

A DISCOVERY METHOD FOR A POWER LINE COMMUNICATION

NETWORK

TECHNICAL FIELD

5 This invention broadly relates to discovery and maintenance protocols for power-line communications (PLC) networks and, more specifically, to discovery methods and apparatus for power line communication networks.

10 BACKGROUND ART

 The following discussion of the background art is intended to facilitate an understanding of the present invention only. The discussion is not an acknowledgement or admission that any of the material referred to is or was part
15 of the common general knowledge as at the priority date of the application.

 Power line communication, or PLC, refers to the practice of using electrical distribution networks for transmitting
20 data between nodes. PLC is typically used as a communication technology to enable sending and receiving of data over existing power cables or lines.

 PLC generally requires a sender node to modulate data to
25 be sent, the sender node then injects the modulated data onto a transmission medium, and a receiver node de-modulates the data to read it. In PLC networks, such a transmission medium is an existing power line network connecting the sender and receiver nodes.

30

SUMMARY OF THE INVENTION

According to a first aspect of the invention there is provided a discovery method for a power line communication network, comprising:

- 5 broadcasting, by a first node, a discovery request to at least one receiver, wherein the discovery request comprises the following information:
- a node address of the first node;
 - a total downward path cost through the first node;
 - 10 a downward hop address associated with the total downward path cost;
 - a total upward path cost through the first node; and
 - an upward hop address associated with the total upward path cost;
 - 15 wherein if the first node is a data concentrator (DC), the total downward path cost and total upward path cost are the values of 0, and wherein the downward hop address associated with the total downward path cost and the upward hop address associated with the total upward path cost are
 - 20 excluded from the discovery request.

According to a second broad aspect of the invention, there is provided a discovery method for a power line communication network, comprising:

- 25 receiving, by a node, a discovery request which comprises a total downward path cost, and a total upward path cost from a transmitter;
- determining, by the node, a total downward path cost from a data concentrator (DC) to the node by using a total downward
 - 30 path cost in the discovery request;
 - storing a node address included in the discovery request and a downward hop cost from the transmitter to the node, the received total downward path cost and the received

total upward path cost from the transmitter into a discovery table of the node.

5 The method may further comprise rejecting, by the node, paths which include the node itself as an upward or downward node in the path.

10 According to a third broad aspect of the invention, there is provided a discovery method for a power line communication network, comprising:

 receiving, by a registering node, a discovery request from a transmitter;

15 choosing the best downward hop address from a stored discovery table, using a total downward path costs from a data concentrator(DC) to the registering node; and

 broadcasting, by the registering node, a registration request;

wherein the registration request comprises the following information:

20 a registering node address;
 a downward hop address;
 a downward hop cost to the registering node; and
 a downward hop's total upward path cost.

25 The registration request may further comprise an upward hop address, an upward hop cost from the registering node to the upward hop address and an upward hop's total upward path cost which are filled by relaying or re-broadcasting nodes.

30 According to a further broad aspect of the invention, there is provided a discovery method for a power line communication network, comprising:

receiving, by a neighbour node, a registration request; and

re-broadcasting the registration request if the neighbour node is in unregistered state or relaying the registration request if the neighbour node is in registered state.

The method may further comprise determining to relay based on a total upward cost of the neighbour node and a downward hop's total upward path cost.

According to a yet further broad aspect of the invention, there is provided a discovery method for a power line communication network, comprising:

transmitting, by a data concentrator (DC), a registration confirm message to at least one registering nodes based on a total upward path costs in registration requests from the at least one registering nodes; and

updating a routing table, by the DC, based on the registration requests from the at least one registering nodes; wherein the registration confirm message comprises the following information:

a registering node address, a downward hop address, a total downward path cost to the registering node, a upward hop address, a total upward path cost from the registering node.

The registering node address may be assigned a shorter address instead of a longer address pre-configured into the registering node.

The method may further comprise receiving, by the at least one registering nodes, the registration confirm message from the DC; and transitioning, by the at least one registering nodes, to registered state.

According to another broad aspect of the invention, there is provided a discovery method for a power line communication network, comprising:

- 5 comparing, by a registered node, a total downward costs through a different downward hop addresses in a discovery table with a current total downward path cost; and
determining, by a registered node, to send a registration request with a newly selected downward hop.

10

According to a further broad aspect of the invention, there is provided a discovery method for a power line communication network, comprising:

- receiving, by a registered node, a polling from a data
15 concentrator(DC); and
transmitting, by the registered node, a response which informs the DC of a newly selected downward hop address.

According to an additional broad aspect of the
20 invention, there is provided a discovery method for a power line communication network, comprising:

- receiving, by a registered node, a discovery request;
determining, by the registered node, a total upward path cost from a node which transmits the discovery request to
25 a data concentrator (DC) if the registered node is replaced as a upward hop for the node; and
deciding to send a registration request to the DC on behalf of the node which transmits the discovery request, if the determined total upward path cost is less than a total
30 upward path cost specified in the discovery request by a certain margin.

The method may further comprise checking, by the registered node, if the upward hop address involves itself and

the upward hop cost has differed above a certain threshold; and transmitting, by the registered node, a registration request to the DC or waiting for a poll, to inform of the changes in upward hop cost.

5

The discovery request may further comprise information associated with multiple downward and upward routes for the first node.

10

The discovery table of the node may further comprise path number information.

The registration request may further comprise path number information.

15

According to another broad aspect of the invention, there is provided apparatus for a power line communication network, comprising:

a power line transceiver;

20

a processor coupled to the power line transceiver; and memory accessible to the processor;

wherein the processor is configured to store a discovery table which comprises a node address and a downward hop cost, the received total downward path cost and the received total upward path cost in the memory for routing packets.

25

The discovery table may further comprise path number information.

30

According to a further aspect of the invention, there is provided apparatus for a power line communication network, comprising:

a power line transceiver;

a processor coupled to the power line transceiver; and

memory accessible to the processor;
wherein the processor is configured to store a routing table
which comprises a downward hop address, a downward hop cost, a
upward hop addresses and upward hop cost for each node.

5

The routing table may further comprise path number
information.

BRIEF DESCRIPTION OF THE DRAWINGS

10

Further features of the present invention are more
fully described in the following description of several non-
limiting embodiments thereof. This description is included
solely for the purposes of exemplifying the present
invention. It should not be understood as a restriction on
the broad summary, disclosure or description of the invention
as set out above.

15

The description will be made with reference to the
accompanying drawings in which:

20

Figure 1 is a diagrammatic representation of a PLC
network with nodes arranged in communication, as indicated by
the arrows;

25

Figure 2 is a diagrammatic representation of the PLC
network of Figure 1, showing the downward and upward links
costs between the respective nodes, along with optimal paths
between nodes;

30

Figure 3 is a diagrammatic representation of the PLC
network of Figures 1 and 2, showing a changed link cost
between node DC and node D, in accordance with a broad aspect
of the invention; and

Figure 4 shows a diagrammatic representation of one example of apparatus for a power line communication network.

5 DESCRIPTION OF EMBODIMENTS

The following description provides examples of different embodiments of methods of discovering and establishing optimal data transmission paths between nodes in power-line communication (PLC) networks. Such PLC networks typically
10 comprise multiple-routing tree networks having asymmetric links between nodes.

Referring now to Figure 1, there is shown an example of a PLC network 10. The network 10 comprises a data concentrator
15 DC (hereafter referenced as 'DC'), which forms the root of the network 10. Further included are nodes A, B, C, D, E and F arranged into a tree network topology via connection paths or links 12, as shown. These paths 12 are typically existing power lines or links of an existing electrical distribution or
20 reticulation network, as is well understood in the art.

In order to form or establish such a single-routing tree network 10, every node in the network 10, including the DC, is pre-assigned a unique address in the network. A node may be
25 assigned a shorter address during a discovery process. However, at network start-up, all the nodes A to F will be in an unregistered state. The DC, being configured as the root of the network 10, will inherently be in a registered state. The manner of forming the network 10 via registration of each node
30 thereon is described below.

As such, in this example, the downward path from the DC to any given node in the network 10 is fully specified by a

network address of the node directly connected to such given node, i.e. the last so-called 'hop' to the given node. In this example, this downward path to the given node is referred to as the 'downward hop address' of the given node.

5

Similarly, an upward path from a given node in the network 10 to the DC is fully specified by the network address of the node directly connected to the given node in the first 'hop' of the given upward path to the DC. In this example, this upward path is referred to as the 'upward hop address' of the given node.

For example, in network 10, the downward path from DC to node F would be from DC to node B to node E to node F, or DC - B - E - F. As a result, the downward hop address of node F would be node E, the last downward 'hop'. This is as a result of the paths to node B and to node E being specified already in the respective network addresses of these nodes. Similarly, the upward hop address of node F would be the address of node D.

It is to be appreciated that the direction of data communication indicated by the network links 12 is determined by the network topology, i.e. the connection characteristics between the nodes. For example, the link 12 between nodes E and F enable unidirectional data transfer only from node E to node F. Alternatively, the network link 12 between nodes B and E allows bi-directional data transfer, as indicated.

In use, the DC periodically broadcasts discovery requests containing a total downward path cost and a total upward path cost of zero value for both to the nodes connected thereto, i.e. nodes A, B and D. Generally, after a node on the network is in a registered state, it also periodically broadcasts

discovery requests which include information of the total downward path cost and total upward path cost to that particular node (described below). Each node is typically configured to broadcast such discovery requests within a specific time slot chosen not to clash with similar broadcast discovery requests from other nodes on the network 10.

In general, a discovery request broadcast by any node on the network 10 includes the following information:

- a) The node address that is transmitting the discovery request;
- b) A total downward path cost through this node (0 if the DC is transmitting);
- c) The downward hop address associated with the total downward path cost (None if the DC is transmitting);
- d) A total upward path cost through this node (0 if the DC is transmitting); and
- e) The upward hop address associated with the total upward path cost (None if the DC is transmitting).

Typically, given the requirements of establishing such a network topology, a node that is unregistered on the network, after receiving a first discovery request, continues to receive discovery requests for a given time.

Each node in the network 10 generally includes a discovery table for storing information received from other nodes in the network. For example, a node (whether in a registered state or not) determines the total downward path cost from the DC to itself using the total downward path cost in the received discovery requests, and stores the node address that transmitted the discovery request together with the downward hop cost from such a transmitter node to itself, the received total downward path cost and the received total

upward path cost from the transmitter node into a discovery table. To avoid any looping paths in the network, a node would reject paths that include itself as an upward or downward node in the path.

5

With reference now to Figure 2, network 10 is shown with the downward and upward links costs between the respective nodes indicated on the links 12. As shown, in this example, at node B, the total downward cost is 3 and the total upward cost is 5.

10

Accordingly, node B's discovery request which it broadcasts over the network 10 contains these values. When node E, for example, receives node B's discovery request, it is able to determine that the downward hop cost from node B to itself is 1 (it does not yet know the upward hop cost from itself to node B). Node E then stores these values into its discovery table.

15

20

By way of illustration, a discovery table for E would be, for nodes B and D:

Downward hop address	Downward hop cost	Received total downward path cost	Received total upward path cost
B	1	3	5
D	3	4	6

After a predetermined period has passed to allow nodes to broadcast their respective discovery requests, an unregistered node can select or choose a best downward hop address from its populated discovery table, using the total downward path costs from the DC to that particular node and subsequently broadcast

25

a registration request via which it registers on the network 10.

5 Accordingly, such a registration request contains the following information:

- a) The registering node address;
- b) The downward hop address;
- c) The downward hop cost to the registering node (this information is typically used for completeness of status information or debugging at the DC, but is not used by the DC, as all downward route decisions are made at the nodes);
- d) The downward hop's total upward path cost;
- e) The upward hop address this request represents;
- 15 f) The upward hop cost from the registering node to the upward hop address; and
- g) The upward hop's total upward path cost;

20 It is to be appreciated that elements e) to g) above are subsequently completed by relaying or re-broadcasting nodes in the network 10.

25 For example, in the case of node E, it chooses node B as the best downward hop. node E then broadcasts a registration request with its own address, node B's address as the downward hop address, 1 as the downward hop cost, and 5 as the downward hop's upward path cost. It does not fill the upward hop details.

30 Neighbouring nodes that receive this registration request may then (or only after consolidating multiple registration requests) relay this information to the DC, or re-broadcast it over the network 10. For example, a registered node (having a known upward path) might relay the information, while an

unregistered node (having no known upward path) might simply re-broadcast the information.

5 Accordingly, a neighbour node may use the downward hop's total upward path cost in the registration request as a guide to determine if it should relay the registration request to the DC. This is generally done to avoid relays through sub-optimal paths.

10 An unregistered neighbour node, when deciding to re-broadcast the registration request will typically:

a) If it is a broadcast registration request with no specified upward hop address, place its own address as upward hop address, put the hop cost from the transmitter node to itself as the upward hop cost in the request,
15 leave the total upward path cost as 0, and re-broadcast the registration request.

b) Alternatively, if it is a broadcast registration request with a specified upward hop address, the neighbour node
20 will add the hop cost from the transmitter node to itself to the upward hop's total upward path cost, and re-broadcast the registration request.

25 As an unregistered node is likely to be downstream, the probability that an unregistered node responds should be minimised. There are, however, rare cases when this may be required for the network to form.

30 In such a case, a registered neighbour node, when deciding to relay the received registration request from the transmitter node, will:

a) If it is a broadcast registration request with no specified upward hop address, place its own address as upward hop address, put upward hop cost from the

transmitter node to itself as the upward hop cost in the request, place its own total upward path cost in the request, and relay the registration request to the DC.

5 b) Alternatively, if it is a broadcast registration request with a specified upward hop address, the neighbour node will add the hop cost from the transmitter node to itself plus its own total upward path cost to the upward hop's total upward path cost in the request, and relay the registration request to the DC.

10

For example, when node E broadcasts its registration request, nodes B and D, as neighbour nodes, receives the request. Based on the downward hop's upward path cost value of 5 in the registration request, node B decides to relay this request to the DC (node B having a total upward cost of 5), placing itself as the desired upward path for node E. Node B also determines that the upward hop cost from node E to itself is 2. So the value 2 is placed in the upward hop cost, and a value of 5 in the upward hop's total upward path cost before relaying it to the DC.

15

20

Similarly, node D decides not to relay this request (node D having an upward path cost of 6). The decision may be probability-based: with a larger difference between the total upward cost of the relaying node and the downward hop's total upward path cost having a lower probability of relaying the request.

25

For example, node F, if not yet registered, when receiving the registration request from node E may re-broadcast the registration request. If it decides to do so, it will put a value of 3 as the upward hop cost, and a value of 0 as the upward hop's total upward path cost. If node D receives this re-broadcasted registration request, it will relay it to

30

the DC by adding its total upward cost of 6 and link cost between node F and itself of 2, to give a value of 8 as the upward hop's total upward path cost.

5 The DC, when receiving the relayed registration requests, stores the downward path (downward hop address) and downward hop cost contained in the registration requests into its own routing table. This downward path information is to be used in source routing to deliver packets to the registering node.

10

In this manner, the DC consolidates all the relayed registration requests from the registering nodes and decides the optimal upward path for each registering node using the total upward path costs in each request.

15

The total upward path costs would be the sum of the upward hop cost and the total upward path cost for that upward hop. If the DC receives a broadcasted (rather than relayed) registration request directly, it adds the upward link cost from the transmitter node of the request to itself, to this total to determine the final total upward path cost.

20

Once the values are stored in the DC's routing table, the DC sends a registration confirmation message through the downward hop address to each node. The registration confirmation message includes information on the optimal upward path (upward hop address) to the registering node, and the total upward path cost to the node, and may include information on a dedicated time slot for the destination node to broadcast discovery requests. The DC also stores the upward hop address and upward hop cost for the registered node in its routing table.

25

30

By way of example, the DC may have the following routing table for the network of Figure 2:

node	Downward hop addresses, and downward hop cost	Upward hop addresses, and upward hop cost
B	DC, 3	DC, 5
E	B, 1	B, 2
D	DC, 4	B, 1
F	E, 3	D, 2
A	DC, 3	DC, 5
C	A, 1	A, 2

Typically, a registration confirmation contains the following information:

- a) The registering node address;
- b) The downward hop address;
- c) The total downward path cost to the registering node;
- d) The upward hop address; and
- e) The total upward path cost from the registering node.

Each node in the network 10, when receiving a registration confirmation message, updates itself to a registered state, and stores the upward hop address, total upward path cost, downward hop address, and total downward path cost; the upward hop address being used by this node when sending packets to the DC. It then begins to broadcast discovery requests, as described above.

A node that is in a registered state generally periodically compares the total downward costs through the different downward hop addresses in its discovery table with its current total downward path cost and can then decide to send a registration request with the newly selected downward hop.

Alternatively, if there is periodic polling from the DC to each node, for example, a “keep alive” poll, the node can inform the DC of the newly selected downward hop address in the response to this poll.

The DC, when receiving the new downward path, can then decide to update its routing table with this new downward path and send a registration confirmation to the node in question. The DC is able to determine from its routing table if this is a valid downward path, as it’s possible to be a looping path. When the node receives the registration confirmation, it changes the total downward path cost appropriately in its discovery request broadcast.

Similarly a registered node, when receiving discovery requests, may determine the total upward path cost from the node that transmitted the discovery request to the DC if it were replaced as the upward hop for that node. If this determined total upward path cost is less than the total upward path cost specified in the discovery request by a certain margin, it can then decide to send a registration request to the DC, on behalf of the node that transmitted the discovery request, so as to change the upward hop for this node to itself.

Also, a registered node, when receiving discovery requests, typically checks if the upward hop address involves itself, and then checks if the upward hop cost has differed above a certain threshold. If it has, it then sends a registration request to the DC, to inform of the changes in upward hop cost, or waits for a “keep alive” poll to inform the DC of the change.

The DC, when receiving this registration request, can then decide to confirm the request by sending the node that transmitted the discovery request a registration confirmation message with the updated upward hop address and total upward path cost. When the node receives the registration confirmation, it updates its routing table with the new values and changes the total upward path cost in its discovery request broadcast.

Changes in total upward cost for a given node may change optimal routing paths for downstream nodes. The DC discovers these changes and informs the affected nodes accordingly either using registration confirmation messages or when performing keep-alive polling.

For example, with reference to Figure 3, the upward link cost between node D and the DC has been changed from a value of 7 to a value of 4. This affects the optimal routing paths through the network.

In such an example, node D would broadcast discovery requests with a total downward path cost of 4 and the old total upward path cost of 6 (via node B). When the DC receives this cost, it compares this with the known upward link cost between node D and itself, which is now 4. If this difference is greater than a certain margin, it then sends a registration confirmation message to node D so that node D changes its routing table to use the DC as its upward hop.

As a result, node D's discovery request broadcast then changes to a total downward path cost of 4 and a total upward path cost of 4. The DC then informs the affected downstream nodes, e.g. node F of the change of its total upward cost, which will now be a value of 6 instead of 8. This can be done

using a new registration confirmation message, or when the DC performs its keep-alive poll of node F.

In networks having multiple (or N) downward and upward routes instead of just one, the DC or a node is able to choose a different path if the application decides that another path to the destination should be attempted if, for example, communication does not seem to work on the original path.

In such an example, the discovery request described above would now include N total downward path costs and N total upward path costs for each relevant node. Similarly, there are now N downward and N upward hop addresses, associated with the total downward and total upward path costs. In contrast with the earlier discovery request described above, the new discovery request will contain the following information:

- a) node address that is transmitting this discovery request;
- b) N total downward paths cost through this node;
- c) N downward hop addresses and path numbers associated with the downward path costs;
- d) N total upward path cost through this node; and
- e) N upward hop addresses and path numbers associated with the upward path costs.

In addition, as there are N downward and upward paths in each discovery broadcast, the path number is also recorded. For example, for the network in Figure 2, if it is assumed that nodes B and E have started broadcasts of discovery requests or packets, broadcasting total downward path costs of 3 and 4, respectively, the discovery table for node D would be, for entries B, E and DC:

Downward	Downward	Received	total	Range	of
----------	----------	----------	-------	-------	----

hop address	hop cost	downward path costs and path number	received upward costs	total path
DC	4	(0, 0)	(0, 0)	
B	2	(3, 0)	(5, 5)	
E	2	(4, 0)	(7, 7)	

The path number, in each case, is 0, as there was only one path for each of them.

5 If $N = 2$, then node D would choose (DC, 0), and (B, 0) as the chosen downward paths as their total downward path costs are the lowest two, being 4 and 5, respectively. This would be placed into the registration request broadcast.

10 In the example described above, an unregistered node would choose N best downward hop addresses from its discovery table and broadcast a registration request. Conversely, in the current example a registration request now contains the following information:

- 15 a) The registering node address;
b) N downward hop addresses and path numbers;
c) N downward hop costs;
d) The range of downward hop upward path costs;
e) The upward hop address this request represents;
20 f) The upward hop cost; and
g) The upward hop's total upward path cost.

As for upward paths, when a node receives a registration request broadcast, it uses the range of upward path costs to
25 determine if it should relay the registration request to the DC, as in the earlier described example. Similarly,

unregistered neighbour nodes receiving registration requests continue to behave as described above.

However, it is to be appreciated that registered neighbour nodes receiving registration requests are handled as above, but that only the lowest total upward path cost is used. The DC will determine the other (N - 1) path costs from its routing table, if required.

Similarly, the DC consolidates all the costs like in the example described above, with the addition of possible multiple paths per relayed request to choose N upward paths. The DC determines N upward paths and costs per request received by looking at its routing table for that node. The DC also validates the paths, places N paths per node in its routing table, and confirms the paths to the registering node.

For example, for N = 2, the routing tables for the DC from the network 10 shown in Figure 2 may look like:

node	Downward hop addresses, path number, and downward hop cost	Upward hop addresses, path number, and upward hop cost
B	(DC, 0, 3), (D, 0, 1)	(DC, 0, 5), (D, 1, 2)
E	(B, 0, 1), (B, 1, 1)	(B, 0, 2), (D, 0, 2)
D	(DC, 0, 4), (B, 0, 2)	(B, 0, 1), (DC, 0, 7)
F	(E, 0, 3), (D, 0, 4)	(D, 0, 2), (D, 1, 2)

If the DC receives a registration request originally from node E but relayed through node B with a lowest total upward path cost of 7, the DC works out that the two possible paths through node B are E - B - DC, which is valid and E - B - D -

B . . . , which is invalid. So the only path cost to consider is 7.

If the DC receives a registration request originally from node E but relayed through node D with a lowest total upward path cost of 8, the DC works out that the two possible paths through node D are E - D - B - DC, and E - D - DC, which are both valid and their associated costs are $2 + 6 = 8$, and $2 + 7 = 9$.

Accordingly, the DC sends a registration confirmation message with the N multiple downward and upward paths and their associated total path costs so that the node may store the confirmed upward paths and downward paths and costs as described above, and perform discovery request broadcasts with the N downward and upward paths and associated total costs.

The maintenance of N optimal downward paths occur as in the example described above, except that the comparison is with N stored downward path costs, and the node sends a registration request if this set of N paths changes.

It is also to be appreciated that the maintenance of optimal upward paths and costs occur as in the above described example, except that:

a) If there are N upward paths in the discovery request, it can decide to send a registration request to the DC if its determined total path is less than the highest total upward path cost in the discovery request.

b) If there are less than N upward paths in the discovery request, the decision can be probability-based, with a larger difference between the determined total upward path cost through itself and the highest total upward

path cost of discovery request resulting in a lower probability of sending a registration request.

Similarly, as described above, the DC discovers which
5 downstream nodes are affected by changes in link costs and informs the downstream nodes accordingly.

For example, in the example above where N = 2
Error! Reference source not found., if the following changes occur in
10 the DC's routing table:

node	Downward hop addresses, path number, and downward hop cost	Upward hop addresses, path number, and upward hop cost
B	(DC, 0, 3), (D, 0, 1)	(DC, 0, 5), (D, 1, 2)
E	(B, 0, 1), (B, 1, 1)	(B, 0, 2), (D, 0, 2)
D	(DC, 0, 4), (B, 0, 2)	(B, 0, 1), (DC, 0, 7) -> (DC, 0, 4), (B, 0, 1)
F	(E, 0, 3), (D, 0, 4)	(D, 0, 2), (D, 1, 2)

The values will accordingly change as follows:

node	Downward hop addresses, path number, and total downward path cost	Upward hop addresses, path number, and total upward path cost
B	(DC, 0, 3), (D, 0, 1)	(DC, 0, 5), (D, 1, 2) -> (DC, 0, 5), (D, 0, 2)
E	(B, 0, 1), (B, 1, 1)	(B, 0, 2), (D, 0, 2) -> (D, 0, 2), (B, 0, 2)
D	(DC, 0, 4), (B, 0, 2)	(B, 0, 1), (DC, 0, 7) -> (DC, 0, 4), (B, 0, 1)
F	(E, 0, 3), (D, 0, 4)	(D, 0, 2), (D, 1, 2)

15

In such an example, the DC will update nodes B and E using either registration confirm messages or through "keep alive" polls. It is to be appreciated that when using the routing table to route packets upwards, the relevant nodes
5 need to specify the route number in addition to the node address.

For example, in the DC's routing table where $N = 2$ above, when node F sends a packet to the DC, it needs to specify
10 either the route (D, 0), or (D, 1). This would result in the route F - D - B - DC, or F - D - DC respectively.

Referring now to Figure 4 of the drawings, there is shown one example of apparatus 20 for a power line communication
15 network 10, as described above. In use, the apparatus 20 connects to a power line 22 to function as a node in a power line communication (PLC) network 10, as described above with reference to Figure 1 to 3.

Typically, the apparatus 20 comprises a power line transceiver 24 via which the apparatus 20 is able to establish an interface 24 with a power line 22. Such an interface 26 is well understood in the art and will not be described in
20 detail.

The apparatus 20 also generally comprises a processor 28 coupled to the power line transceiver 26 by means of a suitable coupling 30. Again, such couplings are well known in the art and will not be described in detail. The processor 28
30 may comprise any suitable processing device, such as a central processing unit or integrated circuit configured for data processing, as is well known in the art. Typically, the

processor 28 comprises some form of arithmetic logic unit (ALU) and processor registers.

5 In this example, the processor 28 further comprises memory 32 which is accessible to the processor 28. The processor 28 is also configured to store a discovery table 34, as described above, which comprises a node address and a downward hop cost, the received total downward path cost and the received total upward path cost in the memory for routing packets. The
10 discovery table 34 typically further comprises path number information.

As described above, the apparatus 20 functions as a node in the PLC network 10 and is subsequently configured to
15 perform the necessary steps and calculations necessary to give effect to the methods described above.

Similarly, in a further example of the apparatus 20, the processor's memory 32 is configured to store a routing table
20 34 which comprises a downward hop address, a downward hop cost, an upward hop addresses and upward hop cost for each node. The routing table 34 typically further comprises path number information in accordance with the methods and requirements of the DC described above.

25 It should be appreciated that the scope of the invention is not limited to the scope of the embodiment described. Various modifications and improvements may be made to the embodiment described without departing from the scope of the invention.

30 Throughout this specification, unless the context requires otherwise, the word "comprise" or variations such as "comprises" or "comprising", will be understood to imply the

inclusion of a stated integer or group of integers but not the exclusion of any other integer or group of integers.

5 It is also to be appreciated that reference to "one
example" or "an example" of the invention is not made in an
exclusive sense. Accordingly, one example may exemplify
certain aspects of the invention, whilst other aspects are
exemplified in a different example. These examples are
intended to assist the skilled person in performing the
10 invention and are not intended to limit the overall scope of
the invention in any way unless the context clearly indicates
otherwise.

CLAIMS

1. A discovery method for a power line communication network, comprising:

5 broadcasting, by a first node, a discovery request to at least one receiver, wherein the discovery request comprises the following information:

 a node address of the first node;

 a total downward path cost through the first node;

10 a downward hop address associated with the total downward path cost;

 a total upward path cost through the first node; and

 an upward hop address associated with the total upward path cost.

15

2. The method of claim 1, wherein if the first node is a data concentrator (DC), the total downward path cost and total upward path cost are the values of 0, and wherein the downward hop address associated with the total downward path cost and
20 the upward hop address associated with the total upward path cost are excluded from the discovery request.

3. A discovery method for a power line communication network, comprising:

25 receiving, by a node, a discovery request which comprises a total downward path cost, a total upward path cost from a transmitter;

 determining, by the node, a total downward path cost from a data concentrator (DC) to the node by using a total downward
30 path cost in the discovery request;

 storing a node address included in the discovery request and a downward hop cost from the transmitter to the node, the received total downward path cost and the received total

upward path cost from the transmitter into a discovery table of the node.

4. The method of claim 3, further comprising rejecting, by
5 the node, paths which include the node itself as an upward or downward node in the path.

5. A discovery method for a power line communication network, comprising:

10 receiving, by a registering node, a discovery request from a transmitter;

choosing the best downward hop address from a stored discovery table, using a total downward path costs from a data concentrator(DC) to the registering node; and

15 broadcasting, by the registering node, a registration request, wherein the registration request comprises the following information:

a registering node address;

a downward hop address;

20 a downward hop cost to the registering node; and

a downward hop's total upward path cost.

6. The method of claim 5, the registration request further comprises an upward hop address, an upward hop cost from the
25 registering node to the upward hop address and an upward hop's total upward path cost which are filled by relaying or re-broadcasting nodes.

7. A discovery method for a power line communication
30 network, comprising:

receiving, by a neighbour node, a registration request;
and

re-broadcasting the registration request if the neighbour node is in unregistered state or relaying the registration request if the neighbour node is in registered state.

5 8. The method of claim 7, further comprising determining to relay based on a total upward cost of the neighbour node and a downward hop's total upward path cost.

9. A discovery method for a power line communication
10 network, comprising:

transmitting, by a data concentrator (DC), a registration confirm message to at least one registering nodes based on a total upward path costs in registration requests from the at least one registering nodes;

15 updating a routing table, by the DC, based on the registration requests from the at least one registering nodes, wherein the registration confirm message comprises the following information:

20 a registering node address;
a downward hop address;
a total downward path cost to the registering node;
an upward hop address; and
a total upward path cost from the registering node.

25 10. The method of claim 9, wherein the registering node address is assigned a shorter address instead of a longer address pre-configured into the registering node.

11. The method of claim 9, further comprising receiving, by
30 the at least one registering nodes, the registration confirm message from the DC; and transitioning, by the at least one registering nodes, to registered state.

12. A discovery method for a power line communication network, comprising:

comparing, by a registered node, a total downward costs through a different downward hop addresses in a discovery table with a current total downward path cost; and

determining, by a registered node, to send a registration request with a newly selected downward hop.

13. A discovery method for a power line communication network, comprising:

receiving, by a registered node, a polling from a data concentrator(DC); and

transmitting, by the registered node, a response which informs the DC of a newly selected downward hop address.

14. A discovery method for a power line communication network, comprising:

receiving, by a registered node, a discovery request;

determining, by the registered node, a total upward path cost from a node which transmits the discovery request to a data concentrator (DC) if the registered node is replaced as a upward hop for the node; and

deciding to send a registration request to the DC on behalf of the node which transmits the discovery request, if the determined total upward path cost is less than a total upward path cost specified in the discovery request by a certain margin.

15. The method of claim 14, further comprising:

checking, by the registered node, if the upward hop address involves itself and the upward hop cost has differed above a certain threshold; and

transmitting, by the registered node, a registration request to the DC or waiting for a poll, to inform of the changes in upward hop cost.

5 16. The method of claim 1, wherein the discovery request further comprises information associated with multiple downward and upward routes for the first node.

10 17. The method of claim 3, wherein the discovery table of the node further comprises path number information.

18. The method of claim 5, wherein the registration request further comprises path number information.

15 19. An apparatus for a power line communication network, comprising:

a power line transceiver;

a processor coupled to the power line transceiver; and

20 memory accessible to the processor, wherein the processor is configured to store a discovery table which comprises a node address and a downward hop cost, the received total downward path cost and the received total upward path cost in the memory for routing packets.

25 20. An apparatus for a power line communication network, comprising:

a power line transceiver;

a processor coupled to the power line transceiver; and

30 memory accessible to the processor, wherein the processor is configured to store a routing table which comprises a downward hop address, a downward hop cost, a upward hop addresses and upward hop cost for each node.

21. The apparatus of claim 19, wherein the discovery table further comprises path number information.

22. The apparatus of claim 20, wherein the routing table
5 further comprises path number information.