



(22) Date de dépôt/Filing Date: 2014/07/04

(41) Mise à la disp. pub./Open to Public Insp.: 2015/09/18

(30) Priorité/Priority: 2014/03/18 (US61/955.018)

(51) Cl.Int./Int.Cl. *H04W 56/00* (2009.01),
H04W 52/02 (2009.01), *H04W 8/22* (2009.01)


(71) Demandeur/Applicant:
SMARTREK TECHNOLOGIES INC., CA

(72) Inventeur/Inventor:
NGUYEN-DANG, THIEN-LY, CA

(74) Agent: ROBIC


(54) Titre : SYSTEME DE RESEAU MAILLE ET TECHNIQUES

(54) Title: MESH NETWORK SYSTEM AND TECHNIQUES



(57) **Abrégé/Abstract:**

The present relates to a node to be integrated in a network of wirelessly connected nodes, to a network system, to method for operating a node of a network, and to a method for communicating within a network. A method is disclosed for communicating within a network of cooperatively synchronized nodes, configured to broadcast data packets to the network during broadcast phases through a flooding mechanism, each data packet comprising hop data. The method comprises, for each broadcast phase, broadcasting a data packet from a source node during a predetermined time slot; and receiving the data packet at one or more destination node during said predetermined time slot. The broadcasting and receiving are repeated according to the hop data, at respective predetermined time slots, wherein each destination node corresponds to a source node in a next execution of broadcasting and the data packet received corresponds substantially to the data packet to be broadcasted in the next repetition.



Abstract

The present relates to a node to be integrated in a network of wirelessly connected nodes, to a network system, to method for operating a node of a network, and to a method for communicating within a network. A method is disclosed for communicating within a network of cooperatively synchronized nodes, configured to broadcast data packets to the network during broadcast phases through a flooding mechanism, each data packet comprising hop data. The method comprises, for each broadcast phase, broadcasting a data packet from a source node during a predetermined time slot; and receiving the data packet at one or more destination node during said predetermined time slot. The broadcasting and receiving are repeated according to the hop data, at respective predetermined time slots, wherein each destination node corresponds to a source node in a next execution of broadcasting and the data packet received corresponds substantially to the data packet to be broadcasted in the next repetition.

MESH NETWORK SYSTEM AND TECHNIQUES

Field:

5

The present relates to network systems and techniques, and more particularly, to a node to be integrated in a network of wirelessly connected nodes, to a network system, to method for operating a node of a network, and to a method for communicating within a network.

10

Background:

The cooperative synchronized networks idea stems from the principle of having a synchronization on the physical layer of a given mesh network, be it via wavelength-level phase synchronization, or via a more easily achievable implementation with synchronization at bit-level in order to enable deployment on existing industry-standard design targets.

Known in the art are contention-enabled technologies such as ZigbeeTM, for example.

Bit-synchronized networks are a special case of antenna diversity found in phased array radar, for example, with less timing constraint, which is enough for ensuring minimal interference between simultaneous packet transmission of multiple node neighbours, the imperfection due to phase delay and multi-path fading being accounted for by FM capture effect. It is very seldom considered as a solution outside of the theoretical realm, and is often only tested in small scale settings where long range delay interference are minimal, or assessed in a pure simulated framework. For instance, the Glossy implementation using DSSS transceivers do this in a setting within a physical boundary that is less than the path delay threshold for which bit-synchronization fails. Extension of the principle to situations where pure bit-synchronization fails, i.e. when path delays exceed

the threshold for proper bit synchronization (such a timing threshold can be defined as corresponding to the $1/(2 * \text{bit-rate})$) using a fallback mechanism has not been assessed in the current art literature.

5 Using such a network, the time taken for a given packet from a source node to flood the whole network can be precisely time-bounded, independently of the density of nodes deployed, in an ideal setting where bit-synchronization holds. This can be used to drastically lower the wake/sleep duty cycle of the network for a given data bandwidth in the case where high node density is required for
10 operating the network. Discounting real-life limitation in synchronization clock precision, the power consumption is in a direct relationship with said duty cycle, and thus can linearly scale down as low as the duty cycle permits.

However, in a real-life implementation setting, the maximum duty cycle is
15 bounded to the worst-case clock tolerances in the manufactured mesh network nodes. As current technology typically allows for clock tolerances between 10 and 100 ppm, this implies that each node clock has to be corrected by a coordinator node in order to correct clock drift periodically in order to maintain the tight tolerances needed for bit-synchronization at medium speeds (for instance, in
20 the order of a few microseconds in the case of a FHSS network at bit-rates in the order of kilobits/sec, and less than 1 microsecond if one were to use a DSSS network such as what is found in IEEE 802.15.4 radio technology). Clock synchronization is a known problem for distributed radio transceiver networks. It is however specially critical in the case of bit-synchronized cooperative networks.

25 Not only must the network have a coordinator node with a reference clock in order to broadcast timing information to the rest of the nodes (be it in a direct or indirect fashion), called in-band clock synchronization, but there must also be a safety mechanism in order to ensure that a given node is awake when this
30 synchronization broadcast arrives. Thus, a safety guard time is usually added in the wake schedule time of any node in order to ensure this. In the case of regular

contention-enabled mesh networks, this guard time has to include the effects of all the following: clock tolerances (drift and jitter), broadcast flood non-linear propagation worst-case time with respect to network size and density, number of retries allowed, etc. and is an ongoing subject of research due to the complexity of contention-enabled broadcast packet propagation.

In the case of a bit-synchronized network, this guard time can be typically linked to the chosen clock manufacturing tolerances. Thus, the state of the art dictates that the minimum safe duty cycle will be limited to the chosen clock relative tolerance (for example, $T=1000$ ppm internal RC oscillator: oscillator tolerance T , desired interval between wake cycles: X sec). Idealized broadcast interval without synchronization guard time: Y sec. Idealized duty cycle (a typical configuration parameter) defined as $D=Y/X$. The required guard time per broadcast cycle is $X*T$. Real duty cycle is thus $R = (Y+X*T)/X = Y/X + T = D + T$. Thus, asymptotic case where D tends to zero, we have minimum duty cycle limited to $R = T$. Thus, in order to accommodate low effective duty cycles, one has to typically use higher-cost crystals, a combination of multiple crystals. This also limits the hardware platforms onto which a cooperative bit-synchronized mesh network can be deployed, and legacy retrofitting is usually impossible, and design required dedicated hardware platforms, lest one needs to cope with less efficient minimum network power consumption, a trade-off that can be unacceptable for long-range, high RF power, long-lasting, low power consumption applications, such as in high-interference locations such as large forest deployments.

25

The alternate method for network-wide clock synchronization is using out-band methods in order to avoid the clock tolerance problem, such as using GPS time base from a GPS unit installed per-node, with implied hefty penalties in node cost, due to raw hardware cost as well as increased integration complexity.

30

The conventional FSK or FSK-like mesh network node electronic design architecture is typically the following: RF node front-end (typically a general purpose RF ASIC with analog front-end circuitry for both transmitter and receiver with a dedicated clock in order to generate local oscillator-derived signals, said ASIC being linked to a "mesh controller" processing unit which contains glue logic interfaces (such as UART, SPI, I2C, CAN, MODBUS etc.) to a host processing unit which is customized according to specific implementation. Thus, there is usually the need for at least 3 dedicated ICs per sensor node, and that is discounting supplementary glue logic circuitry to external sensors in order to link them to said host processing unit. A supplementary technical challenge when employing cooperative bit-synchronized networks is that timing is key, and thus, for single core processors employed as mesh controllers, all interrupts must be typically be deactivated before a timing-critical operation (such as a packet transmit begin trigger signal) occurs. In other terms, running of parallel user-code can possibly jeopardize mesh synchronization integrity between nodes, depending on whether user-code execution overflows into the time slots dedicated for broadcast cycles. This is why current solutions usually employ a supplementary processor core dedicated to user firmware.

Work on optimizing and leveraging a cooperative bit-synchronized network using flooding broadcasts as the sole and prime mechanism of communication has been limited to standard master-slave polling methodologies, where a master (by default, the node chosen as being the coordinator node) polls a node, waits for its answer, then polls the next one in an arbitrary order often predefined from a configuration database (the most known of which is standard round-robin). Moreover, network transit of sensor data, industry-standard mesh networks, is usually limited to collection at the coordinator node location, the synchronization of said collected data to other clients being done using external networks such as IP, telephony or satellite-based networks. A typical example would be a sensor collection system which will have users bound to a mobile device with access to cellular or satellite infrastructure if those users need their sensor data while

roaming in the field. Indeed, in a practical setting, transit of whole network data cannot be feasibly sent to another client node using the same mesh network without heavy power consumption penalties, because for each client, the sensor typically has to actively send the whole of its collected data with said client device
5 using direct client MAC addressing (analogous to IEEE 802.15.4 unicast packet send methods). This is thus not practical in a power-sensitive setting such as for forestry applications, lest costly portable power sources or energy harvesting strategies be used at each deployed node location, because the total network power consumption is typically multiplied by the number of client devices
10 required.

In addition, in a master-slave sensor architecture within a cooperative bit-synchronized mesh network as described above, the coordinator node typically sends a broadcast packet with flooding slotting information in order to time-bound
15 the flooding process to a finite number of hops, and then the requested sensor node being polled or controlled sends back its answer in a similar fashion. This implies that in a whole broadcast flood phase (either "out", i.e. from a coordinator node to the whole network, and "in", i.e. from an end node back to said coordinator), only one data sub-packet can be flooded throughout the whole
20 network at any given time (thus this is very different at the core than ad-hoc mesh routing strategies found within the convention IEEE802.15.4 physical stack). In the case that more than one end node attempt to flood back an answer throughout the whole network (if there are no multiple network operating on different RF channels or hop tables in the case of a FHSS system), either only
25 one packet will reach back to the coordinator node, or even none of them depending on the mesh physical configuration.

This has a corollary: there must be some addressing mechanism put in place in order to ensure that during a broadcast "in" phase, only one end node answers.
30 Typically, a MAC address (or any identification scheme that end nodes agree to) is send by the coordinator node embedded in the data sub-packet (this can be

considered as being part of a 'control' sub-frame in the higher protocol layers according to the OSI model). This means that by its very inception, an end node within such a network must wait for a specific gateway signal before being allowed to send back its data. In the case where a large quantity of nodes are
5 deployed in the field, and for low wake-sleep duty cycles, the reaction time of a node degrades linearly with the number of nodes deployed, using a standard node-balanced polling scheme. This can cause problems when higher rates of polling for specific nodes are required when large networks are deployed that need to save energy in order to last longer.

10

Another corollary is that the communication is done by the very nature of the mesh network in a half-duplex fashion, and thus there will always be a minimal network overhead caused by the requirement of a packet from the coordinator node indirectly containing network sleep synchronization information as well as
15 the request information outbound to a destination end-node. Due to the symmetrical nature of broadcast phases (i.e. phase "out" and phase "in" take the same time as they have to share the same maximum number of hops setting), there is a base information bandwidth reduction factor of 2 in the maximum network speed at which a select end-node can send back a stream of data.

20

Hence, in light of the aforementioned, there is a need for an improved system which, by virtue of its design and components, would be able to overcome some of the above-discussed prior art concerns.

25 **Summary:**

The object of the present invention is to provide a device which, by virtue of its design and components, satisfies some of the above-mentioned needs and is thus an improvement over other related mesh network systems and/or methods
30 known in the prior art.

In accordance with the present invention, the above mentioned object is achieved, as will be easily understood, by a mesh network system such as the one briefly described herein and such as the one exemplified in the accompanying drawings.

5

It is an object of the present to provide a reliable, fast and low power consumption remote data acquisition and device control with multiple channel-sharing clients in an environment with multiple interference-generating physical elements, in the presence of multi-path fading and time delays associated with longer ranges, using an extension of a cooperative bit synchronized frequency-modulated wireless mesh network, and related corresponding manufacturing processes and applications.

Another object is to solve broad engineering requirements besides the use of a bit-synchronized cooperative mesh network using whole network broadcast packet flooding as the principal communication method, for example:

- How to enable deployment of said mesh technology with low duty cycles using existing industry-standard hardware platforms, and how to minimize cost of integration by time-sharing the mesh controller processor with ancillary functionalities;
- How to enable the use of multiple bandwidth-asymmetric clients that communicate via the same sensor network they aim to monitor and control, and that can refresh their databases with the whole network current sensor data without drastically impacting network power consumption;
- How to improve perceived sensor network responsivity via direct interaction between lower protocol levels (according to the OSI model) up to the HMI level
- How to mitigate the bandwidth reduction issue by segmenting a broadcast cycle (network activity cycle between sleep cycles) into more than the

standard 2 broadcast cycles typical of the master-slave half-duplex architecture

- How to improve absolute sensor network responsivity using a novel scheme that we call the "redux phase"

- 5
- How to leverage the low power consumption of the network in order to retrofit duty cycling and auto configuration of a parallel higher-bandwidth and higher-consumption multi-hop network (for instance IEEE802.11s), which can be critical for on-demand mesh wi-fi without power harvesting solutions in place

10

In accordance with an aspect, there is provided a node to be integrated in a network of wirelessly connected nodes, the node comprising:

- a main module comprising a main controller having a main clock, the main controller being operable between an active configuration wherein the main clock oscillates and an inactive configuration wherein the main clock sleeps ;
- 15 - a digital serial bus connected to the main module, the digital serial bus being operable under a communication protocol which is configurable to embed timing information; and
- an external sleep controller being connected to the main module via the
- 20 digital serial bus to operate the main controller between the active and inactive configurations based on said timing information.

In accordance with another aspect, there is provided a network system comprising a plurality of interconnected end nodes as described above, the

25 network further comprising a coordinator node for communicating with the end nodes, the network being configured to broadcast data and synchronize a time base via a flooding mechanism.

In accordance with another aspect, there is provided a method for operating a

30 node adapted to cooperate within a network of nodes, the node comprising a main module and an external sleep controller connected thereto via a digital

serial bus, the main controller having a main clock, the main controller being operable between an active configuration wherein the main clock oscillates and an inactive configuration wherein the main clock sleeps, the method comprising:

- sending from the main controller, timing information embedded in a communication protocol of the digital serial bus, to the external sleep controller, said timing information comprising sleep period data representative of a time period during which the main module is intended to operate in an inactive configuration;
- at the external sleep controller, measuring elapsed time via a clock; and
- when the elapsed time has reached said time period, sending a wake signal from the external sleep controller, over the digital serial bus, to the main module to operate the main module in the active configuration.

In accordance with yet another aspect, there is provided a method for communicating within a network comprising end nodes and a coordinator node, the network being cooperatively synchronized and configured to broadcast data packets during broadcast cycles through a flooding mechanism, each data packet comprising a configuration frame and a data frame, the configuration frame comprising hop data, the method comprising for each broadcast cycle:

- a) during an outbound broadcast phase, sending from the coordinator node a coordinator-originating data packet for polling at least one of the nodes, the data frame of the coordinator-originating data packet comprising address data corresponding to the at least one end node to be polled; and

- b) during an inbound broadcast phase, receiving at the coordinator node a node-originating data packet returned from one of said at least one of the end nodes, the data frame of the node-originating data packet comprising node data from the end node having been polled;

each of said outbound and inbound broadcast phases comprising successive time slots for at least one of said end node to either receive the data packet or rebroadcast the data packet to the network;

the method further comprising:

i) during one of said time slots, receiving at one of the end nodes, one of said data packets; and

ii) during another one of said time slots:

- updating the hop data in the data packet;

5 - rebroadcasting the data packet from the end node; and

- during an inbound broadcast phase, if the end node corresponds to the address contained in the data packet, generating and broadcasting a data packet comprising node data from the end node and corresponding hop data .

10 In accordance with yet another aspect, there is provided a method for communicating within a network comprising end nodes and a coordinator node, the network being configured to broadcast data packets between said nodes during broadcast cycles through a flooding mechanism, consecutive broadcast cycles being separated by a sleep period wherein the end nodes are deactivated
15 for external communication, the method comprising:

- during said sleep period: executing a node-related process, at a main controller of at least one of the end nodes.

In accordance with still another aspect, there is provided a method for
20 communicating within a network of nodes, the network being cooperatively synchronized and configured to broadcast data packets to the network during broadcast phases through a flooding mechanism, each data packet comprising hop data, the method comprising, for each broadcast phase:

a) broadcasting a data packet from a source node during a predetermined
25 time slot;

b) receiving the data packet at one or more destination node during said predetermined time slot; and

repeating steps (a) and (b) according to the hop data, at respective predetermined time slots, wherein each destination node of step (b) corresponds
30 to a source node in a next execution of step (a) and the data packet received in

step (b) corresponds substantially to the data packet to be broadcasted in the next execution of step (a).

The objects, advantages and features of the present invention will become more
5 apparent upon reading of the following non-restrictive description of preferred
embodiments thereof, given for the purpose of exemplification only, with
reference to the accompanying drawings.

Brief description of the drawings:

10

FIG. 1A shows a high level structure of an end node mesh controller, according to an embodiment.

15

FIG. 1B shows a circuit plan of an end node mesh controller, according to an
embodiment.

FIG. 2 shows specific location of event triggers, according to an embodiment.

20

FIG. 3 shows an example of a VM flowchart, according to an embodiment. FIG.
3A, 3B, 3C are enlarged views of portions of FIG. 3.

25

FIG. 4 shows a packet structure and its configuration frame, as well as other
related features, according to an embodiment. FIG. 4A, 4B, 4C, 4D are enlarged
views of portions of FIG. 4.

FIG. 5 is a schematic diagram showing a first polling strategy, according to an
embodiment. FIG. 5A, 5B, 5C, 5D are enlarged views of portions of FIG. 5.

30

FIG. 6 is a schematic diagram showing a second polling strategy, according to an
embodiment. FIG. 6A, 6B, 6C, 6D are enlarged views of portions of FIG. 6.

FIG. 7 is a schematic diagram showing a third polling strategy, according to an embodiment. FIG. 7A, 7B, 7C, 7D are enlarged views of portions of FIG. 7.

5 FIG. 8 is a schematic diagram showing a second polling strategy, according to an embodiment. FIG. 8A, 8B, 8C, 8D are enlarged views of portions of FIG. 8.

FIG. 9 is a flowchart of an exemplary collision management scheme, according to an embodiment. FIG. 9A, 9B are enlarged views of portions of FIG. 9.

10

FIG. 10A, 10B show a flowchart of a RC oscillator-bound sleep controller, according to an embodiment.

FIG. 11A, 11B show a flowchart detailing the main controller (I2C master) side
15 that is connected via I2C bus to an external sleep controller for managing sleep mode, according to an embodiment.

FIG. 12 is a flowchart detailing the operation of an external sleep controller (I2C slave), according to an embodiment.

20

FIG. 13A, 13B show an I2C bus waveform timeline for an external sleep controller communication bus during the sleep initiation phase, according to an embodiment.

25 FIG. 14A, 14B show an I2C bus waveform timeline for an external sleep controller communication bus during the wake-up phase, according to an embodiment.

FIG. 15A is a VM example source code for a simple sensor node (whose main
30 data variables are its own battery state and its radio signal state), according to an embodiment of the VM strategy.

FIG. 15B is a screenshot of a VM tool-chain targeting a mesh controller, according to an embodiment of the VM strategy.

- 5 FIG. 16 shows the structure of the sub-packets encoding node network state, configuration and data variables, within the data payload of a packet, according to an embodiment.

FIG. 17 is a flowchart of a state machine used on the host device connected to a
10 coordinator node for UDP-style polling the network, according to an embodiment.

FIG. 18 is a flowchart of a state machine used on the host device connected to an asymmetric "sniffer" node, according to an embodiment.

- 15 FIG. 19 shows a node registration mechanism using a mobile GPS-enabled device, as well as the possible database synchronization facilities for static base configuration variables, as well as mesh-synchronized data and configuration variables, according to an embodiment.

- 20 FIG. 20 is an example of a possible higher-level protocol on top of a cooperative mesh protocol (using a 1-to-1 acknowledgment scheme), according to an embodiment.

FIG. 21A, 21B show a general principle of acknowledging multiple packets via a
25 single coordinator or end node transmission, according to an embodiment.

FIG. 22A, 22B show a specific case of the multiple acknowledgment strategy which accounts for host-to-main controller communication bus delays, and how it can be used to implement point-to-multipoint handshaking and data transfers for
30 higher-level protocol data chunks segmented from arbitrary data "file" simultaneously between multiple nodes, according to an embodiment.

Detailed description of preferred embodiments:

Broadly, in accordance with an embodiment, there is provided a novel application
5 of a bit-synchronized mesh network leveraging the lower theoretical time bound
of a mesh network packet broadcast which is not limited by conventional
contention effects due to antenna diversity phenomena, as well as leveraging the
capture effect inherent of FM radio distributed networks as a fallback
10 transmission method in the case of heavy multi-path, delay-induced and physical
occlusion-caused interferences, in order to not only enable a standard master-
slave round-robin sensor polling scheme as seen in most distributed sensor
networks, but also to enable multiple bandwidth-asymmetric clients that can
replicate the data received by said master at no supplementary network
bandwidth or speed cost, and thus allows more leeway as for the choice in the
15 optimal power source for each node in terms of cost and maintenance
requirements. In addition, standard flooding scheme is hybridized with a slow-
flooding method called "redux phase", which allows for priority request signals
from end nodes to be flooded back to a coordinator node in a shared, contention-
enabled channel, so as to mitigate the time delay between node polling problem.
20 Moreover, a higher-level integration strategy improving on round-robin polling is
implemented using standard processor I/O interfaces in order to use the full
bandwidth capacity of the network using (but not limited to) a single processor
core that can time-share its computing capacity between mesh network
synchronization and configuration with ancillary functions such as sensor glue
25 code interfacing, thereby lowering integration costs, and allowing advanced
features such as exposing in-band microsecond-level synchronization capabilities
for specific application software living either within the mesh controller itself, or in
an external processing unit. In addition, a design strategy is devised in order to
enable legacy hardware to be bit-synchronized using specialized bus
30 communication strategies in order to interface to an external low-power sleep
controller, or even to use low-precision low-power internal clocks in conjunction

with a regular high-consumption high-precision base clock in order to achieve low duty cycles at low manufacturing cost. Furthermore, a scheme is devised where direct interaction from user interface dictates the polling order of the sensor nodes, in order to maximize network perceived responsivity by the user in a way that is most intuitive and specially suited for deployment using recent touch-enabled HMI paradigms. Finally, a specific implementation of such a mesh network is used to control a parallel high-bandwidth and high-consumption multi-hop network in situations where power is limited to a finite non-rechargeable onboard supply.

TERMINOLOGY

Following are some definitions of terminology used in the present description:

- a node, or sometimes "mote" if it includes a specific sensor implementation, is a network unit comprising of a controller processor and a transceiver, which allows to communicate with its peers in a distributed fashion within a mesh network.
- A coordinator node, also called "gateway", or "master", is the node whose task is controlling network-wide sleep cycles and communication slot assignment for other nodes.
- An end node ("slave") is any other node under the control of the coordinator node.
- A broadcast cycle is defined as a cycle, where packets are flooded throughout the whole mesh network from a coordinator node in what is called a broadcast "out" phase (a "phase" being a subsection of a cycle), and a given end-node floods back an answer throughout the mesh network in what is called a broadcast "in" phase. A standard half-duplex master-slave architecture will have two phases per broadcast cycle: one "out", and one "in". Thus, each network-wide flooding event is called a "phase".

- A packet is defined as a whole unit comprised of digital bit-encoded data sent and received from one transceiver to another. A packet is typically comprised of 2 sub-packets that we call "frames". Typically, a dedicated frame, called "config" frame, is used to transmit permanent mesh configuration data and thus is mainly dependent on the coordinator node configuration. Another dedicated frame, called "data", is used to transmit node addressing (typically a MAC address) and higher-level custom data packets, which depends on the implementation of the protocol layer lying over the physical layer, according to the OSI model. Finally, a frame called the "check" frame, is typically comprised of a packet integrity check method such as CRC. Such a packet structure is analogous to what is defined within the IEEE802.15.4 standard, the exception being that the "config" frame must also hold a critical flood sequence number in order to time-bound a given bit-synchronous packet flooding cycle.
- Each phase can be further subdivided into "slots", which acts akin to TDMA slotting mechanisms found in current literature, with the major difference being the following: the time division access is not used to grant multiple packet data transmission access to the network by multiple end-nodes within the same phase, but rather is employed as a division unit for network-wide flooding of packets containing the same data frame content.
- A cooperative broadcast flooding is defined as a flooding of a specific data "frame" embedded within packets by multiple simultaneous distributed antennas (i.e. network nodes). In the case of a carrier-synchronized flood, the nodes are synchronized at carrier-level, similar to phased-array schemes. In the case of a bit-synchronized scheme, the jitter tolerance of the synchronization is increased to encompass multiple simultaneous, potentially not carrier-synchronized, transmissions of a packet having the same bit-pattern, by said nodes. The bit-synchronization effectiveness threshold is usually defined, according to performance assessment in prior art literature, as roughly half the time for a bit transmission. Both methods

allow synergistic interaction between multiple end nodes, and this synergy is often labelled as "antenna diversity".

- Contention is a synonym for packet "collision". In networks that are not synchronized at the physical level, contention is a major issue affecting power consumption, and measures to mitigate its effect, such as packet acknowledge-retry back-off strategies, is a major ongoing research field.

ACRONYMS

Following are definitions of acronyms appearing in the present description:

ADC

Analog-Digital Converter

ACK

Acknowledge

ALU

Arithmetic Logic Unit

ASIC

Application-Specific Integrated Circuit

CLI

Command-line Interface

CSMA/CA

Carrier sense multiple access with collision avoidance

CTS

Clear-To-Send

DAC

Digital-Analog Converter

5

DB

Database

EEPROM

Electronically-Erasable Programmable Read-Only Memory

10

FCC

Federal Communications Commission

FHSS

15

Frequency Hopping Spread Spectrum

FM

Frequency Modulation

20

FSK

Frequency Shift Keying

GPS

Global Positioning System

25

GUI

Graphical User Interface

HMI

Human-Machine Interface

30

ID

Identification

IEEE

Institute of Electrical and Electronics Engineers

5

I2C

Inter-Integrated Circuit

I/O

10

Input/Output

LPF

Low-Pass Filter

15

LSBit

Least-Significant Bit

MAC address

Media Access Control address

20

MCU

Microcontroller

MIMO

25

Multiple Input Multiple Output

MISO

Multiple Input Single Output

30

NAK

Not Acknowledge

NFC

Near-Field Communication

5

PC

Program Counter

PCB

Polychlorinated Biphenyl

10

RAM

Random-Access Memory

RAMBUF

15

Random-Access Memory Buffer

RC

Resistor-Capacitor

20

RF

Radio Frequency

RISC

Reduced Instruction Set Computing

25

RX

Receive

OOP

30

Object-Oriented Programming

OS

Operating System

OSI model

5 Open Systems Interconnection model

PDU(s)

10 Protocol Data Unit(s). Those are the main exchange protocol data units for a given protocol level N in a multilayered protocol model (such as the OSI model).

PGM

Program

15 **PLL**

Phase-Locked Loop

QoS

20 Quality of Service

SDU(s)

25 Service Data Unit(s). Are the payload(s) of a PDU, transmitted unchanged to a peer within a given protocol level N . In a multilayered protocol model (such as the OSI model), the next protocol level $N+1$ encapsulates its protocol within the SDU of level N .

SFR

30 Special Function Register (usually in the context of microcontroller configuration memory registers)

SIMO

Single Input Multiple Output

SISO

Single Input Single Output

5

SPI

Serial Peripheral Interface bus

SMBus

10

System Management Bus

TCP/IP

Transmission Control Protocol/Internet Protocol

15

TDMA

Time-Division Multiple Access

TRX

Transceiver

20

TX

Transmit

UART

25

Universal Asynchronous Receiver/Transmitter

uC

Microcontroller

30

UDP

User Datagram Protocol

VM

Virtual Machine

5

WAN

Wide-Area Network

In accordance with an embodiment, there is provided a cooperative bit-synchronized flooding scheme, in which a regular broadcast phase is subdivided
10 by a number of non time-overlapping slots corresponding to the maximum number of repeating cycles (a.k.a "hops") that the network allows, in order to time bound packet flooding to strict times. This bounding action is enforced on the end node side via stopping further packet repeating if the packet to be repeated has a current sequence number embedded in the configuration frame which overflows
15 the maximum number of hops set for the whole network. In other terms, the information required to time-bound a packet broadcast flood is embedded as a sequence number in the configuration frame of all packets being broadcasted, and this sequence number can be mapped as an injective function of the number of times a packet has been repeated from the start of the flood, and the end
20 nodes rely on this signal in order to determine when to stop repeating any packet they receive. Every time a packet (which is typically a bit-stream) is repeated, it is so via a variable number of nodes, the number of which depends on network topology, in a bit-synchronized fashion (i.e. with time delay synchronicity better than $1 / (2 * B)$, where B is the bit-rate.

25

The number of time a packet is repeated after a given end node receives its original source can vary from as low as 1, to as high as allowable by the time-bound of any broadcast flood, and those can be spaced with inactive slots in between, mainly in order to diminish the risk of multi-path-related interference
30 due to time delays, among other things. This allows for a certain level of supplementary robustness, as long as the flooding time-bound can allow for such

packet repeated transmissions. Moreover, in the case where bit-synchronization fails due to multi-path fading, to time delays, etc., this allows a system to fall-back to the capture effect inherent of some radio transceiver systems due to their receiver carrier-locking designs. Note that a measure of variability in bit-synchronization jitter is included in order to increase the number of different possibilities of packet capture from different sources (this can be named "capture effect diversity"), and such sources can include (but are not limited by) the following random and pseudo-random sources.

1) Clock synchronization methods

Reference is made to FIG. 1A, 1B, 10A, 10B, 11A, 11B, 12, 13A, 13B, 14A, 14B. The high level structure of an end node mesh controller can be summarized as displayed in Figure 1A, and an example circuit plan is shown in FIG. 1B, whereas more details about the external sleep controller is shown in the FIG. 11A, 11B, 12, 13A, 13B, 14A, 14B.

In contrast with the standard clock synchronization method where the mesh controller IC that interfaces to a RF transceiver and its front-end circuitry uses a high-frequency, high-accuracy (and thus with high power consumption levels) in order to serve the processing power needs and microsecond-level interrupt accuracy needed for a cooperative bit-synchronized mesh network protocol, which disallows for low-power consumption sleep mode, and instead of using a very high quality, high accuracy, low power consumption and thus high-cost, second sleep-mode clock in order to increase the asymptotic effective duty cycle achievable by the system by reducing the guard time per broadcast cycle needed, the approach used, among other things, within the preferred embodiment of the invention, is to use a high-precision master clock in combination with (for medium to low duty cycles requiring an effective duty cycle clock synchronization accuracy of 300 ppm or less) low-cost but high-drift oscillators for managing sleep cycles, such as RC oscillators that can be

implemented at very low cost in ASIC, and are a standard feature in most modern microcontrollers. As those oscillators have typical accuracy an order of magnitude lower than what is required for lower duty cycles (assuming a ppm accuracy as mentioned above, for instance), a method was devised in order to
5 compensate for the lack of accuracy typically caused by thermal noise, temperature variations, etc., which can be summarized as follows.

A clock drift correction calculation mode is occasionally activated, where the RC clock runs while the master high-frequency clock is active, in order to have a
10 timing assessment of the RC clock current offset with respect to the higher-precision master clock. Those values can then be low-frequency filtered (the simplest method being an averaging filter), and a correction factor applied for a finite number of broadcast cycles, in order to account to external influences on a given clock frequency of operation, after which the correction algorithm is ran
15 again.

A trigger for this correction algorithm can be when a sleep-enabled end-node listens to a broadcast "out" packet from its coordinator node, and monitors the parameters which affect the sleep cycle timing. Note that the number of
20 parameters monitor could be zero, in the case where all timing parameters are hard-coded within each node. It calculates the expected sleep cycle delay according to those parameters.

Another trigger is a periodic re-activation of said correction algorithm, where a
25 sleep-enabled end-node uses a counter in order to monitor the number of broadcast cycles elapsed since the last correction cycle has been run. Typically the count threshold is set so that the ratio between the time between successive activations of said correction algorithm, and the time during which said algorithm is ran, is typically equal or lower to the effective desired corrected clock
30 tolerance, in order to minimize the power consumption impact of the correction cycle on overall effective network duty cycle.

If the duty cycle is set at the special case of 100%, only the primary high-frequency clock is used, and the only power-saving features are using mesh controller processor unit idle mode (for powering down any internal peripherals) when time-sharing is inactive and the mesh controller is awaiting a given interrupt from its radio transceiver peripheral.

Else, immediately after, or alternately before, a broadcast cycle, when a trigger activates the correction algorithm, the latter calculates the number of data points needed in order to increase the correction factor precision so that its tolerance become better than the desired corrected clock tolerance value, said number of data points (assuming an averaging window methodology) depending on the jitter of the interrupt signal that captures the length of the average RC clock tick in terms of master clock ticks, due to the combined jitter of the clocks.

15

In a higher-complexity alternate embodiment when the target processor used as mesh controller allows for it, the RC clock drift can be further divided into a random component and a non-random component that can be modeled as being dependent on external factors such as temperature and voltage, in order to pre-compensate for this jitter component and thus allow a lower correction time. Furthermore, the result from this model can be fed to a filtering component, and its nth-order derivatives (with a post-filtering component) used, as a trigger for activating the correction algorithm. This would allow to reduce the frequency of the periodic trigger, while maintaining the same effective accuracy.

25

Another embodiment might use a higher precision low-frequency clock source for managing sleep, either in combination with a high-frequency master clock in the same processor, leveraging the inline clock synchronization facilities used within that processor platform, or in a stand-alone manner, where an internal, PLL-locked higher-frequency clock is periodically activated during wake cycles in order to allow microsecond-level accuracy required for bit-synchronization, while

30

maintaining low power while in sleep mode. This is only possible when one uses a specific hardware platform that have multiple clock interfaces pre-integrated, and do not allow for the mesh controller algorithm to be ported to more generic industry-wide platforms.

5

In the latter case, when a pre-built radio platform with a in-built processor as well as a high speed clock, and when high clock accuracy during sleep is needed in order to reduce the minimum effective duty cycle of the network, is used as basis for retrofitting bit-synchronization facilities in conjunction with ultra-low power sleep modes in between broadcast cycles, the preferred embodiment of the invention then employs a novel synchronization method that can live within any standard synchronous communication bus between the chosen radio platform processor (typically used as mesh controller), and another dedicated IC used for appending low-power sleep mode capabilities, called "external sleep controller", and thus leverages existing bus capabilities of said hardware platform instead of consuming supplementary dedicated pin. This becomes a favourable design proposition when said communication bus is also tied to external sensors, for instance, as it reduces required PCB complexity and thus design time and cost per node.

20

For instance, in the case where a SPI style bus is used, with a set of control pins (typically called NSS pins) in conjunction with synchronous clock and data pins, and in the case where said control pin can deactivate SPI bus at whim, then the data pins can be recycled either while bus is inactive, or when the bus is officially communicating with said sleep controller (in which case even the microsecond-timing of the NSS pin or the clock pin signals can also be leveraged) in order to (1) embed microsecond-level clock timing information within the bus in order to clock synchronize the sleep controller with the master clock within the mesh controller, and (2) as an interrupt wake source from the sleep controller back to the mesh controller. As for transmitting known local clock time base of each

30

controller data back and forth, the regular SPI bus can be used in its regular fashion.

5 In the case where a I2C (or SMBus) style bus is used, with the mesh controller being the I2C bus master, then the clock stretching capabilities of the bus can be leveraged in order to transfer precise interrupt in the same fashion as for the SPI bus, while addressing the sleep controller using its own dedicated I2C address. In the same manner, the equivalent of a start and/or stop signal in the I2C bus can be leveraged by a slave sleep controller, which can pull those lines down to
10 local digital bus ground in order to force a precise wake interrupt to the mesh controller, which can then relinquish bus control after a specified time delay. In this manner, the bus properties stays within specification, as seen by other device users sharing the same bus (the only devices straying from typical bus operation modes being the mesh controller and its external sleep controller).

15

Typically, at least two successive wake signals must be given, the first used to wake the mesh controller, and the second in order for the mesh controller to synchronize its clocks according to its predicted value at wake-up. This is necessary for high-precision microsecond-level synchronization because at mesh
20 controller wake up, there is a variable and high-jitter time delay between the interrupt reception event by said controller, and the moment when its master clock is stable, thus the first interrupt cannot typically be relied on in order to sync the mesh controller clock at wake-up time. It would be possible to disregard this phenomenon and simply use one interrupt, but with an accuracy penalty that
25 would be linked to the accuracy of modeling clock wind-up times for the hardware platform chosen as a mesh controller.

The approach as embodied by both the I2C and SPI approaches above can be generalized to many synchronous bus topologies where signal pins can stray off
30 timing-wise in order to encapsulate clock sync timing information between both

controllers, without affecting the main communication bus protocol specifications as understood by any other hardware sharing that bus.

- 2) Mesh controller time-sharing with sensor peripheral integration firmware
5 within the same processor core

Reference is made to FIG. 2 and 3, as well as to FIG. 11A, 11B, 12, 13A, 13B, 14A, 14B. The specific location of event triggers used in the preferred embodiment can be visualized in figure 2, and an example VM flowchart is
10 provided in figure 3. An example VM client code leveraging this design strategy is shown in FIG. 15A, 15B.

Regular mesh controller architecture is typically using a low cost, low consumption, single-core processor linked to a transceiver ASIC, with must
15 communicate with another host processor in charge of the glue logic between sensor and control interfaces. This is typically due to the strict timing requirements of such a mesh controller depending on the mesh algorithm used, that disallows use of processor cycles while managing timing-sensitive signals, especially when the mesh controller must also already dedicate processor cycles
20 for local bus communication with another client processor while executing mesh-related operations. Moreover, for standard contention-enabled cost-sensitive mesh networks, which usually must also broadcast a synchronization packet in order to in-band synchronize the node clocks together for sleep mode, the unpredictability in flood times for high density of nodes due to contention, which
25 entails use of acknowledge/retry cycles that consume precious time spent in a high power-consumption mode, the mesh controller processor can typically not afford to run client-side custom parallel firmware for specific sensor integration, unless the time taken by said firmware is small enough so as to not affect and impede the operation of the network at low duty cycles. This can be of course
30 ensured in a case by case basis depending on the selected sensor or control external circuitry that are expected to be interfaced to the mesh controller, in

which case custom user sensor glue code is typically designed as being an integral part of the mesh controller firmware itself to ensure well-behavedness, in which case the client portion of the firmware typically be compiled and custom integrated by the IP owners of the mesh controller firmware itself, which doesn't
5 lend well to adoption of such a platform due to a lack of openness. Some specific mesh implementations, such as the popular Digi XBee™ with DigiMesh™ transceiver product line, embed for this purpose a second processor core on the mesh controller PCB dedicated to custom client-side sensor glue-code firmware development.

10

Because the preferred embodiment of the invention employs a cooperative bit-synchronized network, the broadcast flood process becomes predictable and can be time-bounded according to the optimal theoretical relationship between flooding speed and mesh network size. This means that in the case that a
15 network can be operated at a duty cycle less than 100%, there is always a free guaranteed and completely deterministic time slot corresponding to the expected sleep time between broadcast cycles, where processor cycles are typically only used for main communication bus management. This frees it for executing custom client firmware for longer time periods (a feature essential for long ADC
20 acquisition times inherent of many low-power sensors), which can stay well-behaved as long as it respects the guaranteed time bound. It will also ensure that no parallel code execution operates while the mesh-related operations are pending completion, as this would typically involve using interrupts in order to emulate a multi-tasking environment, which would jeopardize the tight timing
25 tolerances required for some critical operations of the bit-synchronized node (such as for transmit begin trigger signals, for instance).

Thus, one can leverage those free processor cycles during this time period while safeguarding the proper operation of mesh-related firmware within the same
30 processor core, using this the following strategy.

The mesh controller sets up event trigger points related to mesh network operation, such as (but not limited to) the following: (1) controller boot-up, (2) entering mesh seek mode, when a node stays in receive mode in order to lock itself and its local clock to the mesh network in which it is intended to operated, 5 (3) leaving mesh seek mode, (4) entering broadcast cycle, (5) leaving broadcast cycle; those event triggers can be activated or deactivated according to a stored configuration data structure residing in the controller non-volatile memory.

10 In a possible embodiment of the invention, those events can also be captured using on-board processor facilities, at a microsecond or sub-microsecond accuracy, and stored in a dedicated section of memory for later use by client code that might need to use a common in-band time base for time-critical sensor operations, such as acoustic phase arrays, for instance.

15 An architecture where user code can only sporadically run as time-limited runtime callback functions after specific events is used (in contrast with continuous user code running), said time-limit depending on the nature of the event. For instance, a time limit as long as the sleep time between broadcast cycles can be used for event 5 (leaving broadcast cycle), can be relatively long for event 1 (boot-up), but 20 should be very short for event 4 (entering broadcast cycle), for instance. This ensures that user code execution be compartmentalized out of the mesh network operation broadcast cycle, therefore interfering with the mesh network precise synchronization. Evidently, for single core processors, this implies that user code cannot safely run as natively compiled code, because then there is no guarantee 25 of the well-behavedness of such user code.

Note that another possible embodiment of the invention can discard the use of time-limits, in which case the network can operate, but with possibility of failure mode due to user-code, depending on the degree of engineering safety that one 30 desires to achieve.

To solve this problem, a dedicated non-volatile memory section is also dedicated to user custom firmware in the form of a virtual machine (VM) bytecode optimized for low-speed low-memory-footprint embedded targets, with the VM architecture being set up so as to trigger separate user code sections according to the mesh controller events as listed above. Note that for single-core targets such as low-cost 8-bit microcontrollers running at typical MHz clock speeds, using a VM usually entails a significant loss in processing power, as a single bytecode opcode being executed can take up to milliseconds to complete (depending on the VM architecture chosen, the most common design choice being to use a RISC VM architecture which sacrifices opcode employment efficiency in order to save precious non-volatile memory space dedicated to VM processor implementation), but due to the deterministic nature of the sleep period between broadcast cycles, such an architecture is usually sufficient to even implement low-efficiency bit-banging of a non-standard sensor peripheral, for instance, all within the sleep cycle between broadcast cycles, even for low maximum number of hops per broadcast phase settings (which decrease the sleep time for a given duty cycle setting).

A known failure mode when using time-shared user-code in the same processor core as the mesh controller, consists of user code overflowing past its allotted execution time slot, which would jeopardize mesh controller timing integrity which is so critical for the good operation of a bit-synchronized network, and can also push the mesh behaviour past its operational parameters as allowed by standard regulating bodies such as FCC. This is one of the reasons why a dedicated user-code core is usually employed in a user-customizable mesh controller as a safeguard measure (short of employing dedicated ASICs for the task). However, the use of a VM architecture allows to address this critical concern for bit-synchronized networks by ensuring that each executed VM bytecode opcode, independently of its nature (be it a ALU operation, an I/O operation, timing peripheral operations, etc.) and expected time of execution completion, be bounded to a deterministic timeout using either a low-level interrupt strategy, or a

more conventional polling strategy for longer custom operations (such as VM delay functions, for instance). Moreover, between each opcode, the processor can then check whether the operation time since the start of a VM execution overflows its allotted time, and then shuts down the VM engine before a
 5 broadcast cycle can begin (and thus operates similarly to a VM watchdog, or a VM hypervisor).

3) Use of multiple bandwidth-asymmetric clients, i.e. upgraded "sniffers"

10 Assuming that the network mode of operation contains at least 1 broadcast out and 1 broadcast in phase (typically symmetric in terms of time taken, and number of effective data bits carried in each data frame of the packets flooded in both directions), this ensures that there is permanent possibility for two-way communication between each end node and the main coordinator node.
 15 However, due to the very nature of a bit-synchronized network, in which every data in transit is flooded to all nodes in order to increase antenna diversity, it is possible to leverage this effect in order to ensure that system-wide database be replicated at different client locations using the same mesh network technology, via sniffing a packet from any location that is connected to the mesh network,
 20 instead of synchronizing them via external network such as cellular towers, for instance, and without hefty power consumption penalties.

Consider a typical sensor management architecture, where a centralized database is managed by a processing unit, which also commands the
 25 coordinator node, and in which every node is registered with the following information (not all-inclusive):

- Configuration variables, which can be either used only by the database management unit, or by the sensor nodes themselves (can contain sensor control variables, node warning/error thresholds, etc.)
- 30 • Data variables (exclusively consists of data send back by end nodes)
- Node network state (ex: online, dead, seeking sync, etc.)

- Node position (this can be neglected if positioning is done for the majority of nodes in an offline fashion, as is the case when someone uses a barcode reading apparatus that contains a GPS unit in order to jot down its position and to register it in database at scan time, in which case
5 dynamic node positioning for specialized sensors or client "sniffer" nodes can be accomplished instead within the data variables above)
- Last good node data reception timestamp

10 In this case, it is desirable to have at least the following sections of the database to be synchronizable via the mesh network at maximum potential speeds, in an UDP fashion (i.e. with no direct acknowledgment mechanisms). This is allowable because sensors are periodically polled, in which case a failure to receive can be corrected after a certain delay corresponding to the time interval between successive polling of the same node address:

- 15 • Configuration variables
- Data variables
- Node network state

20 Less critical sections can be instead synchronized using more conventional methods where the coordinator node sends specific data streams to a node with an acknowledgment scheme, which is less bandwidth efficient, but is sufficient because those sections usually don't need to be updated often.

To do so, a strategy, where all coordinator-bound, user-controlled variables, i.e.
25 those typically modifiable via user HMI, or that need to be resilient to end node failure, such as node network state, is broadcasted to the whole mesh network during the broadcast out phase by the coordinator node. In contrast, all data variables are sent from the specific end node that owns them via a broadcast in phase. Because every packet is of the flooding type, the result is that all variables
30 that need to be synchronized in an efficient manner are available at some point via a packet broadcast flood event, no matter which direction this flood occurs.

Thus, a client user interface, if it is configured with the same database infrastructure as found within the main processing unit that controls the coordinator node, as well as the same software decoding facilities, and is linked via a communication bus to an end node that acts as its network bridge, can then
5 reconstruct the database time-sensitive variables by simply "sniffing", or listening without necessary interaction with the coordinator node or their neighbours, thus allowing for unlimited number of listeners that do not impact mesh performance.

In the case that such a client node needs to send control commands back to the
10 coordinator node (for instance, if a user in the field interacts with an HMI bound to it, such as a smartphone), the default method is for it to wait until a specific outbound request from the coordinator node attempts to access its end node data, in which case it can send a request for priority polling. Then, the coordinator node knows that this specific client need to send database modification requests
15 back to the main coordinator location, and polls the client end node periodically, using a higher-level acknowledgement-based protocol in order to ensure delivery of those updates, until it stops receiving those priority polling requests, in which case it reverts to UDP-style regular sensor node polling. Note that details related to this priority polling scheme will be described in more detail in the following
20 sections.

This is why this strategy is referred to as "asymmetric clients" strategy, because the response time is widely asymmetric (near real-time when a client is listening to broadcast packets from other nodes, and at the whim of the current
25 coordinator node polling strategy if it needs to send back data to the mesh network). This strategy can be furthermore improved using what is named as the "redux" broadcast phase, which will be explained in the next sections..

Thus, the simplest embodiment of such a synchronization scheme between
30 asymmetric clients and a node coordinator is a system in which the software and hardware design of the processing units of both clients and main database

manager linked to the coordinator node, be nearly of the exact same flavour. Thus, in the preferred embodiment of this invention, a smartphone/tablet platform is chosen for the following reasons:

- 5 • Cost-effectiveness with respect to competing platforms with similar processing power;
- Possesses sufficient storage capacity for logging as well as database management, so that it can acts seamlessly as client or main database manager;
- 10 • If a generic wireless standard is used to bridge this computing platform to a mesh controller is used, such as with Bluetooth™, then the same communication method can be used for all devices notwithstanding their nature as a client or database manager;
- Same applies with software: as long as the same version of a software is used, which typically includes low-level mesh control facilities for error
15 checking, database management, and can also leverage the smartphone/tablet natural HMI interface, bit-stream depending on node type decoding is insured to be the same notwithstanding their nature as a client or database manager;
- A smartphone/tablet usually possesses localization functionalities
20 (typically GPS), which can be used for out-band registration of node position at deployment time, therefore avoiding the need of dedicated positioning hardware in each node if such a functionality is desired for its intended application, therefore lowering costs;
- A smartphone/tablet usually possesses a natural registration mechanism
25 via camera vision system (that can be leveraged for barcode reading), which can be leveraged at node deployment time in order to facilitate node address and type registration while in the field, assuming that the required identifier be included with each sensor hardware (typically a barcode);
- A smartphone/tablet possesses all required functionality to sync the
30 databases together, via NFC, wi-fi, via the web, etc. So that database variables that are not synced via network flooding can be ensured to be

the same from device to device in the most simplest embodiment for database synchronization possible.

- Possible upgrades to asymmetric client "sniffers" can also be: the implementation of a sensor-type answer mechanism for those client end nodes, so that a coordinator node can consider them as regular sensors, with useful data such as "sniffer" positional data, and "sniffer" battery state. This allows an operator at the coordinator site to monitor the activity of end users employing those "sniffers", such as for instance an employer monitoring his/her employee activity within a forest plantation such as a sugar bush, for instance.

Note that any HMI does not need to be necessarily on premise with the main coordinator node mesh controller processor: a device could bridge the hardware so that control is done via the cloud using standard web-server technologies.

One possible embodiment leverages generic cloud-based file storage systems in order to enable smartphone/tablet platforms to sync themselves via the cloud, without using application-specific cloud server solutions for this purpose. Again, this leverages the strength of the smartphone/tablet platform, and enables full user mobility, whether it is while roaming within the mesh network, or within access to WAN facilities.

Listed below are example implementations of sensor nodes and their corresponding data/configuration variables in compliance with the strategy mentioned above in order to enable full synchronization of said variables during regular end node polling between a coordinator node and any asymmetric "sniffers". The following examples will be detailed: vacuum sensor, fluid level sensor, and relay multifunction controller.

In order to clarify the context and use case typical of those examples, application of those sensors in the realm of agricultural industries such as sugar making is described as follows: mesh-enabled vacuum and/or temperature sensors are

deployed in the maple orchard, and typically monitors the vacuum level at the end of the collector lines (also known as "master lines"). Those sensors can also be monitoring the vacuum and the positive pressure at the tap locations per tree. This allows to find leaks and to target them in order to raise global vacuum tubing network efficiency, and thus crop yields. Wireless mesh-enabled sensor nodes can also be used to monitor tank levels, to check and control equipment such as reverse-osmosis machines, monitor their electrical status, etc.

For such an application, the environment is relatively hostile to regular energy harvesting strategies such as solar panels. Although those strategies can be used, the high energy-efficiency of the mesh network, for the maximum power output per node, allows use of low-cost primary batteries such as industrial-grade alkaline batteries, and this allows seasonal use for many years in a row, even for a large deployment in the range of hundred of nodes.

VM machines can be implemented for the following sensor examples with the following data/configuration variables:

Vacuum sensor

Data variables (for N pressure ports, implemented for 1 to 3 sensor ports):

- Vacuum (16 bit) x N
- Temperature (16 bit)
- RSSI (8 bit)
- Voltage (8 bit)

Configuration variables:

- Low vacuum error threshold (16 bit) x N
- Low vacuum warning threshold (16 bit) x N (optional)
- Zero calibration (16 bit) x N
- Low temperature warning threshold (16 bit)
- High temperature warning threshold (16 bit)

Level sensor

Data variables

- 5
 - RSSI (8 bit)
 - Type code (4-8 bit)
 - Raw level measurement (8-12 bit)
 - Error code (8 bit)
 - Voltage (8 bit) (optional, for battery operated devices only)
- 10 Configuration variables:
 - Zero calibration (16 bit)
 - Level range (16 bit)
 - Warning high threshold (16 bit)
 - Priority polling threshold (16 bit)
- 15
 - Warning low threshold (16 bit)

Relay/current sensor

Data variables (for L switch inputs, M control channels, N current sensors)

- 20
 - RSSI (8 bit)
 - Control channel current state (1 bit) x M
 - Mechanical switch state (1 bit) x L
 - Local control force database update detection flag (1 bit)
 - Local bypass detection flag (1 bit)
- 25
 - Current sensor channel (16 bit) x N
 - Current change counter (24 bit) x N
 - Mechanical counter (24 bit) x L

Configuration variables (for L switch inputs, M control channels, N current sensors)

- 30
 - Expected (requested) control channel state flag (1 bit) x M

- Reset current counter flag (1 bit) x N
- Reset mechanical counter flag (1 bit) x L
- Current transition hysteresis high threshold (16 bit) x N
- Current transition hysteresis low ratiometric threshold (8 bit) x N
- 5 • Current error threshold high (8 bit) x N
- Current error threshold low (8 bit) x N
- Feature enable bits

For all three sensors, a common set of network states can be used, the currently
 10 selected state for a given node being always sent out via broadcast out from the
 coordinator node as mentioned before. A typical set of network states can be as
 follows:

- OK;
- Dead node (indicates that a given end node has failed to respond to a
 15 coordinator node request after a few retries);
- Network Glitch (indicates a recoverable network error caused by the
 communication bus between host computer and the mesh controller of a
 given coordinator or "sniffer" node);
- Network Critical (indicates a critical error);
- 20 • Seek Mode (indicates that a "sniffer" is currently seeking synchronization
 of a network and is thus not connected to said operating network yet);
- Connected But Silent (indicates that a coordinator is currently
 synchronizing sleep of a network of nodes, but no request for data
 transaction between said coordinator and end nodes is being broadcasted
 25 out);
- Network Contention (used to indicate when an overloaded network results
 in data older than a certain critical time threshold);
- Buffer Overflow (used for gateway or "sniffer" with buffer overflow issues).

30 Note that the relay controller example above needs a data variable called "local
 control force database update detection flag" in order to ensure that any change

in the local state of the end node, such as a change in the "Control channel current state" variable due to local user interaction is not overridden when the corresponding end node receives the current configuration variable "Expected (requested) control channel state flag" , which is broadcasted out by the coordinator node every time the end node is polled and its data variables requested . To do so, the method is as follows:

- The default state is this: anytime the "Expected Control Channel State Flag" is received by the end node, it transfers this variable to "Control Channel Current State" and activates its digital outputs accordingly;
- If local modification of those digital outputs (typically via control switches) is detected, then the node goes in "lock" mode: any configuration variable subsequently sent by the coordinator node is ignored by the end node by default, and the "Local Control Force Database Update Detection Flag" data variable is set and sent back to the coordinator in order to notify the host device controlling that coordinator and holding the main node database that it must update its own configuration variables within its database in order to reflect the changes that was made locally to the end node;
- When the host connected to the coordinator node detects this condition, it updates its configuration variable in its database accordingly;
- The next time the coordinator polls that end node, the "Control channel current state" data variable as stored internally in that end node will match the "Expected (requested) control channel state flag" variable it receives from the coordinator node. It then knows that the database at the coordinator site has correctly updated itself, and thus releases its "lock" and goes back to normal mode.

The locking mechanism described above can be used in a similar fashion for any end node that is controlled by configuration variables and whose state can change locally by some mechanism, such as user interaction.

FIG. 16 shows the packet subdivision for the data/configuration variable paradigm as described above. Reference is made to Fig. 17 and Fig. 18 for the state machine details related to the implementation of such a strategy. Fig. 19
5 shows the high-level details of such a strategy deployed on said Android™ platform.

Some example applications of such a strategy include (but is not restricted to) the following:

10

(1) Generic agricultural and science research applications including:

15

- Maple orchard management, regular industrial monitoring
- HV monitoring of fences
- Monitoring of level in grain/food silos
- General meteorological monitoring (weather stations)
- Etc.

(2) Fleet management

20

It is possible to deploy an array of static end nodes in a city, as well as end nodes in moving vehicles to be tracked (such as a bus fleet) via GPS, which can be especially useful as a low cost alternative to the usual cellular or satellite deployments. The mesh technology would allow fast polling of the whole array of vehicles without unexpected delays due to contention arising from constant
25 changes in end node positions. Static asymmetric "sniffers" can be located in public transportation waiting zones, and, with database replication in place, this will allow efficient informing of the node positions to users of said public transportation system.

30

(3) Monitoring of industrial tubing systems

Sensors deployed in a milieu that can act as acoustic waveguides can leverage the power efficiency of the mesh network. In such a case it is possible to leverage the tubing network for nodes to harvest energy either from an active sound source, or even from the acoustic energy contained in the tubing network from flow of fluids itself.

It is also possible to implement a bit-synchronized network using acoustic transceivers instead of radio transceivers in such a system, the limited bandwidth for transmission channels in such a case being solved by the fact that bit-synchronization allows a large quantity of nodes to use the same transmission channels without undue interference.

(4) Acoustic triangulation using in-band synchronization

Because all sensors are synchronized to a precision higher than half the length of a RF bit transmission itself, which typically is in the order of microseconds, this allows the use of in-band synchronization for acoustic triangulation of signal sources, either via phased array strategies, time of flight measurements, etc.

Examples of use of such a method:

- Gunshot detection and localization for military or police applications;
- Ultrasound leak detection and localization in harsh and wide-ranging environments (such as in a maple orchard);
- Detection and localization of objects or people with acoustics tags (for example, usable in amusement parks, ski centers, and as a safety measure so as to detect and respond to emergency situations), the active acoustic emitter component being situated on the object to be localized in question, with the array of microphones on static end nodes deployed in the field. Alternately, detector microphones with a transceiver can be on the object, with the active acoustic emitter component being on the static end nodes instead..

(5) Drone swarm control

- 5 Because of the predictable nature of the mesh network, as well as the "*redux*" priority polling request capability, this system can be used to connect a swarm of drones and also ensuring high bidirectional responsivity to each individual in a drone swarm.

10 4) HMI retroaction to mesh poll scheduling by coordinator node

In contrast with regular sleep-enabled low-power contention-enabled mesh network architectures which usually attempts to poll multiple sensors in a dedicated time interval after a long period of sleep (for high density of sensors),
15 using a bit-synchronized cooperative architecture allows individual periodic polling of sensors registered in a given database at much closer interval for the same equivalent power consumption. Moreover, said intervals are guaranteed to be of an exact length defined by the bit-synchronized network global configuration parameters such as the maximum hop count, the number of
20 broadcast phases, etc. Therefore, it is possible to emulate near real-time like responsivity of a sensor if needed, by prioritizing its polling time with respect to certain triggers. A possible implementation is as follows:

- Maintain a priority column in the sensor node database, with allows database sorting according to the priority queue number.
- 25 • The polling order is determined by the order in which entries results are listed after a query where the primary sort key is the priority queue number.
- When a node is being polled actively, its corresponding priority queue number is reset to the default value.

Then, the following triggers are set up so that they increment the priority level queue number by a certain magnitude according to the severity ordering of said trigger (depending on specific application requirements):

- Sensor last known data age above a certain threshold that depends on sensor sensitivity to polling speed. For instance, in a specific application where a level sensor monitors a water column that can spill within minutes, this threshold can be set lower in order to ensure that the sensor data be refreshed in a closer fashion than lower-priority sensors such as weather station sensors, for instance.
- User HMI interaction with a specific sensor. For instance, if a smartphone interface is used, the main database management processing unit can detect whether the user has pressed a button in order to see detailed data for a specific sensor, and/or for controlling said sensor or changing its configuration parameters. In this case, the priority queue number is again incremented typically by a higher value than for the trigger above, which results in the network effectively priority polling that sensor in particular. This allows user to interact with that particular sensor as if it was in near-real-time conditions (the responsivity lag being at most $2 \cdot X$, where X is the time interval between 2 successive broadcast cycles).

Moreover, multiple duty cycle network settings can be held by the processor in charge of database management, so that it can switch to higher duty cycle according to specific sensor warning or error conditions, and according to their severity ranking, as well as switch to higher duty cycle if it detect sensor-specific HMI interaction as above. This operates in synergy with the database priority ranking scheme as explained earlier in order to ensure that X be at the lowest allowable setting during sensor-specific user interaction.

An example is as follows: consider a database management unit that holds 2 duty cycles, one labelled "slow", and the other labelled "fast". The default duty cycle setting is "slow". Say a level sensor monitors a water column and detect

that water level is at 75% of full capacity, and sends a warning signal. Then, the duty cycle can be switched to "fast" and will only scale down to "slow" when this signal disappear (to avoid noise-triggering, standard hysteresis methods in the configuration parameters can be used).

5

5) Multi-phase broadcast modes, hot-configurability, and polling speed optimizations

Reference is made to FIG. 4, 5, 6 and 7.

10

For a bit-synchronized network, the bare minimum requirement to synchronize the node local clocks to the coordinator node reference clock is for said coordinator to send a minimum of one packet broadcast out per broadcast cycle (one can call this a "beacon"), which minimally could contain no information other than what is required for the end nodes for time-bounding the packet flooding throughout the network (i.e. at least a sequence number). In such a case, no information can flow back from an end-node to the coordinator node, and thus is of limited utility for sensor polling. However, it is a possible mode of operation that, if there is a method to activate or deactivate it, can allow for the system to stay in "beacon standby", in order to maintain node sync with the network while reducing global power consumption by a factor of 2 when node polling is not needed by user.

In order for this to become useful, a supplementary bit sequence must be included within each "config" frame of each packet broadcasted out and in (both directions are desired so that a node can seek network sync independently of the latter current broadcast phase), encoding the current configuration of the number of broadcast phases desired, among with other dynamic configuration parameters if more configuration malleability is required, including (but not limited to) the number of hops, for instance, and in the case where a FHSS system is used, a mandatory bit sequence encoding the current frequency hop status of the

30

network, the simplest way being to encode the current hop table index of the network within each packet. This allows the whole mesh network to be hot-configurable as soon as a coordinator node packet is broadcasted out and received by all recipients in the network. The packet structure and its configuration frame is shown in Figure 4 (along with the features that will be explained below).

This principle can be further extended so that the "config" sub-frame of any packet encodes the number of phases out, as well as the number of phases in, therefore allowing asymmetric bandwidth capacity according to flood direction. Implemented naively as conventional contention-enabled multi-hop networks might be, this runs the risk, if the number of phases per broadcast is sufficiently high, of bit-synchronization failure due to the local master clock tolerances which varies from node to node. This is the reason why prior art implementations of bit-synchronized networks do not attempt more than transmission of small packet sizes within a maximum of 2 broadcast phases (one for each direction). However this bit-synchronization potential loss of accuracy as the number of phases grows can be alleviated in the following way within each end node.

Because the main time base onto which the whole network latches to is that of the coordinator node, there is an opportunity for using each broadcast phase originating from said coordinator node in order to readjust the local clock of each end-node, but not if the broadcast phase originates from an end-node (because in such case, path delays and accumulated clock drifts.

25

In the case where the accumulated expected drift within a broadcast cycle due to successive broadcast in phases (which do not allow for clock sync reassessment), a maximum bound for the number of such successive broadcast in phases N_l is set according to the expected worst case effective relative clock frequency drift C (defined as absolute frequency drift / frequency), the RF bit-rate B and the time per broadcast phase T_P : $2 * B * T_P * (N_l + 1) < 1 / C$.

30

In the case where higher number of broadcast in phases per broadcast cycle is needed, then there must be a broadcast cycle out interwoven between broadcast cycle in, at a maximum phase spacing corresponding to N_i .

5

For a specific embodiment, if there are more than 1 broadcast out phase used, then for any given measurement index i , the time $T_{IBO}[i]$ between successive broadcast out sync events can be measured, then a correction factor $F[i]$ for the end node clock with respect to its coordinator clock can be approximated as follows: $F[i] = \text{lpf}(T_{IBO}[i] / (N_i[i] + 1))$ value at the time of measurement, where lpf is an implementation-specific low-pass filtering function. Measurements can be free-running between broadcast cycles, with the lpf time constant chosen accordingly, in order to get a noise-reduced assessment of the clock correction factor needed locally for each end node. The standard deviation S_F of this correction factor can be either be qualified offline, or measured in line simultaneously with its actual value, and (assuming the correction factor is applied continuously to the node local clock) its inverse $1 / S_F$ value can be used to replace the $1 / C$ factor mentioned above, which would allow an increase in N_i before bit-synchronization actually fails.

20

In the preferred embodiment of the invention, which uses a FHSS transceiver, the following parameters can be included in the "config" frame in order to enable hot-configurability (as either a table index for saving bit space, or directly as an integer):

25

- network ID
- sequence number
- max sequence number (hop count)
- frequency hop index
- broadcast in count
- broadcast out count
- "redux" phase enable/disable

30

- "redux" phase slot count

Note that the last 2 items will be explained further below. The high-speed master clock error alleviation strategy is shown on Figure 4, at the "regular broadcast level", and all those features are listed in said figure.

Each broadcast phase out is dedicated to a coordinator node transmission, and each broadcast phase in can be used in a TDMA fashion, where a single end-node sends back a packet for its assigned broadcast phase. A standard way of operating in this manner would be to interweave successive broadcast out and broadcast in phases together in order to keep a regular half-duplex architecture, in which one end-node request from the coordinator is sent along with ancillary information (using remaining data bits within the data frame of the packet), and a requested end-node scheduled for reply sends back data in a half-duplex manner.

It is also possible to use a standard SIMO, MISO or MIMO strategy, where a single broadcast out packet from the coordinator contains multiple requests to multiple end node addresses (packet data frame segmentation into independent data sub-frames, in other terms). However, because bit-synchronized networks involves small maximum packet size in order to minimize possible interference due to bit-synchronization jitter between network nodes, this is often not possible unless the number of data bits used per request is small enough.

For instance, for a MISO setting of 1 broadcast out for 4 broadcast in phases, each broadcast in being assigned to a specific node, with a maximum of 20 data bytes in a data frame, and with node addressing taking 4 bytes (a 32.bit MAC address), then there will only be a maximum of 1 byte available sent to a specific end-node, which is often not enough if more than basic requesting is used (as is usually the case for sensor nodes that must also control devices, for instance.

Moreover, supplementary information destined not for one particular node, but system-wide (used for upgraded "sniffers") also requires data frame space.

To solve this data frame space problem, a possible embodiment (assuming to be of the MISO type for explanatory purposes) might use the following strategy, in the case where a coordinator node request packet also contains configuration data to destination end nodes (for instance, error thresholds, output states, etc.).

In preparation for a coordinator node request packet broadcast out phase, assuming that N_i broadcast in phases are used, the coordinator scrolls its node configuration database, which typically consists of node addresses, node sensor type, and individual node sensor configuration parameters. The coordinator fetches the first N_i number of nodes whose node sensor type and configuration parameters match exactly. This ensures that exactly the same configuration data bit stream be generated for all those nodes. In the case where insufficient nodes results from this query, the coordinator can simply leave unused scheduled broadcast in phases empty (by assigning an invalid address node to that phase).

The coordinator node fetches any system-wide data (usually for broadcasting network state to all nodes) that also must be sent, and encodes it in the proper bit-stream format.

The coordinator node encodes the request node addresses, and binds them to specific byte positions within the data frame of the request packet, the exact position of the address bit stream determining which broadcast in reply phase is to be used (this is akin to a TDMA strategy) by the end node possessing that address.

The coordinator node then concatenates the address bit-stream with the common node configuration bit-stream, and then calculates the remaining number of bits available in the data frame.

In a possible embodiment, where the system-wide data bit-stream has variable length, it must also append a length integer to said bit stream, as well as a sequence identifier, so that such data can be segmented and sent via multiple
5 packet transmissions, within the free remaining bits of the data frame of each packet sent (only if the data stream is too wide for the data frame size of a single packet).

In an alternate embodiment, the delimitation of the system-wide data bit-stream
10 can be done using a specific set of reserved bit markers, and the rest of the bit-stream data be encoded in such a way as to not overlap with those reserved bit markers (an example of this would be using BASE64 encoding, with an end-of-line marker, for instance).

15 Using the aforementioned strategy, in the case where the variability between node sensor types within a deployed mesh network is low enough (which is usually the case for the vast majority of mesh network applications), this allows for a maximum speed-up in sensor polling of a factor equal to N_i , with respect to a conventional polling scheme, and allows simultaneous system-wide data
20 transmission from the coordinator node in conjunction with specific end node requests within the same transmitted data frame. See Figure 7 for supplementary information concerning this strategy (note that it assumes that the half-duplex issue is solved, using what is also shown on Figure 6 for the simple case where multi-phase mode is not used).

25

Another issue is that between a request from a coordinator node via a broadcast out phase to a specific end node address, and said end node answer in a broadcast in phase, one has to plan for a non-zero time delay interval between both because in practice all the following events take non-zero time to complete:

- packet reception event in an end-node (this can occur anywhere during a coordinator node broadcast out phase, worst-case will occur at the end of said phase);
- packet information transmission to processing engine (be it an external processor, a communication bus such as UART, or a VM engine);
- processing engine request to sensor hardware delay (due to chosen communication bus, say a I2C bus, or simple digital I/O control pins or analog DAC/ADC interfaces);
- hardware measurement and control delay (for a typical digital sensor, this is related to ADC frequency response and additional processing within the sensor itself);
- data send back from sensor hardware to processing engine;
- processing engine processing and send back to mesh controller for answer packet construction delay.

The usual solution for a bit-synchronized network is usually to simply discard the possibility for multi-phase broadcast cycles and restrict to the standard 1 out / 1 in phase (SISO), so that any variable and unpredictable delay caused from any item listed above, do not cause bit-synchronization failures due to contention caused by an end-node that attempts to flood the network at inappropriate times (i.e. when another is legitimately attempt to flood the RF network). This is usually combined with an approach where a gateway continues to send clock synchronization beacons, but without a data frame payload, while waiting for official receipt of an end node answer, before attempting to send another request to another node, or after a timeout delay. However, using such an approach, the maximum potential polling speed for a given chosen duty cycle is literally halved. (see FIG. 5)

To solve this data frame space problem as well as the reaction delay problem, a polling strategy is via a delayed MISO strategy, meaning that the coordinator node does not receive an expected end node answer in the same broadcast

cycle as that in which the corresponding request has been sent. One might naively assume that this delay corresponds to one sleep interval between broadcast cycles, as seen from the coordinator node standpoint. However, in practice, the delay due to communication between the host processor that manages the sensor database (such as a computer) and the actual mesh controller used as coordinator node also has to be taken into account, in order to make sure that it does not cause similar contention issues as for end nodes. Thus, the minimal interval between a request, and the corresponding expected answer, as seen from the coordinator node standpoint, is 2 sleep intervals between broadcast cycles, and this interval can be increased if needed, depending on the speed of the host processor in charge of database management.

The strategy can be listed step-by-step as follows, with reference to FIG. 6.

15

Every time a broadcast cycle ends, send a marker indicating such event through the communication bus between the controller and the processor in charge of database management (if both are the same processor, then this marker is a software marker) to the latter. As an example, the preferred embodiment uses a default UART with CTS interface using a higher-level command and reply encapsulation of variable type (it can be a bit-stream frame with a start-of-stream marker, length bits, possibly error correction codes, but also ASCII frames, or possibly BASE64 coding, for instance). In such a case, a dedicated bit-stream marker is used to denote the end of broadcast cycle event, and the host processor linked to the controller via UART is thus notified within milliseconds of such an event.

25

The host database manager processor constructs the request data frame that is to be sent to the coordinator node mesh controller. Note that this also causes additional time delays that can be caused by typical computer OS multi-task scheduling, for instance, as well as the query and frame encoding time.

30

The delays above are accounted for and nullified using the following strategy:
when the mesh controller receives the host database encoded data frame, it
stores this, and waits for the next broadcast cycle begin event, thus removing the
5 timing uncertainty.

During that event, the stored data frame is encapsulated in a coordinator packet,
in the manner described in the aforementioned sections, and sent out. Thus
when the broadcast cycle end event occurs, it is guaranteed that the request
10 packet has been broadcasted throughout the whole network. At this time,
assuming packet reception is successful, the end node receives the packet and
decodes its data frame. It then sends out this data to the processor in charge of
controlling and fetching data from the actual sensor circuitry (be it an external
processor, or a VM engine, as used by the preferred embodiment), the latter
15 accomplishes its task, and the resulting answers are sent back to the mesh
controller, all accumulated delays incurred by this process being accounted for
and nullified by ensuring that it runs within the sleep delay between broadcast
cycles (in the case where a multi-core processor or external processor is used as
the glue logic interface, then it is possible to spread the sensor processing delays
20 over more than 1 sleep interval between successive broadcast cycles, but in the
case of a single-core callback-based VM engine, this is often not possible, unless
the user is in charge of manually and tediously managing this delay process).

At the next broadcast cycle event trigger, the end node broadcasts its answer
25 packet back to the coordinator node. Assuming that packet reception by
coordinator node is successful, the data is then transmitted upstream back to the
processor in charge of database management for updating (this is not time
critical because it is a receive only operation that occurs while the broadcast end
marker strategy runs in parallel in order to construct successive requests.

Assuming that all parties involved in a successful polling action (host processor that maintains database, mesh controller used at node coordinator, and end nodes, including all communication buses), as enforced using the above methodology, the sensor polling rate is no more halved with respect to a given
5 duty cycle setting. Thus, effective throughput is increased by a factor of 2.

As for error handling (such as when an end node fails to answer), due to tight restrictions on the timing of the bit-synchronized mesh protocol, its onus thereby lies on the higher-level protocol according to the OSI model. In other terms, the
10 processor in charge of database management (the application-specific host processor) is expected to retry querying a sensor flagged as unresponsive using the same methodology as listed above for a number of times, and then register a sensor node error after the number of retries have unsuccessfully been carried out. In other terms, the strategy is optimal in the case where packet delivery does
15 not have to be assured, and is akin to a UDP communication scheme.

6) "Redux" phase(s)

Reference is made to FIG. 8 and 9. The "redux" strategy is summarized in
20 FIG. 8, and an example of a possible collision management scheme is shown in FIG. 9.

In contrast with the asymmetric client communication scheme as described earlier, the "redux" broadcast phase strategy addresses the inherent limitations of
25 having clients only being able to request for priority polling by the coordinator node after the latter has polled them at least once. To do so, a number of dedicated broadcast phases are set up. Usually 1 single supplementary broadcast in phase, due to the aforementioned reasons, and the following description will assume this for conciseness purposes, although what applies to
30 that single phase can easily be generalized to multi-phase schemes.

Unlike a regular broadcast phase in a broadcast cycle, which are subdivided by the number of slots corresponding to the maximum number of repeating cycles N_H (a.k.a hops) that the network allows in order to time bound packet flooding to strict times, the main innovation behind redux phase, to its core, lies in a simple

5 idea: allow for a setting where the number of redux slots N_{HR} is: $N_{HR} < N_H$ during the redux phase, and ensure completion of a flood within this phase within the currently configured spec N_H by partially flooding within the number of slots in the redux phase, and continue this flooding action during the next, and then the following broadcast cycles, as required, until the total number of redux slots

10 $\text{SUM}(N_{HR}[i]) \geq N_H$, again this novel flooding mechanism being also bounded using a sequence number (called "redux sequence number") which is again an injective function of the number of times a packet in this phase has been repeated. This is therefore a flooding mechanism distributed between broadcast cycles, over a span of multiple sleep intervals, and is henceforth also referred to

15 as the "slow flooding" mechanism, versus the regular "flash flooding" mechanism.

Note that repeated packet repeating strategies as mentioned for regular broadcast phases can also apply, and a redux phase segmentation, in order to minimize the risk of an end node attempting simultaneous redux phase access

20 while a packet is being flooded within this phase, can be firstly applied via a grace period that any end node having just received a redux packet, is forbidden to send another one of its own, and secondly via a contention-mitigation mechanism similar to those used by contention-enabled networks (known and popular examples of which include CSMA/CA often used in IEEE 802.15.4

25 distributed networks, and exponential back-off, often used in high-level networking protocols) for the same purposes. In the preferred embodiment, the grace period is typically simple enough that it can be implemented directly on a mesh controller processor target, if volatile or non-volatile storage space in order to buffer incoming redux message requests is high enough, or if a mechanism to

30 feedback request denial to a host processor communicating with that mesh controller is implemented. As for contention management back-off mechanisms

the preferred embodiment assumes that it is to be implemented by the higher-level protocol using the mesh controller protocol, i.e. by the host (mainly for saving processor time and memory in case of tightly limited embedded targets typically used as mesh controllers).

5

Considering a typical application with the use of $N_{HR} = 1$, which is the minimum redux phase slot count. In this case, a network can reserve this redux phase as a special priority access request channel, which stays unused and free until an asymmetric client end node (a.k.a. "sniffer") desires to request immediate communication with its coordinator node, in which case any node can rely on this special access phase. This allows faster responsivity in asymmetric clients, with a consumption (or bandwidth) penalty that is exactly known and bound, equal (in an ideal implementation) to Y_{REG} / T_{REDUX} , where Y is the time for a regular broadcast cycle, and T_{REDUX} is the supplementary time for the redux phase. In the special case that the slot length T_{SLOT} is similar between all phases, and $T_{SLOT} = T_{REDUX}$, perhaps due to guard times inherent to FHSS or other implementation details being already included in T_{SLOT} (for instance, if there is a 1-on-1 correspondence between a frequency channel FHSS hop and a slot, which is a subset of the general case where multiple slots can be assigned to a specific frequency channel FHSS hop), for N_I broadcast in phases and N_O broadcast out phases, then the idealized power consumption penalty is $N_H * (N_I + N_O) / N_{HR}$, and thus can be potentially an order of magnitude lower than if a regular broadcast phase was assigned to the same task.

Potential applications whose performance within mesh network can be dramatically improved using this strategy include (but are not limited to): control using asymmetric client end nodes, 2-way communication schemes using the mesh (such as text chat, for instance), emergency signal flood-back to coordinator node, etc.

30

7) Parallel high speed network communication and control

For applications off the grid that leverages a distributed mesh network in order to send larger data payload that can be fit within a single mesh network packet (those will be considered for the purpose of this discussion as being "files") than
5 typical sensor information, such as for instance sending camera still shots over the mesh, the bandwidths required can be low enough to even be feasible using a cooperative bit-synchronized mesh network in which typical data frame maximum length is severely limited in order to enforce bit-synchronization despite timing jitter limitations. In such a case, a basic strategy can be implemented
10 where those files can be split, with a unique sequence number encoding the file chunk ID, as well as a supplementary transmission order sequence number if packet ordering is critical, are appended to each chunk so that the resulting combination of bits can be fit within a data frame, and then sent over the network, after a file transfer request packet is sent and acknowledged for. On the receiving
15 end of the transmission, acknowledgment signals can be encoded with the corresponding file chunk ID and transmission order sequence number so that the transmitter knows which chunk has to be sent again. Note that multiple acknowledgements signals can be sent within the same data frame in the case where multi-phase communication is used, in which $N_I \neq N_O$. Note that in a
20 regular communication scheme, the acknowledgment packet is sent back to the sender before the latter attempts further communication, and timeouts are used to determine acknowledgment failure.

However, due to the time deterministic nature of a bit-synchronized cooperative
25 mesh network, such a strategy becomes quite inefficient. Instead, one can leverage the deterministic nature of such a network, and assume that any acknowledgement signal can be encoded by the very presence of an answer packet from the destination node, in the exact broadcast cycle and broadcast phase that such destination node is expected to answer. Therefore, packet
30 timeout cycles can be eliminated from the equation. It is even not necessary to send supplementary acknowledgment information such as the file chunk ID

and/or the sequence order number, because the sender knows exactly for which requested packet sent out the acknowledgment caters for by virtue of the number of broadcast cycles elapsed between the two signals (although this can be optionally kept for redundancy purposes in case host-side time synchronization with the mesh controller fails). In the case of transferring a relatively large file, an error checking mechanism can be appended, not to each packet (because this is pre-included as a base reliability mechanism in the packet construction of the mesh network itself), but rather to the file itself, so that file integrity can be checked at reconstruction on the destination side.

Note that, as shown previously when discussing end node sensor polling strategies with a VM, the actual broadcast cycle interval between packet request from coordinator node, and expected signal answer from queried end-node back at the coordinator node location is at least 2, in the case of a single core mesh controller (but this limitation can be improved by using parallel processing enable hardware). Again, as before, a coordinator node requesting a file from an end node can send requests during each broadcast cycle, and do not need to wait for an acknowledgment signal before sending the next requests in its send pipe, and this is valid notwithstanding the number of nodes being queried (i.e. a single node can be queried multiple times in this fashion).

In the case where such higher bandwidth reliable communication scheme is required not only between a given end node and its gateway, but reliability is also desired for asymmetric client "sniffer" end nodes, then a scheme where those client end nodes are able to send acknowledgment packets for any file, be it from a coordinator node, or from any single other end-node, is needed. In this case, the following optimization strategies can be employed.

Each packet data frame can be separated into sub-frames, and each sub-frame corresponds to a higher-level protocol message (also called "sub-packet" in the code base used for the preferred embodiment), complete with a message ID,

therefore allowing multiple message types to live within a single mesh packet. Some message IDs can be control messages, some can be used as the UDP-like sensor polling protocol as explained previously, some are used to control "file" transfers.

5

Each packet originating from the coordinator node has at least (for the preferred embodiment, exactly) 1 meta-sub-frame, containing the broadcast phase assignment for each end node that must reply, as explained previously (if multi-phase mode is used), as well as a mandatory level-1 sequence number, and a
10 level-2 sequence number. The level-1 sequence number is used to identify and encompass a whole data synchronization cycle between all involved parties (for instance, if there are a fixed number of end node clients and 1 coordinator, each of which has a particular "file" to be broadcasted to all parties ("file" being defined in the most generic way possible as any size-limited bit stream). The level-2
15 sequence number is typically monotonically incremented from a base value for the first broadcast cycle within a level-1 data synchronization cycle, in order to uniquely identify any broadcast cycle that occurs within that data synchronization cycle (so the level-2 sequence is also called the "broadcast cycle ID").

20 Note that for multi-phase broadcasting, the TDMA-like scheme as previously explained has an overhead where addresses are assigned broadcast phase IDs according to their address position in a request packet from a coordinator node. This overhead can be optimized in the case where asymmetric client "sniffer" communication is needed, because typically there are order of magnitudes less
25 "sniffer" end nodes than sensor end nodes. In that case, each asymmetric client can be assigned an ID, either in an out-band fashion, or in-band using specialized higher-level messaging with unique identifiers segregating the type of packet.

30 There are a special subset of messages dedicated to acknowledgment signalling. Those are always prioritized over regular message types. Note that

acknowledgment signals need only be sent for the following cases: coordinator to end node, and end node to end node, and for coordinator packets in a broadcast out phase other than the first (or last, depending on chosen implementation) broadcast out phase, the reason being that the mere fact that an end node
5 replied to a coordinator in its assigned broadcast phase is an implicit acknowledgement of receipt of at least 1 outbound packet from a coordinator node, for bit-synchronized cooperative networks.

All acknowledgement signals leverage the deterministic nature of the bit-synchronized mesh network in order to minimize the number of
10 acknowledgments being sent: instead of acknowledging a node address (or its corresponding reduced-size address, if such a strategy is used, such as the short-addressing vs. 32-bit MAC address scheme used in IEEE 802.15.4) each acknowledgment packet contains a level-2 sequence ID corresponding to the first
15 broadcast cycle that it is attempting to acknowledge, as well as a serialized bit mask, each bit corresponding to a specific broadcast phase. Because specific nodes are bound to a broadcast phase for each broadcast cycle, this allows to acknowledge multiple packets simultaneously, because there is a 1-to-1 mapping for each broadcast phase to a specific node address, as assigned within each
20 packet of a coordinator node, and also allows to broadcast this multi-packet, multiple-destination acknowledgement signal network-wide at all times. Note that the most obvious case would be to have a bit sequence with the exact same length as $(N_o + N_l - 1)$, but in fact this bit sequence can be made longer in order to span multiple broadcast cycles (the level-2 sequence ID allowing any recipient
25 to infer from which starting broadcast cycle one should interpret the acknowledgment bit stream).

Note that the use of a bit mask is not a necessary constraint of such a design, albeit it is highly bandwidth-efficient. In fact, any acknowledgment signalling
30 scheme that does not rely on direct node addressing, but rather by broadcast

phase, due to a 1-to-1 relationship between a phase and a specific end node, is under the umbrella of the present invention.

An example embodiment of such a scheme can be as follows:

5

At the beginning of a level-1 cycle, the coordinator requests the number of bits that each end node client needs to complete this cycle. This is so that the coordinator can schedule the required length of the level-1 cycle.

10 The coordinator monitors and saves every single packet within a level-1 cycle. This is because it is the only device allowed to reschedule a request for a packet resend operation from a specific end node in case of an acknowledgement failure from said end node, and even if an end node did send an acknowledgment network-wide, there is still the possibility that such a flood does not reach the
15 coordinator itself due to network failure. Thus, at the end of a level-1 cycle, when everything else has been sent, the coordinator checks for all missed acknowledgment packets, and explicitly requests the originator packets corresponding to those acknowledgment packets (the coordinator knows which node and with level-2 ID this corresponds to because it keeps a trace of every
20 transaction within the level-1 cycle).

It schedules packet sends from all end node clients in an interlaced way. The simplest methodology is to simply swap from one client end node to the other at each consecutive assigned broadcast in phase. Even if the length of the file that
25 each end node client must send back to the whole network varies, it does not matter in a practical way: when an end node finishes sending its data, it simply uses its assigned phase to send required acknowledgements to data still transiting within the level-1 cycle throughout the network, nothing more.

30 Note that it is also possible to use more than 1 consecutive broadcast phase per client. This has the benefit of allowing end nodes to know whether they can span

their acknowledgment bit mask over more than 1 broadcast cycle, with potential data frame space saving, and thus bandwidth increase (although another possible simpler scheme would be to simply hardcode the consecutive number of broadcast phases per node).

5

An ideal example application of such a scheme is for inter-base communication using a bit-synchronized mesh network. Say one has multiple coordinator nodes managing separate mesh networks on different radio channels (either true RF frequency band channels, or FHSS hop tables, or possibly segregation via a network ID, or possible via an encryption or data whitening scheme, or all possible combinations of the above), and want to join those sub-networks using another mesh network, called the "inter-base" network. Then the "files" that need to be sent are simply the whole database of each of those separate networks, to insure that they be synced at all locations in a distributed way, and those "files" could also include high-bandwidth streams such as low-speed camera streaming, 2-way texting, etc.

Note that such a scheme can also be included within a single mesh network, if strict separation of the data frames of each packet into separate containers employed for both purposes is enforced (the first purpose being the regular UDP-like polling action, the second purpose being "file" streaming to multiple asymmetric client end nodes and/or the coordinator itself). Such a separation can be, at its simplest embodiment, using static bit length containers (due to a packet having a fixed maximum size for a given bit-synchronized cooperative mesh network), or can use more complex variable-size container schemes, in which case the expected container length needed for UDP-like polling action can be predicted by the coordinator node by simply verifying the sensor node type of all sensor nodes deployed, assuming that the container size is fixed for a specific sensor node, a constraint that is easily enforced at a sensor node design level, as a mesh network core feature, as is done in the preferred embodiment of the invention.

Another interesting application of such a scheme (in its simple node to coordinator, without acknowledgement from "sniffers", embodiment, as well as its extended embodiment with "sniffers" included) could also be for network control of a higher bandwidth, lower power efficiency wireless network such as a 802.11 device, for instance. Say one wants to operate a low power sensor network, with the option of activating high-speed camera streaming on demand, without having to rely on rechargeable or costly and potentially maintenance-heavy energy harvesting strategies. In this case, one could link the power pack of the 802.11 device linked to the high-bandwidth requirement equipment to a separate bit-synchronized mesh network whose purpose is to activate, deactivate and configure said 802.11 device (in the preferred embodiment, it is a Linux router running a console-configurable interface, for instance, connected via an adapter to the UART port of the mesh controller). This is possible because the scheme explained above allows for any higher-level protocol to leverage it, including order-sensitive protocols. As such, one can use a "file" as a console command, for instance, and in that case, the mesh network can retrofit configuration-heavy equipment such as 802.11s routers so that their assigned ad-hoc peer address be automatically configured via the alternate, low-power, long lasting cooperative bit-synchronized mesh network, without needing any user interaction whatsoever, making for very easily deployable high-bandwidth mesh networks. Said cooperative bit-synchronized parallel mesh network could also change the assigned RF configuration of those routers depending on how much the router frequency band is crowded, for example.

Reference is made to Fig. 21 and Fig. 22 for possible implementations showing the multi-acknowledgment data sub-frames in action, and Fig. 20 details the possible structure of a basic high-level protocol that allows the implementation of console control of a WIFI802.11s enabled router (based on OpenWRT).

Features of the above-described embodiments may be defined as follows.

(1) An embodiment of the present, uses the FM capture effect as a fallback packet network flooding propagation in case bit-synchronization is in peril due to effects including (but not limited to) path delay, multi-path fading, etc.

(2) Further provided are variations in bit-synchronization timing in order to introduce variability akin to "genetic diversity" between potentially contentious packets out of bit-synchronization, said variations being (but not limited to) of the following natures: random, pseudo-random, etc.

(3) There may also be provided variations in packet RF strength in order to introduce variability akin to "genetic diversity" between potentially contentious packets out of bit-synchronization, said variations being (but not limited to) the following natures: random, pseudo-random, etc.

(4) Item (1) may be combined with item (2) and (3)

(5) In addition to item (3), there is further provided random or pseudo-random source induced from antenna and RF front-end circuitry variable frequency response combined with a FHSS frequency hopping scheme.

(6) In addition to item (2), there is further provided random source induced from mesh controller processor timing jitter, RF front-end timing jitter.

(7) In addition to item (2), there is further provided pseudo-random source induced by the inherent non-linearity of interrupt propagation between the mesh controller processor and the RF front-end circuitry, each of which is operating using clock frequencies with a non-integer ratio.

(8) In addition to item (2), there is further provided leveraging the unpredictability of manufactured clock drifts of the mesh controller and/or the RF front-end circuitry, and its relationship with local variable environmental parameters such as temperature, humidity, etc.

5

(9) In addition to item (2), there is further provided using permutations of combinations of items (6), (7), (8) variability generating schemes.

10

(10) There may be provided a system using a node that sends multiple redundant packets in different time slots in a broadcast phase, all packets carrying the same information except for the required bits that are used to carry slot-dependent information such as a flooding sequence number and any supplementary information that varies according to the current time slot in a flooding broadcast phase (for instance, error correction codes).

15

(11) In addition to item (10), there is further provided regularly spaced redundant packets sends for each node.

20

(12) In addition to item (10), there is further provided redundant packets spaced according to a random or pseudo-random scheme.

25

(13) In addition to item (10), there is further provided redundant packets spacing according to a feedback loop that monitors the success or failure of delivery according to implementation of higher-level protocol used by the network to poll sensors (as a form of inter-communication between different OSI layers, in other terms)

30

(14) There is further provided a system according to item (13) combined with sensor positional information taken using real-time out-band local methods such as an onboard GPS

(15) There is further provided a system according to item (13) combined with sensor positional information taken using locally stored out-band positional information, such as GPS information configuration data being entered into the node at node deployment time by an enabling device that has an onboard position sensor such as (but not limited to) a smartphone

(16) There is further provided a system according to item (13) combined with remotely stored sensor positional information sent from the coordinator node linked to a processing unit with access to sensor database, said positional information being sourced via methods such as (but not limited to) GPS storage in centralized or distributed database at deployment, via a parallel mesh protocol optimized for radio triangulation and neighbour mapping methods, via acoustic phased array deployment synchronized on the mesh network time base, etc.

Other examples and implementations

Various examples of the above-described systems, methods and features for network communication techniques will be described below. It should be understood that the examples described below may be combined with other aspects of the network communication system, method and components as described hereinabove in accordance with various embodiments of the present.

Clock synchronization – node and network

In an embodiment, there is provided a node to be integrated in a network of wirelessly connected nodes. The node comprises a main module, external sleep controller and a digital serial bus connected between the main module and the external sleep controller. The main module comprises a main controller (for example, the above-described "mesh controller") having a main clock (for example, the above-described mesh clock). The main controller is operable between an active configuration wherein the main clock oscillates and an inactive

configuration wherein the main clock sleeps, in order to minimize consumption from the main power source. The digital serial bus is operable under a communication protocol which is configurable to embed timing information. The external sleep controller is operable under a lower power consumption than the main module and operates the main controller between the active and inactive configurations based on the timing information. The main module comprises a communication means connected to the main controller for communicating with other nodes of the network; and a main power source connected to the communication means and the network for supplying power thereto.

The timing information comprises sleep period data representative of a time period during which the main module is intended to operate in an inactive configuration. The external sleep controller comprises an input port connected to the digital serial bus, for receiving the sleep period data; a low frequency clock for measuring elapsed time; a timer connected to the low powered clock for determining when the elapsed time has reached said time period; an output port connected to the digital serial bus to send a signal to the main module to operate the main module in the active configuration.

It is to be understood that the input port and output port may be the same port.

The digital serial bus may be an Inter-Integrated Circuit (I2C) or a Serial Peripheral Interface (SPI) bus, for example.

Network of nodes

A plurality of the above-described end nodes may be interconnected in a network system, the network further comprising a coordinator node for communicating with the end nodes, the network being configured to broadcast data and synchronize a time base via a flooding mechanism. The nodes may be connected in a mesh configuration. Moreover, the network may have a

cooperatively synchronized time base, between the end nodes and the coordinator node.

Sniffer

5

In the described embodiment, the coordinator node comprises a system database. Further, an asymmetric client (or "sniffer") is in communication with one of the end nodes. The asymmetric client comprises a database which replicates at least a portion of the information found in the system database. The
10 asymmetric client is configured to update the replicated database upon receiving data through said flooding mechanism.

Thus, it is to be understood that a client device may be operatively connected to at least one of the end nodes.

15

Parallel network communication and control

The above-described network may be further combined with a secondary network of nodes, connected via the coordinator node. The nodes of the
20 secondary network may be routers which are connected over a high bandwidth network. Further, a computer may be connected to at least one of the nodes of the secondary network.

In addition or alternatively, the secondary network may be another network
25 similar to the previously described one (such as another mesh network, for example). It is to be understood that any suitable number of additional networks (mesh, wifi or other) may be further connected to the first network. It is to be understood that each network may have on or more coordinator node, which may in turn be connected directly or indirectly to the coordinator node(s) of the other
30 networks.

Clock synchronization – method

In accordance with an embodiment, there is provided a method for operating a node adapted to cooperate within a network of nodes. The node comprises a main module and an external sleep controller connected thereto via a digital serial bus, the main controller having a main clock, the main controller being operable between an active configuration wherein the main clock oscillates and an inactive configuration wherein the main clock sleeps. The method comprises:

- sending from the main controller, timing information embedded in a communication protocol of the digital serial bus, to the external sleep controller, said timing information comprising sleep period data representative of a time period during which the main module is intended to operate in an inactive configuration;
- at the external sleep controller, measuring elapsed time via a low frequency clock; and
- when the elapsed time has reached said time period, sending a wake signal (1st wake signal) from the external sleep controller, over the digital serial bus, to the main module to operate the main module in the active configuration.

The sleep period is represented by a number of periods of the low frequency clock. Moreover, the sending of the timing information from the main controller is scheduled to match a start of a period of the low frequency clock. The scheduling comprises delaying the sending of the timing information until the start of a next period of the low frequency clock. This start of the next period is embedded in a communication protocol of the digital serial bus and signaled from the external sleep controller to the main controller. More particularly, the delaying comprises clock stretching and the signaling is made through a clock transition of the external sleep controller.

After sending the wake signal, an initialization signal from the external sleep controller, is further sent over the digital serial bus, to the main module to initialize the main clock (2nd wake signal).

- 5 The timing information sent from the main controller further comprises a time delay to be elapsed between the wake signal and the initialization signal (between the 1st and 2nd wake signals). This time delay is represented by a number of periods of the low frequency clock.
- 10 Moreover, the sleep period data takes into account a guard time to compensate for an error in the low frequency clock. The guard time is calculated according to:

$$\text{guard time} = e \times n + f_{\text{safety}},$$
wherein
 e represents said error (in μsec) of the low frequency clock,
 n represents a number of packet failures ("drops") tolerated, and
 f_{safety} represents a safety factor (in μsec).
- 15

Communication protocol

- 20 In accordance with an embodiment, there is provided a method for communicating within a network. The network comprises end nodes and a coordinator node, the network being cooperatively synchronized and configured to broadcast data packets during broadcast cycles through a flooding mechanism (multiple floodings may be involved).

- 25 Each data packet comprises a configuration frame and a data frame, the configuration frame comprising hop data. The hop data may include a sequence number representing the current number of hops and a hop maximum representing a maximum number of hops for a given broadcast phase. In an alternative embodiment

30

The method comprises for each broadcast cycle:

- a) during an outbound broadcast phase, sending from the coordinator node a coordinator-originating data packet for polling at least one of the nodes, the data frame of the coordinator-originating data packet comprising address data corresponding to the at least one end node to be polled; and
 - 5 b) during an inbound broadcast phase, receiving at the coordinator node a node-originating data packet returned from one of said at least one of the end nodes, the data frame of the node-originating data packet comprising node data from the end node having been polled.
- 10 Each of the outbound and inbound broadcast phases comprises successive time slots for at least one of said end node to either receive the data packet or rebroadcast the data packet to the network. The method further comprises:
- i) during one of said time slots, receiving at one of the end nodes, one of said data packets; and
 - 15 ii) during another one of said time slots:
 - updating the hop data in the data packet (for example, by incrementing or decrementing the sequence number);
 - rebroadcasting the data packet from the end node (if the incremented sequence number is equal or less than the maximum
 - 20 number of hops); and
 - during an inbound broadcast phase, if the end node corresponds to the address contained in the data packet, generating and broadcasting a data packet comprising node data from the end node and corresponding hop data (for example, where the hop data
 - 25 has a reinitialized sequence number).

Each time slot has a duration ranging between about 1 μ sec and about 100 μ sec. It is to be understood that the time slots may have variable durations depending on the particular implementation.

In accordance with an embodiment, each of said outbound and inbound broadcast phases comprises a set number of time slots ranging between 2 and 31. It is to be understood that in accordance with alternate embodiments, a broadcast phase may have any suitable number of time slots.

5

Each of the outbound and inbound broadcast phases further comprises a frequency change guard time period. The frequency change guard time period may be set to a period set up to about 10 μ sec, for example. In accordance with an embodiment, the guard time is set to 1 μ sec.

10

In other words, there is provided a method for communicating within a network of nodes. The network is cooperatively synchronized and configured to broadcast data packets to the network during broadcast phases through a flooding mechanism, each data packet comprising hop data. The method comprises, for each broadcast phase:

15

- 1) broadcasting a data packet from a source node during a predetermined time slot; and
- 2) receiving the data packet at one or more destination node during said predetermined time slot.

20

The method further repeats steps (1) and (2) according to the hop data, at respective predetermined time slots, wherein each destination node of step (2) corresponds to a source node in a next execution of step (1) and the data packet received in step (2) corresponds substantially to the data packet to be broadcasted in the next execution of step (1).

25

The hop data may represent a maximum number of hops and a current hop number. In this case, the method further comprises at each destination node of step (2), updating the current hop number in the data packet. The repeating is thus executed until the current hop number has reached the maximum number of hops.

30

The current hop number corresponds to a sequence number and the updating comprises incrementing the sequence number; the method further comprises before step (1), initializing the sequence number.

5

The nodes comprise end nodes and a coordinator node. Each broadcast phase comprises: at least one outbound broadcast phase wherein the coordinator node broadcasts a coordinator-originating data packet directed to one or more of said end node; and at least one inbound broadcast phase wherein a node-originating data packet is received at the coordinator node.

10

The method further comprises broadcast cycles, where each broadcast cycle comprises successive ones of the at least one outbound broadcast phase and of the at least one inbound broadcast phase, consecutive broadcast cycles being separated by a sleep period.

15

The coordinator-originating data packet of the outbound broadcast phase comprises address data representing to one or more end node to be polled and each node-originating data packet of the outbound broadcast phase comprises node data from each end node having been polled.

20

In accordance with an embodiment, each node comprises a storage for storing configuration data, wherein the data packets comprise updated configuration data, the method further comprising after step (b), updating at each destination node, the corresponding configuration data based on the updated configuration data of the data packet in order to synchronize the configuration data of ends nodes in the network

25

Multi-phase broadcast modes

30

In the above-described method, the broadcast cycle may comprise multiple inbound broadcast phases, i.e. step (b) of the afore-mentioned method is executed multiple times sequentially, in respective inbound broadcast phases, within a given broadcast cycle.

5

This feature may be useful for segmenting larger volume elements to be transmitted across the network. In such a case, the data frame of each of the node-originating data packet of the said multiple inbound broadcast phases of step (b) comprises a segmented portion of a data element, to be reconstituted at the coordinator node.

10

In some cases, data frame of the coordinator-originating data packet of the outbound broadcast phase may comprise address data associated to a plurality of said end nodes (for example, to poll several nodes with one polling message); the method then comprises, for each of the plurality of end nodes, broadcasting a data packet in respective ones of said multiple inbound broadcast phases.

15

In other embodiments, the broadcast cycle may comprise multiple outbound broadcast phases, for example to transfer data to an asymmetric client, as previously explained.

20

Acknowledgement of receipt

In other embodiments, the data packet broadcasted by a source node is addressed to a plurality of destinations nodes, the data packet having data set for acknowledging receipt, each destination node being associated to a predetermined position in the data set, the method further comprising for each of the destination nodes of step (b), updating the received data packet at the corresponding predetermined position to indicate an acknowledgement of receipt.

25

30

In some embodiments, a high-level communication protocol is wrapped in bounded bit streams called "files" and segmented in chunks in order to be broadcastable via the lower-level protocol. Each chunk is identified by an unique sequence number so as to enable "file" reconstruction at receiver site. Each
5 chunk is sent within free space available in a data payload of the packet to be broadcasted by a given node. Acknowledgement may be sent which directly contains information pertaining to the chunk that has been acknowledged, such as (but not limited to) the chunk sequence number. Acknowledgement may be sent which contains information pertaining to a specific broadcast cycle and
10 broadcast phase where said chunk was broadcasted, so as to uniquely identify it via the structure of the broadcast phases themselves. Moreover, acknowledgement of an end node for a coordinator originating packet may be implicit to that end node by virtue of simply replying to said coordinator packet at the requested broadcast cycle and broadcast phase, said request being sent at
15 all times during broadcast out phases by said coordinator node.

The identification of a specific broadcast cycle and/or phase combination may be sent as a number identifier. Alternatively, the identification of a specific broadcast cycle and/or phase combination may be sent as a bitmask, where the position of
20 a bit correlate directly to the position of a given broadcast phase. The method may allow for acknowledgement verification and network health checking. Moreover, acknowledgement signals may be combined with chunks so as to be sent in the same packet.

25 Also, mutiple acknowledgements may be sent simultaneously in the same packet. Simultaneous acknowledgement may be used in order to achieve handshaking between a set of nodes, whether a coordinator node or an end node, so as to enable robust point-to-mutipoint communication between all nodes within the set. The number of end nodes in the set may match exactly the
30 number of broadcast in phases per cycle, and each of those end nodes are then assigned a specific broadcast "in" phase within a broadcast cycle, so that any of

those end nodes can acknowledge all transactions from the broadcast cycle preceding it by a known number of broadcast cycles.

Priority management in poll scheduling by coordinator node

5

In accordance with an embodiment, each node address is associated to a priority queue number in a storage. The storage is preferably located at the coordinator node. The priority queue numbers represent a polling order for the coordinator node. The address in the coordinator-originating data packet of the outbound
10 broadcast phase corresponds to the end node having a highest priority based on the priority queue number. After the outbound broadcast phase, the priority queue numbers are reassigned.

The reassigning may comprises:

- 15
- lowering a priority of the end node having been polled;
 - maintaining a higher priority level of critical ones of said end nodes; and/or
 - maintaining a higher priority level of predetermined ones of said end nodes for a period of time.

20 The reassigning may be initiate further to receiving a priority change request at the coordinator node, comprising a node address and a corresponding new priority queue number, in which case the reassigning comprises updating in the storage the priority queue number associated to the node address, with the new priority queue number.

25

The priority change request may be initiated by an end node, for example a "sniffer" node which may be made via a user interface of a client application.

Alternatively, the priority change request may be initiated in response to a
30 detected trigger event. A trigger event may include: a warning or error threshold detected on a given end node; an old timestamp condition on a given end node;

and/or a user interacting with a gauge on a user interface connected to the coordinator node (in which case the priority mode stays active for a while, and then stops after a while if the client application remains idle).

5 Asymmetric client

In addition, the coordinator node a system database, and at least one of the end nodes of step (i) is an asymmetric client (or "sniffer") comprising a replicated database, i.e. at least a portion of the content of the replicated database is
10 substantially sourced from the system database at the coordinator node. The method further comprises after step (i) updating, via the asymmetric client, the replicated database with content of the data packet. The content of the replicated database comprises at least one group among: configuration variables, data variables and node network state information.

15

The method may thus further include sending a command from the asymmetric client directed to the coordinator node, for example to obtain specific content to be updated in the replicated database.

20 Redux phase

The broadcast cycle further comprises a redux phase. The redux phase includes one or more time slot for an end node to either receive the data packet or broadcast the data packet to the network, for a partial flooding to the network, so
25 as to split the flooding over multiple broadcast cycles.

An end node (for example, a "sniffer") may initiate communication with the coordinator node by broadcasting a data packet during the redux phase.

30 Thus, the redux phase may be used by the "sniffer" to request priority (i.e. sending a priority change request to the coordinator node). More particularly, in

one embodiment, the coordinator may send a flag in its outbound broadcast to indicate that it gives permission to a particular sniffer node of a specific address to request priority at this time. In an alternative embodiment, the network is configured such that the sniffer or end node is able to do send a priority change request by its own initiative (unsupervised redux slot), in which case the coordinator receiving such a request automatically reassigns the priority scheduling.

Sleep

10

Consecutive broadcast cycles are separated by a sleep period, during which end nodes are deactivated for external communication (i.e. between nodes of the network). The sleep period may have a duration of up to about 1 min.

15 Mesh controller time-sharing with sensor peripheral integration firmware within same processor core

In accordance with an embodiment, each end node comprises a main controller. The method further comprises executing a node-related process, at the main controller of at least one of the end nodes, during the sleep period.

Each end node of further comprises a non-volatile memory having stored thereon virtual machine (VM) code means being configured to execute the node-related process (i.e. a process that is internal to the node). A node-related process may be triggered by at least one of the following events: booting-up the main controller, activating a network seeking mode, deactivating a network seeking mode, entering a broadcast cycle and leaving a broadcast cycle. The virtual machine code means may be further configured to limit execution time of the node-related process based on the duration of the sleep period. For example, the virtual machine code means may include a hypervisor to limit the execution time based on a comparison of a predicted execution time with remaining sleep time.

Possible user processes employed within VMs include:

- 5 • Fetching analog sensor data interfaced to the available signal conditioning hardware on the mesh controller processor (for example: onboard internal ADCs, digital inputs, timer capture, etc.). This potentially allows to capture sensor data from standards such as 0-10V, 4-20 mA, with minimal signal conditioning circuitry (typically a resistor or a voltage divider is the only thing required for such applications).
- 10 • Sending analog output signals to external circuitry directly from available hardware on the mesh controller processor (for example: onboard internal DACs, digital outputs, timer compare/PWM generators).
- Controlling and acquiring data from I2C/SMBus enabled sensors.
- Controlling and acquiring data from SPI enabled sensors.
- 15 • Controlling and acquiring data from custom digital protocol sensors (ex: 1-Wire) via custom bit-banging protocols implemented in the user VM code.
- Doing basic digital signal conditioning operations (typically mathematical operations for converting a raw sensor data into the proper format for transmission via radio).
- 20 • Capturing the precise timing of the mesh network, and transmitting that information to a coprocessor so as to in-band synchronize the timers of external devices (this is especially useful for applications that need a distributed common time base without having to rely on costly solutions such as individual GPS modules per node). This feature is typically
- 25 implemented for time of flight triangulation, such as for gunshot localization via a distributed network of wireless acoustic sensors.
- Reading and reconfiguring the mesh controller's own internal configuration registers (this also potentially enables meta-programming).
- Storing information in accessible internal RAM/ROM memory sections of
- 30 the mesh controller (typically for storing non volatile sensor factory calibration data).

- Communicating event information to external circuitry.
- Communicating internal mesh controller information to external circuitry without using standard communication bus (typically a serial port).

5 Parallel network communication and control

In some embodiments, the network may be combined with a secondary network of nodes (or more additional network(s)), connected to the coordinator node. The method comprises operating from the coordinator node, the secondary network
10 between a powered configuration wherein the nodes of the secondary network are connected to a power source, and a sleep configuration wherein the nodes of the secondary network are disconnected from the power source. A data element from the secondary network may be encapsulated in said data packets being broadcasted in the network of end nodes.

15

In accordance with one embodiment, there is provided a method for communicating within a network comprising end nodes and a coordinator node, the network being configured to broadcast data packets between said nodes during broadcast cycles through a flooding mechanism, consecutive broadcast
20 cycles being separated by a sleep period wherein the end nodes are deactivated for external communication. The method comprises during said sleep period, executing a node-related process, at a main controller of at least one of the end nodes.

25 The secondary network may be a high bandwidth network, thereby allowing streaming of sources such as camera feeds unrestricted by the primary network bandwidth limitations. Alternatively, the secondary network may be a IEEE802.11 network. Still alternatively, the secondary network may be a cooperative mesh network. It is to be understood that multiple of the above-mentioned examples of
30 secondary networks including a combination of the above-examples or the like may be connected to the main network.

Auto-configuration of the hardware platforms responsible for managing said secondary network may be achieved using the primary network. In such a case, network configuration may be accomplished via wrapping of a high-level protocol with handshaking within the communication protocol paradigm of the primary network in order to change configuration data of said secondary network hardware platform such as, but not restricted to, the following: peer addresses, network channel, network bandwidth, network encryption.

- 10 The hardware potentially operating the secondary network may be leveraged in order to locally interface with compatible sensors and controllers. The hardware may be a network router is used as the main secondary network is a wifi-enabled router. In addition, legacy gauges and instrument panels are retrofitted via a camera digitizing their state and encapsulating the resulting data in packets
15 broadcasted in the primary network as a regular end node sensor.

Embodiments of the present may apply in the areas of forestry, military, industrial and construction applications.

- 20 The above-described embodiments are considered in all respect only as illustrative and not restrictive, and the present application is intended to cover any adaptations or variations thereof, as apparent to a person skilled in the art. Of course, numerous other modifications could be made to the above-described embodiments without departing from the scope of the invention, as apparent to a
25 person skilled in the art.

Claims:

1. A node to be integrated in a network of wirelessly connected nodes, the node comprising:

- 5 - a main module comprising a main controller having a main clock, the main controller being operable between an active configuration wherein the main clock oscillates and an inactive configuration wherein the main clock sleeps;
- 10 - a digital serial bus connected to the main module, the digital serial bus being operable under a communication protocol which is configurable to embed timing information; and
- an external sleep controller being connected to the main module via the digital serial bus to operate the main controller between the active and inactive configurations based on said timing information.

15

2. A node according to claim 1, wherein the main module comprises:

- a communication means connected to the main controller for communicating with other nodes of the network; and
- 20 - a main power source connected to the communication means and the network for supplying power thereto.

20

3. A node according to claim 1 or 2, wherein the timing information comprises sleep period data representative of a time period during which the main module is intended to operate in an inactive configuration wherein the external sleep controller comprises:

25

- an input port connected to the digital serial bus, for receiving the sleep period data;
- a low frequency clock for measuring elapsed time;
- a timer connected to the low powered clock for determining when the elapsed time has reached said time period;

30

- an output port connected to the digital serial bus to send a signal to the main module to operate the main module in the active configuration.

4. A node according to any one of claims 1 to 3, wherein the digital serial bus comprises an Inter-Integrated Circuit (I2C).
5. A node according to any one of claims 1 to 3, wherein the digital serial bus comprises a Serial Peripheral Interface (SPI) bus.
6. A network system comprising a plurality of interconnected end nodes according to any one of claims 1 to 5, the network further comprising a coordinator node for communicating with the end nodes, the network being configured to broadcast data and synchronize a time base via a flooding mechanism.
7. A network system according to claim 6, wherein the end nodes are connected in a mesh configuration.
8. A network system according to claim 6 or 7, wherein the synchronizing of the time base is cooperative between the end nodes and the coordinator node.
9. A network system according to any one of claims 6 to 8, wherein the coordinator node comprises a system database the network system further comprising an asymmetric client being in communication with one of the end nodes, the asymmetric client comprising a replicated database said asymmetric client being configured to update the replicated database upon receiving data through said flooding mechanism.
10. A network system according to any one of claims 6 to 9, further comprising a client device being operatively connected to at least one of the end nodes.

11. A network system according to any one of claims 6 to 10, in combination with a secondary network of nodes, connected to the coordinator node.

5 12. A network system according to claim 11, wherein the nodes of the secondary network are connected over a high bandwidth network.

13. A network system according to claim 11 or 12, wherein nodes of the secondary network comprise routers.

10 14. A network system according to any one of claims 11 to 13, further comprising a computer being connected to at least one of the nodes of the secondary network.

15 15. A method for operating a node adapted to cooperate within a network of nodes, the node comprising a main module and an external sleep controller connected thereto via a digital serial bus, the main controller having a main clock, the main controller being operable between an active configuration wherein the main clock oscillates and an inactive configuration wherein the main clock sleeps, the method comprising:

20 - sending from the main controller, timing information embedded in a communication protocol of the digital serial bus, to the external sleep controller, said timing information comprising sleep period data representative of a time period during which the main module is intended to operate in an inactive configuration;

25 - at the external sleep controller, measuring elapsed time via a clock; and

- when the elapsed time has reached said time period, sending a wake signal from the external sleep controller, over the digital serial bus, to the main module to operate the main module in the active configuration.

30 16. A method according to claim 15, wherein the sleep period is represented by a number of periods of the low frequency clock.

17. A method according to claim 16, wherein the sending of the timing information from the main controller is scheduled to match a start of a period of the low frequency clock.
- 5 18. A method according to claim 17, wherein the scheduling comprises delaying the sending of the timing information until the start of a next period of the low frequency clock.
- 10 19. A method according to claim 18, wherein said start of the next period is embedded in a communication protocol of the digital serial bus and signaled from the external sleep controller to the main controller.
- 15 20. A method according to claim 18 or 19, wherein said delaying comprises clock stretching.
21. A method according to any one of claims 15 to 20, further comprising after sending the wake signal:
- sending an initialization signal from the external sleep controller, over the
- 20 digital serial bus, to the main module to initialize the main clock.
22. A method according to claim 21, wherein the timing information sent from the main controller further comprises a time delay to be elapsed between the wake signal and the initialization signal.
- 25 23. A method according to claim 22, wherein the time delay is represented by a number of periods of the low frequency clock.
24. A method according to any one of claims 15 to 23, wherein the sleep period
- 30 data takes into account a guard time to compensate for an error in the low frequency clock.

25. A method according to claim 24 wherein the guard time is calculated according to: $\text{guard time} = e \times n + f_{\text{safety}}$, wherein e represents said error of the low frequency clock, n represents a number of packet failures tolerated, and

5 f_{safety} represents a safety factor.

26. A method for communicating within a network comprising end nodes and a coordinator node, the network being cooperatively synchronized and configured to broadcast data packets during broadcast cycles through a flooding mechanism, each data packet comprising a configuration frame and a data frame, the configuration frame comprising hop data, the method comprising for each broadcast cycle:

10

a) during an outbound broadcast phase, sending from the coordinator node a coordinator-originating data packet for polling at least one of the nodes, the data frame of the coordinator-originating data packet comprising address data corresponding to the at least one end node to be polled; and

15

b) during an inbound broadcast phase, receiving at the coordinator node a node-originating data packet returned from one of said at least one of the end nodes, the data frame of the node-originating data packet comprising node data from the end node having been polled;

20

each of said outbound and inbound broadcast phases comprising successive time slots for at least one of said end node to either receive the data packet or rebroadcast the data packet to the network;

the method further comprising:

25

i) during one of said time slots, receiving at one of the end nodes, one of said data packets; and

ii) during another one of said time slots:

- updating the hop data in the data packet;

- rebroadcasting the data packet from the end node; and

30

- during an inbound broadcast phase, if the end node corresponds to the address contained in the data packet, generating and broadcasting a data packet comprising node data from the end node and corresponding hop data.

5

27. A method according to claim 26, wherein each time slot has a duration ranging between about 1 μ sec and about 100 μ sec.

10

28. A method according to claim 26 or 27, wherein each of said outbound and inbound broadcast phases comprises a set number of time slots ranging between 2 and 31.

15

29. A method according to any one of claims 26 to 28, wherein each of said outbound and inbound broadcast phases further comprises a frequency change guard time period.

30. A method according to claim 29, wherein the frequency change guard time period is set to a period set up to 10 μ sec.

20

31. A method according to any one of claims 26 to 30, wherein the broadcast cycle comprises multiple inbound broadcast phases.

25

32. A method according to claim 31, wherein the data frame of the coordinator-originating data packet of the outbound broadcast phase of step (a) comprises address data associated to a plurality of said end nodes; the method further comprising, for each of the plurality of end nodes of step (a), executing step (b) sequentially, in respective ones of said multiple inbound broadcast phases.

30

33. A method according to claim 32, wherein the data frame of each of the node-originating data packet of the said multiple inbound broadcast phases

of step (b) comprises a segmented portion of a data element, to be reconstituted at the coordinator node.

- 5 34. A method according to any one of claims 26 to 33, wherein the broadcast cycle comprises multiple outbound broadcast phases.
35. A method according to any one of claims 26 to 34, further comprising:
- in a storage, associating to each node address a priority queue number, the priority queue numbers representing a polling order;
10 wherein the address in the coordinator-originating data packet of the outbound broadcast phase of step (a) corresponds to the end node having a highest priority based on the priority queue number.
- 15 36. A method according to claim 35, further comprising after step (a), reassigning the priority queue numbers in the storage.
37. A method according to claim 36, wherein said reassigning comprises at least one of:
- lowering a priority of the end node having been polled in step (a);
20 - maintaining a higher priority level of critical ones of said end nodes; and
- maintaining a higher priority level of predetermined ones of said end nodes for a period of time.
- 25 38. A method according to any one of claims 35 to 37, further comprising:
- receiving a priority change request at the coordinator node, comprising a node address and a corresponding new priority queue number;
- updating in the storage the priority queue number associated to the node address, with the new priority queue number.
- 30 39. A method according to claim 38 wherein the priority change request is initiated by an end node.

40. A method according to claim 38 or 39, wherein the priority change request is initiated in response to a detected trigger event.
41. A method according to claim 40, wherein said trigger event comprises at least one of: a warning or error threshold detected on a given end node; an old timestamp condition on a given end node; and a user interacting with a gauge on a user interface connected to the coordinator node.
42. A method according to any one of claims 26 to 41, wherein the coordinator node comprises a system database, and at least one of the end nodes of step (i) is an asymmetric client comprising a replicate database, the method further comprising after step (i):
- updating, via the asymmetric client, the replicated database with content of the data packet.
43. A method according to claim 42, wherein the content comprises at least one group among: configuration variables, data variables and node network state information.
44. A method according to claim 42 or 43, further comprising sending a command from the asymmetric client directed to the coordinator node.
45. A method according to any one of claims 26 to 44, wherein the broadcast cycle further comprises a redux phase, the redux phase comprising one or more time slot for an end node to either receive the data packet or broadcast the data packet to the network, for a partial flooding to the network, so as to split the flooding over multiple broadcast cycles.
46. A method according to claim 45, wherein an end node initiates communication with the coordinator node by broadcasting said data packet during the redux phase.

47. A method according to any one of claims 26 to 46, further comprising a sleep period between two consecutive broadcast cycles, during which end nodes are deactivated for external communication.

5

48. A method according to claim 47, wherein the sleep period has a duration of up to about 1 min.

10

49. A method according to any one of claims 47 or 48, wherein each end node of step (i) comprises a main controller, the method further comprising:
- executing a node-related process, at the main controller of at least one of the end nodes, during the sleep period.

15

50. A method according to claim 49, wherein each end node of step (i) further comprises a non-volatile memory having stored thereon virtual machine code means being configured to execute the node-related process.

20

51. A method according to claim 49, wherein the node-related process is triggered by at least one of the following events: booting-up the main controller, activating a network seeking mode, deactivating a network seeking mode, entering a broadcast cycle and leaving a broadcast cycle.

25

52. A method according to claim 50 or 51, wherein the virtual machine code means is further configured to limit execution time of the node-related process based on the duration of the sleep period.

30

53. A method according to claim 52, wherein the virtual machine code means comprises a hypervisor to limit the execution time based on a comparison of a predicted execution time with remaining sleep time.

54. A method according to any one of claims 26 to 53, wherein the network is combined with a secondary network of nodes, connected to the coordinator node, the method comprising:
- operating from the coordinator node, the secondary network between a powered configuration wherein the nodes of the secondary network are connected to a power source, and a sleep configuration wherein the nodes of the secondary network are disconnected from the power source.
55. A method according to claim 54, wherein a data element from the secondary network is encapsulated in said data packets being broadcasted in the network of end nodes.
56. A method for communicating within a network comprising end nodes and a coordinator node, the network being configured to broadcast data packets between said nodes during broadcast cycles through a flooding mechanism, consecutive broadcast cycles being separated by a sleep period wherein the end nodes are deactivated for external communication, the method comprising:
- during said sleep period: executing a node-related process, at a main controller of at least one of the end nodes.
57. A method according to claim 56, wherein each end node further comprises a non-volatile memory having stored thereon virtual machine code means being configured to execute the node-related process.
58. A method according to claim 57, wherein the node-related process is triggered by at least one of the following events: booting-up the main controller, activating a network seeking mode, deactivating a network seeking mode, entering a broadcast cycle and leaving a broadcast cycle.

59. A method according to claim 57 or 58, wherein the virtual machine code means is further configured to limit execution time of the node-related process based on the duration of the sleep period.

5

60. A method according to claim 59, wherein the virtual machine code means comprises a hypervisor to limit the execution time based on a comparison of a predicted execution time with remaining sleep time.

10 61. A method according to any one of claims 56 to 60, wherein each data packet comprises a configuration frame and a data frame, the configuration frame comprising hop data, the method further comprising, for each of said broadcast cycle:

- 15 a) during an outbound broadcast phase, sending from the coordinator node a coordinator-originating data packet for polling at least one of the nodes, the data frame of the coordinator-originating data packet comprising an address for each end node to be polled; and
- b) during an inbound broadcast phase, receiving at the coordinator node a node-originating data packet returned from one of said at least one
20 of the end nodes, the data frame of the node-originating data packet comprising node data from the end node having been polled;

each of said outbound and inbound broadcast phases comprising successive time slots for at least one of said end node to either receive the data packet or rebroadcast the data packet to the network;

25 the method further comprising:

- i) during one of said time slots, receiving at one of the end nodes, one of said data packets; and
- ii) during another one of said time slots:
- updating the hop data in the data packet;
 - 30 - rebroadcasting the data packet from the end node; and

- during an inbound broadcast phase, if the end node corresponds to the address contained in the data packet, generating and broadcasting a data packet comprising node data from the end node and corresponding hop data.

5

62. A method for communicating within a network of nodes, the network being cooperatively synchronized and configured to broadcast data packets to the network during broadcast phases through a flooding mechanism, each data packet comprising hop data, the method comprising, for each broadcast phase:

10

a) broadcasting a data packet from a source node during a predetermined time slot;

b) receiving the data packet at one or more destination node during said predetermined time slot; and

15

repeating steps (a) and (b) according to the hop data, at respective predetermined time slots, wherein each destination node of step (b) corresponds to a source node in a next execution of step (a) and the data packet received in step (b) corresponds substantially to the data packet to be broadcasted in the next execution of step (a).

20

63. A method according to claim 62, wherein the hop data represents a maximum number of hops and a current hop number, the method further comprising:

25

- at each destination node of step (b), updating the current hop number in the data packet;

wherein said repeating is executed until the current hop number has reached the maximum number of hops.

30


64. A method according to claim 63, wherein the current hop number corresponds to a sequence number and said updating comprises

incrementing the sequence number, the method further comprising before step (a), initializing the sequence number.

- 5 65. A method according to any one of claims 62 to 64, wherein the nodes comprise end nodes and a coordinator node and wherein each broadcast phase comprises:
- at least one outbound broadcast phase wherein the coordinator node broadcasts a coordinator-originating data packet directed to one or more of said end node; and
 - 10 - at least one inbound broadcast phase wherein a node-originating data packet is received at the coordinator node.
- 15 66. A method according to claim 65, comprising broadcast cycles, each broadcast cycle comprising successive ones of said at least one outbound broadcast phase and of said at least one inbound broadcast phase, consecutive broadcast cycles being separated by a sleep period.
- 20 67. A method according to claim 66, wherein the coordinator-originating data packet of the outbound broadcast phase comprises address data representing to one or more end node to be polled and each node-originating data packet of the outbound broadcast phase comprises node data from each end node having been polled.
- 25 68. A method according to any one of claims 62 to 67, wherein the data packet broadcasted at step (b) is addressed to a plurality of destinations nodes, the data packet having data set for acknowledging receipt, each destination node being associated to a predetermined position in the data set, the method further comprising for each of the destination nodes of step (b), updating the received data packet at the corresponding predetermined
- 30 position to indicate an acknowledgement of receipt.

69. A method according to any one of claims 62 to 68, wherein each node comprises a storage for storing configuration data, wherein the data packets of step (a) comprise updated configuration data, the method further comprising after step (b), updating at each destination node, the
5 corresponding configuration data based on the updated configuration data of the data packet in order to synchronize the configuration data of ends nodes in the network.

FIGURE 1A - TYPICAL END NODE BLOCK SCHEMATIC






FIGURE 2 - VM ENGINE TRIGGER POINTS





FIGURE 3 - EXAMPLE END NODE CALLBACK VM ENGINE FLOWCHART



Note: Buffer R is used as a temporary volatile memory for the VM engine, and can be used as shared buffers for other operations not running in parallel with said engine (such as UART CTS-controlled buffers, for instance) in order to reduce memory footprint

Example special functions executable by user (non-exhaustive list):

GetExecType(): returns eventTypeId number

Send(var len): sets vm_regBufferCount = len, thus triggering a send operation, either via local comm. Bus, or via RF packet reply

Delay(var len): this should be converted at bytecode compile time to multiple static calls to static delay function, so that the VM hypervisor can check for time overflows

GetBuffer(): set of accessor functions for R

SetBuffer(): set of mutator functions for R

CopyBuffer(r1, r2, len): copy len bytes from location R[r2] to buffer location R[r1]

GetRegisterXXX(): set of accessors for configuration register in memory set XXX (can be volatile, non-volatile...). In the preferred embodiment, there are 3 such sets: RAM, RAMBUF, EEPROM. EEPROM is the non-volatile holder, RAMBUF is the temporary structure that the user is allowed to write to, and RAM is the effective configuration used by the mesh controller.

SetRegisterXXX(): set of mutators for configuration register

TransferConfigXXX(): set of functions to transfer of all configuration register from RAMBUF to destination XXX. If XXX=RAMBUF, then transfer from RAM->RAMBUF

GetPinIn(ID): get state of I/O pin ID. Note that allowable ID is defined by VM engine (as a sandbox security measure, not all pins of the hardware uC is typically allowed).

GetPinDir(ID): get direction of I/O pin ID

SetPinOut(ID, val): set state of I/O pin ID

SetPinDir(ID, dir): set direction of I/O pin ID

GetPerReg(reg): accessor for a SFR ("peripheral" uC-specific register) (reg is mapped to specific uC's own SFRs, with optional read/only flags also mapped, in order to generate a sandbox in which user code can control certain native SFRs, but not all of them, for mesh operation safety purposes)

SetPerReg(reg): same as above, but mutator

I2C_Stop(): master I2C stop signal

I2C_ReadAck(): master I2C read and ACK signal send

I2C_ReadNak(): master I2C read and NAK signal send

I2C_Start(addr_rw): master I2C start command, with addr_rw is I2C address bit-combined with read-write flag (LSBit)

I2C_Write(val): master I2C write command




FIGURE 3A




FIGURE 3B

Note: Buffer R is used as a temporary volatile memory for the VM engine, and can be used as shared buffers for other operations not running in parallel with said engine (such as UART CTS-controlled buffers, for instance) in order to reduce memory footprint

Example special functions executable by user (non-exhaustive list):

GetExecType(): returns eventTypeID number

Send(var len): sets vm_regBufferCount = len, thus triggering a send operation, either via local comm. Bus, or via RF packet reply

Delay(var len): this should be converted at bytecode compile time to multiple static calls to static delay function, so that the VM hypervisor can check for time overflows

GetBuffer(...): set of accessor functions for R

SetBuffer(...): set of mutator functions for R

CopyBuffer(r1, r2, len): copy len bytes from location R[r2] to buffer location R[r1]

GetRegisterXXX(...): set of accessors for configuration register in memory set XXX (can be volatile, non-volatile...). In the preferred embodiment, there are 3 such sets: RAM, RAMBUF, EEPROM. EEPROM is the non-volatile holder, RAMBUF is the temporary structure that the user is allowed to write to, and RAM is the effective configuration used by the mesh controller.

SetRegisterXXX(...): set of mutators for configuration register

TransferConfigXXX(): set of functions to transfer of all configuration register from RAMBUF to destination XXX. If XXX=RAMBUF, then transfer from RAM->RAMBUF

GetPinIn(ID): get state of I/O pin ID. Note that allowable ID is defined by VM engine (as a sandbox security measure, not all pins of the hardware uC is typically allowed).

GetPinDir(ID): get direction of I/O pin ID

SetPinOut(ID, val): set state of I/O pin ID

SetPinDir(ID, dir): set direction of I/O pin ID

GetPerReg(reg): accessor for a SFR ("peripheral" uC-specific register) (reg is mapped to specific uC's own SFRs, with optional read/only flags also mapped, in order to generate a sandbox in which user code can control certain native SFRs, but not all of them, for mesh operation safety purposes)

SetPerReg(reg): same as above, but mutator

I2C_Stop(): master I2C stop signal

I2C_ReadAck(): master I2C read and ACK signal send

I2C_ReadNak(): master I2C read and NAK signal send

I2C_Start(addr_rw): master I2C start command, with addr_rw is I2C address bit-combined with read-write flag (LSBit)

I2C_Write(val): master I2C write command

FIGURE 3C

FIGURE 4 - EXAMPLE MESH CONTROLLER PHYSICAL LEVEL PROTOCOL




FIG. 4A

EXAMPLE MESH CONTR

Config frame lev

Data frame len	Packet type ("air cmd" or "bus")	nS (max hop count)	Duty cycle	SEQ Slot #
----------------	----------------------------------	--------------------	------------	------------

Packet level

sync bits (PLL lock)	preamble bits	Config frame (can be encrypted)
----------------------	---------------	---------------------------------

FHSS level

FHSS slot				FHSS slot			
FHSS freq. change guard time	Slot 21	Slot 22	Slot 23	Slot 24	FHSS freq. change guard time	Slot 25	Slot 26

Regular broadcast phase level

Phase OUT 1	Phase in 1	Phase in 2	Phase OUT 2
-------------	------------	------------	-------------

Local master clock sync

Local master clock

FIG. 4B

OLLER PHYSICAL LEVEL PROTOCOL

rel

Bcast out count	Bcast in count	NWK ID	nSR enable	nSR (redux)	FHSS Hop table index

Config frame var

nS = 5
Duty cycle = 12.5% (1/8)
Bcast out count = 2
Bcast in count = 4
nSR enabled = true
nSR = 2

Static var

Min interval between retransmits = 2
Max retransmits per phase = 3

Data frame (can be encrypted)	Error check

FHSS slot				
Slot 28	FHSS freq. change guard time	Slot 29	Slot 30	Slot 31
Slot 32				

FHSS slot				
Slot 30	Slot 31	Slot 32		

Phase in 3	Phase in 4
---------------	---------------




FIG. 4C





FIG. 4D

FIGURE 5 - POLLING STRATEGY #1 - REGULAR HALF-DUPLEX



CA 02856027 / 2014-07-04

COORDINATOR NODE																		
Host communication bus																		
Req. ascend ↓ ↑ Ans. ascend ↑ Req. regular ↓ ↑ Ans. regular ↓ Req. ascend ↑ Ans. ascend ↓ Req. regular ↓ ↑ Ans. regular ↑																		
Holding buffer																		
Mesh protocol																		
Phase OUT 1	Phase in 1	Idle	Phase ↑OUT 1	Phase in 1	Idle	Phase ↓OUT 1	Phase in 1	Idle	Phase ↓OUT 1	Phase in 1	Idle	Phase ↑OUT 1	Phase in 1	Idle	Phase ↓OUT 1	Phase in 1	Idle	Phase ↑OUT 1
Req. ascend 1 Req. regular 2 Req. ascend 3																		
END NODE 1																		
Host communication bus																		
Holding buffer																		
Mesh protocol																		
Phase OUT 1	Phase in 1	Sleep	Phase ↓OUT 1	Phase in 1	Exec. ascend 1	Sleep	Phase ↑OUT 1	Phase in 1	Sleep	Phase ↓OUT 1	Phase in 1	Sleep	Phase ↑OUT 1	Phase in 1	Sleep	Phase ↓OUT 1	Phase in 1	Sleep
Req. ascend 1 Ans. ascend 1																		
END NODE 2																		
Host communication bus																		
Req. regular ↑ ↓ Ans. regular ↓ Req. regular ↓ Ans. regular ↑																		
Holding buffer																		
Mesh protocol																		
Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase ↓OUT 1	Phase in 1	Sleep	Phase ↑OUT 1	Phase in 1	Sleep	Phase ↓OUT 1	Phase in 1	Sleep	Phase ↑OUT 1	Phase in 1	Sleep	Phase ↓OUT 1
Req. regular 2 Ans. regular 2																		
END NODE 3																		
Host communication bus																		
Holding buffer																		
Mesh protocol																		
Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Exec. ascend 3	Sleep
Req. ascend 3																		

NOTE:

NOTE. We see that the req. 3 from host connected to coordinator did not come in time, and thus there is a break in the polling sequence that makes for a supplementary lag equal to the sleep cycle time. Because this is a regular half-duplex, not speed optimized, there is no risk of buffer overrun between either host or mesh controller. If that was not the case, without proper host synchronization, then the host might erroneously consider that an answer from an end-node matches an incorrect expected slot, and there is also a risk of buffer overrun.

FIG. 5A

POLLING STRATEG




FIG. 5B

Y #1 - REGULAR HALF-DUPLEX

LEGEND

Receive

Transmit

TOR NODE

Ans.
regular

Req.
airmid

2

3

Req.
airmid
3

Idle	Phase OUT 1	Phase in 1	Idle	Phase OUT 1	Phase in 1	Idle

ODE 1

Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep

ODE 2

Ans.
regular

2

Ans.
regular

2									
Mesh protocol									
Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1 ↓ Phase in 1
END N									
Host communication bus									
Holding buffer									
Mesh protocol									
Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1 Phase in 1

NOTE:

We see that the req. 3 from host connected to coordinator did not come in time, and thus there is a break in the polling sequence that makes for a supplementary lag equal to the sleep cycle time. Because this is a regular half-duplex, not speed optimized, there is no risk of buffer overrun between either host or mesh controller. If that was not the case, without proper host synchronization, then the host might erroneously consider that an answer from an end-node matches an incorrect expected slot, and there is also a risk of buffer overrun.

Fig. 5C


		2							
Sleep	Phase OUT 1	Phase in 1 ↑	Sleep	Phase OUT 1	Phase in 1	Sleep	Phase OUT 1	Phase in 1	Sleep

ODE 3										

FIG. 5D

FIGURE 6 - POLLING STRATEGY #2 - PARALLEL SISO

CA 02856027 2014-07-04




NOTE:


The safest way to implement the host communication bus is to disable it for the whole duration of a broadcast cycle (for a coordinator), and from the first used broadcast in phase to the end of its broadcast cycle (for an end node). A dead zone can be appended before the expected time where disabling needs to be in place, in order to compensate for host-side time imprecision. For instance, using a UART port, one could employ its CTS pin in order to block host UART transmission to the mesh controller. In this case, because the CTS pin would not release the port until the broadcast cycle end marker is sent back to the host via the port, a host can use detection of 2 consecutive broadcast end markers before its own UART transmit buffer is empty, for instance, as a symptom of an overrun, which can then be corrected for.

As for the end node side, same security measures can be implemented if the end node uses the host communication bus. In the case of use of a hypervisor-enabled VM engine instead, it is completely safe, because the VM engine can shut down a VM program if it goes beyond its allotted time slot in-between broadcast cycles.

F1G.6A

- POLLING STRA





NOTE:

The safest way to implement the host communication bus is to disable it for the whole duration of a broadcast cycle of its broadcast cycle (for an end node). A dead zone can be appended before the expected time where disabling imprecision. For instance, using a UART port, one could employ its CTS pin in order to block host UART transmission: release the port until the broadcast cycle end marker is sent back to the host via the port, a host can use detection buffer is empty, for instance, as a symptom of an overrun, which can then be corrected for.

As for the end node side, same security measures can be implemented if the end node uses the host communication completely safe, because the VM engine can shut down a VM program if it goes beyond its allotted time slot in-b-

FIG. 6C

		2		2	
Sleep		Phase OUT 1	Phase in 1	↓Phase OUT 1	Phase in 1
				Exec. aircmd 2	
		Sleep			Sleep
		Phase OUT 1	Phase in 1	↑Phase in 1	Sleep

ODE 3

		Ans. aircmd 3		Req. aircmd 3	
Exec. aircmd 3		Sleep	Phase OUT 1	Phase in 1	↓Phase OUT 1
			Phase in 1	Sleep	Phase in 1
					Exec. aircmd 3
					Sleep

cle (for a coordinator), and from the first used broadcast in phase to the end
y needs to be in place, in order to compensate for host-side time
ssion to the mesh controller. In this case, because the CTS pin would not
in of 2 consecutive broadcast end markers before its own UART transmit

ation bus. In the case of use of a hypervisor-enabled VM engine instead, it is
etween broadcast cycles.

Fla. 6D

FIGURE 7 - POLLING STRATEGY #3 - PARALLEL MISO (ex. Bcast. In count = 2)
(same as MIMO with 1 coordinator request per broadcast cycle, no end node bus shown for simplicity purposes)


[illegible]

FIG. 7A

POLLING STRATEGY #3 - PARALLEL

(same as MIMO with 1 coordinator request per broadcast cycle, not 2)

COORDINATOR








FIG. 7B



- MISO (ex. Bcast. In count = 2)
end node bus shown for simplicity purposes)





TOR NODE



ODE 1



ODE 2



ODE 3

														Req aircmd	3			Ans aircmd	3
Mesh protocol																			
Phase OUT 1	Phase in 1	Phase in 2	Sleep	Phase OUT 1	Phase in 1	Phase in 2	Sleep	↓Phase OUT 1	Phase in 1	Phase in 2	Exec aircmd	3	Slp.	Phase OUT 1	Phase in 1				
END N																			
Holding buffer																			
														Req aircmd	4			Ans aircmd	4
Mesh protocol																			
Phase OUT 1	Phase in 1	Phase in 2	Sleep	Phase OUT 1	Phase in 1	Phase in 2	Sleep	↓Phase OUT 1	Phase in 1	Phase in 2	Exec aircmd	4	Slp.	Phase OUT 1	Phase in 1				
END N																			
Holding buffer																			
														Req aircmd	5				
Mesh protocol																			
Phase OUT 1	Phase in 1	Phase in 2	Sleep	Phase OUT 1	Phase in 1	Phase in 2	Sleep	Phase OUT 1	Phase in 1	Phase in 2	Sleep	↓Phase OUT 1	Phase in 1	Phase in 1	Phase in 1				

FIG 7C

			Req aircmd 3			Ans aircmd 3		
--	--	--	--------------------	--	--	--------------------	--	--

Phase in 2	Sleep	↓ Phase OUT 1	Phase in 1	Phase in 2	Exec aircmd 3	Slp.	Phase OUT 1	Phase in 1	Phase ↑ in 2	Sleep	Phase OUT 1	Phase in 1	Phase in 2	Sleep
---------------	-------	---------------------	---------------	---------------	---------------------	------	----------------	---------------	--------------------	-------	----------------	---------------	---------------	-------

ODE 4

			Req aircmd 4			Ans aircmd 4		
--	--	--	--------------------	--	--	--------------------	--	--

Phase ↑ in 2	Sleep	Phase OUT 1	Phase in 1	Phase in 2	Sleep	↓ Phase OUT 1	Phase in 1	Phase in 2	Exec aircmd 4	Slp.	Phase OUT 1	Phase ↑ in 1	Phase in 2	Sleep
--------------------	-------	----------------	---------------	---------------	-------	---------------------	---------------	---------------	---------------------	------	----------------	--------------------	---------------	-------


ODE 5

			Ans aircmd 5			Req aircmd 5			Ans aircmd 5		
--	--	--	--------------------	--	--	--------------------	--	--	--------------------	--	--

Phase in 2	Exec aircmd 5	Slp.	Phase OUT 1	Phase ↑ in 1	Phase in 2	Sleep	↓ Phase OUT 1	Phase in 1	Phase in 2	Exec aircmd 5	Slp.	Phase OUT 1	Phase in 1	Phase ↑ in 2	Sleep
---------------	---------------------	------	----------------	--------------------	---------------	-------	---------------------	---------------	---------------	---------------------	------	----------------	---------------	--------------------	-------


FIG 7D

FIGURE 8 - POLLING STRATEGY #4 - PARALLEL MISO



(comm. bus ignored for simplicity purposes) (ex. Bcast. In count = 2, Ns = 4)

[illegible]



TOR NODE

[illegible]

ODE 1

		Ans aircmd 1				Req aircmd 1			
Phase in 2	EXEC air- cmd 1	slp	Phase OUT 1	Phase in 1	Phase ↑ in 2	Sleep ↓ 3	Phase OUT 1	Phase in 1	Phase in 2

ODE 2

[illegible]

ODE 3




FIG 8C

			Req aircmd 3			Ans. aircmd 3									
Phase in 2	↑ 3	Sleep	↓ OUT 1	Phase in 1	Phase in 2	Exec air- cmd 3	Slp.	Phase OUT 1	Phase in 1	Phase ↑ in 2	Sleep	Phase OUT 1	Phase in 1	Phase in 2	Sleep
ODE 4															
				Req aircmd 4				Ans. aircmd 4							
Phase ↑ in 2	↑ 3	Sleep	Phase OUT 1	Phase in 1	Phase in 2	↑ 3	Sleep	↓ OUT 1	Phase in 1	Phase in 2	Exec air- cmd 4	Phase OUT 1	Phase ↑ in 1	Phase in 2	Sleep
ODE 5															
			Ans. aircmd 5			Req aircmd 5			Ans. aircmd 5						
Phase in 2	↑ 3	Exec air- cmd 5	Slp.	Phase OUT 1	Phase ↑ in 1	Phase in 2	Sleep	↓ OUT 1	Phase in 1	Phase in 2	Exec air- cmd 5	Phase OUT 1	Phase in 1	Phase ↑ in 2	Sleep

FIG 8D

FIGURE 9 - EXAMPLE BASIC REDUX PHASE ACCESS FLOWCHART

A very basic contention management mechanism for the redux slot (demonstrates the general idea)

This assumes that redux phase is only used for priority access requests. All other transactions should be done using regular broadcast phases to limit contention possibilities

Basic examples of improvements: a stacking system to store pending packets while back-off is in force





FIG. 9A

- EXAMPLE BASIC REDUX PHASE ACCESS FLOWCHART

A very basic contention management mechanism for the redux slot (demonstrates the general idea)

This assumes that redux phase is only used for priority access requests. All other transactions should be done using regular broadcast phases to limit contention possibilities

Basic examples of improvements: a stacking system to store pending packets while back-off is in force



ASYNCHRONOUS
PARALLEL EVENT: End node
to gateway priority access
needed

Calculate whether packet must be sent according to $F(\text{SEQ_R} - \text{FIRST_SEQ_R})$, where F is a 1 to 1 function mapping a number to a boolean value

A typical pseudo-code function is
$$F(x) = \lfloor (x \% \text{SKIPCOUNT}) - \text{SKIPCOUNT} \rfloor \leq 0$$

or
$$F(x) = \lfloor (x \% \text{SKIPCOUNT}) - \text{SKIPCOUNT} \rfloor < \text{MAX_COUNT}$$




FIG. 9B

RC SLEEP MODE
EXAMPLE SIMPLE IMPLEMENTATION





FIG. 10A



This example implementation is optimized for low-power microcontrollers with 16-bit counters or better.

Note: MeshClock is a 48 bit timer, usually implemented as a combination of 3 x 16 bit (or more) timers, that resets to 0 every broadcast cycle. Its value must be corrected according to the sync signal received by an end node from a gateway.

All time units are in terms of MeshClock, except for those in RTC ticks (i.e. for variables labeled RTCTICKS)

All calculations are fixed-point, with what that implies for arithmetic operations.

Hardware events can be captured with this timer (or timer combo) and the time of occurrence within 1 instruction cycle can be determined.

The function MeshTimedActions_GetLastAsyncCaptured() is our example function that fetches the last MeshClock value corresponding to the last captured event, precise to 1 instruction cycle.

The function RTCTicks() gets the current RTC value in RTC ticks.

This implementation checks to ensure that enough safeguard time is available before triggering sleep mode with RC clock as RTC as main sleep clock, so that there is no time desynchronization in the case that the sleep time is dangerously close to the expected main clock stabilization time, as well as an user-defined minimum (RTC_MIN_ACCEPTABLE_RTCTICKS).

It is assumed that RTC is a 16-bit timer. In case of a 32-bit timer or better, then the RTC overflow portion of the flowchart can most definitely be omitted.






FIG 10B

EXTERNAL SLEEP CONTROLLER (MESH MCU
SIDE, I2C master)





FIG. 11A








FIG. 11B

TIMING DIFFERENTIAL



#19.12



EXTERNAL SLEEP CONTROLLER
(External controller side, I2C slave)

FIG. 13A

Byte ID assignment
ADDR = 0x48
CMD_CLKSYNC = 0x35
CMD_STARTSLEEP = 0x69

Nomenclature
E_TMR: External sleep controller timer

PART 1: !




FIG. 13B

EXTERNAL SLEEP CONTROLLER

1: SLEEP INITIATION BY MASTER (MESH CONTROLLER)




FIG. 14A

Byte ID assignment
ADDR = 0x48

Nomenclature
E_TMR: External sleep controller timer

EXTERNAL S
PART 2: WAKEUP BY S





FIG 14B

LEAP CONTROLLER SLEEP SLAVE (EXTERNAL SLEEP CONTROLLER)



Source Code Exhibit 2 - Example VM code

FIGURE 15A

SOURCE CODE EXHIBIT - VM Client

```

#include "refdes.evi"
#define I2C_ADDR 0x91

//A mesh VM can be used to create a user-space sandbox in which a user (i.e. integrator)
//can safely implement application-specific code within the same mesh controller, without
//supplementary hardware needed.

//This is an example of a repeater node implementation, which only return 2 bytes when receiving
//a request for information from the coordinator nodes: RSSI (8bits), and voltage (8bits)

//at bootup, a special code living within the sleep controller allows user to change
//rf channel via a button, using a specific button-pressing sequence.

//Custom VM function executed at mesh controller bootup
function exec_bootup(){
    //Insert specific boot-up code here
}

//Custom VM function executed when end node mesh controller receives an "air cmd", which is
//a command from coordinator node to execute a VM function over the air
function exec_aircmd(){
    local V;

    //time to read voltage using I2C bus
    if(I2C_Start(ADDR) ){ //failure
        V=0;
    }
    else{
        V = I2C_ReadNak();
    }
    I2C_Stop();


    SetBuffer(0,GetRSSI(),1); //first byte to be sent is RSSI
    SetBuffer(1,V,1); //2nd byte to be sent is Voltage fetched from I2C bus
    Send(2); //send 2 bytes back to gateway
}

//In this particular implementation of a VM, the same VM codebase is executed with the same entry point
//independently of the event chosen. The last event type triggered is stored in the mesh controller
//memory, and can be fetched back with VM high-level language .GetExecType() function
function main()
{
    local execType;

    execType = GetExecType();
    if(execType==MESHEXECTYPE_BOOTUP_bm){
        exec_bootup();
    }
    else if(execType==MESHEXECTYPE_AIRCMD_bm){
        exec_aircmd();
    }
}

```


FIGURE 16 - ASYMMETRIC SNIFFER ENABLED NETWORK IMPLEMENTATION – PACKET CODEC




Note: Varia are optional bitstream spaces in the remaining available space within a packet that can accommodate a slower, higher-level protocol, such as for instance, encoding a 2nd binary message using BASE64 and split and spliced into multiple subsequent packets, with ASCII characters 13 or 10 to denote message end marker, for instance.

A sniffer can use the varia spaces to signal a priority request, or alternately such a flag can be embedded into said sniffer data variables (as such a sniffer node is also an end node like any other sensor end node).

FIGURE 17 - ASYMMETRIC SNIFFER ENABLED NETWORK IMPLEMENTATION –
GATEWAY STATE MACHINE – UDP STYLE POLLING


(simplified, no handling of pathological cases, no handling of dynamic mesh network configuration changes, HIGH-LEVEL PROTOCOL (1) mode summarized)



Notes:

- (1) HIGH-LEVEL PROTOCOL can refer to any higher-layer protocol that can be used for large file transfers between nodes, or VM firmware update over-the-air, for instance
- (2) In case of timeout, front-end is flagged that there is a hardware connection problem.

FIGURE 18 - ASYMMETRIC SNIFFER ENABLED NETWORK IMPLEMENTATION – END
NODE SNIFFER STATE MACHINE – UDP STYLE POLLING
(simplified, dynamic config checking enabled, HIGH-LEVEL PROTOCOL (1) summarized)







FIGURE 20 - Simple High-Level Protocol Model stack example for single gateway-node (Master-Slave) communication

	Data unit	True layer level	Code-Name	Function
Host layers (high-level protocols)	Data (File)	5+	5+. Custom	Custom layers can be superimposed on top of the previous layers. For instance, one can decide to emulate a full OSI stack by replicating the OSI layer from the OSI Layer 2 (Data Link Layer), with its full 64-bit MAC addressing scheme. Another example, in the case where only pure Master-Slave connectivity between a gateway and a specific end node is required (i.e. no end node to end node message routing is required, which would entail an intermediary layer to emulate full peer-to-peer mode), and without strict requirement for QoS (Quality of Service) such as found in TCP/IP, it is possible to send binary blobs such as camera images in this fashion, and also possible to emulate a terminal link with a command-line interface between the gateway and the end node. A typical use-case of this methodology can be embodied by a Command-line Interface (CLI) enabled router connected to a mesh controller, which can be remote controlled by the host connected to the gateway via CLI commands. In this case, the type of CLI stream as implemented in a typical computer Operating System (OS) can be encoded within the FILE_TYPE header of the previous protocol layer, in order to distinguish between CLI stream type, for instance (in/out streams, versus error streams, is a typical example).
	Segments (Chunks)	4. Message Layer	4. Smartrek RF Level 4	A layer in which a long bit-stream (also called "file") of any kind is segregated by type via a header type identifier FILE_TYPE. Post-pended to the PDU, i.e. in its footer, is a checksum/error-correction code in order to further safeguard the payload against corruption.
	Packet (Datagram)	3. Packet Segmentation Layer	3. Smartrek RF Level 3	A layer dedicated to segmenting longer message into shorter bitstreams that can fit within the payload of the precedent layer. Message ordering is enforced using a message sequence identifier in the PDU, called MSG_SEQ, and another sequence identifier, called "frame sequence ID", or FRAME_SEQ, is used to order the parts of a message that was splitted in order to fit individual RF packets (maximum of splitted components is thus dependent on the bit length of FRAME_SEQ). This layer can be designed to be symmetrical, i.e. the same protocol can be used for a gateway to end node message, or vice versa. The PDU also contains a PREAMBLE identifier, with identify the function of the PDU sent, with minimally includes: Frame Data Carry function (a part from a higher-level message is encapsulated in such a frame), Frame Count function (identifies the expected number of parts splitted from a given source message), Frame Cancel function (flags destination that a current send procedure has been interrupted), Buffer Flush function (to reset the buffers of the higher-level protocol, and thus reset the state machines used for protocol handling) and Frame Acknowledge function (used to answer back, i.e. acknowledge, a Frame Data Carry or a Frame Count function sent by the other peer). The typical behavior after a source peer sends a long message segmented into chunks in this fashion to another destination peer, is for the latter to send an acknowledgement every time it receives a message from the source peer, and accumulate the chunks in an ordered fashion inside of a holding structure (such as a tree map, in one possible embodiment). The source peer accumulates the acknowledgement PDUs from the destination peer, and at the end of its current message send cycle, it checks to see if any acknowledgement from the destination is missing. If so, it builds a list of chunks to be re-sent to destination, and re-sends them. It typically repeats this procedure until all chunks have been received, or after repeated failed deliveries condition is detected, in which case it considers its message delivery to have failed, and notifies its host of the condition (a simpler basic implementation can neglect this timeout condition, and any message send with failed acknowledgement reception by sender will trigger a sender retry until its protocol buffers are forcibly cleared by host, for instance).
				A mesh data connection with network dynamic parameters for broadcast cycle configuration, channel, FHSS hop table (if applicable to embodiment), error correction code, and data payload. The data payload (a.k.a. the SDU) is considered as being the "data frame" of the physical (a.k.a. media) layer, and the rest of the information sent is considered as being part of the "config frame" of said layer (by default contained in the header of the PDU, the footer containing the error correction code). Moreover, a MAC addressing scheme is typically implemented here, with optional multi-phase feature. Redux priority access requests signals are directly addressed by a given host to this layer.
Media layers (low-level protocols)	Bit/Frame	2. Physical Data Link	2. Smartrek RF Level 2	A cooperative bit-synchronized data connection for mesh networks, the standard of which has 1 gateway connected to multiple end nodes. Each transmission from a node to another is occurring within a <i>broadcast cycle</i> . Typical embodiment uses a FHSS FSK transceiver with a super-heterodyne receiver, and with a method for controlling sleep cycles to ensure that all end nodes are synchronized with the main gateway every broadcast cycle.
	Bit	1. Physical Raw	1. Smartrek RF Level 1	

FIG. 21A


General principle for multi-




-acknowledgement scheme

FIG 21B


Simplified multi-acknowledgement scheme, generic ideal case
(no delay between gateway requests and end node response)
(timeline as shown from the coordinator standpoint) for 1 coordinator G# and 4
end-nodes 1#, 2#, 3#, 4#



gment scheme, with realistic compensation for processing delays between a coordinator
nd the corresponding answer (timeline as shown from the coordinator standpoint)



Example special-case multiple acknowledgment scheme leveraging implicit end-r broadcast cycle delay, settings: $B_0 = 1$, $B_1 = 3 = 3$ end nodes (labeled 1



SUB-PACKET LEGEND	
B-(i#, j#, k#)	Coordinator command telling end-nodes which one is allowed to communicate during which phase. A B-(i#, j#, k#) command result in an answer at (X, Y+2, 0) by node i#, at (X, Y+2, 1) by j#, and at (X, Y+2, 2) by k# (as seen by coordinator node upon for communication prior to beginning this communication protocol mode).
Q-i#	Coordinator request for end node i send-back of their requested number of assigned transmission p
R-n	End node i# answer to Q-i# query, with n being the number of packet per level-1 cycle requested
P(X, Y, Z)	Custom payload (can consist of multiple sub-packets as implemented by the higher-level protocol u with an answer expected to occur at (X, Y + 2, Z)
A-bxxx	Sequence of acknowledgment sub-packets (content not explicitly listed) using bit-mask strategy (x being a bit). Those a explicit acknowledgement needed for a coordinator source packet). If reception at coordinator site oft that acknowledge for all phases that was sent by an end node at (X, Y - 2, ...), which thus corresponds to a prior coordinator order at (X, Y
A-(X, Y, Z) A-(X, Y, Z+1) ...	Decoded acknowledge sub-packet content listing

Because of the broadcast cycle delays due to host bus transmission of all nodes, the minimum number of broadcast cycles per synchronization cycle corresponds to the number requi any Q-i# requests sent at the beginning of said synchronization cycle, as well as for any end node to receive the acknowledge signals for all previous cycles. In case of acknowledgem chunk in any of the next "files" to be sent in its sequence of files to be segmented and sent via synchronization cycles.

Note that for this special case, network throughput is maximized if the "file" data sent per node per synchronization cycle are close to be of the same size. It is possible to use those "fi condition.

FIG. 22A

node to coordinator answers as acknowledgments for coordinator packets, 2 1#, 2#, 3#) and 1 coordinator (labeled G#) using high-level protocol




FIG. 22B


- GATEWAY AND SNIFFER DATABASE SYNCHRONIZATION

Legend

- > Node state and configuration variables
- > Data variables

Example of cloud server :
Using Google Drive the devices connected to the gateway sends XML containing Base Configuration of the gateway and then another devices fetches this data to put it in the Base Configuration of the "Sniffer".

Cloud-synchronized web client



Cloud-synchronized native client



Base configuration	Current State
MAC address	Node state (from gateway)
Type of device	Configuration variables (from gateway)
GPS coordinates	Data variables (from node)
Etc.	

Node registration procedure

QR Code on each node physical package



Scanning it provides information about the device

Use of a vision recognition device for barcode recognition (ex.: camera)

The device used to scan the QR code provides GPS data

- MAC address
- Bluetooth® MAC address (if applicable)
- Type of the device
- Serial (if applicable)

Node now registered in database with proper map positioning

Base configuration	Current State
MAC address	Node state (from gateway)
Type of device	Configuration variables (from gateway)
GPS coordinates	Data variables (from node)
Etc.	



Coordinator

Data/configuration variables and node state are synchronized on each broadcast cycle via flooding throughout network of nodes

Node

Node

Node

Node

Node replying

Sniffer



Base configuration	Current State
MAC address	Node state (from gateway)
Type of device	Configuration variables (from gateway)
GPS coordinates	Data variables (from node)
Etc.	

Same data but not synchronized unless the user asks for it. It is then synchronized via higher-level protocols, NFC, WIFI or cloud server.