

(51) International Patent Classification:  
**G01S 5/14** (2006.01) **H04B 7/26** (2006.01)(21) International Application Number:  
PCT/AU2011/000252(22) International Filing Date:  
7 March 2011 (07.03.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
2010900961 9 March 2010 (09.03.2010) AU(71) Applicant (for all designated States except US): **COMMONWEALTH SCIENTIFIC AND INDUSTRIAL RESEARCH ORGANISATION** [AU/AU]; Limestone Avenue, Campbell, Australian Capital Territory 2602 (AU).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **JURDAK, Raja** [LB/AU]; 212/26 Felix Street, Brisbane, Queensland 4000 (AU).(74) Agent: **WATERMARK PATENT AND TRADE MARK ATTORNEYS**; Level 2, 302 Burwood Road, Hawthorn, Victoria 3122 (AU).


(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: METHOD AND SYSTEM FOR ENERGY-EFFICIENT GPS LOCALISATION



(57) **Abstract:** A method (400) of estimating a location of a mobile node(100) in a network comprising a plurality of similar nodes. The mobile node(100) comprises a GPS receiver (114) and a radio communications transceiver (112) having a variable communications range for communication with proximate nodes. The method (400) includes activating (402) the GPS receiver (114) to determine an initial location estimate of the mobile node (100) at a first time instant, before deactivating (404) the GPS receiver (114) throughout a time interval subsequent to the first time instant. An ongoing estimate of uncertainty of location of the mobile node (100) during the time interval in which the GPS receiver (114) is deactivated is computed (406) based upon an estimated motion of the mobile node (100). The node communicates (410,412) via the radio communications transceiver (112) with one or more proximate nodes within a selected contact range, and receives (414) estimated location and uncertainty information from each such proximate node. The node (100) then updates its estimate of uncertainty of location based upon the received location and uncertainty information and the selected contact range.

**METHOD AND SYSTEM FOR ENERGY-EFFICIENT GPS LOCALISATION****FIELD OF THE INVENTION**

The present invention relates generally to localisation, or position estimation, of mobile nodes, such as mobile computing and/or sensor nodes. More particularly, the invention is directed to methods and apparatus for locating mobile nodes by fusion of GPS positioning measurements with contact logging in a network of mobile nodes.

**BACKGROUND OF THE INVENTION**

There are many circumstances in which it may be desirable to estimate and/or monitor the location of people, animals and/or objects using portable mobile computing or sensor devices. In one particular application of such technology, which provides (without limitation) a preferred embodiment of the present invention described herein, the outdoor tracking of livestock, such as cattle, may be used to monitor the position of animals, for example to ensure that they remain within an intended grazing zone, or alternatively outside an exclusion zone. In this application, the cattle may be fitted with "smart collars" that contain wireless sensor nodes and GPS modules. However, other applications include the tracking of people within public spaces and buildings, such as hospitals, museums and so forth, by way of mobile devices such as cellular phones, PDAs, and the like, or the tracking of tagged objects, such as containers within shipping yards, warehouses, and similar locations.

In recent times, the global positioning system (GPS) has become a widespread technology for localisation of people, animals, vehicles and other objects. The GPS system is able to provide localisation to a high degree of accuracy, eg typically within five metres using widely available commercial civilian receivers. However, a significant disadvantage of the GPS positioning system in mobile applications is the substantial power consumption associated with identifying, tracking, receiving and processing the GPS satellite signals. In portable devices, this power requirement leads to a rapid depletion of batteries. By contrast, in the case of the exemplary application of cattle tracking, animals may be relatively free to roam over a large and remote area, and may only be herded in for health checks, treatment and other procedures at relatively long intervals, such as six months. For reasons of cost, practicality and animal

welfare, the batteries powering the mobile sensor nodes (*ie* smart collars) are desirably of a minimum size. Similar considerations apply to mobile devices such as smart phones and PDA's, for which it is an inconvenience to the user if the batteries require unduly frequent recharging.

One solution to the battery life problem in GPS positioning devices is duty cycling of the GPS receiver. This refers to a process whereby a GPS receiver is activated only periodically, in order to obtain a position fix, but remains inactive for a majority of the time. While this extends battery life, a problem which arises in relation to continuously mobile nodes is the increasing uncertainty of location between GPS measurements.

It has been recognised in the prior art that within networks, or clusters, of mobile devices, measurements of relative locations of neighbouring devices may be used to reduce uncertainty. Communication or ranging between proximate devices in networks or clusters is known as "contact logging". In a publication to Chan *et al*, "Collaborative Localisation: Enhancing Wi-Fi – Based Position Estimation with Neighbourhood Links in Clusters", *Pervasive*, 2006, pages 50 to 66, simple ranging is used to improve accuracy of location estimation within an in-building Wi-Fi network, which suffers from inherent inaccuracies due to interference and other effects. US patent application publication no. 2009/0212995, to Wu *et al*, published on 27 August 2009, discloses a system in which mobile nodes share location estimates within a GPS-denied environment, in order to improve overall positioning accuracy. European patent application no. EP 2084553, in the name Massachusetts Institute of Technology, also published at WO 2008/115209 on 25 September 2008, discloses a method in which a network of mobile nodes share probability distribution information of their respective location estimates, in order to improve the position estimation of each node.

However, there remains a need for improved techniques for reducing uncertainty of location of nodes which use GPS duty cycling for periodic absolute positioning. The present invention seeks to address this need based upon a fusion of GPS measurements with contact logging, and with a particular goal of providing a means to control a trade-off between power consumption (and hence battery life) against positional uncertainty.

**SUMMARY OF THE INVENTION**

In one aspect, the present invention provides a method of estimating a location of a mobile node in a network comprising a plurality of similar nodes, wherein the mobile node comprises a GPS receiver and a radio communications transceiver having a variable communications range for communication with proximate nodes, the method comprising the steps of:

- activating the GPS receiver to determine an initial location estimate of the mobile node at a first time instant;

- deactivating the GPS receiver throughout a time interval subsequent to the first time instant;

- computing an ongoing estimate of uncertainty of location of the mobile node during said time interval based upon an estimated motion of the mobile node;

- communicating via the radio communications transceiver with one or more proximate nodes within a selected contact range, wherein said communicating includes receiving estimated location and uncertainty information from each said proximate node; and

- updating said estimate of uncertainty of location of the mobile node based upon said received location and uncertainty information and said selected contact range.

Advantageously, embodiments of the present invention enable reductions in the uncertainty of location of a mobile node between activations of a GPS receiver based upon the presence and corresponding location/uncertainty information of similar mobile nodes within a selected contact range. In contrast with prior art contact logging methods, embodiments of the present invention utilise the capacity of a radio communications transceiver to operate over a variable range (for example by varying transmission power and/or receiver sensitivity) such that the contact range is neither fixed nor arbitrary, but rather may be specifically controlled such that, in combination with location information available from other nodes within the contact range, this information may be used in improving location estimation and reducing location uncertainty.

In the step of computing an ongoing estimate of uncertainty of location, the estimated motion of the mobile node may be an assumed speed of the mobile

node, and uncertainty is thereby computed as the assumed speed multiplied by the elapsed time since the first time instant. In a simple implementation, the assumed speed may be a fixed value, such as a maximum expected speed of the mobile node, an average speed, or a typical speed. Alternatively, the assumed speed may be estimated from a measured motion of the mobile node during a recent prior time interval. In yet another alternative implementation, the assumed speed may be based upon a mathematical and/or statistical model of mobile node motion using measured motion of the mobile node during one or more prior time intervals.

In a particularly preferred embodiment, the contact range is selected based upon a statistical model of distance between mobile nodes in the network. Such a model may include a statistical distribution of distances between mobile nodes. Furthermore, the contact range may be dependent upon an instantaneous estimate of uncertainty of the location of the mobile node.

In one particularly advantageous implementation, the contact range is selected so as to maximise a statistical measure comprising a product of probability that at least one other node is within the contact range, and a probability that said at least one other node has a lower uncertainty in its estimated location than that of the mobile node.

It will be appreciated that optimisation of the contact range in order to enable a substantial reduction in positional uncertainty allows for a greater time interval between successive activations of the GPS receiver, for a given maximum allowable positional uncertainty, or "absolute acceptable uncertainty" (AAU). In turn, this facilitates a reduction in power consumption by the GPS receiver, and extended battery life in mobile nodes.

As such, in preferred embodiments the method includes a further step of re-activating the GPS receiver to determine a new location estimate at the end of the time interval. In order to achieve the aforementioned advantages in terms of power consumption and battery life, the duration of the time interval is dependent upon accumulated uncertainty of location, and in particular the GPS receiver is preferably activated when the computed estimate of uncertainty reaches or exceeds a predetermined AAU. The AAU may be a fixed value, or alternatively it may be location and/or time dependent. For example, in the case of a livestock

monitoring application it is particularly advantageous to reduce the AAU as an animal approaches a perimeter or boundary of an exclusion zone, to ensure that it does not enter the exclusion zone prior to a subsequent GPS positioning measurement. Conversely, when an animal is located at a great distance from any critical boundary or threshold, a larger AAU may advantageously be employed in order to reduce the frequency of GPS measurement, and hence further increase battery life.

Notwithstanding the aforementioned advantages of utilising a variable AAU, in some circumstances it may be acceptable to set the duration of the time interval in accordance with a predetermined (*ie* fixed) GPS activation duty cycle. Through the fusion of GPS measurements with contact logging using a variable contact radius, the GPS activation duty cycle may be longer than would be acceptable in the absence of such contact logging, again resulting in a reduction in power consumption, and improvement in battery life.

It should be noted that embodiments of the invention that employ a fixed assumed speed for uncertainty estimation, in combination with a fixed AAU, are equivalent to a system with a fixed predetermined duty cycle. That is, an appropriate fixed duty cycle may be determined by suitable selection of assumed speed and AAU. Conversely, embodiments employing variable speed estimation and/or a variable AAU will exhibit a variable GPS duty cycle, *ie* activation of the GPS receiver for positioning measurements “upon demand”.

In another aspect, the present invention provides a mobile node adapted to estimate its location when deployed in a network comprising a plurality of similar mobile nodes, the mobile node comprising:

- at least one microprocessor;

- a GPS receiver operatively coupled to the microprocessor;

- a radio communications transceiver operatively coupled to the microprocessor, the radio communications transceiver having a variable communications range, for communication with proximate nodes in said network;
- and

- at least one storage medium operatively coupled to the microprocessor, the storage medium containing executable instruction code which, when executed

by the microprocessor, cause the microprocessor to implement a method comprising the steps of:

- activating the GPS receiver to initiate a GPS location measurement;
- receiving an estimate of the mobile node location from the GPS receiver at a first time instant;

- deactivating the GPS receiver throughout a time interval subsequent to the first time instant;

- computing an ongoing estimate of uncertainty of location of the mobile node during said time interval based upon an estimated motion of the mobile node;

- determining a contact range for communication with proximate mobile nodes in the network;

- communicating via the radio communications transceiver with one or more proximate nodes within said contact range, wherein said communicating includes receiving estimated location and uncertainty information from each said proximate node; and

- computing an updated estimate of uncertainty of location of the mobile node based upon said received location and uncertainty information and said contact range.

In yet another aspect, the invention provides a method of estimating the locations of a plurality of mobile nodes which comprise a network of nodes adapted to communicate via respective radio communications transceivers having variable communication range, wherein one or more said mobile nodes is a GPS-equipped node comprising a GPS receiver, the method comprising the steps of:

- each of said plurality of mobile nodes maintaining a respective ongoing estimate of uncertainty of node location, based upon a corresponding estimated motion of the mobile node;

- each of said one or more GPS equipped nodes activating its GPS receiver to measure a respective node location estimate at a corresponding time instant, updating its ongoing estimate of uncertainty in accordance with said measurement, and deactivating its GPS receiver throughout a subsequent time interval;

each of said plurality of nodes communicating via its respective radio communications transceiver with proximate nodes within a selected contact range, wherein said communicating includes receiving estimated location and uncertainty information from each said proximate node; and

each of said plurality of nodes updating its respective estimate of uncertainty of node location based upon location and uncertainty information received from one or more proximate nodes, and its selected contact range.

Advantageously, the method enables relatively accurate position estimates obtained by GPS-equipped nodes to be shared among proximate nodes within the network, such that the operation and power consumption of GPS receivers may be reduced while maintaining an overall desired limitation on the accumulated location uncertainty of all nodes within the network. Furthermore, in embodiments in which all nodes are GPS-equipped, the method may be inherently fair, in the sense that requirement to perform GPS measurements is, on average, shared among the nodes in an unbiased manner. The reason for this is that, in accordance with preferred implementations of the method, a node which has just performed a GPS measurement will have a minimum uncertainty of location, while corresponding proximate nodes will be able to reduce their estimates of location uncertainty, albeit to a lesser degree. Due to the higher uncertainty in location of the proximate nodes, a fair and unbiased implementation of the method will result in a different node next activating its GPS receiver and, on average, in each proximate node performing an equal proportion of GPS measurement, each measurement being effectively shared with other proximate nodes.

Further preferred features and advantages of the invention will be apparent to those skilled in the art from the following description of a preferred embodiment of the invention, which should not be considered to be limiting of the scope of the invention as defined in the preceding statements, or in the claims appended hereto.

It will be understood that the word “comprises” and grammatical variations thereof, when used in this specification, is taken to specify the presence of stated features, integers, steps or components, but does not preclude the presence or addition of further features, integers, steps, components or groups thereof.

**BRIEF DESCRIPTION OF THE DRAWINGS**

A preferred embodiment of the invention is described with reference to the accompanying drawings, in which like reference numerals refer to like features, and wherein:

Figure 1 is a block diagram illustrating the major components of a mobile node according to a preferred embodiment of the invention;

Figure 2 is a schematic diagram illustrating a model of estimated location uncertainty according to the preferred embodiment;

Figure 3 is a further schematic diagram illustrating a general principal of sharing location information between mobile nodes in a network, according to the preferred embodiment;

Figure 4 is a flow chart illustrating a method of estimating a location of a mobile node in a network comprising a plurality of similar nodes, according to the preferred embodiment;

Figure 5 is a flow chart illustrating GPS duty cycling according to a preferred embodiment of the invention;

Figure 6 is a flow chart illustrating opportunistic contact logging according to the preferred embodiment;

Figure 7 is a graph illustrating a probability density function of relative distances between mobile nodes in an exemplary application of the invention;

Figure 8 is a graph illustrating a relationship between the number of mobile nodes in a network and the expectation of reducing location uncertainty as a function of a selected contact range, in the exemplary application;

Figure 9 is a graph illustrating the relationship between contact duration and distance in the exemplary application;

Figures 10(a) and 10(b) are graphs illustrating the impact of contact logging on GPS duty cycle in the exemplary application;

Figure 11 shows graphs illustrating the effect of selection of contact range for a fixed maximum target uncertainty, in the exemplary application;

Figure 12 shows graphs illustrating the effect of selection of contact range for a variable maximum location uncertainty, according to the exemplary application;

Figure 13 is a flow chart illustrating a method of adaptive duty cycling according to the preferred embodiment;

Figures 14(a) and Figure 14(b) show graphs illustrating the effect of preference for accuracy or energy minimisation, respectively, in an adaptive duty cycling scheme according to the preferred embodiment; and

Figure 15 is a block diagram illustrating a location management model for setting of contact radius and GPS duty cycle according to an embodiment of the invention.

### **DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS**

A preferred embodiment of the invention, as described herein, is directed to an outdoor location monitoring application for tracking cattle using “smart collars” that comprise mobile nodes embodying the present invention, each having a GPS receiver and a radio communications transceiver having a variable communications range for communication with proximate nodes (*i.e.* nearby smart collars). The goal in this application is to monitor cow position, and to ensure that each cow remains outside an exclusion zone, in effect implementing a “virtual fence”. The present disclosure is directed to the position monitoring aspect of this application, and accordingly details of the mechanism employed in order to influence cattle behaviour to remain outside the exclusion zone are not discussed herein. It will therefore be appreciated that the principles of the present invention are not restricted to the application of cattle monitoring, and may equally be applied in alternative embodiments to other applications, such as: the tracking of people within public spaces and buildings, such as hospitals, museums and so forth, by way of mobile devices such as cellular phones, PDAs, and the like; the tracking of tagged objects, such as containers within shipping yards, warehouses, and similar locations; and indeed to a variety of other location tracking applications in which similar mobile node technology may be employed.

Within the exemplary application of cattle tracking, a desired target node lifetime is around six months, which is the interval at which animals are brought in for health checks, treatment and sorting. Achieving this lifetime is a challenge, because the GPS module in each mobile node has substantial power requirements. Constant operation of the GPS receiver is not feasible, due to

rapid depletion of the battery. Battery size is limited by considerations such as animal welfare.

Duty cycling of the GPS receiver is one known method to reduce battery depletion, and thereby increase node lifetime. Periodic activation of the GPS receiver results in intermittent position fixes, between which time instants uncertainty of node position increases. A primary objective of the present invention is to reduce the growth in positioned uncertainty during periods of GPS receiver deactivation, in order to improve node lifetime and/or to reduce positional uncertainty, at least on average, for a given GPS duty cycle period.

Figure 1 is a block diagram illustrating a mobile node 100 embodying the present invention, and suitable for use in the abovementioned exemplary application. The mobile node 100 is generally an embedded microprocessor based system, that is sufficiently lightweight, and compact in size to be fitted to a smart collar.

The mobile node 100 includes a microprocessor 102, which is interfaced via an address/data bus 106 to a number of peripheral devices. Peripheral devices specifically relevant to the location monitoring application are shown in Figure 1, however it will be appreciated that in general additional peripheral interfaces may be provided, such as serial communications ports and the like, that have been omitted from the figure to improve clarity. It should also be appreciated that in some embodiments various components of the embedded system may be implemented using a microcontroller, or by a custom or semi-custom "system on chip" integrated circuit, such that the various components illustrated in Figure 1 need not be discrete devices, and may be elements of one or more integrated circuits.

The microprocessor 102 is interfaced via the bus 106 to a non-volatile memory/storage device 108, which in the exemplary application is typically a solid state device, such as read only memory (ROM), flash memory, or the like. The storage device 108 contains programs and data required for the operation of the mobile node 100, including various software components implementing functions according to a method embodying the present invention, as discussed in greater detail below.

The mobile node 100 further includes one or more additional storage devices 110, in the form of volatile memory, such as random access memory (RAM) interfaced via the address/data bus 106, for containing transient program instructions and data relating to operation of the mobile node 100.

The mobile node 100 further includes a radio communications transceiver 112, which has a variable communications range selectable under control of the microprocessor 102. The variable communications range may be implemented via a facility of the transceiver 112 to control signal transmission power and/or signal reception sensitivity.

The mobile node 100 further includes a GPS positioning receiver 114, which is operable under control of the microprocessor 102 to be activated as required, in order to obtain a GPS positioning fix (*i.e.* a measurement of location utilising the global positioning system), and to be deactivated, in order to conserve power, at other times.

The entire mobile node 100 is portable, and is powered by a battery 118.

At any given time during operation of the mobile node 100, the non-volatile 108 and volatile 110 storage devices collectively include program instructions and/or data 116, 117 that are executed and processed by the microprocessor 102 in order to implement various features of the present invention, as described in greater detail below with reference to the remaining drawings. In general, these features include operation of the radio communications transceiver 112 to select a communications range and communicate with proximate mobile nodes, and operation of the GPS receiver 114 to obtain a GPS position fix when required and to deactivate the receiver 114 in order to conserve power at other times. The program instruction code also includes instructions enabling calculation of a desired communications range for the transceiver 112, and for processing data received from the GPS receiver 114 and the radio transceiver 112, along with other relevant information, in order to compute and maintain an ongoing estimate of uncertainty of location of the mobile node 100.

Turning now to Figure 2, there is shown a schematic diagram 200 illustrating a model of location uncertainty according to a preferred embodiment of the invention. At some particular time instant, a GPS measurement is taken resulting in an accurate estimate of location 202. For present purposes, a

simplifying assumption is made that a GPS measurement produces a precise position, with no associated uncertainty. However, it will be appreciated that the actual uncertainty associated with a GPS measurement (typically in the order of five metres for civilian commercial receivers) may readily be incorporated into the uncertainty model of the present invention.

During a time interval subsequent to the time instant at which the GPS measurement is taken, the GPS receiver 114 is deactivated. Throughout this time interval the mobile node 100 may be in motion, and accordingly uncertainty in location increases. According to an appropriate estimate of speed of the mobile node, and assuming that direction of motion is random and unknown, at some later time the uncertainty may be represented by a circle 206 of appropriate radius, centred upon the initial measurement location 202. A maximum acceptable uncertainty, referred to as “absolute acceptable uncertainty” (AAU), is represented by the solid circle 208. In accordance with the preferred embodiment of the present invention, it is an objective to activate the GPS receiver 114 and obtain a subsequent location measurement before the uncertainty reaches the AAU, while avoiding unnecessary activation of the GPS receiver 114, that may result in unnecessary battery depletion. If the modelled or estimate of speed of the mobile node 100 exceeds the actual speed of the node, then the subsequent GPS measurement location 204 will be within the circle 208, representing the AAU. On the other hand, if the modelled estimate of speed of the mobile node 100 underestimates the actual speed then it is possible that the subsequent GPS measurement location 210 will be outside the radius of the circle 208, which is considered to be an error.

In accordance with the preferred embodiment of the invention, a primary performance metric is the ratio of errors to the total number of GPS measurements. It should be appreciated that the invention does not seek to estimate actual position during time intervals in which the GPS is deactivated, but rather to control the probability that each subsequent GPS measurement occurs within a defined radius, *ie* the AAU, of the previous measurement.

Figure 3 is a schematic diagram illustrating a model 300 of a general principle of collaborative localisation implemented by embodiments of the present invention. In particular, the model 300 represents a first node, the last GPS

position fix of which was at a location 302, and which has since moved to a new location 303. A second node has a last GPS position fix at a location 304, and for simplicity it is assumed that the node has not moved since this fix. The current uncertainty in location of the first node is represented by the dashed circle 306. This uncertainty is relatively large, as compared with the current uncertainty in location of the second node, represented by the smaller dashed circle 308. This may be, for example, because the second node has more recently obtained a GPS position fix than the first node, and/or because the estimated speed of the first node is greater than that of the second node.

The first node operates its radio transceiver 112 in order to communicate with all proximate nodes within a selected range (*ie* radius), represented by the solid circle 310. Since the majority of the circle 308 lies within the circle 310, there is a high probability that the second node is within this effective contact range. If so, then the first node contacts the second node, and the second node transmits its last known location 304, and current uncertainty estimate 308, to the first node. The first node is then able to calculate that its current location can be no further from the second node's last known location 304 than a distance equal to the sum of the second node's uncertainty 308 and the contact range. If this uncertainty is less than the first node's current uncertainty estimate 306 (as it is in the example 300), then the first node is able to improve its estimate of uncertainty in accordance with the further dashed circle 312. This reduction in uncertainty will generally enable the first node to increase the time until next activation of its GPS receiver 114.

Turning now to Figure 4, there is shown a flowchart 400 illustrating the preferred general method for estimating a location of a mobile node in a network of such nodes, such as a herd of cattle fitted with smart collars including nodes 100. At step 402, the GPS receiver 114 is activated, in order to obtain a GPS position fix of the current location of the node. The time taken to acquire a fix (*ie* the GPS lock time) is somewhat variable, including deterministic and random components depending upon the time for which the GPS receiver has been deactivated, as well as factors such as interference, received signal strength, and so forth. In practice, the total time in the exemplary application required to obtain GPS lock varies between a few seconds, and (rarely) in excess of two minutes.

In any event, once a GPS position fix has been obtained, the GPS receiver is deactivated at step 404.

Steps 406 to 414 shown in the flowchart 400 comprise a loop that is repeatedly executed during the time that the GPS receiver is inactive. It will be appreciated that the flowchart 400 represents an exemplary algorithm, and that other implementations and variations are possible, and such algorithms as would be apparent to persons skilled in the art are all within the scope of the present invention. In the algorithm represented by the flowchart 400 it is assumed that the steps 406 to 414 are performed periodically, in accordance with a uniform time unit.

At step 406, an uncertainty estimate is computed, corresponding with the end of the current time unit. At step 408, a comparison is performed to determine whether the current computed uncertainty (*ie* as will apply at the end of the current time unit, and preferably also taking into account an anticipated GPS lock time) will result in the total uncertainty exceeding the AAU. If so, then control is returned to step 402, wherein the GPS receiver is reactivated in order to obtain a new position fix.

Otherwise, control passes to step 410, in which an appropriate contact radius is selected.

At step 412 the node transceiver 112 is activated, with its communication range set to the selected contact radius, and an attempt is made to contact any other nodes within range. If no other nodes are within range before the end of the current time unit, control is returned directly to step 406.

If one or more proximate nodes is contacted, then at step 414 communication is initiated with each such node in order to obtain the proximate node's location and uncertainty estimates, which may be used to update the current node's uncertainty estimate. Further details of this process are described below, with reference to Figure 6.

Turning now to Figure 5, there is shown additional detail of the process for computing an uncertainty estimate at step 406. It is assumed that the algorithm 400 maintains a record of the GPS state, having at least two relevant values, namely "lock" indicating that a GPS position fix has just been measured, or "off"

indicating that the GPS receiver has been deactivated, and the most recent position fix is no longer valid (*ie* at least one time unit has since passed).

At step 502, a check is performed to determine whether the current GPS state is “lock”. If so, then the current position and uncertainty are updated at step 504. In particular, the current position is set to the GPS position fix value, the uncertainty is set to a minimum value, corresponding with the GPS accuracy (possibly being zero if it can be considered negligible). Optionally, a record is kept of the present time as the most recent lock time, and the GPS state is reset to “off”.

At step 506, a current speed estimate ( $S_c$ ) is computed. In accordance with three exemplary implementations, one of the following speed models is employed:

- a static model, based on speed statistics and judgement;
- a dynamic model, which is based on a last observed speed of the mobile node, computed from the difference in GPS co-ordinates over the last time interval; and
- a probabilistic model, which is based upon the last observed speed (as in the dynamic model) and a state model of the mobile node, which takes into account historical information regarding typical behaviour of node motion (*ie* known cattle behaviour).

Table 1 below summarises the three exemplary speed models. In the static model, a constant speed is assumed, which may be, for example, a conservative speed estimate such as maximum speed, an average speed estimate, a typical speed estimate, or some other estimate based upon general historical data and appropriate judgement. The dynamic speed model estimates the current speed either as the most recently observed speed over the last GPS fixed time interval, or a fixed speed (based on speed statistics and judgement as in the static model), whichever is the larger. The probabilistic model is a simple Markov model of cow speed, in which the animal has a slow moving and a fast moving state. The estimated speed is generally the last observed speed, as for the dynamic model, but at each time unit the speed is adjusted towards a constant speed (similar to the static model) in accordance with a first order filter having a time constant which is a function of the probability that the speed will

transition out of the initial state (*ie* the slow or fast moving state). The general reason behind this model is that, as the time since the last speed observation grows, its significance decreases until at some point, and in the absence of any better information, the estimate reverts to a constant assumed speed.

| <i>Static</i>   | <i>Dynamic</i>                                                     | <i>Probabilistic</i>                                                                                                                                                                |
|-----------------|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $S_e = \bar{s}$ | if ( $s(t) > \bar{s}$ )<br>$S_e = s(t)$<br>else<br>$S_e = \bar{s}$ | $i = t - \text{lastlocktime}$<br>if ( $i == 0$ )<br>$S_e = s(t)$<br>if ( $S_e > \bar{s}$ )<br>$P = t_{22}$<br>else<br>$P = t_{11}$<br>else<br>$S_e = P \times S_e + \bar{s}(1 - P)$ |
| (a)             | (b)                                                                | (c)                                                                                                                                                                                 |

Table 1: Exemplary node speed models

At step 508, the uncertainty is increased in accordance with the computed speed estimate, multiplied by the time unit.

Figure 6 is a flow chart showing greater detail of the contact logging and uncertainty update step 414. It is assumed that there are N proximate nodes within the contact radius, each of which is identified by an index i. The aim of the process 414 is to determine a minimum uncertainty  $u_{\min}$  and a corresponding last position estimate  $l_{\min}$  which are initialised at step 602 to be the present node's current uncertainty and position.

At step 604, the node uses its communications transceiver 112 to request uncertainty and position information from the next proximate node, i. At step 606, the received uncertainty and location information  $u_i$   $l_i$  are received. At step 608 a test is performed to determine whether the current minimum uncertainty  $u_{\min}$  is greater than the sum of the contact radius R and the uncertainty  $u_i$  of the proximate node i. In the case of a positive result, the node updates its current position and uncertainty information so that the position is set to be the received last known position  $l_i$  of the proximate node, and the uncertainty  $u_{\min}$  to the sum of  $u_i$  and the contact radius R.

At step 612, a check is performed to determine whether there are more proximate nodes (*i.e.*  $i$  is incremented and compared with  $N$ ). If so, then control passes back to step 604, otherwise the procedure terminates.

It should be noted that the flow chart in Figure 6 represents a conservative implementation of an uncertainty estimation/improvement algorithm. In particular, use of the contact radius  $R$  at steps 608 and 610 inherently reflects a worst case assumption, *i.e.* that the proximate node is at the edge of the contact radius. Alternative, less conservative, assumptions may be utilised. For example, assuming a uniform distribution of location, the expected distance to any node within the contact radius  $R$  is  $R/2$ , and this value could therefore be used in the calculations at steps 608 and 610. Alternatively, a statistical model based upon measured internode distances may be employed (*eg* the statistical data shown in Figure 7, and discussed in greater detail below, could be utilised). Finally, if the transceivers 112 in each of the mobile nodes are capable of performing range estimation, for example via RSSI, an improved estimate of the actual distance between proximate nodes can be obtained.

An appropriate method for selecting the contact radius at step 410 will now be described. In general, it may be observed that the contact radius is desirably chosen so that one or more proximate nodes are within range, enabling the exchange of uncertainty information. Furthermore, the presence of such nodes is only of value to the local node if at least one of the proximate nodes has a lower uncertainty. Conversely, the use of a large contact radius in an endeavour to maximise the number of proximate nodes becomes counter-productive, since the contact radius itself contributes directly to the uncertainty in location of the local node based upon information received from a proximate node (*i.e.* at step 608).

By a simple probabilistic argument, it can be shown that the optimum contact radius is that which maximises a quantity  $P_{use}$  defined by the following equation, wherein  $\hat{r}$  is the estimated optimum contact radius:

$$P_{use} = [1 - (1 - P(\hat{r}))^{(N-1)}] \times \frac{U(t)}{AAU}$$

$$\text{where } U(t) = \hat{r} + (t - t_{contact}) \times \bar{s}$$

$$\text{s.t. } U(t) < AAU$$

More particularly, in order to implement the condition that the uncertainty  $U(t)$  remains less than the AAU, the quantity  $P_{use}$  should be integrated over time, in the interval following activation of the GPS receiver until the expected time at which uncertainty reaches the AAU. That is, an optimum value for  $R$  (ie  $R = \hat{r}$ ) maximises the following integral:

$$\int_0^{\frac{AAU - \hat{r}}{\hat{r}}} P_{use} dt$$

In order to obtain the relevant probability that at least one node is within a given contact radius  $R$ , the probability distribution  $P(r)$  is required. Figure 7 shows graph 700 in which a histogram 706 represents distance between each pair of cows in a herd of 35 cattle as a proportion of time (ie probability) during a two day observation period. The x-axis 702 is distance, while the y-axis 704 is proportional to probability density. The histogram 706 therefore has the approximate form of the actual probability density function for internode distances in the herd of cattle. As can be seen, the internode distance is concentrated in the range of 0 to 20 metres, and suggests that setting contact radius at or below 20 metres may be highly beneficial for cooperative localisation, since the probability of internode distance being less than 20 metres is nearly 73%. Calculations of the integral above have therefore been performed using contact radius values of 5 metres, 10 metres, 20 metres and 30 metres, based upon the data shown in Figure 7, and for between 2 and 10 mobile nodes. The results of the calculations are shown in the graph 800 in Figure 8.

In Figure 8, the x-axis 802 of the graph 800 represents the number of mobile nodes, while the y-axis 804 represents the value of the integral above. The computed values for contact radius of 5 metres, 10 metres, 20 metres and 30 metres are shown by the curves 806, 808, 810 and 812 respectively. The results show a high dependence of the optimum contact radius upon the number of nodes, and in particular that larger numbers of nodes correspond with smaller optimal contact radius. For example, a 30 metre contact radius is never optimal, even with only two nodes, and there is little difference between a choice of 10 or 20 metres in contact radius in a two-node network. Between three and seven nodes, a 10 metre contact radius is optimum out of the four calculation sets, while

for eight or nine mobile nodes the results for both 5 and 10 metre contact radius are similar. For 10 or more nodes it is apparent that a 5 metre contact radius is optimal although the benefit over a 10 metre contact radius is likely to be minimal.

A further factor that must be taken into account, particularly with the use of a small contact radius, is contact duration. That is, if mobile nodes remain in close proximity, yet more in and out of contact range (which may happen for a small contact radius), then there is a risk of highly variable neighbour information resulting in reduced benefits from contact logging. Figure 9 is a graph 900 showing contact duration of the 35 cattle in the measured herd, wherein the x-axis 902 is the relative distance between a pair of cows, and the y-axis 904 is the time spent within the corresponding distance. The curve 906 represents the average contact duration, while error bars 908 are indicative of the variance in the data set. It is clear from the error bars that the contact durations are highly variable, however any two cows will stay within 10 metres' contact for nearly 100 seconds on average, and within 20 metres' contact for more than 6 minutes on average.

Figures 10(a) and (b) are graphs illustrating the impact of contact logging on GPS duty cycling for two nodes. Figure 10(a) illustrates a case in which each node independently tracks its uncertainty estimate, and acquires a GPS position fix whenever uncertainty approaches the AAU, *ie* equivalent to the method 400 shown in Figure 4, without implementation of steps 412 and 414. More particularly, the graph 1000 shows time on the x-axis 1002, and uncertainty on the y-axis 1004. The trace 1006 represents the uncertainty for node 1, while the trace 1008 represents the uncertainty for node 2. In each case, the uncertainty increases according to the relevant speed estimate until the AAU is approached, at which time a GPS position fix is obtained, and the uncertainty returns to zero. This results in five fixes for node 1 (*eg* at time instant 1010), and four fixes for node 2 (*eg* at time instant 1012).

This is compared with the graph 1020 in Figure 10(b), in which contact logging is employed with a contact radius of 10 metres. Again, the x-axis 1022 is time, and the y-axis 1024 is uncertainty. Trace 1026 represents uncertainty for node 1, and trace 1028 represents uncertainty for node 2. The uncertainty for node 2 approaches the AAU at time instant 1030, at which time it activates its

GPS receiver and obtains a position fix. Accordingly, at this instant 1030 the uncertainty at node 2 returns to 0. Node 1, which is in the vicinity of node 2, can then rely on the latest GPS fix to reduce its uncertainty to the 10 metre contact radius. A feature of this contact logging strategy is its inherent fairness, since nodes that have recently acquired a fix will have a smaller uncertainty estimate than their neighbours, such that another proximate node will be next to reach its maximum allowable uncertainty, and activate its GPS receiver at a subsequent time. For example, in the graph 1020 in Figure 10(b), node 1 reaches maximum uncertainty at subsequent time instant 1032, at which point it activates its GPS receiver to obtain a position fix, of which node 2 is then able to take advantage, now reducing its uncertainty to the 10 metre contact radius. This alternation continues, such that within the exemplary time window node 2 obtains a GPS position fix three times, while node 1 does so twice. Both nodes have therefore reduced the number of activations of their GPS receivers, and indeed the total number of GPS activations has been reduced from nine, in the absence of contact logging to five, with the use of contact logging.

Turning now to Figure 11, there are shown graphs 1100 and 1110, summarising the GPS duty cycle and error rate respectively. More particularly, the graph 1100 shows contact radius on the x-axis 1102, and GPS duty cycle (*ie* total proportion of time for which the GPS receiver is active) on the y-axis 1104. Similarly, the x-axis 1112 of graph 1110 is contact radius, while the y-axis 1114 is the error rate, expressed as a percentage of actual occasions upon which a GPS fix finds a node to have moved a greater distance than the AAU from the previously measured location (*ie* cases in which the calculated uncertainty has underestimated the true distance travelled). The two traces 1106 and 1116 represent the use of a static speed model, while traces 1108 and 1118 represent a dynamic speed model (see Table 1).

The results in Figure 11 demonstrate that for both speed models a minimum GPS duty cycle (*ie* minimum energy consumption) is achieved for a contact radius of 10 metres, whereas error rate is continuously improving with increasing contact radius. There is little difference in performance between the static and dynamic models in terms of optimum energy consumption, however the dynamic model has significantly improved error performance.

Recalling that the tracking data upon which the results in Figure 11 are based represent a herd of 35 cattle, it might be expected that the optimum contact radius would be 5 metres as indicated by the results in the graph 800 of Figure 8 for large numbers of mobile nodes. However, as previously discussed the benefit of a smaller contact radius is offset by variability in contact duration.

The results in Figure 11 are for a static (*ie* fixed) AAU. Figure 12 shows similar results when a dynamic AAU is utilised. In particular, these results represent a scenario in which cattle are to be tracked in order to implement a “virtual fence”, which enables the use of a larger AAU (and therefore less frequent GPS position fixes) for nodes located at a distance from the virtual fence line. As individual nodes approach the virtual fence, the AAU is reduced accordingly, to maximise the probability of obtaining a further GPS fix prior to crossing of the fence line. Assuming that animals within the herd spend a majority of time away from the fence line, it is expected that further reductions in GPS duty cycle may be achieved.

As in Figure 11, the two graphs 1200, 1210 in Figure 12 show contact radius on the x-axis 1202, 1212, and GPS duty cycle and error rate on the respective y-axis 1204, 1214. Traces for static speed model 1206, 1216 and dynamic speed 1208, 1218 are also shown.

In the case of a dynamic AAU, the use of a dynamic speed model results in a substantial reduction in GPS duty cycle 1208, as compared with a static speed model 1206. Assuming that each node spends a majority of its time distant from the virtual fence line, the corresponding dynamic AAU is large, such that the dynamic speed model is able to benefit from averaging actual node movement over long time periods. In this regard, it should be recognised that nodes change direction as well as speed over time, although the different speed models have no explicit directional dependency. There is no marked minimum in GPS duty cycle for the dynamic AAU case, as there was in the static case, this behaviour now being dominated by the fact that the AAU is large for a majority of the time, and thus increasing the contact radius has little impact on the frequency of GPS activation. Nonetheless, the full benefit of contact logging is achieved at a contact radius of 10 metres. In some implementations it is undesirable to utilise a larger context radius than necessary, due to potential increases in power

consumption from the use of higher radio transmit power and/or receiver sensitivity.

As with the case of static AAU, the graph 1210 demonstrates that increasing contact radius can improve error rate in the case of a dynamic AAU, for both static and dynamic speed models. The results in Figures 11 and 12 demonstrate that it is in general possible to minimise GPS duty cycle (*ie* energy consumption) by suitable selection of contact radius for any given speed model and selected AAU strategy. At the same time, it is always possible to reduce the error rate by increasing contact radius, and/or via selection of a different speed model or AAU strategy. Accordingly, it must always be possible to trade off GPS duty cycle and energy consumption in order to obtain an improved error rate. Figure 13 shows a flowchart 1300 of an adaptive duty cycling algorithm that may be employed in conjunction with the overall method 400 in order to allow a controlled trade off between energy consumption and error rate.

The inputs to the algorithm 1300 are the most recent and the previous GPS position fix coordinates. That is, the algorithm 1300 is preferably executed after each GPS fix, following step 402 in the method 400. At step 1302, the algorithm calculates the distance  $d$  between the most recent GPS coordinates and the previous GPS coordinates. At step 1304 this distance is compared with the AAU, and if it is greater than the AAU an error is recorded at step 1306.

At step 1308, an average error rate,  $p$ , is calculated (*ie* total number of recorded errors divided by total number of GPS fixes). This actual error rate is compared, at step 1310, with a target error rate  $p^*$ .

If the actual error rate exceeds the target error rate, then step 1312 is executed. This step assumes that the operator of the mobile network has indicated a preference either for accuracy, or for achieving a specific energy target. If the objective is to improve accuracy, then a default speed estimate used in the speed model is increased, which will increase the subsequent frequency of GPS fixes, thereby reducing the error rate. Alternatively, if a residual energy measure (*eg* remaining battery life) exceeds a target value at the current time, the default speed estimate may also be increased, since there is capacity to increase the frequency of GPS fixes, regardless of the operator's accuracy preference.

In the alternative situation (*ie* the actual error rate does not exceed the target error rate) control passes to step 1314, which determines whether the actual error rate is less than the target error rate (as opposed to being equal with it). If so, then there is capacity to allow the error rate to increase, and accordingly the default speed estimate is decreased at step 1316.

Advantageously, the amount by which the speed is increased in step 1312 may be different from the amount of decrease in step 1316. In particular, it is beneficial to set a larger magnitude of speed increase in step 1312 than the magnitude of speed decrease in step 1316, such that the system responds rapidly to an excessive error rate, with a slower energy recovery.

Figures 14(a) (b) show results for residual energy (measured as remaining battery charge in mAhr) and corresponding error rate using the algorithm 1300 with an error rate target of 5%. In Figure 14(a) the user preference is for accuracy, whereas in Figure 14(b) the user preference is for target energy consumption.

Turning first to Figure 14A(a), the graphs 1400 show time (in hours) on the respective x-axis 1402, 1410 and residual energy 1404 and error rate 1412 on the y-axis. The trace 1406 represents the target residual energy required for the battery to last for the duration of the exemplary 12 hour period. The trace 1414 represents the target error rate of 5%. As can be seen, with the user preference set to accuracy both the measured error rate 1416 and the real error rate 1418 generally converge to the target value. (Note that the measured rate indicates the node's online measurements, while the real error rate is the ground truth value from the tracking data). However, in order to achieve the target accuracy the actual energy consumption, as represented by the residual energy trace 1408, exceeds that required in order to achieve the 12-hour target, and accordingly, the battery is exhausted at around 10 hours.

The graphs 1420 in Figure 14(b) again show time in hours on the respective x-axis 1422, 1430 and residual energy 1424 and error rate 1432 on the y-axis. The target residual energy is represented by trace 1426, while the target error rate is represented by trace 1434. With user preference set to energy, it can be seen that the actual energy usage shown by the trace 1426, closely tracks the target, such that the desired 12-hour battery life is achieved. By comparison,

however, the measured error rate 1436 and real error rate 1438 both substantially exceed the target 1434.

In addition to the formal optimisation of the contact radius (via maximisation of  $P_{use}$ ), each mobile node may apply a number of additional criteria, based upon heuristics and/or practical limitations of the monitoring application, in order to adjust the contact radius. Exemplary additional criteria include:

- **Size of monitored area** – the size of the monitored area sets realistic upper bounds on the maximum contact radius that can be set. In addition, the ratio of the current contact region (determined by the current contact radius) and the monitored area represents the proportion of the area covered with a node's current radio signal. If the node's GPS duty cycle is higher than the target GPS duty cycle, then the node can increase its contact radius to try to cover a larger portion of the total monitored area. However, increasing the contact radius to cover too large a proportion of the total monitored area could provide little extra benefit. Similarly, if a node determines that its GPS duty cycle is lower than the target duty cycle, then the node can afford to reduce its contact radius, thereby reducing the proportion of the total monitored area covered by contact beacons.
- **Number of nodes in the deployment** – the number of nodes in the deployment enables a node to determine the proportion of nodes covered within its current contact radius, and to determine when it is beneficial to increase or decrease that proportion from the perspective of localization accuracy and energy efficiency. More specifically, a node can compare its current number of neighbours to the total number of nodes in the deployment. This is the current proportion of nodes covered by the node's current contact beacons. If the node's GPS duty cycle is higher than the target GPS duty cycle, then the node can increase its contact radius to try to cover a larger proportion of the nodes in the deployment. This should only take place if the predicted increase in covered nodes, based on the statistical distribution of separation distances between the nodes (*eg* as represented by the

histogram 706 in Figure 7), is significantly higher than an appropriate threshold. If a node predicts that an increase in its contact radius will not result in a significant increase in the proportion of covered neighbours, then it should not change its contact radius.

- **Current node velocity** – the node's current velocity may be used in predicting how the node location may change in the near future. A node can use its current speed to determine  $P_{use}$  for various contact radii, and select the contact radius that maximizes the area under the  $P_{use}$  curve in order to maximize the chance of useful contact with neighbours. This process can be repeated dynamically throughout a deployment and processed locally at each node.
- **Proximate node velocity** – tracking the node's velocity along with that of proximate nodes (within the current contact radius) may be used to provide the node with a prediction of relative location of proximate nodes, and corresponding neighbourhood size in the near future. For instance, if most of a node's neighbours have a velocity of zero, and the node itself has a velocity of zero, then it can predict that the current neighbours will continue to be neighbours in the near future. The speed/velocity of a node relative to its neighbours therefore enables prediction of the number of neighbours in the next period. If a node predicts that its number of neighbours will decrease based on its speed relative to its neighbours, then it can choose to increase its contact radius pre-emptively to counter this effect. Conversely, if the node predicts that its neighbourhood size will be greater than required in coming periods (*eg* if the GPS duty cycle is below target), then the node can reduce its contact radius resulting in a reduction of neighbourhood size.
- **Clustering patterns** – mobile entities (*eg* cattle within a herd) may tend to exhibit distinguishable clustering patterns (*ie* historical patterns of relative location and movement) that may be learned by an online algorithm to refine the prediction of future neighbourhood size, which in turn may be used to adapt the contact radius. The online learning algorithm may operate as follows. Each node maintains a record of

the proportion of time that it has been in contact with any neighbour, and ranks these neighbours accordingly. Based on these ranks, each neighbour is assigned a reliability measure, and the reliability measure can change dynamically during the deployment. When a node periodically considers whether to change its contact radius, it uses the reliability scores in its forecasting of neighbourhood size. For instance, if a node has five neighbours, with three highly reliable and two unreliable neighbours, and all nodes are moving at the same speed, then the node can safely predict that it will have three neighbours in the next period for contact logging purposes. If its current contact radius with five neighbours is borderline in terms of achieving the target GPS duty cycle, it can pre-emptively increase its contact radius in anticipation of its neighbourhood size decrease.

- **Terrain map/proximity of obstacles** – an *a priori* known terrain map can be loaded into each node's memory 108 to enable the node to determine that obstacles are nearby and to pre-empt loss of connectivity with neighbours by increasing its contact radius, and/or to decrease contact radius when it is in obstacle-free terrain to reduce its uncertainty from neighbour position estimates. For instance, a node that is moving into an elevated region of the deployment area may start to lose connectivity with some neighbours on lower lying areas, due to antenna plane mismatches. In anticipation of this, the node can increase the transmit power on its transceiver to increase the likelihood that it will log contact with enough neighbours. Similarly, if a node is entering an area with heavy foliage, it may need to significantly increase its transmit power to counter the attenuation effect of the trees. Alternatively, if the node has below target GPS duty cycle, it can maintain its current transmit power as it enters the dense foliage area, which results in an effective decrease in its contact radius.
- **Environmental conditions** – air temperature, air humidity, rain and other environmental conditions can affect radio propagation and coverage of contact logging signals, and corresponding adjustment of contact radius may be employed to maintain a stable neighbourhood

size in the presence of changing environmental conditions. Apriori developed models for signal propagation in relation to various environmental conditions can serve as an additional input into a node's decision on setting its transmit power for a target contact radius. For instance, as humidity increases, radio signals propagate more poorly. As a result, maintaining a stable contact radius requires an increase in the transmit power. If it happens that the node had decided to reduce its contact radius when it detects increased humidity, it may decide to maintain its transmit power at current levels. Similar adjustments for other environmental factors are also possible.

Figure 15 is a block diagram 1500 illustrating the manner in which the above criteria feed into the setting of the contact radius and GPS duty cycle. The user policy 1502, terrain model 1504, environmental model 1506, and deployment features 1508 serve as pre-deployment inputs to the location management model (LMM) 1510, which duty cycles the GPS and sets the contact radius 1512. The LMM 1510 also relies on dynamic inputs that change during the deployment, including neighbourhood information 1514 and node information 1516 to determine the GPS duty cycle and the contact radius 1512 for the next period. These settings impact the energy consumption of the node, as determined by an energy model 1518, which feeds back the updated energy budget into the LMM 1510.

The above criteria may be implemented by piggybacking information on radio beacon messages (*eg* for current node speed information), the provision of onboard sensors (*eg* for monitoring environmental conditions) and/or by incorporating locally available information into a database held in node memory 108 (*eg* a terrain map). At any time instant, each mobile node can determine the best contact radius to use in order to maximum the chance that there will be a neighbour in range with a useful location estimate, in order to maximise deactivation periods of the node's GPS receiver.

In one embodiment, dynamic setting of contact radius is implemented via a central coordinator, such as a network gateway, to compile all available information from nodes periodically, and to determine an optimal network wide contact radius for maximizing the average  $P_{\text{use}}$  over the network. In an alternative

embodiment, each node is configured to select its own optimal contact radius in collaboration with its neighbours. This may be achieved through message exchanges and periodic negotiations among nodes for dynamic setting each of their optimal contact radii.

While the foregoing description has covered various exemplary features of a preferred embodiment of the invention, it will be appreciated that this is not intended to be exhaustive of all possible functions and features provided within various embodiments of the invention. For example, the various different speed models, AAU strategies, and adaptive duty cycling algorithms may be combined in a variety of ways in order to create embodiments offering different energy and accuracy trade offs, as may be appropriate for different applications (*eg* for tracking of people, shipping containers, and so forth). Furthermore the speed models and other strategies disclosed herein are intended as examples of possible approaches, and a range of alternative models, algorithms and strategies may be derived while still employing the general principles of the invention. The scope of the invention is thus not to be limited to the preferred embodiment as described herein, but rather is as defined by the claims appended hereto.

CLAIMS:

1. A method of estimating a location of a mobile node in a network comprising a plurality of similar nodes, wherein the mobile node comprises a GPS receiver and a radio communications transceiver having a variable communications range for communication with proximate nodes, the method comprising the steps of:

activating the GPS receiver to determine an initial location estimate of the mobile node at a first time instant;

deactivating the GPS receiver throughout a time interval subsequent to the first time instant;

computing an ongoing estimate of uncertainty of location of the mobile node during said time interval based upon an estimated motion of the mobile node;

communicating via the radio communications transceiver with one or more proximate nodes within a selected contact range, wherein said communicating includes receiving estimated location and uncertainty information from each said proximate node; and

updating said estimate of uncertainty of location of the mobile node based upon said received location and uncertainty information and said selected contact range.

2. The method of claim 1 wherein, in the step of computing an ongoing estimate of uncertainty of location, the estimated motion of the mobile node is an assumed speed of the mobile node, and uncertainty is computed as the assumed speed multiplied by the elapsed time since the first time instant.

3. The method of claim 2 wherein the assumed speed is a fixed value.

4. The method of claim 2 wherein the assumed speed is estimated from a measured motion of the mobile node during a recent prior time interval.

5. The method of claim 2 wherein the assumed speed is based upon a mathematical and/or statistical model of mobile node motion using measured motion of the mobile node during one or more prior time intervals.
6. The method of claim 1 wherein the contact range is selected based upon a statistical model of distance between mobile nodes in the network.
7. The method of claim 6 wherein the contact range is further dependent upon an instantaneous estimate of uncertainty of the location of the mobile node.
8. The method of claim 7 wherein the contact range is selected so as to maximise a statistical measure comprising a product of probability that at least one other node is within the contact range, and a probability that said at least one other node has a lower uncertainty in its estimated location than that of the mobile node.
9. The method of claim 6 wherein including a step of adjusting the selected contact range in accordance with one or more of the following criteria: an overall geographical size of the network; a total number of nodes in the network; a current communications range of the radio communications transceiver; a current node velocity estimate; one or more proximate node velocity estimates; one or more known node clustering patterns; one or more characteristics of local terrain; one or more measured environmental conditions and/or changes therein; and a distance from a predetermined virtual or real geographic boundary.
10. The method of claim 1 including a further step of re-activating the GPS receiver to determine a new location estimate at the end of the time interval.
11. The method of claim 10 wherein the GPS receiver is activated when the computed estimate of uncertainty reaches or exceeds a predetermined value AAU.
12. The method of claim 10 wherein said predetermined value of uncertainty is location and/or time dependent.

13. A mobile node adapted to estimate its location when deployed in a network comprising a plurality of similar mobile nodes, the mobile node comprising:

at least one microprocessor;

a GPS receiver operatively coupled to the microprocessor;

a radio communications transceiver operatively coupled to the microprocessor, the radio communications transceiver having a variable communications range, for communication with proximate nodes in said network; and

at least one storage medium operatively coupled to the microprocessor, the storage medium containing executable instruction code which, when executed by the microprocessor, cause the microprocessor to implement a method comprising the steps of:

activating the GPS receiver to initiate a GPS location measurement;

receiving an estimate of the mobile node location from the GPS receiver at a first time instant;

deactivating the GPS receiver throughout a time interval subsequent to the first time instant;

computing an ongoing estimate of uncertainty of location of the mobile node during said time interval based upon an estimated motion of the mobile node;

determining a contact range for communication with proximate mobile nodes in the network;

communicating via the radio communications transceiver with one or more proximate nodes within said contact range, wherein said communicating includes receiving estimated location and uncertainty information from each said proximate node; and

computing an updated estimate of uncertainty of location of the mobile node based upon said received location and uncertainty information and said contact range.

14. A method of estimating the locations of a plurality of mobile nodes which comprise a network of nodes adapted to communicate via respective radio communications transceivers having variable communication range, wherein one

or more said mobile nodes is a GPS-equipped node comprising a GPS receiver, the method comprising the steps of:

- each of said plurality of mobile nodes maintaining a respective ongoing estimate of uncertainty of node location, based upon a corresponding estimated motion of the mobile node;

- each of said one or more GPS equipped nodes activating its GPS receiver to measure a respective node location estimate at a corresponding time instant, updating its ongoing estimate of uncertainty in accordance with said measurement, and deactivating its GPS receiver throughout a subsequent time interval;

- each of said plurality of nodes communicating via its respective radio communications transceiver with proximate nodes within a selected contact range, wherein said communicating includes receiving estimated location and uncertainty information from each said proximate node; and

- each of said plurality of nodes updating its respective estimate of uncertainty of node location based upon location and uncertainty information received from one or more proximate nodes, and its selected contact range.

1/11




Figure 1






Figure 2



3/11




Figure 4

4/11




Figure 6

5/11




Figure 7




Figure 8

6/11




Figure 9




Figure 10(a)

7/11




Figure 10(b)




Figure 11

8/11




Figure 12

9/11






Figure 13






Figure 15

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2011/000252

## A. CLASSIFICATION OF SUBJECT MATTER

Int. Cl. *G01S 5/14* (2006.01) Int. Cl. *H04B 7/26* (2006.01)

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

## B. FIELDS SEARCHED

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

EPODOC, WPI, TXTE, USPTO, INTERNET (estimate, location, GPS, wake, sleep, speed, transceiver, range and like terms)

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                                                        | Relevant to claim No.        |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| X<br>Y    | JIANG, L. et al. "SenSearch: GPS and Witness Assisted Tracking for Delay Tolerant Sensor Networks" AD HOC, MOBILE AND WIRELESS NETWORKS Lecture Notes in Computer Science, 2009 Vol. 5793/2009 pp 255-269<br>Whole document, especially Sections I, II, III, IV (A, B, E), figure 1<br>Abstract, figure 6 | 1, 10, 12<br>2-5, 11, 13, 14 |
| Y         | US 6581546 B1 (DALLAND et al.) 24 June 2003<br>Abstract, column 1 lines 55-64, column 2 lines 41-67, column 4 lines 5-12, column 9 lines 34-67, column 11 lines 14-67, column 15 line 34 to column 18 line 59, figures 9, 10 and 12                                                                       | 13,14                        |
| Y         | WO 2006/099632 A2 (QUALCOMM INCORPORATED [US/US]) 21 September 2006<br>Abstract, paragraphs 9, 59-61, 71, 72, 77, 80-85 and 92-96, figures 3, 6, 7 and 11, claims 7-11                                                                                                                                    | 2, 4, 5,<br>11,13,14         |
| Y         | SIKKA, P. et al. "Wireless sensor devices for animal tracking and control" Proc. of the 29 <sup>th</sup> Annual IEEE International Conference on Local Computer Networks (LCN'04), 16 November 2004<br>Abstract, Sections 1, 2.2, 3.3                                                                     | 13,14                        |

☒ Further documents are listed in the continuation of Box C☒ See patent family annex

|                                                                                                                                                                         |                                                                                                                                                                                                                                                  |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| * Special categories of cited documents:                                                                                                                                |                                                                                                                                                                                                                                                  |
| "A" document defining the general state of the art which is not considered to be of particular relevance                                                                | "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention                                              |
| "E" earlier application or patent but published on or after the international filing date                                                                               | "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone                                                                     |
| "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art |
| "O" document referring to an oral disclosure, use, exhibition or other means                                                                                            | "&" document member of the same patent family                                                                                                                                                                                                    |
| "P" document published prior to the international filing date but later than the priority date claimed                                                                  |                                                                                                                                                                                                                                                  |

Date of the actual completion of the international search  
17 May 2011Date of mailing of the international search report **26 MAY 2011**

Name and mailing address of the ISA/AU

AUSTRALIAN PATENT OFFICE  
PO BOX 200, WODEN ACT 2606, AUSTRALIA  
E-mail address: pct@ipaustalia.gov.au  
Facsimile No. +61 2 6283 7999

Authorized officer

**DALE E. SIVER**AUSTRALIAN PATENT OFFICE  
(ISO 9001 Quality Certified Service)  
Telephone No : +61 2 6283 2196

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2011/000252

| C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT |                                                                                                                                                                                                                                                                   |                       |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Category*                                             | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                | Relevant to claim No. |
| Y                                                     | SWAIN, D.L., et al. "Using high fix rate GPS data to determine the relationships between fix rate, prediction errors and patch selection" Biological Modelling 212 (2008) pp 273-279<br>Abstract, Sections 2.1, 2.2, 2.3, equation 5 and Section 3.1, Conclusions | 2, 3, 5, 13,14        |

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/AU2011/000252**

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

| Patent Document Cited in<br>Search Report                                                                                 |            | Patent Family Member |            |    |            |
|---------------------------------------------------------------------------------------------------------------------------|------------|----------------------|------------|----|------------|
| US                                                                                                                        | 6581546    | AU                   | 2003215152 | WO | 03069363   |
| WO                                                                                                                        | 2006099632 | AU                   | 2006225107 | CA | 2601203    |
|                                                                                                                           |            | US                   | 2006211430 | US | 7702370    |
|                                                                                                                           |            |                      |            | EP | 1859294    |
|                                                                                                                           |            |                      |            | US | 2011070896 |
| Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001. |            |                      |            |    |            |
| END OF ANNEX                                                                                                              |            |                      |            |    |            |