An Overview of Positioning by Diffusion

Michael Spratt
Mobile Systems and Services Laboratory
HP Laboratories Bristol
HPL-2001-207
September 9th , 2001*

positioning, position based services, location, bluetooth, mobile services

Positioning by Diffusion is a method by which mobile and static devices can calculate their positions, by means of exchanging information over short range wireless links. It is intended for deployment within shopping malls, corporate buildings, theme parks etc. The implementation is in software, and does not impose extra hardware requirements on mobile or static devices. Another objective is to minimize the system administration and deployment costs, by minimizing the required number of devices which have their position programmed into them – devices we define as Access Points – for the system to operate. The positioning system works instead by seamless inter-working of the algorithm across Access Points, static devices which learn their position rather than having it pre-programmed, and mobile devices.

* Internal Accession Date Only Approved for External Publication
© Copyright Hewlett-Packard Company 2001

An Overview of Positioning by Diffusion

Michael Spratt, Hewlett-Packard Laboratories Bristol.

Abstract

Positioning by Diffusion is a method by which mobile and static devices can calculate their positions, by means of exchanging information over short range wireless links. It is intended for deployment within shopping malls, corporate buildings, theme parks etc.. The implementation is in software, and does not impose extra hardware requirements on mobile or static devices. Another objective is to minimize the system administration and deployment costs, by minimizing the required number of devices which have their position programmed into them – devices we define as Access Points – for the system to operate. The positioning system works instead by a seamless inter-working of the algorithm across Access Points, static devices which learn their position rather than having it pre-programmed, and mobile devices.

1 Introduction

This paper is a general overview of Positioning by Diffusion, a method for mobile and static devices to calculate their position from communications with other devices, across short-range wireless links. It is aimed at an in-building environment, and is characterized by low complexity of implementation, and low overhead in system administration.

Positioning within buildings is a relatively new idea, in contrast to positioning outside of buildings (often called 'Location' in this context), which is fairly well established. Several Location techniques have been standardized for use in conjunction with mobile cellular technology, and the much more accurate GPS is also becoming much more widely used. However, these technologies will not work well within buildings, and the cellular technologies do not have sufficient accuracy. Consequently new ideas have to be sought, based on short range wireless technology.

Some of the applications of cellular and GPS Location will also be appropriate within buildings. For example, positioning in the event of a request for help in an emergency is a service relevant to shopping malls, airports, hospitals and so on. Other classes of application, such as enabling visitors within corporate buildings quickly find facilities and services, or to automatically locate the nearest suitable printer for example, or allowing hot-desking roving workers to be easily located by co-workers, are applications specific to the indoor as opposed to the out of doors environment.

One option for in-building positioning systems is to densely cover the area with fixed devices which have been programmed to hold their position – devices which we term Access Points. However, deploying many Access Points involves a lot of administration in terms of programming into each Access Point its position, as well as installation and maintenance overhead.

The alternative approach described in this paper minimizes the number of Access Points required for the system to operate, by taking advantage of the existence, in the near future, of the multitude of short-range wireless enabled stationary devices. For example, in the corporate environment PCs, printers, scanners, photocopiers etc. are in the near future likely to be enabled with short-range wireless connectivity. Such devices are termed Static Devices.

The Positioning by Diffusion approach is characterized by a seamless inter-working of the positioning algorithm across Access Points, Static Devices as described, and Mobile Devices. The actual movement of Mobile Devices is used to diffuse position information from Access Points – which have had their position programmed into them – to Static Devices which the former passes close to. Subsequently, other Mobile Devices passing close to Static Devices can obtain this stored position information to perform positioning calculations. The position data originally emanating from a small number of Access Points can be combined to perform positioning calculations.

The Positioning by Diffusion is also seamless, in that – given a small number of Access points - it can allow communicating Static Devices in the absence of Mobile Devices, to perform position calculations. In a similar way, in the absence of Static Devices, Mobile Devices can often perform reasonable position estimates, by exchanging among themselves position data, which was originally sourced at the Access Points.

Positioning by Diffusion is also characterized by low device complexity, in that a device enabled with short-range wireless – Bluetooth for example – will implement the Positioning by Diffusion in a small software module, and will not require any extra hardware. In particular, no changes to antennae are required.

The Positioning by Diffusion is aimed mainly at in-building deployment, where GPS cannot operate well, if at all. Positioning by Diffusion could be deployed within corporate buildings, or government or administrative buildings. It could also be usefully deployed in shopping malls, hospitals, airports, tourist attractions and theme parks, city centers and university campuses, with short range wireless coverage.

The typical applications will vary from deployment to deployment. Within a corporate building there is an increasing tendency for people to 'hot-desk', in other words come into the building several times per week, sit in a different location each time, and plug in a laptop. A positioning system could allow them to be quickly found by colleagues. Another application of positioning is asset tracking of electronic office equipment. Similarly, within the medical, engineering and scientific environment, the tracking of expensive equipment is important.

Within the mobile communications industry, a key marketing aim has been the personalization of the user experience. Within a shopping mall for example, the ability to position a Mobile Device allows such personalized service as sending the mobile user text or audio information about new products, prices, discount offers etc. which are in the user's immediate vicinity.

Within campuses, corporate buildings, tourist attractions, theme parks etc. an obvious application of positioning is allowing visitors to navigate through an unfamiliar environment. There are indeed many other positioning applications, and additionally many other services which can be enhanced by incorporating the positioning information.

Another important application of positioning technology is as the prerequisite first stage in Geographic Routing, particularly for Ad Hoc Networks [2], [3]. [3] and [4] report positioning algorithms, which assume both continuous connectivity between adjacent devices and accurate distance measurement based on hardware such as Time of Arrival. The assumptions for Positioning by Diffusion are more general, in that a more dynamic environment is assumed in which connectivity is not always continuous, due to the movement of Mobile Devices, and extra hardware to very accurately measure distance is not assumed to be available.

A couple of points about the algorithm are relevant. The (probably novel) approach of using the actual movement of mobile users to in effect carry out diffusion of position information to Static Devices opens up the possibility of using the geometric method, described in this paper, of calculating the position estimate. The method is probably an optimal calculation, given the system assumptions, of the position estimate.

The Positioning by Diffusion uses a slightly non-obvious way of conveying positioning information from device to device. Initially the best idea appears to be for devices to send to other devices the best estimate of their position. There are several drawbacks with this, one being that it may not cope well with very heavily asymmetric patterns of movement of mobile users and uneven distributions of Access Points. Another drawback is that if the output of the positioning calculation for one device is used as the input to the positioning calculation for another device, any inaccuracies in the algorithm will tend to build up cumulatively within the distributed system of calculations. Instead, in Positioning by Diffusions, Static Devices and Mobile Devices will communicate to each other the raw information about the Access Points, to allow calculations based directly on the source, rather than processed data.

Bluetooth ([5] and [6]) is an example of a short-range wireless network over which Positioning by Diffusion could operate, and indeed Positioning by Diffusion has been submitted to the Bluetooth Local Positioning Working Group as an algorithm for implementing the positioning. However, Position by Diffusion is in general independent of the short-range wireless technology, and could be envisaged to operate in combination with various other short-range wireless networks.

2 A Generic Overview of the Operation of the Algorithm

In this section, a high-level summary is given, of how the Positioning by Diffusion algorithm operates, focusing in a sense on the 'Diffusion' aspect of the algorithm. As discussed before, Access Points have an inherent knowledge of their position, and may well have been programmed by an administrator to hold accurate position information.

Since there is some administrative overhead in ensuring the Access Points each have their position accurately programmed into the device, a motivating assumption has been that there should be a requirement for only minimal number of Access Points to make the system operate. Access Points may be stand-alone devices, or the Access Point functionality could be incorporated into another device, such as a device allowing short-range wireless access to LAN or internet services.

Static Devices participating in the algorithm are defined as short-range wireless equipped devices which are generally static while operating (e.g. PCs, most printers, multifunctional photocopiers etc.), and which do not have inherent position information, but which can estimate their position as a result of the Positioning by Diffusion Algorithm. Static Devices can in effect hold positioning information they have accumulated as a result of the algorithm, over a relatively long period of time.

Positioning information can be transferred in the following ways :

1. From Access Point to Mobile Device
2. From Access Point to Static Device
3. From Static Device to Static Device
4. From Static Device to Mobile Device
5. From Mobile Device to Static Device
6. From Mobile Device to Mobile Device

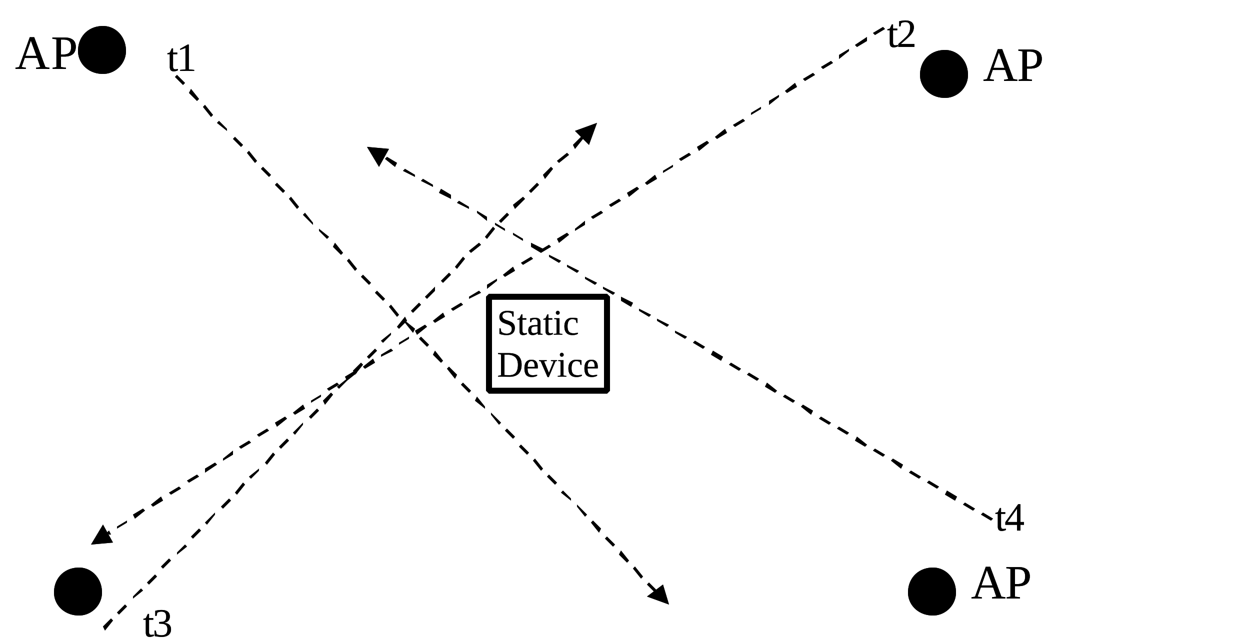

Figure 1. The dotted lines show the path of 4 Mobile Devices, which pass near to the respective Access Points at times t1, t2, t3 and t4, and progress towards to the Static Device, communicating position information with the latter. The times t1, t2, t3 and t4 could be separated from each other by hours or even days. Subsequently, this information can be communicated to another Passing Mobile.

.

A number of situations which can typically arise, are outlined as follows.

a) Initially Static Devices have no positioning information. A Mobile Device passes

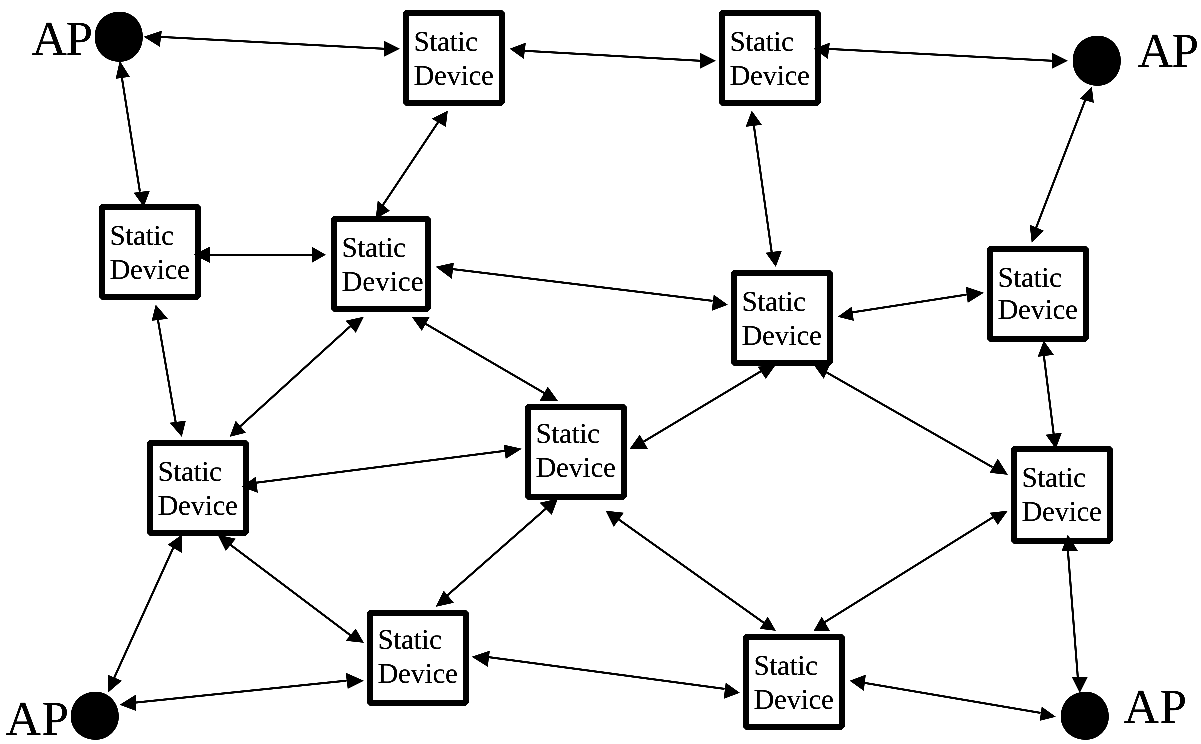

Figure 2. Static Devices communicate positioning information to neighbours within short-range wireless range. This allows position estimates to be made, even in the absence of Mobile Devices.

an Access Point and receives positioning information from the Access Point. A short time later it passes a Static Device, and so is able to pass on the positioning information, sourced at the Access Point, to the Static Device. Similarly, other Mobile Devices will first pass an Access Point, then the Static Device, again passing further positioning information to the Static Device. The Static Devices gain progressively more positioning information as time goes on. Some time later – possibly many hours or even days later – a further Mobile Device passes such a Static Device, and the positioning information is passed from the Static Device to the Mobile Device.

In general, the Static Devices can obtain positioning information from passing Mobile Devices, which can subsequently be used by another Mobile Device. Another way in which positioning information may be diffused in this situation is by a Mobile Device receiving positioning information from another Mobile Device.

b) It is possible that in some situations, there may be no Mobile Devices. Provided the Static Devices form a contiguous network (i.e. successive devices in range of each other), and there are a number of Access Points around the edge of the contiguous network, then positioning information can be sourced from the Access Points, and diffuse among the Static Devices. This allows the Static Devices to make a reasonable estimate of their position.

   It is of course also possible for a particular Static Device to make use of some position information passed from nearby peer Static Device(s), and also to make use of positioning information communicated from a passing Mobile Device as described in a).

c) It is possible that in some situations, there may be no Static Devices. A Mobile Device in range of one or more Access Points can estimate its position by positioning information passed from the Access Point(s) to the Mobile Device. The Mobile Device can also receive positioning information from other Mobile Devices. Even when Mobile Devices are some way from Access Points, a group of Mobile Devices, which have come from different directions (and have communicated with different Access Points) can diffuse their information to each other, and be in a position to perform reasonably good position estimates.

d) Another possible situation is where the Access Points are much more densely situated, and so Mobile and Static Devices may be able to communicate with multiple Access Points. Algorithms need to be able to operate in this situation, but, owing largely to the cost involved, Access Points are much more likely to be deployed more sparsely, and this paper focuses on this latter situation.

The operation with very small and very low power Mobile Devices, which will have to communicate less frequently, needs considering. In the sort of situation described in a), there are relatively few Access Points, but a number of Static Devices. The Static Devices can obtain positioning information only from passing low power Mobile Devices which happen to be actively communicating at that point.

So Static Devices still build up positioning information over time as outlined in a), albeit at a slower rate. Subsequent Mobile Devices may then obtain positioning information from a nearby Static Device, possibly as a result of a user or application actively requesting a position estimate.

3 A Look at the Geometry of the Algorithm

3.1 A Simple Example

In this section, some simple situations are outlined, in which knowledge accumulated by short-range wireless transmission can be used to determine the region in which a Mobile Device is likely to be located. In the simplest case, the region will be defined as the area within a circle, but in more complicated cases, a somewhat more complex region is derived.

Suppose a Mobile Device A can communicate with an Access Point AP1. In general, a maximum distance of transmission, *MD*, can be determined. This might be a maximum range of the particularly wireless technology. Alternatively there may be a maximum range for the particular circumstances. For example, the receiver may use the strength of received signals, possibly compared with the received signal strength of other recently received transmissions, to get a better estimate of maximum distance.

For simplicity of explanation, a general maximum distance for the particular wireless technology will be used in the description that follows.

a) The simplest case is when the Mobile Device A can communicate with Access Point AP1. It is clear that $Dist(A, AP1) \leq MD$, where *Dist* is the geometric distance between the 2 points within the brackets. Clearly, the equation means that A is located within a circle centered at AP1, of radius *MD*. To express this more generally, A's *Region of Possible Positions* is the circle centered at AP1, with radius *MD*.

b) If a second Mobile Device B is in range of A, while A is also within range of AP1, then $Dist(B, AP1) \leq 2MD$, in other words, B's Region of Possible Positions is within a circle of radius 2*MD* from AP1.

c) It is assumed that within a typical environment for which the algorithm is envisaged – e.g. corporate building, shopping mall, etc., there will be a practical maximum limit of pedestrian speed above which it is very unlikely mobile users, and hence Mobile Devices, will travel. This maximum speed is abbreviated to *MS*. Considering again Mobile Device A, suppose A loses touch with Access Point 1, and the time of its last communication with AP1 is t_1.
Then at a later time t, $Dist(A, AP1) \leq MD + (t - t_1)MS$
i.e. A's Region of Possible Positions is within a circle of $MD + (t - t1)MS$ of AP1.

   This circle is clearly expanding in size as time progresses, the radius increasing at a rate *MS* per unit time. The expanding circle implies that in the absence of further information, there is less certainty as time progresses of where exactly the Mobile Device is positioned.

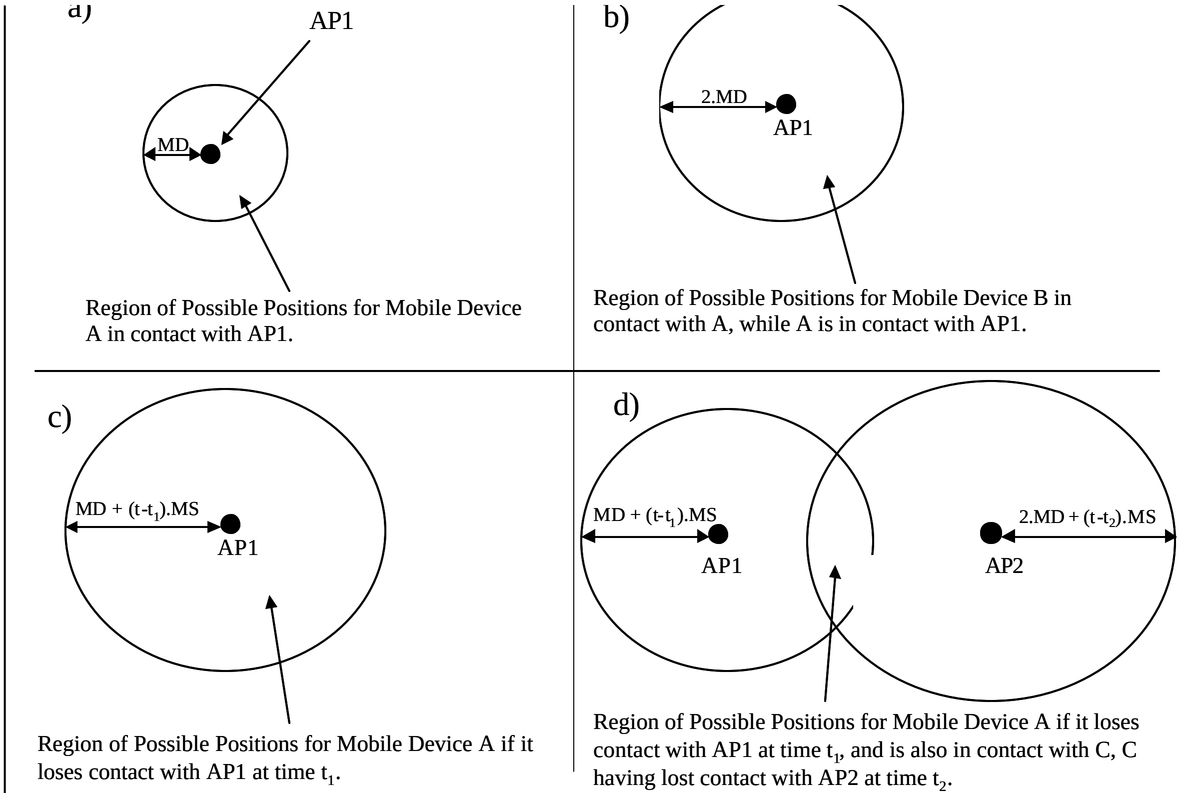

Figure 3. Some Examples of the Region of Possible Positions in Various Circumstances

d) Suppose further that Mobile Device C last communicates with Access Point AP2 at time t_2, and then communicates with Mobile Device A at time t. Clearly, $Dist(C, AP2) \leq MD + (t - t_2)MS$, and since C and A are now communicating, $Dist(A, AP2) \leq 2MD + (t - t_2)MS$ i.e. A is within a circle of $2MD + (t - t_2)MS$ of AP2.
However, it is also known from c) that $Dist(A, AP1) \leq MD + (t - t_1)MS$
Therefore, from the information available, it can be deduced that A's Region of Possible Positions is within the intersection of the two circles, one centered at AP1 and the other at AP2.

It is also clear that if Mobile Device A communicates with a further device D, which recently communicated with an Access Point AP3, Mobile Device A can also deduce it is within a third circle centered at AP3, in addition to the two circles mentioned. So A's Region of Possible Positions is within the intersection of three circles. Figure 4 in fact shows such a situation.

It is possible to generalize about the Region of Possible Positions, using a) to d) above. Point a) indicates that while a Mobile Device can communicate with only one Access Point, its Region of Possible Positions is a circle of radius *MD*, centered at the Access Point. As is shown in b) when the Region of Possible Positions is communicated from one device to another, its size expands by *MD* in all directions. From c) one sees that in the absence of further positioning information, the Region of Possible Positions for a Mobile Device expands at a rate *MS* per unit time. Finally, d) shows that if positioning information pertaining to more than one Access Point is obtained by a Mobile Device, the Region of Possible Positions will be the intersection of two or more circles.

3.2 The Protocol for Diffusing and Updating Position Information

The observations of the previous section are used when defining the protocol for diffusing the position information. The information is conveyed in *Position-Distance Pairs*. A Position-Distance Pair has two fields, position and distance, and is expressed as {Position,Distance}. The Position field will in general be a coordinate in 2 dimensions, so the Position-Distance Pair in its longer form could be written as {($Position_x$,$Position_y$), Distance}.
When two devices communicate, the Position-Distance Pair is in effect copied from one device to another, and many such operations cause it to be diffused among devices.
As the Position-Distance Pair is copied or diffused, it is manipulated according to the following rules. The rules are framed in a way which can be implemented on Mobile and Static Devices, and also take into account that inter-device transmissions will generally be in the form of a discrete event, when a packet or frame is sent between devices.

1. A Position-Distance Pair is generated by an Access Point, and initially, the Distance Value is set to 0. The Position field is set to be (the x- and y-coordinates of) the position of the Access Point.
2. The Position field does not alter as the Position-Distance Pair diffuses from device to device.
3. When a copy of the Position-Distance Pair is received by a Mobile Device or a Static Device, the Distance field is increased by *MD*. As discussed previously, *MD* may be a general maximum distance for a particular short-range wireless technology, or may be a maximum distance of transmission based on current and local factors, as determined by the receiving device.
4. Each device has a small piece of memory for storing Position-Distance Pair information. When a Position-Distance Pair is received by a device, subsequent to increasing the Distance field by *MD* as just described, the receiving device carries out the following calculation. If the receiving device does not already have in its memory a Position-Distance Pair with the same Position field value as the incoming Position-Distance Pair, then the incoming Position Distance Pair is stored in the device's memory. If the receiving device does already have in its memory a Position-Distance Pair with the same Position field value as the incoming Position-Distance Pair, then if the stored Distance value is higher than the incoming Distance value, the stored Distance value is overwritten by the incoming Distance value.
5. During the time a Position-Distance Pair is held in a Mobile Device's memory, the Distance field value is increased by *MS* per unit time. This could be done on a continuous basis, or a more efficient implementation could be to use an internal clock, and to do the appropriate increase (based on elapsed time) on the Distance

value(s) in memory, the next time it is utilized in a calculation or communicated to another device. This increase of *MS* per unit time is not carried out on Static Devices.

As can be seen, a Position-Distance Pair is sourced at an Access Point. While held within a Mobile Device, its Distance value increases at a constant rate *MS*, due to the fact that the Mobile Device may move, and so its Region of Possible Positions increases. Also, when inter-device transmission takes place, a copy of the Position-Distance Pair is transferred to the new device, with an increase in Distance value, owning to the fact that the Regions of Possible Positions has increased by *MD*.

The protocol for individual Position-Distance Pairs has been described above. In practice, a group of more than one Position-Distance Pairs will be transferred from one device to another, with each individual Position-Distance Pair being processed according to the above rules. The transfer of a group of more than one Position-Distance Pairs being transferred from one device to another implies that a Region of Possible Positions more complex than just a circle is being transferred. In general, a 2-way exchange of Position-Distance Pairs between communicating devices is most effective, but there may be some dependencies on the underlying short-range wireless technology.

Within Mobile Devices, Position-Distance Pairs do not exist indefinitely. They may be superseded by Position-Distance Pairs pertaining to the same position, but with a lower Distance Value (see 4.). There is also a mechanism described later which will cause less useful Position Distance Pairs to be eliminated.

From an intuitive point of view, as the Distance value of a Position-Distance Pair increases over time, the Position-Distance Pair becomes less useful, since the Region of Possible Positions defined becomes larger i.e. the knowledge of the position of the device which can be derived becomes more vague. The Distance value can be seen to be inversely related to the confidence level associated with the Position-Distance Pair. Intuitively it can be seen as well that faster moving Mobile Devices will tend to convey more useful information to other devices. By the time a slower moving device reaches a peer device and transfers positioning information, the Distance value(s) of the Position-Distance Pair(s) transferred will have increased, thus rendering the positioning information less useful, and more likely to be eliminated in the mechanism described in point 4., described above.

A point related to the last paragraph is that the size of the Region of Possible Positions is an indication of the confidence level of a position estimate. A large Region of Possible Positions indicates less certainty or confidence of the exact position of a Mobile or Static Device. A confidence measure can therefore be calculated, based on the area or dimensions of the Region of Possible Positions, and this can be output to an application in parallel with any actual estimate of position.

Within Static Devices, the Distance Value of the Position-Distance Pair is not increased in proportion to *MS*, as is done in Mobile Devices, and in this sense, the confidence level associated with the Position-Distance Pair does not diminish over time. Position-Distance Pairs are stored over a much longer timescale than within Mobile Devices. In some senses, the Static Device acts as a 'memory' for Mobile Devices. A Mobile Device conveys positioning information to the Static Device which at some time later is used by another Mobile Device. Since the Position-Distance Pair information stored by a Static Device may be used by multiple Mobile Devices, there is a case for a refinement within the Static Device, in which rather than the lowest Position-Distance Pair being used for position estimates and passed on to Mobile Devices, a 5% percentile value over the distribution of Distance values received, is instead used. This, or similar refinements being analyzed, would have the effect of eliminating rogue Position-Distance Pairs, such as might occur very occasionally, if for example a carrier of a Mobile Device runs rather than walks, and exceeds the assumed maximum speed of *MS*. In any event, it is pragmatic to age out positioning data in Static Devices after a certain period.

4 The Calculation of the Best Estimate

As a result of the protocol described in the last section, a device will have a number of Position-Distance Pairs in its memory, from which it can carry out an estimate of its position. Several algorithms are possible, varying from the simple and less accurate to the somewhat more complex and more accurate, and three such algorithms are outlined here.

4.1 Estimation by Geometric Calculation

The most intuitive position estimation algorithm, consisting of a few simple steps, is now described. The simplest case is when a device has just one Position-Distance Pair in its memory. In this case, the device only has information pertaining to one Access Point, and so the position of that Access Point is the device's best estimate for position.

If there are more than one Position-Distance Pairs in memory, then in general, the Region of Possible Positions is defined as the intersection of the areas of several circles. A special case is where one circle pertaining to a Position-Distance Pair is entirely within all of the other circles pertaining to the other Position-Distance Pairs. A simple arithmetic check will determine this. In this case, the situation is similar to that above, and the best estimate is simply the position of the Access Point at the center of the circle which is inside all of the others.

In general, the estimation algorithm will have a number of Position Distance Pairs which correspond to a set of circles, each centered at an Access Point. The Position and the Distance fields of the Position Distance Pair define the circle center and circle radius respectively. The Region of Possible Positions is the area formed by intersection of the circles. In the case of 3 intersecting circles, the Region of Possible Positions could be loosely described as a triangle with curved edges. For more than 3 intersecting circles, the Region of Possible Positions could be described as a polygon with curved edges.

The best calculation of position would be to evaluate the center of gravity of the Region of Possible Positions. A pragmatic and fairly straightforward estimate of this is simply to determine the mean position of the points of corners of the (curved) triangle or polygon forming the Region of Possible Positions.

This is done in a simple two-stage process. Firstly the set of all of the intersections of pairs of circles is calculated. The coordinates of the intersects can be calculated using a simple quadratic formula. Secondly a subset of *interior intersections* is created in the following way. If an intersection point of a pair of circles is inside all of the other circles (i.e. all of the circles not involved in the intersection), then it is defined as being in the subset of interior intersections. Otherwise, it is not included in the subset of interior intersections. This evaluation involves a simple calculation of the distance from the point of intersect to the center of a circle (the Position field of the Position Distance Pair), and comparing it to the circle radius (the Distance field of the Position Distance Pair).

To take a simple example of the two stage process involved, assume a device has 3 Position-Distance Pairs in its memory, as shown in Figure 4. The intersection of the 3 circles (which defines the Region of Possible Positions) creates a roughly triangular shape (with curved rather than straight edges), bounded by intersection points a, b and c.

In the first stage of the process, each pair intersection point of the (respective pairs of) circles is calculated, so in this case the coordinates of the 6 points a, b, c, d, e and f are calculated . In the second stage, each of these 6 points is evaluated to see if it is inside the remaining circle not involved in the intersection. For example, point a is the intersection of the two circles centered at AP1 and AP3, and it is inside the remaining circle centered at AP2. Point a is therefore an interior intersection. If one looks at point d, formed by the intersection of the circles centered at AP1 and AP2, it is outside the remaining circle, the circle centered at AP3. Point d is therefore not an interior intersection.

Once this simple calculation is made for each of the 6 intersect points, 3 intersections - a, b and c – are found to be interior intersections. The mean of these 3 points of interior intersection is defined as the estimate of the position of the device. A slightly enhanced version of this calculation is presented in the next section.

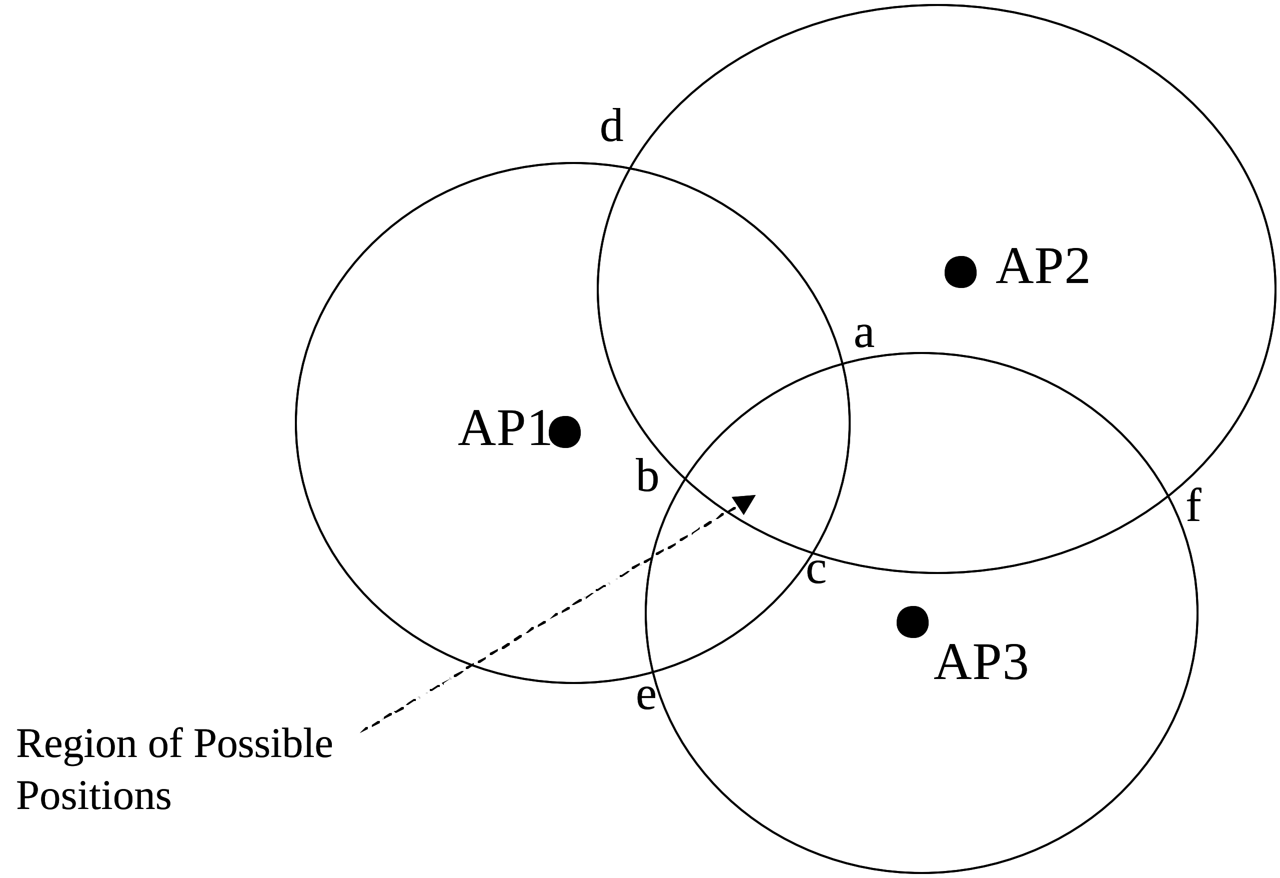

Figure 4. The Region of Possible Positions formed by the Intersection of Circles

A simple way of minimizing the number of Position-Distance Pairs transmitted across the short-range wireless links is inherent in this method. In the previous paragraphs, the stage of creating the subset of interior intersections has been described. A simple way of keeping the most useful Position-Distance Pairs and discarding the less useful ones, is simply to keep only those Position-Distance Pairs which participated in the creation of interior intersections. Any Position-Distance Pair, which has not participated in the creation of the subset of interior intersections, is marked as not to be sent to peer devices across the short-range wireless link in subsequent transmission.

4.2 Estimation by Geometric Calculation – Enhanced Version

The accuracy of position estimation of the Estimation by Geometric Calculation method outlined in 4.1 can be improved with the following enhancement. The basic method is very similar to that of 4.1, but a post-processing enhancement to the basic algorithm can improve the results.

Some motivation for the enhancement is given first of all. One approximation used in the Estimation by Geometric Calculation is calculation of the center of gravity of the polygon, by taking the mean of the interior intersection points at its corners. It may be that the corners are grouped very irregularly. For example, if several successive sides of the polygon are very short, several corners will all be located close together, and this may cause the estimation by taking the mean of the position of the corners to lead to a somewhat biased result.

It is to be noted that the effect of any such bias described would be minimal if the size of the polygon (or of the Region of Possible Positions) were small. Also, other approximation effects will tend to diminish, with a smaller size of Region of Possible Positions. A useful strategy is therefore to reduce the size of the Region of Possible Positions before carrying out the averaging calculation.

This is carried out as follows. Firstly a copy of the original Distance Values of the Position Distance Pairs is stored in memory. (This has to be done, because when Position Distance Pairs are subsequently sent to other devices, they have to be as originally defined, not in their altered state after the following operation.) Then, at each iteration, if the size of the Region of Possible Positions is above a target size, each of the Distance values is multiplied by a factor k (of less than one), so that all of the Distance values are reduced proportionally. From a geometric point of view, the circles centered on the respective Access Points all are reduced by the same ratio. This process is carried on until a target size is reached. Then the same calculations as in the Estimation by Geometric Calculation is made, but now using a potentially much reduced Region of Possible Positions.

The target size is best framed in terms of the size of the area bounded by the Access Points the device is aware of through the Position fields of the Position-Distance Pairs it holds in memory. Since the target size calculation is simply used to determine how many iterations of the multiplication by k are carried out, great accuracy in determining the target size is not necessary. Therefore, the size of the area bounded by the Access Points could be estimated by taking the smallest rectangle which would contain all of the Access Points, which can be calculated in a simple way. Similarly the size of the Region of Possible Positions can be estimated by framing it with the smallest possible rectangle. The target size could therefore be some fraction of the size of the smallest rectangle enclosing the Access Points, and the iterations reducing the size of the Region of Possible Positions can continue until the smallest rectangle enclosing it is less than the target.

During the course of reducing the size of the Region of Possible Positions, the number of points in the set of interior intersections may reduce i.e. some circles may cease to be involved in its definition. The efficient implementation of the iteration process involves some finesse, as it encompasses the geometry involved as well as iterative convergence processes.

4.3 Estimation by Numerical Calculation

Another possible method of performing a calculation of a best estimate of position is through a simple direct numerical calculation, rather than performing a computational algorithm. The input of the calculation is the Position-Distance Pairs the device has in its memory.

An obvious candidate is a weighted average estimate of the form

$$x\_est = \frac{c_1 \times AP1_x + c_2 \times AP2_x + c_3 \times AP3_x + ... + c_N \times APN_x}{c_1 + c_2 + c_3 + ... + c_N}$$

$$y\_est = \frac{c_1 \times AP1_y + c_2 \times AP2_y + c_3 \times AP3_y + ... + c_N \times APN_y}{c_1 + c_2 + c_3 + ... + c_N}$$

where APi_x and APi_y are the x-coordinate and the y-coordinate of Access Point i. The weighting factors are determined in terms of the Distance values of the respective Access Points' Position-Distance Pairs, and in the environments tested, the most promising simple definition appeared be $c_i = \frac{1}{d_i^{\,2}}$

The chief motivation for this approach is the ease of implementation.

5 A Discussion on where Access Points should be situated.

It is now relevant to have a brief general discussion about the best placing of Access Points, the pre-programmed source of position information. Often, the Access Point will be incorporated in a general wireless to wired network bridging device, allowing for example mobile users to access the internet or corporate LAN across the short-range wireless link. If this is the case, there may be less choice as to the placing of the device. However, where there is some element of choice as to where Access Points should be situated, it is useful to have some guidelines.

The description of diffusion of position information in a) of Section 2 discussed how the initial stage for position information to be diffused among Static and Mobile devices is often for Mobile Devices to pass an Access Point, and a short time later transfer position information to a Static Device, or possibly a Mobile Device. To maximize the diffusion of position information from the Access Points, it is often logical for Access Points (the source of the position information) to be placed at entrances to an area where positioning is to take place. For example, in corporate buildings, Access Points should be placed where people first enter a particular floor i.e. at stairs, elevator exits etc. As people arrive for work for example, and leave the stairs or lift exit, they will then tend to diffuse the position information while walking to their working location. Similarly if positioning is to be implemented in a public building such as a shopping mall, it is useful to position Access Points at the entrances.

It is also appropriate to position Access Points in situations where there is a heavy flow of Mobile Devices passing. In the example of pedestrian streets within a shopping mall, Access Points could be positioned at street intersections, which also has the benefit of partitioning the area into straight sections of pedestrian streets. However, it is clear that in such situations Static Devices (which learn position information rather than having to be programmed) could also be placed there in such a situation. A Static Device which lies somewhere between Access Points can be expected to accumulate enough position information to be able to estimate its position, and to enable passing Mobile Devices to do likewise.

In general, a useful strategy is to place the Access Points at the extremes of the space in which the positioning is implemented, and to allow Static Devices situated in the interior of the space, to participate in the way described. Focusing for example on a floor of a corporate building, the Access Points could in particular be situated particularly in the corners, and on the edge walls, and as well at the stairs and lift exits, as discussed.

6 A Comparison between the Algorithms

6.1 The Simulation Environment

In order to be able to examine the characteristics of the various algorithms with respect to realistic numbers of Mobile Devices, a simulation environment was developed. The code was written in Java, and a display of the movement of Mobile Devices, together with the configuration of Access Points and Static Devices could be observed in real time on the monitor. Also, long simulation runs could be carried out, in order to examine the characteristics of the algorithms.

A good degree of flexibility was obtained by allowing the definition of the movement patterns of Mobile Devices to be defined completely independently of the location of Static Devices, both of these being completely independent of the definition of the location of the Access Points.

6.2 The Simulation Experiments and Results

A number of experiments are reported, which give an insight into the various algorithms. A rectangular space was used, which could represent a storey of a large corporate building, or a large space in a commercial public building. The dimensions of the rectangular space was 400 pixels in the east-west direction, and 800 in the north-south dimension. The maximum distance of transmission of the short-range wireless transmission was initially assumed to be 60 pixels. (To convert this to a real-life situation with meters, one might divide these numbers by 4, the simulation would in effect be of a 200 meters by 100 meters space, and a wireless technology with a maximum range of 15 meters. However, the results are described in terms of pixels, not in terms of meters.)

Mobile and Static Devices attempt to communicate with another device within range every 5 seconds, and if successful, they can exchange Position-Distance Pair information.

6.2.1 Experiment 1 : 72 Static Devices and 8 Mobile Devices

The intention of this experiment was to simulate a corporate environment, in which there are relatively few Access Points, but many Static Devices, such as PCs and printers. There were 6 Access Points, one in each of the four corners the other two were placed halfway along the two longer sides.

The 72 Static Devices were evenly distributed throughout the rectangular space, in a 6 x 12 grid, but were just far enough apart that they could not communicate with each other. They therefore relied on the movement of Mobile Devices to obtain position information. The 8 Mobile Devices represent a small number of personnel who at any one time might be moving around the floor of the corporate building. The movement was semi-random, with the person coming into the space through an entrance (representing a lift exit for example).

The mean error in calculating the position of Static Devices, and of Mobile Devices, for the three algorithms discussed, is shown in the table below. The table shows the mean errors after 25,000 cycles. One cycle in the simulation corresponds to the one second in

	Static Devices	*Mobile Devices*
Estimation by Geometric Calculation	18.36 (4.6 meters)	33.09 (8.3 meters)
Estimation by Geometric Calculation – Enhanced Version	13.87 (3.5 meters)	27.89 (7.0 meters)
Estimation by Numerical Calculation	45.84 (11.5 meters)	51.05 (12.8 meters)

Table 1 The results of Experiment 1, with 72 evenly distributed Static Devices and 8 Mobile Devices

the situation being simulated, so 25,000 cycles approximately corresponds to around 7 hours, i.e. a little less than a working day. As discussed before, if the figures below reported in terms of pixels, are divided by 4, this will approximately correspond to meters in the real situation, and these figures are given in brackets in the tables.
One can see that the Enhanced Version of the Estimation by Geometric Calculation gives a significant gain in accuracy as compared to the original version. Also, the very simple Estimation by Numerical Calculation performs significantly worse, as one might expect from such a simple algorithm. Interestingly, the figures for Estimation by Numerical Calculation hardly improve at all with a longer simulation time, while the results for the

first two algorithms do show a gradual improvement as the simulation progresses. For example, after a million cycles, for the Enhanced Version of the Estimation by Geometric Calculation, the Static Devices have a mean positioning error of 7.6 pixels, and the Mobile Devices have a mean error of 23.5 pixels.

6.2.2 Experiment 2 : 4 Static Devices and 80 Mobile Devices

In this experiment, the same overall rectangular space was used, and the same positioning of the Access Points. However, rather than a corporate environment with large numbers of Static Devices, a public building or tourist attraction was assumed, with a large number of Mobile Devices moving around, but from time to time semi-randomly changing direction. The experiment had 80 Mobile Devices moving within the rectangular space. It was also assumed there could well be a small amount of infrastructure, and so 4 Static Devices were also placed within the space, well separated from each other and from the Access Points.

The simulation results converge to remain almost stationary after a certain point, and the mean error in positioning estimate after 400,000 cycles are given.

	Static Devices	*Mobile Devices*
Estimation by Geometric Calculation	17.75 (4.4 meters)	78.38 (19.6 meters)
Estimation by Geometric Calculation – Enhanced Version	17.75 (4.4 meters)	65.27 (16.3 meters)
Estimation by Numerical Calculation	48.50 (12.1 meters)	70.29 (17.6 meters)

Table 2 The results for Experiment 2, with 4 Static Devices and 80 Mobile Devices

This is rather a different scenario from the previous experiment, in that the 6 Access Points and 4 Static Devices provide a much sparser coverage for Mobile Devices. Consequently, the mean error for the Mobile Devices is higher. The relatively scarce availability of positioning information means that the disadvantages of the Estimation by Numerical Calculation are less pronounced. The Estimation by Geometric Calculation in the absence of the enhancement will in some situations be performing calculations with relatively large Regions of Possible Position, and some of the issues outlined in Section 4.2 arise. The benefits of using the Enhanced Version of the Estimation by Geometric Calculation is seen in the results.

6.2.3 Experiment 3 : 72 Static Devices

Although a positioning system operating in the absence of Mobile Devices would be an unusual situation, this experiment was carried out to illustrate that the algorithms can operate in this environment.

The same rectangular space and the same placing of Access Points as in Experiment 1 were defined. Also, the 72 Static Devices were distributed in the same way as Experiment 1. There were no Mobile Devices in this experiment. A slight increase in the wireless transmission distance allowed the Static Devices to communicate with their nearest neigbour, and also for the Access Points communicate with the nearest Static Devices. (In Experiments 1 and 2, no such communication could take place, and so the Static Devices relied entirely on Mobile Devices for their position information.

	Static Devices
Estimation by Geometric Calculation	50.97 (12.7 meters)
Estimation by Geometric Calculation – Enhanced Version	24.78 (6.2 meters)
Estimation by Numerical Calculation	47.33 (11.8 meters)

Table 3 The results for Experiment with 72 Static Devices in the Absence of Mobile Devices

The results show that the algorithms still operate reasonably well in this environment, with the Enhanced Version of the Estimation by Geometric Calculation able to perform best.

7 Conclusions

This paper has provided an overview of the basic method of Positioning by Diffusion, a method of performing positioning based on information transferred across short-range wireless links. It utilizes the motion of mobile users who carry Mobile Devices, to diffuse, or spread positioning information through the network. The paper reports on some simulations which have been performed to highlight pertinent features of the performance of 3 methods of performing the final positioning calculation.

An aim of Positioning by Diffusion is to have a positioning system which minimizes the requirement for Access Points, which are devices which have to be programmed with accurate position information from which other devices – Static Devices and Mobile Devices – derive their own position estimate. Clearly, minimizing the required number of Access Points minimizes system administration and installation requirements.

Position by Diffusion, which is implemented in devices as software, and does not require any extra hardware, operates in a seamless way. Assuming there are a minimal number of Access Points, it can operate with a network of just Static Devices, just Mobile Devices or a mixture of both.

The initial step, before the position calculation is made, is to propagate or diffuse Position-Distance Pairs among Mobile and Static Devices. This approach avoids the feeding of the output of one estimate calculation, as a parameter into the input of another estimate calculation, the latter approach tending to make any approximation effect in the calculations cumulatively build up.

The results of 3 different methods of performing the final position calculation have been simulated in this paper :

1. Estimation by Geometric Calculation
2. Estimation by Geometric Calculation – Enhanced Version
3. Estimation by Numerical Calculation

A summary of observations from the simulations is that 3. can be described as a fairly crude method and very simple to implement, but is unable to make really accurate estimates, even when the quality of the information available is good. Method 2. of the list above is an enhancement of 1., and 2. performs sufficiently better than 1. to justify the enhancement. Further refinements to the basic algorithms presented here are being investigated, to further enhance the algorithm performance.

In general the results indicate that in the scenarios simulated Positioning by Diffusion is accurate enough for the majority of in-building positioning applications, but there will be some applications where very high accuracy is necessary, and positioning based on special hardware additions to the wireless communications systems would then be required.

Position by Diffusion has been submitted for consideration to the Bluetooth Local Positioning Working Group, which has been charged with standardizing a local positioning technology for Bluetooth. However, Position by Diffusion is not actually specific, and can potentially operate in conjunction with other short-range wireless technologies.

References

[1] http://www.navcen.uscg.gov/gps/default.htm US Coast Guard Navigation Center GPS Web pages.

[2] J Li , J Jammotti, D S J De Couto, D R Karger, and R Morris. A Scalable Location Service for Geographic Ad Hoc Routing. ACM Mobicom 2000, Boston, MA.

[3] C Savarese, J Rabaey, and J Beutel. Locationing in Distributed Ad-Hoc Wireless Sensor Networks. Proc ICASSP2001, Salt Lake City, Utah.

[4] S Capkun, M Hamdi and J-P Hubaux. GPS-free Positioning in Mobile Ad-Hoc Networks. Proc. of 34th Hawaii International Conference on System Science, 2001.

[5] B A Miller and C Bisdikian. Bluetooth Revealed. Prentice-Hall, 2001.

[6] J Bray and C F Sturman. Bluetooth, Connect Without Cables. Prentice Hall PTR, 2001.