transmission of a discovery pulse in the next cluster of slots of the route discovery area of the frame. Control then returns to 504.

- [0138]** In FIG. 20B, example forwarding operation of a node using the frame definition of FIGs. 15A-15D is shown. Control begins with the pulse forwarding area of each frame at 604. Control determines whether an event has been detected by the node and, if so, control transfers to 608; otherwise, control transfers to 612.

- [0139]** At 608, control selects a best hop for transmission of the pulse toward the sink. This may be done by choosing the neighbor with the earliest associated time. Control continues at 616, where control schedules transmission of a pulse in the slot corresponding to the neighbor selected in 608 during the cluster of

slots corresponding to the present node within the pulse forwarding area of the following frame. Control then returns to 604.

[0140] At 612, control determines whether a pulse has been received. If so, control transfers to 620; otherwise, control transfers to 624. At 620, control
5 determines whether the slot in which the pulse was received corresponds to the present node. If so, control continues at 628; otherwise, control ignores the pulse and transfers to 624.

[0141] At 628, control notes the cluster in which the pulse was received, which indicates the original source node of the event resulting in that pulse. Control
10 continues at 632, where a next hop is chosen to forward the pulse toward the sink. Control continues at 636, where control schedules transmission of a pulse in the slot corresponding to the neighbor selected in 632 during the cluster of slots corresponding to the original source node determined in 628 of the following frame. Control then returns to 604. In 624, control determines whether
15 the end of the final cluster of slots of the pulse forwarding area of the frame has been reached. If so, control ends; otherwise, control returns to 604.

MODIFICATIONS

Frame Definition

[0142] FIG. 21A shows an example of a frame definition including modifications
20 made to the frame definition of FIGs. 15A-15D. The frame includes a downlink (away from the sink) area including a synchronization area shown in more detail in FIG. 21B and a discovery area shown in more detail in FIG. 21C. Note that the letter M is used to indicate the number of nodes in FIGs. 21A-21C.

[0143] The frame also includes an uplink (toward the sink) area including a
25 control area for energy management, a localization area for pulse forwarding, a protection area for avoiding false positives, and a flag area for avoiding routing loops. The frame also includes a response area to increase reliability by providing a mechanism for a node to determine whether a transmitted pulse has been received.

[0144] Within each slot, the pulse is transmitted at the midpoint of the slot. This is to account for the fact that relative clock drift between nodes may occur in either direction. With the pulse targeted to the middle of the slot, clock drift and propagation delay that is less than half the size of the slot will not cause the pulse to encroach into an adjacent slot.

[0145] In FIG. 21B, the synchronization area is shown as having a number (N_{sync}) of pulse slots. In a simple case, N_{sync} is equal to 1, meaning that a single synchronization pulse is defined for the frame. When N_{sync} is greater than 1, each of the pulse positions will either include a synchronization pulse transmitted by the sink or will not contain a pulse.

[0146] In normal operation, each of the pulse positions may be occupied by a pulse. The sink may need to communicate data to the entire network at once, in which case some of the pulses may be absent. This provides a binary codeword for the sink to use when communicating with the entire network. In various implementations, the first pulse and last pulse within the synchronization area may always be transmitted. Therefore, the pulses in the middle can be used to transmit binary information. In these implementations, the sink can transmit $2^{N_{\text{sync}} - 2}$ different messages per frame.

[0147] Although not shown in FIG. 21A, the binary codeword in the synchronization area may be used to indicate whether various portions of the frame will be omitted. For example only, one message transmitted in the synchronization area may instruct the nodes that discovery will be skipped for the current frame. The remaining areas of the frame may then shift earlier in time to occupy the nonexistent discovery area. In other implementations, the definition of the frame is invariable to prevent nodes from using an incorrect frame definition if they do not correctly receive the binary code in the synchronization area.

[0148] In FIG. 21C, the discovery area is seen to have a slot for a synchronization pulse from the sink as well as slots for each of the nodes. The sink can initiate route discovery by transmitting a pulse in the sink slot. This pulse is then received by neighboring nodes and in the following frame those

nodes will transmit their own pulses within the respective slots of the discovery area.

[0149] Note that between the sink slot and the node slots are delay intervals.

These intervals are chosen to allow a node's transceiver to transition from

5 transmitting a pulse to receiving a pulse. For example, assume that node 1 has received a discovery pulse in the previous frame and therefore schedules transmission of a pulse in slot 1 of the discovery area of the current frame. After transmitting the pulse in slot 1 of the discovery area, node 1 should immediately change to listening in case node 2 transmits a pulse in the second slot. This
10 delay interval is referred to as transceiver turnaround time and, if measured in units of slots, may be indicated as γ slots in length. Therefore, as indicated in FIG. 21A, the length of the discovery area is equal to one slot (for the sink) and M slots (one for each node) plus γ times M slots (the delay interval prior to each node slot) this can be written as $(\gamma+1)*M+1$.

15 **[0150]** The localization area is similar to the pulse forwarding section described above. Each node has a corresponding group of slots and an event that originated in a certain node is always transmitted within the group of slots corresponding to that node regardless of which node is currently transmitting the pulse. The slots within each group of slots indicate the destination (next hop) of
20 the pulse. Therefore, there are M groups of slots corresponding to the M cells within which events may be detected. In each group of slots there are $M+1$ slots indicating the number of destinations (M nodes + 1 sink) for each pulse to be transmitted.

Route Diversity

25 **[0151]** As described above, each node determines a next hop toward the sink based on the routing table. For example, the neighbor whose route discovery pulse was received first is likely the best (shortest) route back to the sink. To improve transport reliability, a node may transmit a pulse to multiple next hops. This is referred to as route diversity.

30 **[0152]** A variable δ defines the number of next hops to which a node will transmit a pulse. In the examples discussed above, δ is equal to 1, meaning that

each pulse is only transmitted to a single next hop. If δ is greater than 1, the pulse will be forwarded to the best δ next hops. Because pulses that reach the sink are all received by the sink within the same group of slots, the sink is able to determine where the event occurred regardless of the route the pulse took in reaching the sink. As δ is increased, reliability is increased incrementally while energy use also increases.

Energy Management

[0153] Each node may alternate between three states: transmission, listening, and sleeping. During route discovery and pulse forwarding, each node begins in the listening state. During route discovery, a node switches from listening to transmitting for the slot assigned to that node in the discovery area. The node then transitions back to the listening state for the remainder of the discovery area.

[0154] During pulse forwarding (the localization area of the frame), the node enters the transmission state during the group of slots corresponding to that node. The node enters the listening state for the other groups of slots of the localization area. If the node has not transmitted a pulse during the localization area, the node may enter the sleep state during the response area. If the node does not have a pulse scheduled for transmission, the node may enter the sleep state for a new group of slots corresponding to that node in the localization area.

[0155] The control area of the uplink section of the frame has been added to allow nodes to enter the sleep state during portions of the localization area. The control area has M slots, one for each of the nodes in which an event can be detected. When a node is planning to transmit a pulse in the localization area corresponding to node X, a pulse is first transmitted in slot X of the control area. This pulse indicates to receivers that a pulse will be transmitted in the Xth group of slots of the localization area, and therefore the receivers should remain in the listening state during that group of slots.

[0156] The control area is therefore a cumulative representation of all the groups of slots in which a transmission will be made. Therefore, any slots of the control area in which a pulse was not received indicates to the node that the

sleep state can be entered for the corresponding groups of slots. In addition to conserving energy by allowing intraframe sleeping, some false positive pulses will be filtered out. A false positive present in any of the groups of slots where a corresponding transmission is not announced in the control area will not be
5 received because the nodes will be in a sleep state during those groups of slots.

Compression

[0157] As a power consumption enhancement, nodes may be programmed to delay transmitting a pulse for a randomly selected number of frames. Meanwhile, if a node detects that another node within the same cell has already transmitted
10 a pulse indicating detection of an event, the detecting node recognizes that a pulse transmission is no longer necessary. This optimization will mean, depending on the number of nodes in a cell and the maximum delay interval, that most pulses detected by multiple nodes will only be transmitted by a single one of the nodes.

[0158] When a node detects multiple events in quick succession (e.g., detecting a moving target multiple times while it is within the detection range), it generates multiple pulses, each of which are independently forwarded to the sink. Temporal pulse compression can be used to remove this redundancy by ensuring that a node generates at most one pulse within a certain time duration,
20 which is referred to as a Deactivation Period. Temporal compression is also applicable while pulse forwarding so that a node forwards at most one pulse originated from the same event area within a Deactivation Period.

False Positives

[0159] For additional protection against false positive pulses, a protection area
25 R slots long is added to the frame. Each node that is transmitting a pulse in the localization area transmits one or more pulses in the protection area. The pulses in the protection area form an R-bit protection code. Because the transmissions of nodes within a transmission area overlap, the resulting pulses in the protection area are a logical OR of the R-bit protection codes generated by each node.

[0160] A receiver can calculate the R-bit protection codes that should have been transmitted based on pulses received in the localization area, and if a

logical OR of the expected protection codes does not match the sequence of pulses actually received in the protection area, an error is declared. The error means that one or more pulses were lost and/or one or more false pulses were detected. As a result, all pulses received in the localization area are ignored

5 based on the error.

[0161] The algorithm for determining the R-bit protection code can be as simple as, for each pulse to be transmitted in the localization area, taking the node number from which that pulse originated and performing a modulo R operation on that number. The resulting remainder from the modulo R operation

10 determines a single bit to be set in the R-bit protection code.

[0162] The R-bit protection code is transmitted using a pulse to indicate a set bit, and an absence of a pulse to indicate a cleared bit. When a node is transmitting only a single pulse in the localization area, the node will similarly transmit only a single pulse within the protection area. Protection against a false

15 positive pulses is increased as the ratio of R to number of nodes increases.

Acknowledgements

[0163] When a node transmits a pulse, there is no way for that node to know whether the pulse was received correctly. To provide this feedback, a response area has been added to the frame of FIG. 21A. The response area is divided into

20 a reception sub-area and a source sub-area. When a node correctly receives a pulse in a frame, the node will schedule two pulses for transmission in the upcoming response area.

[0164] The first pulse is in the reception sub-area, and indicates the number of the node that correctly received that pulse. The second pulse is transmitted in

25 the source sub-area, and corresponds to the originating node where an event was detected that lead to the received pulse. The reception sub-area includes a slot for each of the nodes that might receive a pulse, as well as the sink. The source sub-area includes M slots, one for each node that could have detected an event and originated a pulse. In various implementations, the number of slots in

30 the source area is increased to one more than M.

[0165] When a sender detects expected pulses in the reception sub-area and the source sub-area indicating that the desired recipient correctly received the pulse, the sender may assume that the pulse is on its way to the sink. However, if a response is not received, the sender will retransmit the pulse in the following
5 frame. After a certain number of retransmissions, the node may stop attempting to transmit the pulse in order to conserve energy. For example, the intended recipient node may have been damaged or depleted its power. Eventually another route discovery process will commence, and assuming there remains a route from the sink to the node, the node will be able to determine a new best
10 route back to the sink.

Prioritized Routing

[0166] Instead of waiting for another route discovery process to be completed, nodes may use a prioritized routing strategy, where the first attempt of a transmission is made to the best next hop. If the best next hop does not provide
15 a successful response, the node can next attempt a retransmission to the second-best next hop.

[0167] In a case where the diversity δ is greater than 1, this prioritized routing is performed in groups of δ next hops. In other words, the first attempted transmission is made to the δ best next hops. If no response is successfully
20 received, the following δ next hops will be used for the following transmission. Once all of the potential hops in the routing table are exhausted, the attempted transmissions will wrap back to the best next hop.

[0168] For example only, consider a node that has five potential next hops in its routing table, in increasing distance from the sink, and δ is three. The node will
25 first attempt transmission of a pulse to the first three next hops. If no response is successfully received, the node will then attempt retransmission using the fourth and fifth hops in the routing table as well as the first hop. If a response is still not successfully received, the node will attempt to retransmission to the second, third, and fourth hops.

Loop Avoidance

[0169] In FIG. 22, an example of a difficult routing situation is shown. At the moment in time depicted in FIG. 22, cells n_1 , n_2 , n_3 , n_4 , and n_5 are experiencing a fault, which may be temporary or permanent. Cell m is surrounded by cells n_1 - n_5 .

- 5 Cell m may have previously had routes to the sink through a number of cells, including cell n_1 and cell n_2 . Similarly, cell n_6 , which is not in a fault condition, may have a best route to the sink via cell m .

[0170] If cell m detects an event and needs to transmit a pulse to the sink, cell m transmits a pulse to the best next hop, which may be cell n_1 in this case.

- 10 However, because cell n_1 is experiencing a fault, no successful response is returned to cell m . Eventually, using prioritized routing, cell m will transmit the pulse to cell n_6 .

[0171] When cell n_6 forwards that pulse, cell n_6 may attempt to forward that pulse back to cell m , as cell m may be the best next hop toward the sink. Cell n_6 is unaware of the failure of cell n_1 . Cell m , however, will not provide a successful response to the pulse because cell m is in a transmit state when cell n_6 forwards the pulse. Further, cell m may be programmed not to respond to pulses resulting from events that originated in cell m .

- 20 **[0172]** Cell n_6 , therefore, uses prioritized routing and eventually attempts to transmit the pulse to cell k . Cell k forwards the pulse to the best next hop, which in the case may be back to cell n_6 . This creates a routing loop, preventing the pulse from reaching the sink. To resolve some of these loops, a flag area is added to the uplink area of the frame in FIG. 21A.

- 25 **[0173]** The flag area includes a single slot that is used as follows. In the routing table of a node (in this case, cell n_6), the node groups next hops into two groups, one that is closer to the sink than n_6 and one group that is further from the sink than n_6 . When cell n_6 is going to transmit a pulse in the localization area, cell n_6 determines whether the destination (in this case, cell k) is further from the sink or closer to the sink.

- 30 **[0174]** In FIG. 22, cell k was likely further from the sink according to the routing table. Because a pulse is being transmitted to a node further from the sink, a

pulse is also transmitted within the flag area slot. When the diversity number δ is greater than 1, if any of the destinations for the pulse are further from the sink than the transmitter, a pulse is transmitted during the flag area slot.

[0175] When cell k receives the pulse from cell n_6 , with the flag being set by the transmitted pulse in the flag area, cell k will examine previous pulse transmission records. Each node may maintain pulse transmission records from within a specified period of time, which may be defined in units of frames. When a node receives a flagged event pulse, the node determines whether that pulse had already been sent to any of the next hops. If so, the node forwards the pulse to a different next hop or next hops. Otherwise, the node operates in the default manner, which is to initially send the pulse to the best next hop.

[0176] In the example of FIG. 22, cell k determines whether a pulse originating in cell m has been transmitted within the predetermined period. In the case of a routing loop, the answer will be yes because the pulse will have been transmitted from cell k to cell n_6 . Because the answer is yes, cell k transmits the pulse to a further one of the next hops (in this case, cell p). This allows the pulse to escape from the routing loop, and in the example of FIG. 22, reach the sink along a new route.

[0177] When the route diversity number δ is greater than 1, a node determines whether a pulse has been sent to the first group of δ next hop cells, and if so, forwards the pulse to the following group of δ next hop cells. Eventually, the route discovery process should establish new routing table entries in which cell m recognizes the best next hop to the sink is through cell n_6 , cell n_6 recognizes that the best next hop is cell k, and cell k recognizes that the best next hop is cell p.

Methods

[0178] In FIG. 23, an example method of performing route discovery based on the frame definitions in FIGs. 21A-21C is shown. Control begins at 704, where control determines whether a pulse has been received. If so, control transfers to 708; otherwise, control remains at 704. At 708, control determines whether the transmitter of the discovery pulse is a predefined neighbor of the node. If so,

control transfers to 712; otherwise, control returns to 704. In various implementations, neighbors may not be predefined, and therefore 708 is omitted, and control proceeds directly from 704 to 712.

[0179] At 712, control adds an entry into a routing table based on the transmitter of a discovery pulse. The entry also includes a time stamp where earlier time stamps indicate more direct routes to the sink. Control continues at 716, where control determines whether a discovery pulse has been transmitted within a predetermined number of prior frames. For example only, the predetermined number may be named T_{DISC} , which is equal to the number of frames between the sink's initiation of each discovery process.

[0180] At 716, if the node has already transmitted a discovery pulse within this window, control returns to 704; otherwise, control transfers to 720 to schedule transmission of the discovery pulse. The discovery pulse will be scheduled for transmission in the discovery area of the next frame. Control then returns to 704.

[0181] FIGs. 24A-24B together describe an example method of event reporting and forwarding according to the frame definitions of FIGs. 21A-21C. The method is described for a node labeled as node X. Control begins at 804, where a transmission queue is cleared. The queue is an array of tuples, with each tuple including a source node and a retransmission count.

[0182] Control continues at 808, where if an event is detected, control transfers to 812; otherwise, control transfers to 816. At 812, the event that was detected at node X is added to the queue. The source of the event is therefore X and the retransmission count is initially zero. Control then continues at 816. At 816, control determines whether there are any events in the queue. If so, control transfers to 820; otherwise, control transfers to 824 in FIG. 24B via the circled letter A.

[0183] At 820, the next event in the queue to be transmitted has a source stored in variable S. Therefore, at 820, control schedules transmission in slot S of the control area. This indicates to other nodes that they should be in a listening state during that cluster of slots of the localization area to receive the forwarded pulse.

[0184] Control continues at 828, where control determines whether the event that is about to be transmitted is within a recent history data structure or if the event as received was flagged. If so, control transfers to 832; otherwise, control transfers to 836. A flag indicates that a routing loop may be present, and
5 therefore, instead of a standard next hop, an alternative next hop should be used to retransmit the pulse.

[0185] This may be accomplished by, for example, increasing the retransmission count. For example, the retransmission count may cause the node to iterate through a list of next hops. With a retransmission count of zero, a
10 pulse may be transmitted to the best next hop. When the retransmission count is one, the pulse will be retransmitted to the second-best next hop. Therefore, if the history indicates that a pulse relating to the event was successfully transmitted to the best next hop as well as to the second-best next hop, control may set the retransmission count to two, which will cause the third-best next hop to be tried.

[0186] Even if the event transmission is not stored in recent history, if the event was received with a flag, control may set the retransmission count to one to attempt to avoid creating a routing loop. In various implementations, adjusting the retransmission count may be performed only when both the event is present in recent history and the event had been flagged. In any case, control then
15 continues at 836.

[0187] At 836, control determines the next hop, labeled as H, based on the retransmission count. Although a single next hop is shown in the flowcharts, when the route diversity parameter δ is greater than 1, a pulse may be transmitted to multiple next hops at once. For example only, with a
25 retransmission count of zero and a δ of two, the pulse may be transmitted to the best two next hop entries in the routing table. With a retransmission count of one, the pulse may be transmitted to the third and fourth entries in the routing table.

[0188] Control continues at 840, where control schedules transmission of a
30 pulse in slot H of cluster S of the localization area. Control continues at 844, where control calculates a protection code and schedules transmission of that protection code following the localization area.

[0189] Control continues at 848, where control determines whether next hop H is located further away from the sink than is node X. For example only, this may be determined if the earliest entry in the routing table, with a next hop of H, is later than the best next hop in the routing table. If H is in a direction that is away
5 from the sink, control transfers to 852 to schedule transmission of a pulse in the flag area of the frame. In either case, control continues at 856.

[0190] At 856, if an indication of successful response has been received that matches the transmitted pulse, control transfers to 860; otherwise, control transfers to 864. At 860, control removes the event from the queue and
10 continues at 868, where the successful transmission of the event, along with the retransmission count at which success was achieved, is added to the history structure. Control then continues at 824 in FIG. 24B. At 864, control increments the retransmission count for the event, and control continues at 824 in FIG. 24B.

[0191] At 824, control determines whether one or more pulses were detected in
15 the control area of the frame. If so, control transfers to 872; otherwise, control transfers to 876, where the transceiver can sleep during the upcoming localization area because no reception is expected. If a transmission is scheduled during one of the clusters, that control area will have a pulse and the transceiver will be in transmit mode during those clusters of the localization area.
20 Control then continues at 808 in FIG. 24A via the circled letter B.

[0192] At 872, control sleeps during clusters in the localization area that were not indicated by corresponding pulses in the control area. Control continues at 880, where if one or more events are received in the localization area, control transfers to 884; otherwise, control continues at 808 in FIG. 24A.

[0193] At 884, control calculates protection codes for each of the events that were received, and creates a combined calculated protection code based on the logical OR of the individual protection codes. Control continues at 888. If the calculated protection code matches the received protection code, control continues at 892; otherwise, control returns to 808 of FIG. 24A.

[0194] At 892, control selects a first event of the received events, and continues
30 at 896. At 896, control determines whether the next hop specified in the selected

event is equal to X. If not, the event is discarded and control continues at 900; otherwise, control transfers to 904. At 904, control determines whether the source of the event is equal to X. If so, the pulse has returned in a loop back to node X, and is therefore discarded as control transfers to 900. If the source of the event is other than node X, control transfers to 908.

[0195] At 908, control sends an indication of a successful response for the selected event. Control continues at 912, where the received event is added to the queue for forwarding. Control continues at 916, where if a flag accompanied the event, control transfers to 920; otherwise, control returns to 808 in FIG. 24A. At 920, control adds a flag attribute to the event in the queue, and returns to 808 in FIG. 24A.

[0196] At 900, control determines whether there are additional received events to analyze. If so, control transfers to 924; otherwise, control returns to 808 in FIG. 24A. At 924, control selects the next one of the received events and continues at 896.

POINT-TO-POINT ROUTING

Actuators

[0197] FIG. 25 is an example wireless sensor network, including both sensor nodes and actuator nodes. In FIG. 25, each cell is shown as being a sensor cell containing only sensor nodes, or an actuator cell containing only actuator nodes. However, in various implementations, cells may include both sensor nodes and actuator nodes. According to the previously-described cellular routing, each sensor cell would communicate information to the sink, and the sink would send commands out to actuators. In other words, there would be no direct communication between sensors and actuators.

[0198] There may be applications where minimal hierarchical coordination and processing is required, meaning that sensors may be able to communicate directly to actuators. For example only, in a fire suppression system, fire detection sensors may communicate directly to sprinkler system actuators. The latency involved and potential disruption due to communication or device failure

when sending data to the sink and then relaying commands back out to actuators might make the sink-oriented architecture less desirable in these applications. See also Q. Huo, B. Dong, and S. Biswas, "Self-Organized Pulse Switching for Binary Sensing and Actuation," Proceedings of the SPIE, Volume 8753 (2013), the entire disclosure of which is hereby incorporated by reference into this application.

Frame Definition

[0199] In FIGs. 26A-26D, an example frame format is discussed that permits one cell to communicate directly to another cell without relying on the sink – i.e., bypassing the sink unless the sink lies on a best path to the other cell. FIGs. 26A and 26B are alternative frame definitions that are alternately used based on instructions from the sink.

[0200] The frames in both FIGs. 26A and 26B begin with the same three slots. The first slot is a synchronization pulse for synchronizing all of the nodes in the network. The next two slots allow the sink to instruct the nodes whether the frame of FIG. 26A (a discovery frame) or the frame of FIG. 26B (an event reporting frame) will be used. The discovery frame allows for route discovery from each node to every other node. The event reporting frame allows nodes to report events to each other as well as to the sink. In various implementations, including the one shown in FIG. 26A, the discovery frame may also include an event reporting area so that events can report it even during a discovery process.

[0201] Following the synchronization slot is a start discovery slot, which is then followed by a stop discovery slot. When the sink transmits a pulse in the start discovery slot, the nodes are instructed to use the discovery frame definition up until the frame in which the sink sends a pulse in the stop discovery slot. In other words, a pulse in the start discovery slot indicates that the discovery frame of FIG. 26A will be used while a pulse in the stop discovery slot indicates that a slot in the event reporting frame of FIG. 26B will be used.

[0202] In some implementations, a single pulse will be sent in the start discovery slot to indicate that the discovery frame of FIG. 26A will be used until

the nodes are instructed otherwise. A node that, for some reason, does not properly receive the start discovery pulse may not know to switch from the event reporting frame to the discovery frame. Therefore, in other implementations, the sink transmits a pulse in the start discovery slot during each frame that is defined as a discovery frame and transmits a pulse in the stop discovery slot for each frame defined as an event reporting frame.

[0203] The sink initiates discovery periodically. The frequency of route discovery trades off rapid adaption to changing conditions against increased energy usage in performing route discovery. FIG. 26C shows an example of the discovery area of FIG. 26A. For a network having M cells (each cell having one or more sensors and/or one or more actuators) there are M clusters of M slots each.

[0204] In between each slot is a delay interval that allows node transceivers to switch from transmitting to receiving. When measured in increments of a slot, this delay is referred to as γ . Route discovery is essentially performed by having each node broadcast a signal unique to that node. As the unique signal propagates throughout the network, every other node can determine the direction of arrival of that signal, and therefore determine a route back to that node.

[0205] Because of the time multiplexing implemented in the discovery area of FIG. 26C, all of these unique signals can be broadcast simultaneously and can propagate simultaneously throughout the network. Using cell 1 as an example, cell 1 transmits a pulse in slot 1 of cluster 1 of the discovery area. Cluster 1 is dedicated to the unique signal from cell 1, and each node that relays that signal from cell 1 will relay the signal in cluster 1 of the discovery area of each frame.

[0206] When cells that are neighbors of cell 1 detect the pulse in cluster 1, the neighboring cells recognize that the pulse indicates a route to cell 1. The slot within the first cluster indicates what that route is. Because these cells are direct neighbors of cell 1, the pulse is in slot 1 and the neighbor cells recognize that a route to cell 1 is simply through cell 1. For example, cell 5 (a neighbor of cell 1 in the example of FIG. 25) will now have a routing table entry indicating that a pulse could be sent to cell 1 where the next hop is cell 1 itself.

[0207] In the discovery area of the following frame, cell 5 will forward on a pulse in cluster 1 indicating a route to cell 1. Node 5 will place the pulse in the fifth slot of the first cluster. In this way, when a neighbor of cell 5, such as cell 8, receives this pulse, cell 8 can determine that the route to cell 1 includes a next hop of cell 5. Each of the other clusters corresponds to the remaining nodes, and because each cluster is dedicated to routes to a specific node in each frame, any given node may identify a number of routes to other nodes as well as advertise a number of routes to various nodes.

[0208] In FIG. 26D, the event reporting area of FIG. 26B is shown in detail.

Because of the discovery process, each node knows one or more next hops through which to send a pulse to a specified destination. The event reporting area includes a localization area, which includes a forwarding sub-area, a destination sub-area, and a source sub-area.

[0209] When a node sends or forwards an event pulse the forwarding sub-area indicates the next hop to which the pulse is being sent. The destination sub-area indicates the final destination of the event pulse, and the source sub-area indicates the node that originated the event pulse. To allow communication to the sink, the destination sub-area and the forwarding sub-area can each be increased by one slot. Similarly, to allow for transmissions by the sink, the source sub-area can be increased by one slot.

[0210] As a specific example of forwarding an event, consider node 5 transmitting an event to node 9 using a next hop of node 8, where the event originated at node 1. Node 5 would therefore transmit a pulse in the 8th slot of the forwarding sub-area, transmit a pulse in the 9th slot of the destination sub-area, and transmit a pulse in the 1st slot of the source sub-area.

[0211] In other words, to transmit an event, a node transmits three pulses with only one pulse in the forwarding sub-area, only one pulse in the destination sub-area, and only one pulse in the source sub-area. Nodes cannot transmit more than one event at a time because then two pulses would be present in each sub-area and there is no way of indicating which pulse in the forwarding sub-area corresponds to which pulse in the destination sub-area, and similarly which pulse in these areas corresponds to which pulse in the source sub-area. As a result, if

nodes in close proximity each attempt to transmit an event a collision occurs, as discussed in more detail in FIG. 27A.

[0212] As shown in FIG. 26D, the event reporting area includes a response area. The response area is used by nodes who detected pulses in the
5 localization area to respond as to whether an event was successfully received. A collision sub-area includes a slot that can be used by a receiving node to indicate that a collision has occurred. A node determines that a collision occurred if more than one pulse is received in any of the forwarding sub-area, the destination sub-area, or the source sub-area of the localization area. If no
10 collision is detected, the node does not transmit a pulse in the type sub-area.

[0213] If a node has successfully received an event during the localization area, the node will transmit corresponding pulses in the response area. Specifically, the source sub-area of the response area indicates the originating source of the event that has just been received, the destination sub-area indicates the ultimate
15 destination of the event that has been successfully received, and the reception area indicates the node number that has successfully received the event.

[0214] When a transmitting node receives a response matching the transmitted event, further retransmission is stopped. If no response is received, the node will attempt to retransmit, potentially using prioritized routing as described above to
20 try different next hop nodes. However, if a transmitting node receives a pulse in the collision sub-area of the response area, this indicates that a collision has occurred and therefore the transmitting node will back off before attempting to retransmit.

[0215] This backoff may be randomized so that if multiple nodes are backing off
25 from a collided transmission, the nodes are less likely to attempt to retransmit at the same time. The backoff interval can be expressed in number of frames and can be as low as one, indicating that retransmission will be attempted in the following frame, or up to an upper limit. The upper limit may be increased if more than one collision has occurred consecutively. Increasing the upper limit
30 decreases the chances that yet another collision will occur.

Collisions

[0216] FIG. 27A is a graphical example of an example collision occurring. Cell 3 is attempting to transmit an event that originated in cell 3. For this reason, cell 3 transmits a pulse in the third slot of the source sub-area. The destination of the event is cell 5, and therefore cell 3 transmits a pulse in the fifth slot of the destination sub-area. The next hop in the routing table of cell 3 for a destination of cell 5 is cell 2. Therefore, cell 3 transmits a pulse in the second slot of the forwarding sub-area.

[0217] Meanwhile, cell 4 has detected an event, and therefore transmits a pulse in the fourth slot of the source sub-area. The destination of the event is cell 1, and therefore cell 4 transmits a pulse in the first slot of the destination sub-area. The next hop in the routing table of cell 4 for a destination cell of cell 1 is cell 2. Therefore, in the forwarding sub-area, cell 4 transmits a pulse in the second slot.

[0218] If cell 3 and cell 4 are each transmitting their respective events in the same frame, cell 2 receives a combination (essentially a logical OR) of the two events. In other words, a single pulse is received in the second slot of the forwarding sub-area. However, pulses are received in both the third and fourth slots of the source sub-area. Further, pulses are received in the first and fifth slots of the destination sub-area.

[0219] As a result, although cell 2 recognizes that at least two events were forwarded to cell 2, cell 2 is unable to identify which pulse in the source sub-area matches which pulse in the destination sub-area. Cell 2 therefore recognizes that a collision has occurred, and in the response area following the localization area transmits a pulse in the collision sub-area to indicate that a collision has occurred. Cells 3 and 4 will therefore choose a random number of frames to back off before retransmitting their respective events.

[0220] In various implementations, a collision may not be declared if multiple pulses are received in only one of the sub-areas. In other words, if only a single pulse is received in the other two sub-areas, the receiving node can uniquely identify the two events that were received simultaneously. For example only, see FIG. 27B, where both cells 3 and 4 are transmitting events to cell 5 via cell 2.

[0221] Therefore, cell 2 receives a single stacked pulse in the second slot of the forwarding sub-area and a single stacked pulse in the fifth slot of the destination sub-area. Meanwhile, in the source sub-area, cell 2 receives pulses in slots 3 and 4. Cell 2 can infer that two events, each having a destination of cell 5 and each being sent to cell 2, have corresponding sources of cell 3 and cell 4. This assumes that neither the pulse in cell 3 nor the pulse in cell 4 is a false pulse and also assumes that a second pulse was not lost in the forwarding sub-area and that a second pulse was not lost in the destination sub-area. To help justify these assumptions, protection mechanisms such as are described above, including a protection area, may be implemented.

Methods

[0222] FIG. 28A is an example method for performing route discovery using the frames defined in FIGs. 26A-26D. Control begins at 1004, where if a start discovery pulse has been received, control transfers to 1008; otherwise, control remains at 1004. At 1008, control schedules a discovery pulse transmission. The current node is labeled in FIG. 28A as node X, and therefore the pulse transmission is scheduled to occur in slot X of cluster X of the discovery area of the frame.

[0223] Control continues at 1012, where an array named sent[] is cleared. The sent[] array tracks whether a discovery pulse has been sent during the present discovery phase for a given source node. Control continues at 1016, where control determines whether the discovery pulse has been received. If so, control transfers to 1020; otherwise, control transfers to 1024. At 1020, the slot during which the discovery pulse was received is stored as variable Y and the cluster in which the discovery pulse was received is stored in variable Z.

[0224] At 1020, control determines whether node Y is a neighbor of node X. If so, control transfers to 1028; otherwise, control ignores the discovery pulse and transfers to 1024. At 1028, control adds a routing table entry for node Z with the next hop being node Y. Control transfers to 1032, where if a discovery pulse has already been sent for node Z, control transfers to 1024. If, however, a discovery pulse has not yet been sent for node Z, control transfers to 1036 where the sent[] array is updated to indicate the transmission of the discovery pulse.

[0225] Control continues at 1040, where a pulse is scheduled for transmission in slot X of cluster Z. Control then continues at 1024. At 1024, if a stop discovery pulse has been received, control transfers to 1004; otherwise, control returns to 1016.

5 **[0226]** FIG. 28B is an example method of event reporting for a node (labeled node X) using the frames defined in FIGs. 26A-26D. Control begins at 1104, where a queue is cleared. Each element of the queue is a tuple of source and destination. The source is the original source of the event that resulted in a pulse being transmitted while the destination is the ultimate destination to which the
10 event is being reported.

[0227] Control continues at 1108, where if an event is detected by node X, control transfers to 1112; otherwise, control transfers to 1116. At 1112, control determines a destination node (stored in variable D) to which to report the event. For example only, an event can be reported to the sink, which by convention in
15 this example is a value of zero. Control continues at 1120, where an event is added to the queue with a source of X and a destination of D. Control then continues at 1116.

[0228] At 1116, control determines whether an event has been received. If so, control transfers to 1124; otherwise, control continues at 1128. At 1124, control
20 determines whether a collision has occurred while the event was being received. If so, control transfers to 1132, where control transmits a pulse in the collision sub-area of the event reporting area of the frame. Control then continues at 1128. If no collision occurred, control transfers to 1136.

[0229] If the forwarding node specified in the received event is equal to node X,
25 control continues at 1140; otherwise, control discards the event and continues at 1128. At 1140, control determines whether the source node of the received event is X. If so, that event originated at the present node and some form of loop has occurred. Therefore, if the source of the event is equal to X, control discards the event and transfers to 1128; otherwise, control continues at 1144.

30 **[0230]** At 1144, control transmits an indication of a successful response, and continues at 1148. At 1148, control determines whether the destination of the

event is node X. If so, the event has reached its ultimate destination, and control continues at 1152; otherwise, control transfers to 1156. At 1152, the event is processed, which may include operating an actuator. Control then continues at 1128. At 1156, the event has not yet reached its final destination, and therefore
5 the event is added to the queue. Control then continues at 1128.

[0231] At 1128, control determines whether there is an event in the queue. If so, control transfers to 1160; otherwise, control returns to 1108. At 1160, control determines whether a back-off interval is ongoing. If so, control is waiting to transmit due to a previous collision, and control therefore returns to 1108;
10 otherwise, control transfers to 1164. At 1164, control determines a next hop for the next event in the queue, and at 1168, control transmits the event with the next hop specified in the forwarding sub-area.

[0232] Control continues at 1172, where if a collision indication has been received, control transfers to 1176; otherwise, control transfers to 1180. At 1180,
15 control determines whether a response that matches the transmitted event has been received. If so, control transfers to 1184 and removes the event from the queue. Control then continues at 1108. If a matching response was not received, the event remains in the queue and control returns to 1108.

CONCLUSION

20 **[0233]** The foregoing description is merely illustrative in nature and is in no way intended to limit the disclosure, its application, or uses. The broad teachings of the disclosure can be implemented in a variety of forms. Therefore, while this disclosure includes particular examples, the true scope of the disclosure should not be so limited since other modifications will become apparent upon a study of
25 the drawings, the specification, and the following claims. As used herein, the phrase at least one of A, B, and C should be construed to mean a logical (A or B or C), using a non-exclusive logical OR. It should be understood that one or more steps within a method may be executed in different order (or concurrently) without altering the principles of the present disclosure.

30 **[0234]** In this application, including the definitions below, the term module may be replaced with the term circuit. The term module may refer to, be part of, or

include an Application Specific Integrated Circuit (ASIC); a digital, analog, or mixed analog/digital discrete circuit; a digital, analog, or mixed analog/digital integrated circuit; a combinational logic circuit; a field programmable gate array (FPGA); a processor (shared, dedicated, or group) that executes code; memory
5 (shared, dedicated, or group) that stores code executed by a processor; other suitable hardware components that provide the described functionality; or a combination of some or all of the above, such as in a system-on-chip.

[0235] The term code, as used above, may include software, firmware, and/or microcode, and may refer to programs, routines, functions, classes, and/or
10 objects. The term shared processor encompasses a single processor that executes some or all code from multiple modules. The term group processor encompasses a processor that, in combination with additional processors, executes some or all code from one or more modules. The term shared memory encompasses a single memory that stores some or all code from multiple
15 modules. The term group memory encompasses a memory that, in combination with additional memories, stores some or all code from one or more modules. The term memory may be a subset of the term computer-readable medium. The term computer-readable medium does not encompass transitory electrical and electromagnetic signals propagating through a medium, and may therefore be
20 considered tangible and non-transitory. Non-limiting examples of a non-transitory tangible computer readable medium include nonvolatile memory, volatile memory, magnetic storage, and optical storage.

[0236] The apparatuses and methods described in this application may be partially or fully implemented by one or more computer programs executed by
25 one or more processors. The computer programs include processor-executable instructions that are stored on at least one non-transitory tangible computer readable medium. The computer programs may also include and/or rely on stored data.

CLAIMS

What is claimed is:

1. A sensor network comprising:
 - a sink coupled to a substrate and configured to transmit a periodic
 - 5 ultrasonic pulse on the substrate; and
 - a plurality of sensor nodes, wherein a first sensor node of the plurality of sensor nodes is coupled to the substrate and is configured to:
 - receive the periodic ultrasonic pulse from the substrate,
 - synchronize an internal clock of the first sensor node to the sink
 - 10 based on the periodic ultrasonic pulse,
 - selectively detect an event in a region surrounding the first sensor node, and
 - in response to detecting the event, transmit a first ultrasonic pulse toward the sink on the substrate.
- 15 2. The sensor network of claim 1 wherein:
 - the periodic ultrasonic pulse has a first transmission range,
 - the first ultrasonic pulse has a second transmission range, and
 - the first transmission range is greater than the second transmission range.
- 20 3. The sensor network of claim 2 wherein:
 - the sink transmits the periodic ultrasonic pulse using a first voltage,
 - the first sensor node transmits the first ultrasonic pulse using a second voltage, and
 - the first voltage is greater than the second voltage.
- 25 4. The sensor network of claim 3 wherein the first voltage is approximately double the second voltage.
5. The sensor network of claim 1 wherein a second sensor node of the plurality of sensor nodes is configured to:
 - receive the first ultrasonic pulse, and

in response to receiving the first ultrasonic pulse, transmit a second ultrasonic pulse toward the sink.

6. The sensor network of claim 1 wherein the first sensor node is configured to:

5 determine an arrival time of a first mode of the periodic ultrasonic pulse,
determine an arrival time of a second mode of the periodic ultrasonic pulse,

estimate a propagation delay of the periodic ultrasonic pulse from the sink to the first sensor node based on the arrival time of the first mode and the arrival
10 time of the second mode.

7. The sensor network of claim 6 wherein the first sensor node is configured to adjust synchronization of the internal clock of the first sensor node based on the propagation delay.

8. The sensor network of claim 6 wherein:
15 the first mode is a symmetric mode, and
the second mode is an asymmetric mode.

9. The sensor network of claim 1 wherein:
the sensor network is spatially divided into a plurality of cells, and
each sensor node of the plurality of sensor nodes is assigned to one of
20 the plurality of cells.

10. The sensor network of claim 9 wherein each sensor node of the plurality of sensor nodes is assigned to one of the plurality of cells during a pre-programmed setup phase of the sensor network.

11. The sensor network of claim 1 further comprising an actuator node
25 including a physical actuator.

12. The sensor network of claim 11 further comprising a plurality of actuator nodes including the actuator node.

13. The sensor network of claim 11 wherein the physical actuator operates in response to a command from the sink.
14. The sensor network of claim 11 wherein the physical actuator operates in response to data received from one or more of the plurality of sensor nodes.
- 5 15. The sensor network of claim 14 wherein the data bypasses the sink and the actuator node receives the data directly from the one or more of the plurality of sensor nodes.
16. The sensor network of claim 1 wherein the first sensor node includes an ultrasonic transducer coupled to the substrate.
- 10 17. The sensor network of claim 16 wherein the ultrasonic transducer includes a piezoelectric element.
18. The sensor network of claim 16 wherein the ultrasonic transducer is placed between a first side of the substrate and a first magnet, and wherein a second magnet is located on an opposite side of the substrate.
- 15 19. The sensor network of claim 18 wherein the ultrasonic transducer is secured to the first magnet with an epoxy.
20. An apparatus comprising:
a frame synchronization circuit configured to identify a start of a frame in response to receipt of a global synchronization pulse; and
20 a pulse transmission control circuit configured to, in response to detection of an event by a sensor associated with the apparatus, transmit a first pulse toward a first destination to signal detection of the event,
wherein a temporal position of the first pulse transmission within the frame indicates a spatial position of the event, and
25 wherein the temporal position of the first pulse transmission within the frame also indicates a next hop toward the destination to which the first pulse is directed.

21. The apparatus of claim 20 wherein a portion of the frame is divided into clusters of slots, and wherein each of the clusters corresponds to a distinct spatial position of an event.
22. The apparatus of claim 21 wherein, within each of the clusters, each of
5 the slots corresponds to a distinct next hop.
23. The apparatus of claim 20 further comprising a discovery circuit configured to determine next hops to a plurality of destinations including the first destination.
24. The apparatus of claim 20 further comprising a pulse forwarding circuit
10 configured to:
 receive a second pulse from a second apparatus,
 determine a spatial location of a second event in response to a temporal position of the second pulse in the frame, and
 wherein the pulse transmission control circuit is configured to, in response
15 to the second pulse, transmit a third pulse toward the first destination,
 wherein a temporal position of the third pulse transmission within the frame indicates a spatial position of the second event, and
 wherein the temporal position of the third pulse transmission within the frame also indicates the next hop toward the destination to which the third pulse
20 is directed.
25. The apparatus of claim 20 wherein the first pulse is an ultrasonic pulse and wherein the global synchronization pulse is an ultrasonic pulse.
26. A sensor node comprising:
 the apparatus of claim 25; and
25 an ultrasonic transducer.
27. The sensor node of claim 26 wherein the ultrasonic transducer comprises a piezoelectric element.
28. The sensor node of claim 26 wherein the ultrasonic transducer comprises a piezoelectric wafer active sensor.

29. The sensor node of claim 26 further comprising a sensor configured to identify changes in an environment proximate to the sensor node.
30. The sensor node of claim 26 further comprising:
a clock circuit configured to generate a periodic waveform;
5 a timer circuit configured to output an enable signal for a predetermined length of time;
a pulse multiplexer configured to provide an output based on the periodic waveform while the enable signal is being output by the timer circuit; and
a transmit/receive multiplexer configured to selectively provide the output
10 to an ultrasonic transducer.
31. The sensor node of claim 30 further comprising:
an amplifier circuit configured to generate an amplified signal by amplifying a received signal from the ultrasonic transducer; and
an envelope circuit configured to generate a processed signal by applying
15 an envelope function to the amplified signal, wherein timing of received pulses is determined based on the processed signal.




FIG. 1
Prior Art




FIG. 2
Prior Art




FIG. 3




FIG. 4




FIG. 5A




FIG. 5B




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12A





FIG. 12B






FIG. 13A




FIG. 13B




FIG. 13C









FIG. 14

18/39





Node	Neighbors	Reception Times
Sink	-	-
Node 1	-	-
Node 2	S	T_0
Node 3	S	T_0
Node 4	-	-
Node 5	-	-

FIG. 16A





FIG. 16B

Node	Neighbors	Reception Times
S	2	T_1
	3	T_1
Node 1	2	T_1
Node 2	S	T_0
	3	T_1
Node 3	S	T_0
	2	T_1
Node 4	3	T_1
Node 5	-	-





FIG. 16C

Node	Neighbors	Reception Times
S	2	T_1
	3	T_1
Node 1	2	T_1
Node 2	S	T_0
	3	T_1
	1	T_2
Node 3	S	T_0
	2	T_1
	4	T_2
Node 4	3	T_1
Node 5	4	T_2






FIG. 16D

Node	Neighbors	Reception Times
S	2	T_1
	3	T_1
Node 1	2	T_1
Node 2	S	T_0
	3	T_1
	1	T_2
Node 3	S	T_0
	2	T_1
	4	T_2
Node 4	3	T_1
	5	T_3
Node 5	4	T_2



Node	Neighbors	Reception Times
S	2	T_1
	3	T_1
Node 1	2	T_1
Node 2	S	T_0
	3	T_1
	1	T_2
Node 3	S	T_0
	2	T_1
	4	T_2
Node 4	3	T_1
	5	T_3
Node 5	4	T_2




FIG. 17A





FIG. 17B

Node	Neighbors	Reception Times
S	2	T_1
	3	T_1
Node 1	2	T_1
Node 2	S	T_0
	3	T_1
	1	T_2
Node 3	S	T_0
	2	T_1
	4	T_2
Node 4	3	T_1
	5	T_3
Node 5	4	T_2




FIG. 18






FIG. 19

27/39

**FIG. 20A**

**FIG. 20B**





FIG. 21C







FIG. 22

31/39

**FIG. 23**

32/39












FIG. 25






FIG. 27A






FIG. 27B

**FIG. 28A**

39/39

**FIG. 28B**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2014/067609

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G08B 9/00 (2015.01)

CPC - G08B 25/14 (2014.12)

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) - G08B 1/08, 9/00, 29/00 (2015.01)

USPC - 340/286.02, 511, 531, 539.1, 539.22; 370/282; 367/137

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
CPC - G08B 25/14, 17/06, 25/10 (2014.12) (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Orbit, Google Patents, ProQuest

Search terms used: sensor, nodes, sink, substrate, internal clock, multihop, transducer, ultrasonic, pulse

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2010/0027379 A1 (SAULNIER et al) 04 February 2010 (04.02.2010) entire document	1-19
Y	US 2011/0285516 A1 (RITTER) 24 November 2011 (24.11.2011) entire document	1-19
Y	US 8,499,632 B1 (IHN et al) 06 August 2013 (06.08.2013) entire document	8
Y	US 2006/0178156 A1 (KIM) 10 August 2006 (10.08.2006) entire document	9-10
Y	US 2004/0100394 A1 (HITT) 27 May 2004 (27.05.2004) entire document	11-15
Y	US 2009/0059842 A1 (MALTSEFF et al) 05 March 2009 (05.03.2009) entire document	17
Y	US 2010/0328075 A1 (RAHAMIM et al) 30 December 2010 (30.12.2010) entire document	18-19
Y	US 2005/0134574 A1 (HILL) 23 June 2005 (23.06.2005) entire document	19
A	US 2008/0048605 A1 (FRANKE) 28 February 2008 (28.02.2008) entire document	1-19
A	US 2011/0268139 A1 (CARACAS et al) 03 November 2011 (03.11.2011) entire document	1-19
A	US 2008/0048905 A1 (MCEWAN) 28 February 2008 (28.02.2008) entire document	1-19

☐

Further documents are listed in the continuation of Box C.

☐

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

16 March 2015

Date of mailing of the international search report

13 APR 2015

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2014/067609

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:
See last page

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-19

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2014/067609

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I, claims 1-19, drawn to a sensor network.

Group II, claims 20-31, drawn to an apparatus.

The inventions listed as Groups I-II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons: the special technical feature of the Group I invention: a sink coupled to a substrate and configured to transmit a periodic ultrasonic pulse on the substrate; and a plurality of sensor nodes as claimed therein is not present in the invention of Group II. The special technical feature of the Group II invention: a frame synchronization circuit configured to identify a start of a frame in response to receipt of a global synchronization pulse; and a pulse transmission control circuit as claimed therein is not present in the invention of Group I.

Groups I and II lack unity of invention because even though the inventions of these groups require the technical feature of detecting an event by a sensor and transmitting a pulse, this technical feature is not a special technical feature as it does not make a contribution over the prior art.

Specifically, US 2008/0048905 A1 (MCEWAN) 28 February 2008 (28.02.2008) teaches detecting an event by a sensor and transmitting a pulse (a precision pulse detector, comprising a transceiver 12 which transmits pulses and receives echo, or range, pulses in response to a timing control circuit 40, Para. 27 and Fig. 1).

Since none of the special technical features of the Group I or II inventions are found in more than one of the inventions, unity of invention is lacking.