(12) UK Patent Application (19) GB (11) 2500609 (13) A

(43) Date of A Publication 02.10.2013

(21) Application No: 1205269.2

(22) Date of Filing: 26.03.2012

(71) Applicant(s):
Ubiquisys Limited
(Incorporated in the United Kingdom)
Windmill Hill Business Park, SWINDON, Wilts, SN5 6QR, United Kingdom

(72) Inventor(s):
Andrea Giustina
Cristavao Da Silva

(74) Agent and/or Address for Service:
Haseltine Lake LLP
Redcliff Quay, 120 Redcliff Street, BRISTOL, BS1 6HU, United Kingdom

(51) INT CL:
H04W 36/04 (2009.01)

(56) Documents Cited:
EP 2352332 A1
EP 2343920 A1
EP 2326122 A1
WO 2011/133921 A1
WO 2011/020488 A1
WO 2010/149829 A1
WO 2010/033729 A2

(58) Field of Search:
INT CL **H04W**
Other: **Online: EPODOC, WPI**

(54) Title of the Invention: **Mobile communications network**
Abstract Title: **Handover of user equipment from a large cell to a small cell operating with the same physical cell identity as other small cells**

(57) A cellular telecommunications network comprises a plurality of large cells (e.g. macrocells) wherein the coverage area of at least one of the large cells includes a plurality of small cells, each operating with a respective Physical Cell Identity (PCI) that does not unambiguously identify each of the small cells (i.e. the small cells may operate with the same Physical Cell Identity (PCI) as other small cells within the coverage area of the large cell). Hence, on determining that handover of user equipment is required from a serving large cell to one of the small cells included within its coverage area, a handover request is sent from the serving large cell identifying the Physical Cell Identity (PCI) of one of the small cells and including additional information to allow the network to reduce ambiguity in an identity of the small cell. The additional information may comprise information relating to a location of the user equipment within the large cell, a timing offset between transmissions of the large cell and the small cell, or neighbour cells that can be detected by the user equipment.

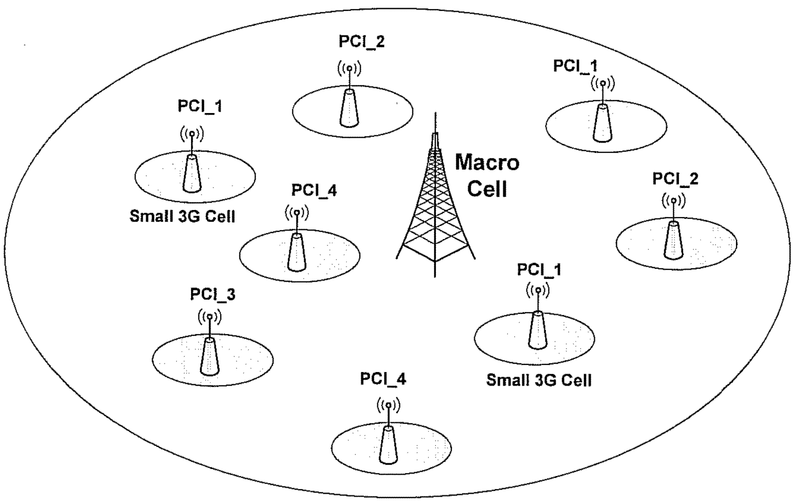

Figure 2

GB 2500609 A

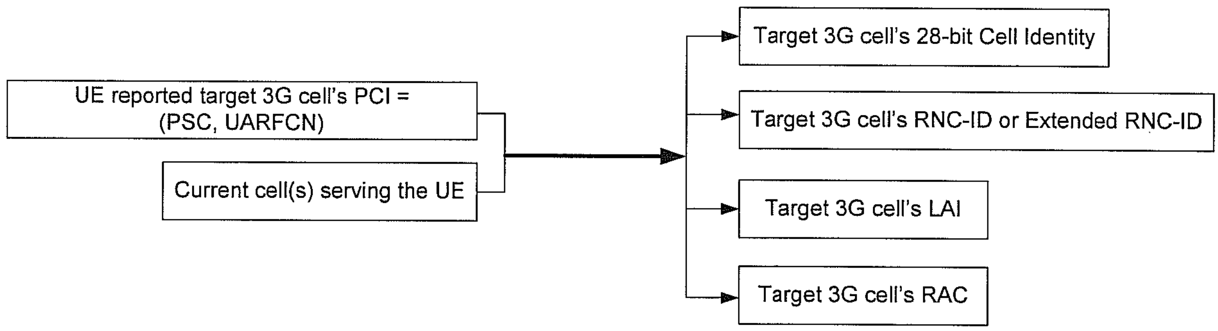

Figure 1

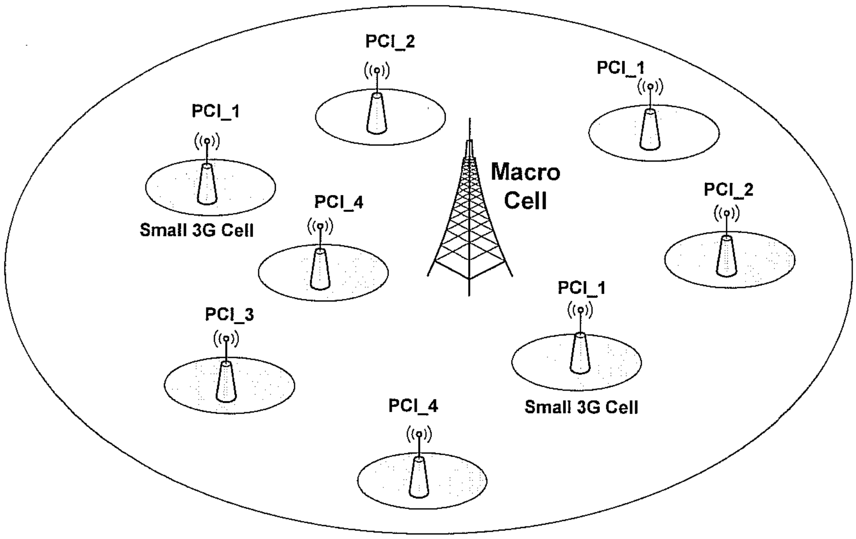

Figure 2

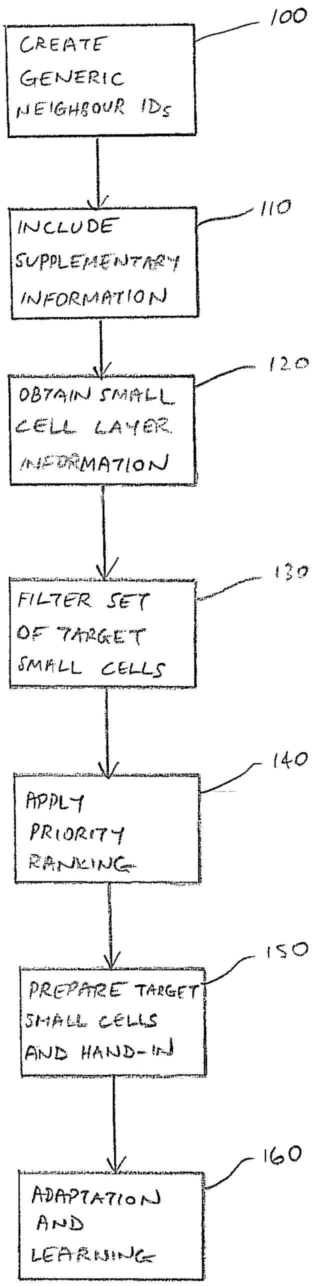
CREATE GENERIC NEIGHBOUR IDs
100
INCLUDE SUPPLEMENTARY INFORMATION
110
OBTAIN SMALL CELL LAYER INFORMATION
120
FILTER SET OF TARGET SMALL CELLS
130
APPLY PRIORITY RANKING
140
PREPARE TARGET SMALL CELLS AND HAND-IN
150
ADAPTATION AND LEARNING
160

Figure 3

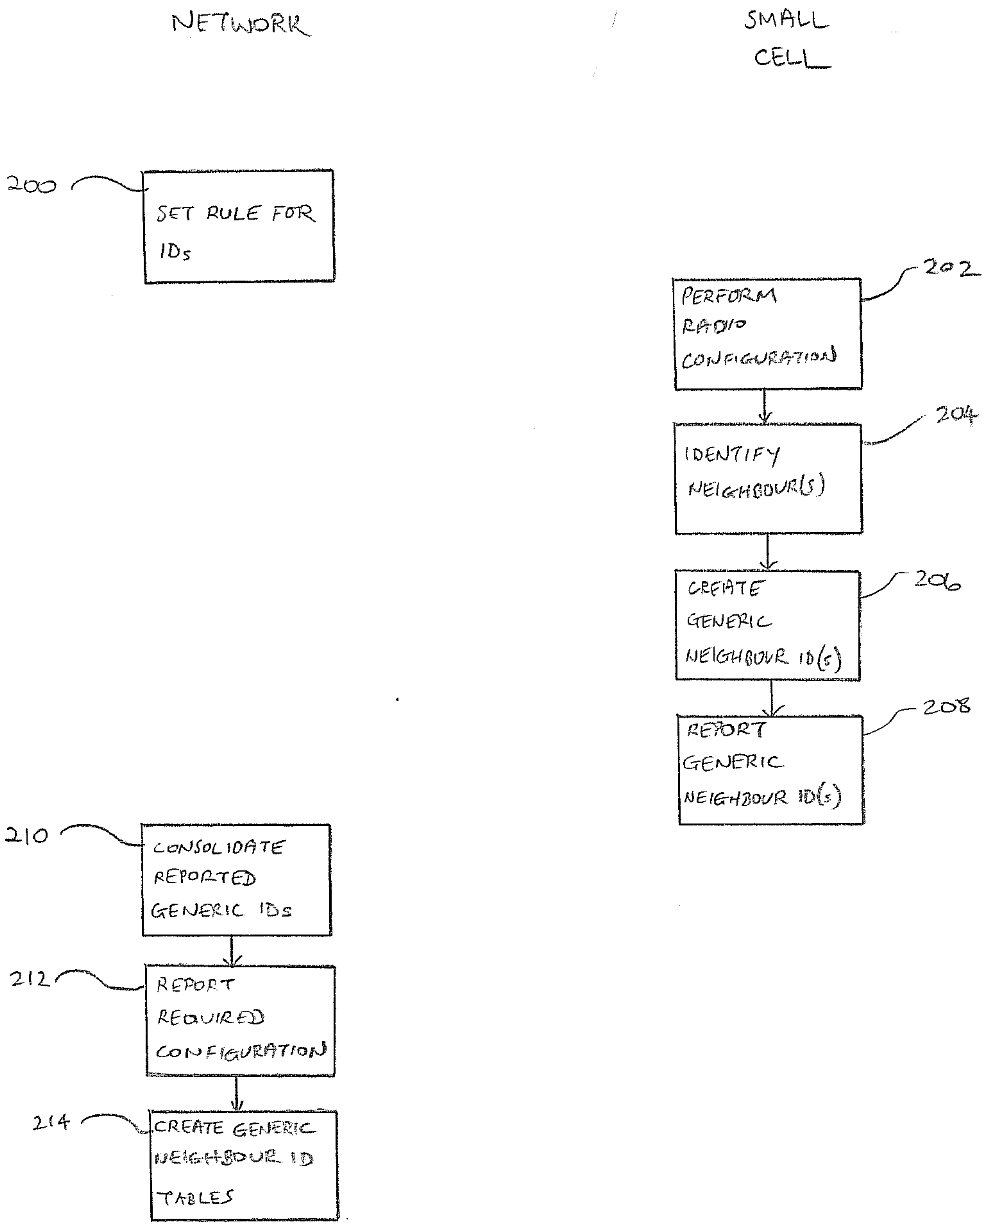
NETWORK
SMALL CELL
200
SET RULE FOR IDs
202
PERFORM RADIO CONFIGURATION
204
IDENTIFY NEIGHBOUR(S)
206
CREATE GENERIC NEIGHBOUR ID(S)
208
REPORT GENERIC NEIGHBOUR ID(S)
210
CONSOLIDATE REPORTED GENERIC IDs
212
REPORT REQUIRED CONFIGURATION
214
CREATE GENERIC NEIGHBOUR ID TABLES

Figure 4

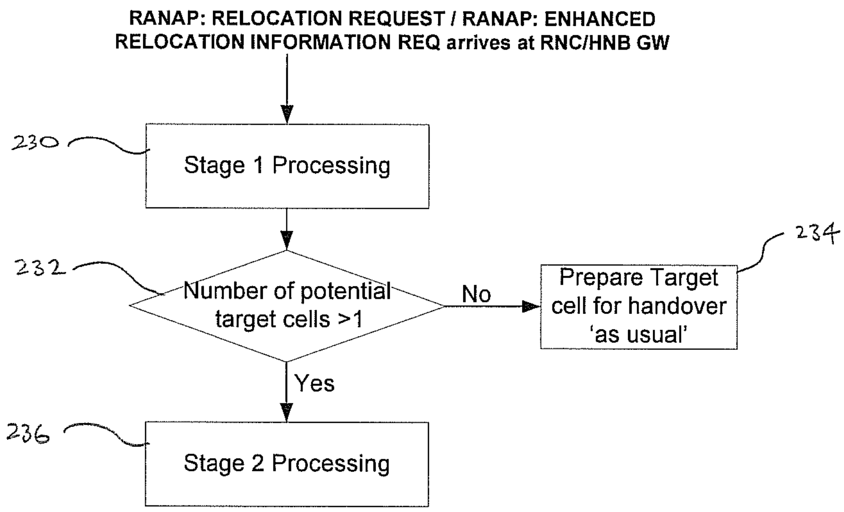

Figure 5

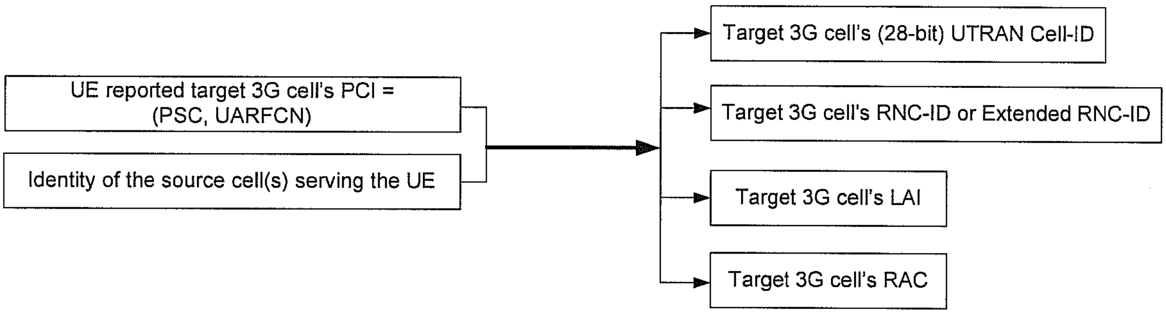

Figure 6

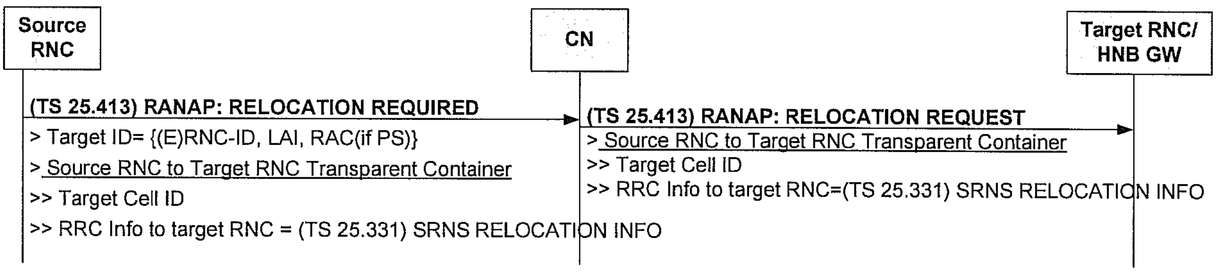

Figure 7

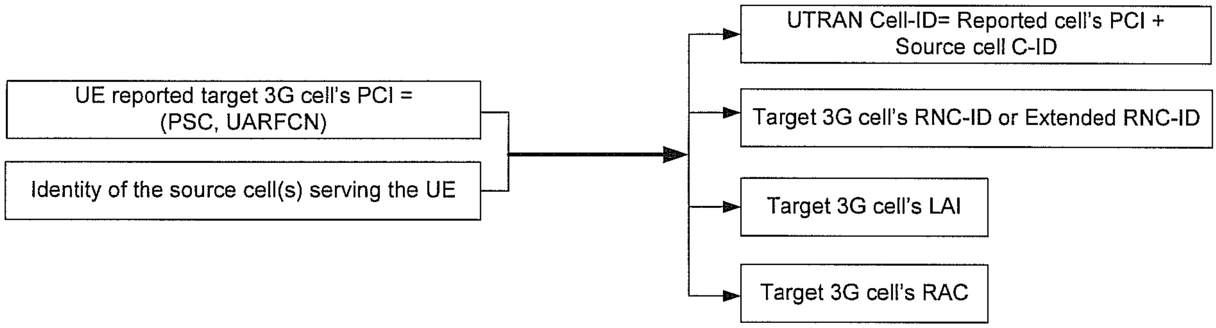

Figure 8

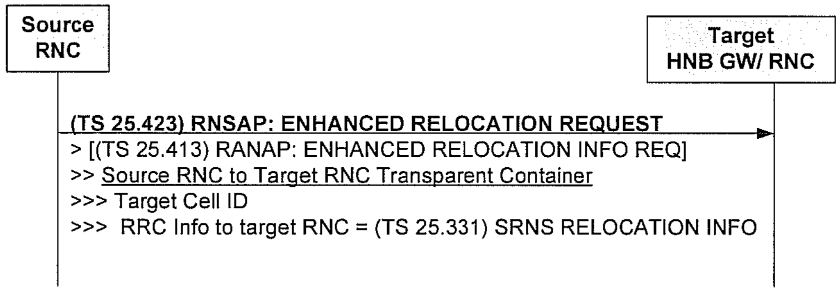

Figure 9

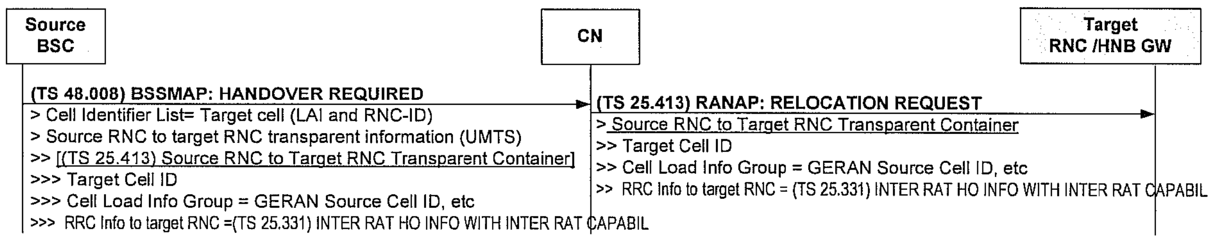

Figure 10

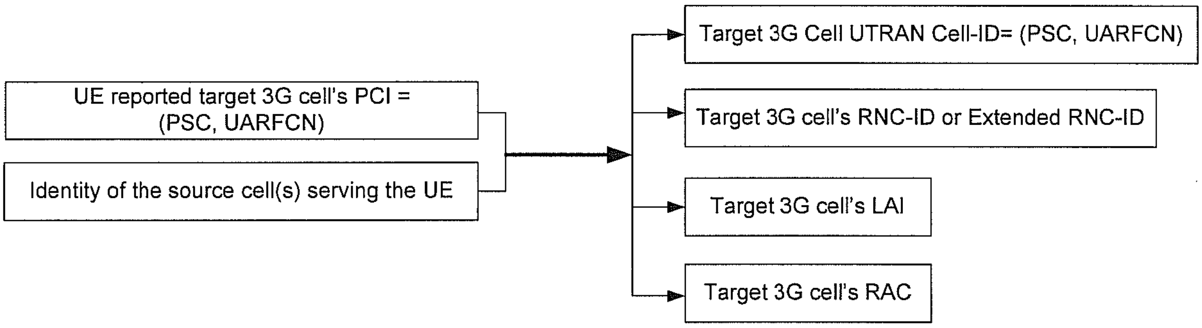

Figure 11

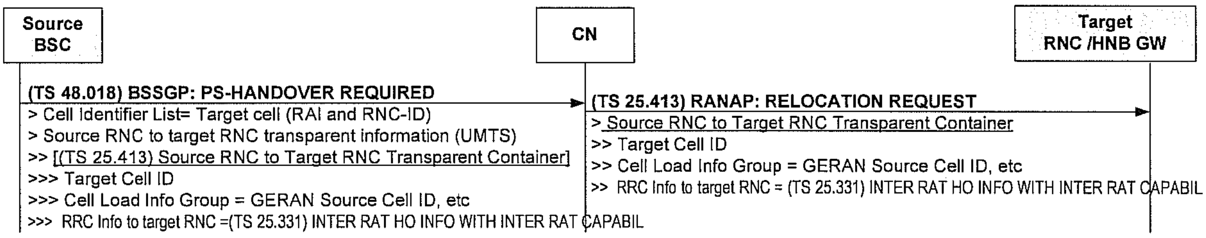

Figure 12

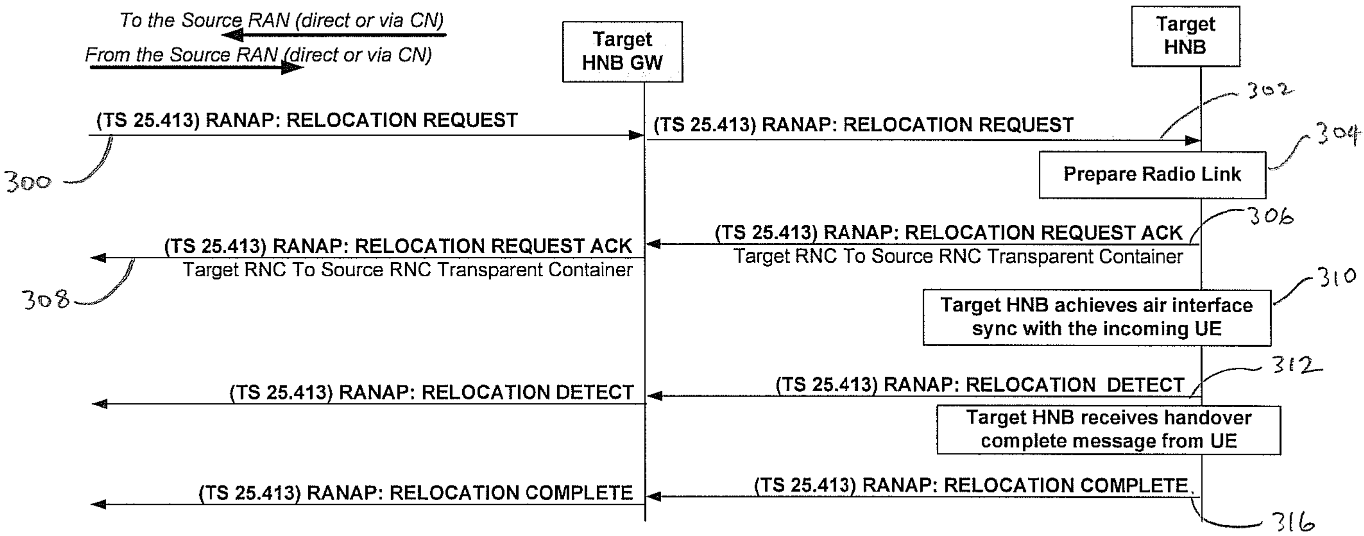

Figure 13

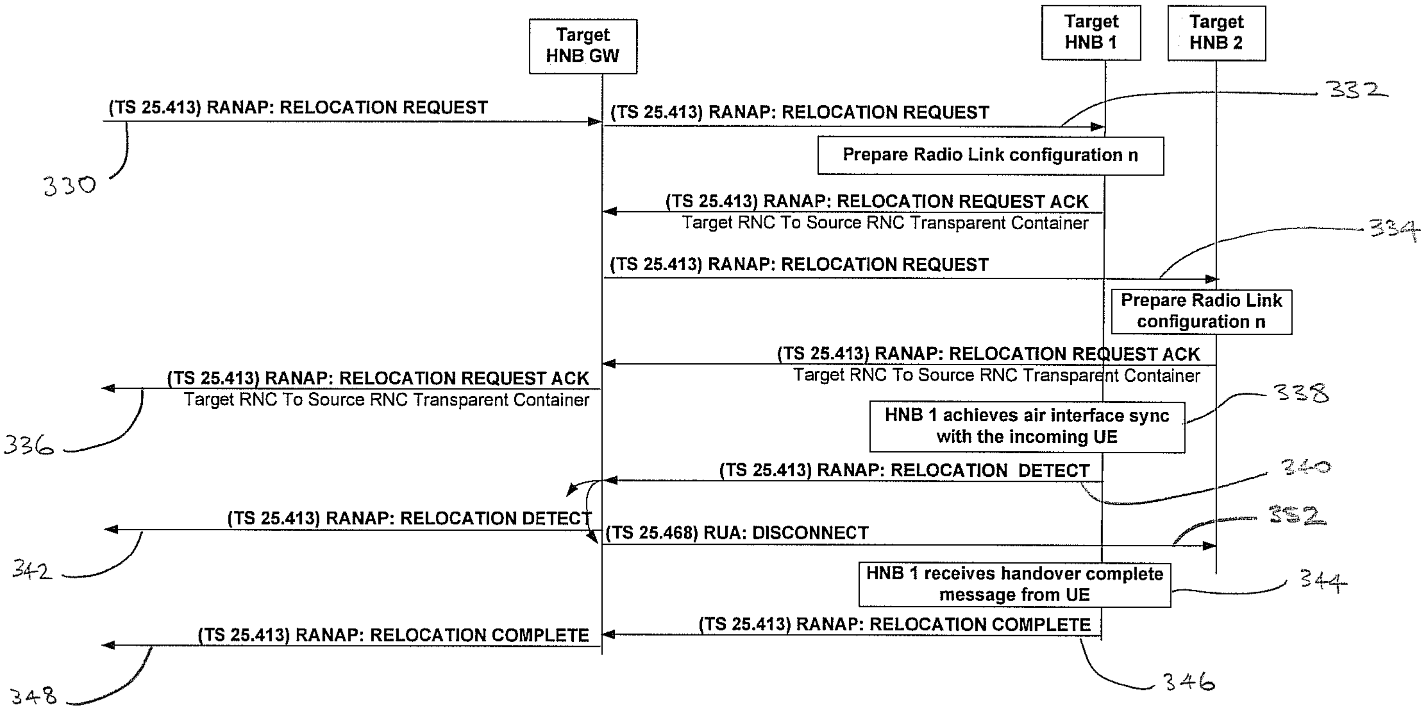
Target HNB GW
Target HNB 1
Target HNB 2
(TS 25.413) RANAP: RELOCATION REQUEST
330
(TS 25.413) RANAP: RELOCATION REQUEST
332
Prepare Radio Link configuration n
(TS 25.413) RANAP: RELOCATION REQUEST ACK
Target RNC To Source RNC Transparent Container
(TS 25.413) RANAP: RELOCATION REQUEST
334
Prepare Radio Link configuration n
(TS 25.413) RANAP: RELOCATION REQUEST ACK
Target RNC To Source RNC Transparent Container
(TS 25.413) RANAP: RELOCATION REQUEST ACK
Target RNC To Source RNC Transparent Container
336
HNB 1 achieves air interface sync with the incoming UE
338
(TS 25.413) RANAP: RELOCATION DETECT
340
(TS 25.413) RANAP: RELOCATION DETECT
342
(TS 25.468) RUA: DISCONNECT
352
HNB 1 receives handover complete message from UE
344
(TS 25.413) RANAP: RELOCATION COMPLETE
346
(TS 25.413) RANAP: RELOCATION COMPLETE
348

Figure 14

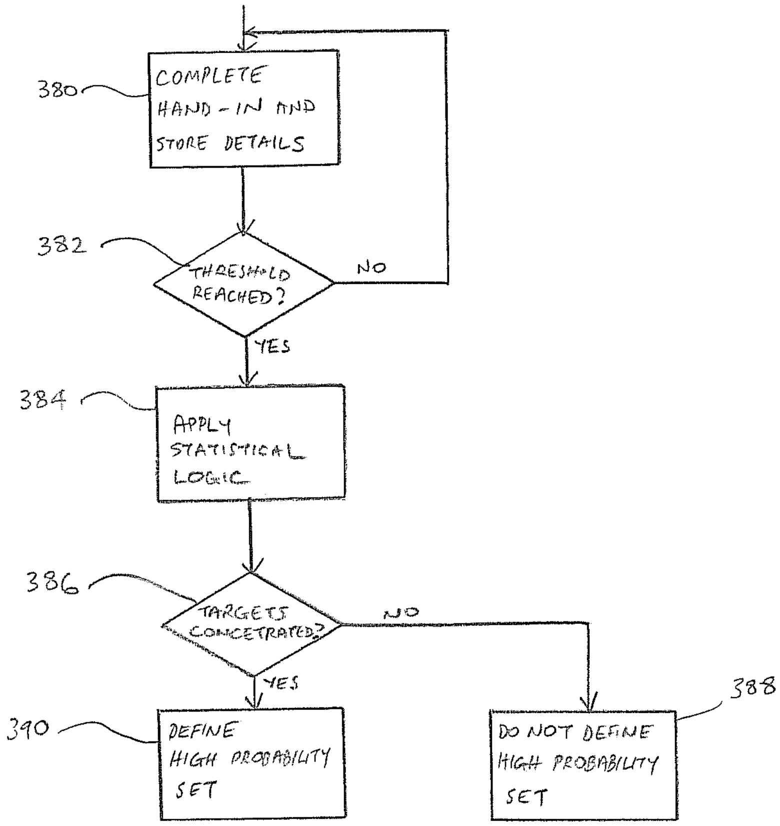

Figure 15

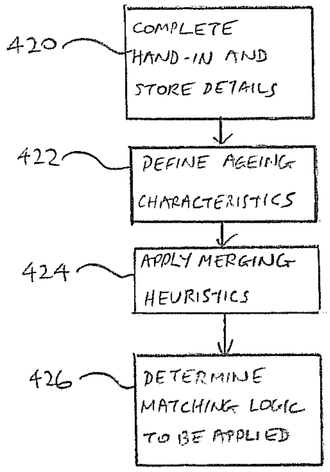

Figure 16

MOBILE COMMUNICATIONS NETWORK

This invention relates to a cellular mobile communications network, and to basestations for use in such a network.

In a cellular network, a user equipment device operates by establishing a connection with a basestation of the network. Each basestation has a respective coverage area, and the network operator typically aims to ensure that the user equipment device can establish a connection with one of the basestations, wherever it might be.

Thus, as a user equipment device moves from the coverage area of one basestation into the coverage area of another basestation, it is necessary for its connection with the first basestation to be handed over to the second basestation. Generally, it is necessary to identify the second basestation unambiguously, before the user equipment connection with the first basestation is handed over to the second, target basestation.

In a traditional cellular network, the coverage areas, or cells, of the basestations are all of approximately similar sizes, and the handover takes place as a user equipment device reaches a boundary between two cells.

In such a network, it is possible for the second basestation to be identified unambiguously, for example by means of the carrier frequency and scrambling code being used by the second basestation in the case of a Universal Mobile Telecommunications System (UMTS) network using Wideband Code Division Multiple Access (WCDMA) in the air interface. Any given user equipment device will probably be able to detect signals transmitted by only one basestation having any specific combination of carrier frequency and scrambling code.

However, increased demand for network capacity can cause congestion in such networks. One solution to this problem is the use of small cells and heterogeneous networks. Thus, basestations with much smaller coverage areas have been introduced, for example to provide extra network capacity in locations with large numbers of users; or to provide network coverage within buildings where signal reception might otherwise be a problem; or to provide network capacity within particular business or residential premises, for the occupiers of those premises.

In such a network, there might be a large number (perhaps in the tens, hundreds, or even thousands) of small cells within one of the larger cells (referred to as macrocells).

In a heterogeneous network of this type, the number of combinations of carrier frequency and scrambling code is not sufficient for the combination of carrier frequency and scrambling code used by one basestation to be unique, even amongst the small cells within one of the macrocells.

As a result, there is an issue of the interworking between the small cells and the Macro Layer (ML), and one particular problem concerns hand-in, i.e. how to perform a handover from the ML cells to small cells. Some user equipment devices complying with more recent releases of the 3GPP standards are able to identify a target basestation by means of its unique cell identifier, but these user equipment devices are expected to remain relatively rare for some years, and so it is necessary to allow for successful handover from a macrolayer basestation to a particular small cell basestation, even when the user equipment device is unable to report the target cell identity, and so the ML cell cannot distinguish the actual target of the handover.

Hand-in from the Macro Layer to the small cells layer is not well supported in current technology implementation (before Release 9 in UMTS UEs and networks). Essentially, the hand-in issues are aspects of the “one-to-many” relationship that exists between the large cells making up the Macro Layer and the small cell layer.

Traditional handover techniques require the full identification of the target cell (here called *unambiguous hand-in*). In current technology implementation, this is based on the macro RNC using simple UE measurements (UARFCN and PSC) in conjunction with the identity of the current service macrocell, to identify the target cell. Release 9 UEs will report the full target cell ID, but their penetration is still minimal and a solution to this problem must be valid for all UEs including legacy ones.

In a heterogeneous network deployment, the Macro Layer neighbour list is soon exhausted by adding just a few small cells in the same area, and so there is no possibility to scale beyond a dozen or so small cells per macro cell, while the operator requirement is to get to hundreds if not thousands of small cells in the area covered by a macro cell (in particular in the case of 2G macro cells, which typically cover an area greater than 3G macro cells).

For any small cells deployment to scale to more than a dozen per macro cell (including very large 2G macro cells) and keep hand-in capability with the current UE population, the solution must allow for *ambiguity* of the handover target cells, i.e. multiple small cells in the same ML area (under the same ML cell coverage) reuse the same UARFCN and PSC and are hence indistinguishable just by using the pre-R9 UE reports.

Small cells can form multiple layers, outdoor and indoor, open mode and closed mode, from “larger” small cells layer to “smaller” small cells and femtocells. Seamless mobility must be achieved between all of these layers and also between one of these small cells layers and the ML. The same techniques applied to achieve seamless mobility and hand-in between the ML and a small cells layer must be applicable to the hand-in from a “larger” small cells layers to a “smaller” small cells layers or between peer small cells. Furthermore, to make it scalable into the highest numbers, the small cells layers must achieve seamless mobility within and between layers with very limited engineering needs, i.e. the small cells layer should preferably self-configure and also produce automatic reports on how to simply reconfigure the ML RNCs or “larger” small cells layers to allow for the seamless mobility.

In 3GPP radio access networks, i.e. a GSM EDGE Radio Access Network (GERAN), a UMTS Radio Access Network (UTRAN) or an Extended UMTS Radio Access Network (E-UTRAN), the smallest logical component is a cell. In every 3GPP Public Land Mobile Network (PLMN), each cell belonging to each RAT (Radio Access Technology) is identified by a logical cell identifier that takes a value that is unique in the PLMN. 3G cells can be located in a Node B or in a Home Node B.

In the specific case of the UTRAN, there is an extra degree of complexity in that each cell is actually associated with two logically independent globally unique 28-bit cell identifiers, namely the UTRAN Cell-ID defined in 3GPP TS 25.401 and the Cell Identity IE as defined in 3GPP TS 25.331. In typical networks, these two identifiers are set to the same value, but they do not necessarily have to coincide.

The UTRAN Cell-ID identifies the cell in the context of signalling between UTRAN nodes and is not visible to the UE. This identity is defined in 3GPP TS 25.401, as a 28-bit field composed of two logical components: the identity of the RNC that controls the cell plus a C-ID value which identifies the cell within that Controlling RNC. Typically, the RNC-ID is a 12-bit quantity which makes the C-ID a 16 bit quantity. However, in 3GPP Release 7 the concept of 'Extended RNC-ID' was introduced and, in UTRANs using this option, each RNC is identified by a 16-bit Extended RNC-ID, which means that the C-ID is reduced to a 12-bit range. Thus (28-bit) UTRAN Cell-ID = (12-bit) RNC-ID + (16-bit) C-ID or (28-bit) UTRAN Cell-ID = (16-bit) Extended RNC-ID + (12-bit) C-ID.

The 28-bit Cell Identity IE is defined in 3GPP TS 25.331 and broadcast by each cell in SIB 3. This is the logical cell identity that is visible to the UE. In a typical UTRAN the operator will set it to the UTRAN Cell-ID, but the only standard-based requirement on this value is that it be unique within the PLMN.

Before 3GPP Release 9 (R9), a UE was unable to report the Cell Identity IE broadcast by a measured 3G cell while in dedicated mode and there was no scope for the RAN serving the UE (GERAN, UTRAN, E-UTRAN) to request the UE to do so. In R9 this functionality was introduced in order solve the problem of how to uniquely identify a target 3G cell (for the purposes of handover) from UE measurement reports when the target 3G cell has not been deployed in coordination with the Source RAN.

In handover to RNC/HNB GW procedures the Source RAN must provide the unique identifier of the target 3G cell in the handover request message it sends toward the target RNC/HNB GW.

Typically the handover is triggered at least in part by UE measurements reports on the quality of the target 3G cell. However, in pre-R9 UEs, in these measurement reports the UE identifies the target 3G cell by what is often called the cell's Physical Cell Identity (PCI), which is the combination of the UTRAN frequency operated by the cell (identified by its UARFCN) and the Primary Scrambling Code (PSC) operated by the cell; which takes a value in the 0-511 range.

Traditional 3G deployments are characterised by two important properties. Firstly, every time a new 3G cell is introduced in the network it is allocated a locally unique

PCI, where the term ‘locally unique’ means unique within the coverage area of the small number of cells of each RAN from which handover to that new 3G cell may take place. Secondly, such networks follow a coordinated deployment model whereby, every time a new 3G cell is introduced in the network, its identifiers (i.e. PCI, UTRAN Cell-ID, RNC-ID/Extended RNC-ID, Location Area Identifier (LAI), and Routing Area Code (RAC)) are added to the neighbour relations table defined in the source RAN for any of its cells from which handover to the new 3G cell is to be supported. In a deployment of this type, as illustrated in Figure 1, when a UE operating in dedicated mode in one or more cells in the source RAN reports the PCI of a 3G cell to the Source RAN, and the RAN knows the current serving cell(s) of the UE, it is able to derive the 3G cell ‘s UTRAN Cell-ID, RNC-ID/Extended RNC-ID, LAI, and RAC.

Thus, the 3GPP handover procedure requires a full engineering in the ML of the target cells. However, operators are currently using a very different deployment model for what are currently called 3G small cells, which range from the residential Home Node B cell to picocells or microcells. These cells are expected to become widely deployed and carry a very considerable fraction of the UTRAN’s traffic in the near future. It is thus very important to support handover to these cells.

The problem is that the small cell deployment model breaks both of the above listed characteristics of the traditional 3G deployment model, namely the characteristics which in those traditional deployments enable the source RAN to unique identify a target 3G cell for handover purposes without requiring the UE to report the 28-bit Cell Identity IE broadcast by the 3G cell. First, the small cells are deployed in an uncoordinated fashion, i.e., when one of these cells is deployed at a certain location the operator will not generate a new cell-specific entry to the neighbour relations table that the Source RAN holds for each of its cells. On the contrary, the objective is to make the addition and removal of such cells as invisible as possible to the overlapping source RANs so as to minimise the operating expense. It is disadvantageous for the operators to have to go beyond the provisioning of a small number of generic entries into the neighbour relations table for each source cell, into which all small 3G cells deployed under the source cell must somehow fit. Second, local uniqueness of the PCIs of the small cells will often be logically impossible. That is, the number of these small cells under a source RAN cell might well be smaller than the number of PCIs that the operator has reserved to be used by the deployment.

Thus the target cell identification problem in handover from a source RAN to this type of 3G cells has two aspects. Firstly, the entries in the neighbour relations tables at the Source RAN that point to the small 3G cells cannot contain any information pertaining to a specific small 3G cell logical 28-bit cell identity, which means that the Source RAN handover request message cannot convey that information. Even if, somehow, the Source RAN handover request message can be made to provide the same information that is used by the Source RAN in traditional deployments to identify a target cell (i.e. the PCI of the target cell and the identity of the source cell(s)), this information will not be enough to allow the Target system to determine the 28-bit target cell identity, because of the fact that the PCI is no longer locally unique in the small 3G cell deployments.

The present invention seeks to address this problem without using the Release 9 solution that is based on the UE reporting (to the Source RAN) the unique Cell Identity IE that each small 3G cell broadcasts. This is because, currently, there are few commercial UEs supporting this Release 9 feature, and it is clear that UEs that do not support this feature will be in use for many years. Thus, an alternative solution must be provided to the users of those UEs to enjoy similar quality of experience. Moreover, the feature also requires an upgrade to Release 9 of the Source RAN which, in the case of the GERAN, may not be widely implemented.

According to the present invention, there are provided methods of operation of a telecommunications network. Specific aspects of the invention define methods as performed in small cell basestations, methods as performed in core network nodes such as gateway nodes, small cell basestations themselves, and computer readable media containing code for causing the small cell basestations to perform the methods as defined in the other aspects.

For a better understanding of the invention, and to show how it may be put into effect, reference will now be made, by way of example, to the accompanying drawings, in which:-

Figure 1 illustrates the information available in a conventional deployment of 3G cells;

Figure 2 shows an example of a telecommunications network operating in accordance with the present invention;

Figure 3 is a flow chart, illustrating a method in accordance with an aspect of the invention;

Figure 4 is a flow chart, illustrating in more detail a part of the method of Figure 3;

Figure 5 is an illustration of a part of the process shown in Figure 3;

Figure 6 illustrates the information available in a deployment of cells;

Figure 7 illustrates a relocation procedure;

Figure 8 illustrates the information available in an embodiment of the invention;

Figure 9 illustrates a further relocation procedure;

Figure 10 illustrates a still further relocation procedure;

Figure 11 illustrates the information available in an embodiment of the invention;

Figure 12 illustrates a still further relocation procedure;

Figure 13 illustrates a still further relocation procedure;

Figure 14 illustrates a still further relocation procedure;

Figure 15 is a flow chart, illustrating a method in accordance with an aspect of the invention; and

Figure 16 is a flow chart, illustrating a further method in accordance with an aspect of the invention.

The invention is described herein for the purposes of illustration in the context of a network operating in accordance with 3rd Generation Partnership Project (3GPP) standards. However, the same principles can be used in networks made up of basestations operating in accordance with some or all of 2G, 3G, WiFi, LTE or other technologies.

Figure 2 shows a small part of a heterogeneous network, for the purposes of illustration only, with multiple small cells within the coverage area of one macrolayer (ML) cell. Each small cell has a Physical Cell Identity (PCI), which is the combination of the UTRAN frequency operated by the cell (identified by its UARFCN) and the Primary Scrambling Code (PSC), but the number of available PCIs, PCI_1, PCI_2, PCI_3, PCI_4 is smaller than the number of small cells within the macro cell, and so there is more than one small cell having each available PCI.

In such a heterogeneous network, small cells can be deployed in the ratio of hundreds if not thousands to one ML cells. In other networks in accordance with the invention, there are different small cell layers, and there might for example be hundreds of small cells in each ML cell, and hundreds of femtocells within some of the small cells. In such a network, it is necessary to support mobility within and between the multiple small cell layers, and in particular to allow seamless service continuity between small cell layers and the ML, including allowing for hand-in from one layer to a layer of smaller cells while requiring limited ML configuration and no special features in the ML.

In such a network, there will be a very high frequency and primary scrambling code (PSC) reuse. In addition, the UE population will include legacy (that is, pre-R9) UEs, and so the mobility cannot depend on the use of procedures specified only in R9.

The present invention relates to various methods performed in a telecommunications network, and in particular to methods performed in small cell basestations. Typically, such a basestation includes a processor and memory for storing software for causing these methods to be performed, and aspects of the invention relate to the parts of the overall method performed in the basestations themselves, and to the computer readable media containing the code for causing the basestations to perform the methods.

Figure 3 is a flow chart, illustrating a method in accordance with an aspect of the invention. The method shown in Figure 3 will refer to a specific situation, in which the originating cell is a macrolayer (ML) cell. However, as mentioned above, there might be multiple small cell layers, in which case the method described herein is equally applicable to the hand-in from a layer of relatively large small cells to a layer of relatively small small cells, such as femtocells, and the description should be understood as equally relevant to that situation, except where specifically mentioned otherwise.

The method starts at step 100, in which generic neighbour IDs are created in the ML. As described in more detail below, the generic neighbour IDs point to the small cells layer, and identify univocally the originating cell and the radio characteristics of the target cells. The small cells can then self-assign to a generic neighbour ID and the ML configuration can be automated.

Step 110, which is optional, enables the originating cell to supplement the mandatory generic neighbour ID information with additional information, such as location, timing offset, UE measurements (radio signature), neighbour IDs measured by the small cell itself, handover priority, etc.

In step 120, information is obtained from the small cell layer to improve the precision of the target filtering.

In step 130, when a hand-in request is received in the small cell layer, filtering techniques are applied as described in more detail below to reduce the set of target small cells, possibly to one single target, based on target cell radio characteristics, open/closed mode access, any supplemental information provided and the handover history.

In step 140, optionally, a priority ranking within the set of target cells is applied.

In step 150, the femtocell gateway (FGW) prepares the set of target small cells for receiving the hand-in.

In step 160, when the hand-in has been completed, there is a process of continuous adaptation and learning, storing the hand-in probability and the hand-in signature (described in more detail below) for better targeting in future.

Figure 4 is a flow chart, illustrating in more detail the processes carried out in step 100 of the method of Figure 3. As described above, the aim here is to create generic neighbour IDs in the ML that point to the small cells layer, identifying univocally the originating cell and the radio characteristics of the target cells. The small cells can self-assign to a generic neighbour ID and the ML configuration can be automated.

Thus, in step 200, the network sets a rule for generating generic neighbour IDs. The generic neighbour ID will include a part of the identity of the ML cell, and so the rule will define this. For example, the rule might have the effect of using one generic neighbour ID for all small cells sharing the same radio configuration (i.e. UARFCN and PSC) under the same ML cell. For smaller scale deployments, or for further minimising the ML reconfiguration efforts, a generic neighbour ID might be used across a group of ML cells instead of just one ML cell. Thus, the rule might be to allow the creation of the same generic neighbour ID across a Location Area Code (LAC), Routing Area Code (RAC) or Service Area Code (SAC) area.

In step 202, the small cell performs its radio configuration, and chooses its UMTS Terrestrial Radio Access (UTRA) Absolute Radio Frequency Channel Number (UARFCN) and its Primary Scrambling Code (PSC).

In step 204, the small cell identifies its neighbours, in particular its neighbour ML cells(s) (and its higher layer small cell neighbours in the case where there are multiple layers of small cells). Multiple ML cells overlap in most areas, and so the small cell will often identify more than one neighbour ML cell. This identification can take place by direct radio measurements including detecting the cell ID transmitted by the ML, or indirectly via reports from neighbouring small cells that have detected ML cells, or via interactions from a management system or FGW after the small cell has reported its current location (which might be self-detected, e.g. via GPS, or via radio measurements or based on the known locations of neighbour small cells). The small cell also identifies its small cell neighbours to allow mobility within the small cell layer, but this is less relevant for the issue of hand-in.

In step 206, the small cell assigns itself to the relevant generic neighbour ID(s). Each generic neighbour ID is created according to the rule created in step 200, by combining the relevant parts of the cell ID of the ML cell identified in step 204 (which might be the Cell ID, or might be the LAC, etc) with the UARFCN/PSC chosen by the small cell in step 202.

In step 208, the small cell reports the full list of generic neighbour IDs created in step 206, together with additional information about the ML cell-ID, the small cell UARFCN/PSC, the small cell location, etc. The report is sent to a server in the network, for example the HMS or the FGW or a separate management support server.

In step 210, the server in the network consolidates the generic neighbour IDs that it has received from the small cells and, in step 212, it automatically creates a periodic report of the new configuration needed in the ML RNCs (or in the relevant small cell layers), identifying the ML cell-ID associated with each generic neighbour ID. It will be apparent that, when the number of small cells deployed reaches a certain level, a report from a newly configured small cell report may not trigger any reconfiguration in the macrolayer, because there may already exist a small cell that has identified the same ML cell as a neighbour and that is using the same UARFCN/PSC combination. The server that collects the generic neighbour IDs also creates automatically a periodic configuration report for the server that applies the filtering in step 130.

The server may keep the full table of generic neighbour IDs for correlation, or may delegate it to the ML RNC management systems and/or to the small cell management system. In some embodiments, the server that collects the generic neighbour IDs can be integrated with or into the management systems of the ML RNCs and/or of the server that applies the filtering in step 130, so that the reports it creates are automatically configured into such systems.

Based on the reports received, the macrolayer radio network controller (ML RNC), or equivalent node, is able to create a generic neighbour ID table in step 214. An example of such a table is given below, in a case in which there are two frequencies (F_2 and F_3) assigned to small cells, with eight scrambling codes (SC_1, SC_2, …, SC_8) assigned to F_2 and five scrambling codes (SC_1, …, SC_5) assigned to F_3, and two femtocell gateways (FGW_1 and FGW_2).

Thus, this table indicates, when it becomes necessary to perform a hand-in from a macrolayer cell to a small cell, the target cell information that would be required in any potential Handover Request (HO_REQUEST) message, including the identity of the respective FGW RNC and the generic neighbour ID, which is constructed, e.g., as <ML cell ID (without MNC/MCC)> <2 bits to identify the frequency> <4 bits to identify the SC>

Inputs			Outputs	
ML cell ID	Target small cells radio characteristics		Target cell information (for the HO_REQUEST message)	
	UARFCN	PSC	FGW RNC ID	Generic neighbour ID
ML_cell_1	F_2	SC_1	FGW_1	<ML_cell_1> <F_2> <SC_1>
		SC_2	FGW_1	<ML_cell_1> <F_2> <SC_2>
		——	——	——
		——	——	——
		——	——	——
		——	——	——
		——	——	——
		SC_8	FGW_1	<ML_cell_1> <F_2> <SC_8>
	F_3	SC_1	FGW_2	<ML_cell_1> <F_3> <SC_1>
		——	——	——
		——	——	——
		——	——	——
		SC_5	FGW_2	<ML_cell_1> <F_3> <SC_5>
ML_cell_2	F_2	SC_1	FGW_1	<ML_cell_2> <F_2> <SC_1>
		SC_2	FGW_1	<ML_cell_2> <F_2> <SC_2>
		——	——	——
		——	——	——
		——	——	——
		——	——	——
		——	——	——
		SC_8	FGW_1	<ML_cell_2> <F_2> <SC_8>
	F_3	SC_1	FGW_2	<ML_cell_2> <F_3> <SC_1>
		——	——	——

		——	——	——
		——	——	——
		SC_5	FGW_2	<ML_cell_2> <F_3> <SC_5>
ML_cell_3	——	——	——	——
——				
——				
——				
ML_cell_120	F_3	SC_5	FGW_2	<ML_cell_120> <F_3> <SC_5>

As discussed above, when it is determined that a handover is required, and the required handover is to a small cell, the originating cell is not able to identify the target cell unambiguously, but is only able to identify the target cell by means of the generic neighbour ID. There may be many small cells that share the same generic neighbour ID. The optional step 110 in the method of Figure 3 therefore involves providing supplementary information, that can be applied to filter the possible targets.

Thus, when sending the handover request, the originating cell can supplement the mandatory generic neighbour ID information with as much additional information as possible. This additional information might not be sufficient to allow the intended target cell to be identified unequivocally, but it might be sufficient to identify a group of cells that contains the intended target cell, or it might be sufficient to determine whether each of the cells sharing the generic neighbour ID is more or less likely to be the intended target cell. This can be described as partially identifying the target cell.

For example, the supplementary information might relate to location. As another example, the supplementary information might relate to a timing offset, for example between the timing of the transmissions by the originating cell and by the target cell, as detected by the UE. As a further example, the supplementary information provided by the originating small cell might contain details of radio measurements made by the UE, such as the cell IDs of other macrolayer or small cells whose transmissions can be detected by the UE (the radio signature), as this information will help to identify the location of the UE.

As a further example, the supplementary information can include information (such as cell IDs) about neighbour cells detected by the small cell itself. Thus, where the small cell can see and detect the pilot signal of neighbour cells and/or can capture cell-IDs measured by R9+ UEs, it can correlate them to the radio UARFCN/PSC measurement. Target cell-IDs are reported as a pre-restricted choice to the target small cell controller. Of course, this applies only to the case where the small cell detects neighbour ambiguity, i.e. that multiple small cells surrounding it share the same UARFCN/PSC.

As a still further example, the originating cell can detect how urgent the handover is, based on the set of measurements that it has received from the UE that is to be handed over. For example, the originating cell can determine how rapidly the path loss is changing and the current quality of the connection, and can set a priority level to be used in the controller of the target small cell, for example by applying serial retries only to lower priority handovers.

Some of the supplementary information can be derived in the originating cell itself, while some can be generated by means of a suitable client running on the UE to be handed over.

As mentioned above, in step 120 of the process in Figure 3, information is obtained from the small cell layer to improve the precision of the target filtering.

For example, information from small cell layer can be used to improve the precision of the target filtering. One possibility is to use small cell self-reported radio information, such as measured neighbour cell IDs, UARFCN/PSC, power levels and path losses of the other ML cells/small cells seen. Another possibility is to report a measured timing offset with ML cells, for very precise timing offset matching. A further possibility is to report location.

As mentioned above, in step 130 of the process in Figure 3, a hand-in request is received in the small cell layer, and the hand-in request can identify the target cell only by means of the generic neighbour ID, which might be shared by a large number of small cells. The FGW or small cells controller then applies filtering techniques to reduce the set of target small cells, potentially to one single target in some cases.

When the originating layer is a “larger” small cells layer, this layer can provide additional information on top of the target UARFCN/PSC but also apply techniques to further reduce the number of targets in the “smaller” small cell layer, i.e. by identifying the most likely target cell-IDs out of direct radio measurements of the pilot of such targets. In any case, these techniques cannot, and in typical deployments will not, guarantee the elimination of the target ambiguity but rather work on reducing such ambiguity.

The filtering can be based on the target cell radio characteristics, whether the potential target cells sharing the generic neighbour ID are open mode or closed mode cells, and based on any supplemental information provided and the handover history.

Summary:

1. Use the generic neighbour ID specified by the originating ML cell/group to identify a set of small cells (part of the manual or self-assign process).
2. For closed access mode, apply the UE authorisation rules to reduce the target set before applying the other techniques. In case of a mix of closed mode and open mode potential targets, apply the authorisation filtering only to the closed mode ones only and leave all of the open mode ones in the target group. Authorisation checks may be applied to the type of service currently in use by the UE (e.g. blocking a multi-RAB HO if not supported in the small cells layer).
3. Optionally, use the UE location where known to further restrict the small cell target set.
4. Optionally, use the UE radio signature to further restrict the target set. This is similar to location, but can simply use the raw UE radio measurements as reported to the macro RNC.
5. Optionally, apply heuristics based on history to subdivide each small cell sets into two groups: a group where frequent hand-ins are attempted, and another where there are none or rare. Potentially we could also self-identify the true *mobility hotspots* (which will need multiple landing channels).

Figure 5 is an illustration of the target filtering process that takes place in steps 130 and 150 of the process shown in Figure 3.

Thus, in step 230 a first processing stage is carried out, which consists of trying to identify the correct uncoordinated target 3G cell from the contents of the RANAP: RELOCATION REQUEST/ RANAP: ENHANCED RELOCATION INFORMATION REQUEST message received by a RNC/HNB GW based on the contents of those messages that were generated by the Source RAN.

It is then determined in step 232 whether the first processing stage has reduced the original set of target cells, all identified by the same generic neighbour ID, to a single target small cell. If so, the process passes to step 234, and the handover preparation can take place in the same way as in the case of a coordinated deployment. If it is found in step 232 that the first stage filtering has failed to reduce the set of targets to a single target, then the process passes to a second stage processing in step 236. As described in more detail, the second stage processing uses a combination of heuristics and parallel handover preparation techniques to ensure that the UE eventually is handed over to the correct target 3G cell

In order to describe the invention it is convenient to explain how a Source RAN to RNC/HNB GW handover is currently implemented in what we have described above as a “traditional” 3G deployment’, i.e., a coordinated deployment where new 3G cells are planned into the existing overlapping RANs and where each 3G cell is allocated a locally unique PCI.

Case 1: Handover from Source RNC to Target RNC or HNB GW

This corresponds to either an inter-RNC handover or a RNC to HNB GW handover. In the case of inter-RNC handover, technically called SRNS Relocation, there are currently two supported methods, namely the Legacy SRNS relocation procedure (see section 7.11.1.2.1 ‘Using SRNS Relocation scheme’ of 3GPP TS 25.931 v9.0.0) and the Enhanced SRNS relocation procedure (see section 7.11.1.2.2 ‘Using Enhanced SRNS Relocation’ of 3GPP TS 25.931 v9.0.0).

In the case of RNC to HNB GW handover, the standards currently only support the Legacy procedure because no Iur interface between RNC and HNB GW has been introduced. However it is likely that this will be changed and, in any case, there is nothing stopping the deployment of an Iur interface between a RNC and a HNB GW, provided that the HNB GW performs its role of appearing to the rest of the PLMN as a regular RNC. This means that the Extended SRNS relocation procedure is also relevant when the target 3G cell is a HNB.

Case 1A: Handover based on the Legacy SRNS relocation procedure

When a UE reports the PCI of a 3G cell, the Source RNC checks its neighbour relations tables and, with knowledge of the identity of the source cell, is able to derive the (28-bit) UTRAN Cell-ID, RNC-ID or Extended RNC-ID, LAI and RAC of the target cell, as shown in Figure 6.

Figure 7 shows the currently defined procedure (section 7.11.1.2.1 'Using SRNS Relocation scheme' of 3GPP TS 25.931 v9.0.0). In order to request the Target RNC/HNB GW to prepare the target 3G cell, the Source RNC sends a RANAP: RELOCATION REQUIRED message to each CN domain for which the UE currently has a Iu-signalling connection. The message always contains a Target ID IE and a Source RNC to Target RNC Transparent Container IE, and this in turn always contains the RANAP Target Cell ID IE and the RRC Information to Target RNC container. In the case of a 3G to 3G handover, this contains the SRNS RELOCATION INFO container whose contents are defined in 3GPP TS 25.331.

The information in the Target ID consists of the target cell's RNC-ID / Extended RNC-ID and LAI if the CN domain is CS, or the target cell's RNC-ID / Extended RNC-ID, LAI and RAC if the CN domain is PS.

From the information contained in the Target ID IE, the CN domain that receives this message can derive the transport layer address of the Iu interface that connects it to the target RNC/HNB GW and as such uses that Iu-interface to send a RANAP: RELOCATION REQUEST message to that target RNC/HNB GW. This message contains the unchanged Source RNC to Target RNC Transparent Container IE with the Target Cell ID IE and SRNS RELOCATION INFO container.

When the target RNC or HNB GW receives this message it uses the globally unique UTRAN Cell-ID value in the Target Cell ID IE to uniquely identify the Node B /HNB that controls the target cell and communicates with it (via the Iu/Iuh interface) to prepare the necessary resources to receive the incoming UE.

If we now consider a situation where the target 3G cell was deployed in an uncoordinated way, it is easy to see that this procedure would break down, because the Target Cell ID IE populated by the Source RNC will not contain the UTRAN Cell-ID of the target 3G cell, since the neighbour relation entry that was used to by the Source RNC is a generic entry to be re-used by multiple un-coordinated 3G cells.

According to embodiments of this invention, the Source RNC can provide two types of information, namely: the current location of the UE relative to the source RAN, in this case the GERAN; and the PCI of the uncoordinated 3G cell that was reported by the UE in the Source RNC to Target RNC Transparent Container IE to help the Target RNC/HNB GW identify the uncoordinated target 3G cell reported by the UE.
The simplest embodiment does not does not require any changes in the operation of the Source RNC or the RANAP protocol. It affects only how the neighbour relation tables in the Source RNC are provisioned. While this is different from traditional coordinated 3G deployments, but requires only the same level of O&M effort.

The Source RNC to Target RNC Transparent Container IE to be generated by the Source RNC contains two key information fields for this invention, namely the Target Cell ID IE (as discussed above), and the UTRAN - Radio Network Temporary Identifier (U-RNTI) IE (present in the SRNS RELOCATION INFO container). The U-RNTI is a quantity allocated by the Source RNC that uniquely identifies the UE in the UTRAN while in RRC connected mode. Crucially, among other information it always contains the Source RNC's RNC-ID or Extended RNC-ID (See 3GPP TS 25.331).

Thus without any changes to current operation the Source RNC is already providing its RNC-ID or Extended RNC-ID in the U-RNTI IE. That means that the 28 bit Target Cell ID IE does not also need to carry that information in order for the Source RNC to Target RNC Transparent Container IE to carry the UTRAN Cell ID of one of the of the cells that is serving the UE in the Source RNC. Therefore, when the UTRAN is operating 12-bit RNC IDs, this means that the 28-bit Target Cell ID IE has 12 bits free after being used to carry the 16 bit C-ID. Similarly, when the UTRAN is operating 16-bit Extended RNC IDs, this means that the 28-bit Target Cell ID IE has 16 bits free after being used to carry the 12 bit C-ID.

Even 12 bits is enough to carry the PCI reported by the UE because this is composed of (PSC, UARFCN) where the PSC has a 9-bit range (0-511) and the number of 3G carrier frequencies used by an operator is unlikely to exceed a 3-bit range, i.e., 8. In fact, the number of PSCs typically reserved for un-coordinated 3G cells is much smaller than the 512-sized logical range, so effectively 12 bits is more than enough.

Thus in this embodiment, the operator will reserve a certain number of PCIs to be used by the un-coordinated 3G deployment. It will then configure special neighbour relations in each Source RNC for these PCIs. Thus, as shown in Figure 8, when the Source RNC decides to perform a handover to one of these cells it generates a Source RNC to Target RNC Transparent Container IE, where the Target Cell ID IE contains the C-ID of one of the source cells plus the UE-reported PCI.

When the Target RNC/HNB GW receives the Source RNC to Target RNC Transparent Container IE, it uses the Target Cell ID IE plus the U-RNTI IE to derive the UTRAN Cell-ID of the source cell, plus the PCI of the target 3G cell as reported by the UE. It will then input this information in a database, which will output one or more candidate target 3G cells. This completes the first stage processing at step 230 in Figure 5.

If a single cell is output, then handover preparation towards that cell takes place and there is no need for further processing to support the handover. Otherwise, second stage processing will be required.

In another embodiment, the Source RNC is configured to encode additional information about source cells, in order to aid the target RNC/HNB GW to determine the target 3G cell but without changes to the RANAP protocol.

At a minimum, the operator will configure the SRNS RELOCATION INFO container in the Source RNC to Target RNC Transparent Container IE to contain the UE's neighbour cell lists just prior to the handover, despite the fact that the presence of this information is optional. A more involved implementation would see the bit space of the neighbour cell list IE being used to carry information like the Geolocation of the source cell(s).

In another embodiment of this invention the Source RNC generates a Source RNC to Target RNC Transparent Container IE carrying additional proprietary IEs to encode additional information about source cells in order to aid the target RNC/HNB GW to determine the target 3G cell, e.g., Geolocation of the source cell(s).

Case 1B: Handover based on the Enhanced SRNS relocation procedure

This case does not require any additional handling relative to Case 1A described above, because the relevant information transfer from the source RNC to the Target RNC/HNB GW is still performed using the same Source RNC to Target RNC

Transparent Container IE. The only difference is that it, in this case, the Source RNC to Target RNC Transparent Container IE is sent directly from the source RNC to the Target RNC/HNB GW via the Iur interface between them, instead of being routed via the CN.

As shown in Figure 9, the handover is prepared by the Source RNC sending a RNSAP: ENHANCED RELOCATION REQUEST message (as defined in 3GPP TS 25.423) directly to the Target RNC/HNB GW via the Iur-interface. That message contains a RANAP: ENHANCED RELOCATION INFORMATION REQUEST (defined in 3GPP TS 25.413), which in turn carries the Source RNC to Target RNC Transparent Container IE with the same components as in the case of the Legacy SRNS Relocation procedure, in particular the Target Cell ID IE and the SRNS RELOCATION (which carries the UE's U-RNTI).

Case 2: Handover from Source BSC to Target RNC or HNB GW

In this case the method follows the same principle as in Case 1 above, i.e. to try and enable the Target RNC/HNB GW to identify a reported un-coordinated 3G cell by having the source system provide two types of information to the Target RNC/HNB GW (via the Source RNC to Target RNC Transparent Container), namely the current location of the UE relative to the source RAN, in this case the GERAN; and the PCI of the uncoordinated 3G cell that was reported by the UE.

While it is not possible in this case to use the detailed mechanism described in Case 1, which can be implemented without proprietary protocol changes, there is an alternative mechanism that provides the same functionality.

Case 2A: Handover of CS domain service from GERAN A/Gb mode to UTRAN

In 3GPP PLMNs, the rule adopted for handover-related communications between network nodes that do not speak the same (protocol) language is that the source/sender of the information adapts to the target/receiver. In the case of handover from GERAN A/Gb mode to UTRAN, this means that, although the Source BSC in the GERAN normally does not use the RANAP protocol, since this is the protocol used by the UTRAN the Source BSC must provide the necessary handover information to the UTRAN via a RANAP message carried to the CN via what is called the Source to Target RNC transparent container.

Figure 10 shows the relevant IEs of the signalling involved in the initial communication between the Source BSC and the Target RNC/HNB GW in this case. As defined in 3GPP TS 48.008 (and shown in Figure 10), in the case of handover to a 3G cell the Source BSC needs to generate a BSSMAP: HANDOVER REQUIRED message carrying what is called the “Source RNC to target RNC transparent information (UMTS)”. This container carries another container which is the already discussed RANAP Source RNC to Target RNC Transparent Container IE, (as defined in 3GPP TS 25.413).

The simplest embodiment does not require any additional capabilities in the Source BSC or proprietary changes to the current protocols. In this embodiment, the operator will reserve a certain number of PCIs to be used by the un-coordinated 3G deployment. The operator will then configure the BSCs to always include the optional Cell Load Information Group IE in the ‘Source RNC to Target RNC Transparent Container IE’ that it generates when it triggers handover to one of the reserved PCIs. As defined in 3GPP TS 25.413, this IE will contain the Source Cell Identifier which provides the unique cell identity of the source GERAN cell in the GERAN.

The operator will also configure special neighbour relations in each Source BSC for these reserved PCIs. As shown in Figure 11, the 28 bits of the Target Cell ID IE in the Source RNC to Target RNC Transparent Container IE are used to encode the PCI (PSC, UARFCN) of the target 3G cell. Thus, when the Target RNC/HNB GW receives the Source RNC to Target RNC Transparent Container IE, it can use the Target Cell ID IE and the Cell Load Information Group IE to learn the PCI and source cell identity reported by the UE.

It will then input this information in a database, which will output one or more candidate target 3G cells. This completes the first stage processing at step 230 in Figure 5.

If a single cell is output, then handover preparation towards that cell takes place and there is no need for further processing to support the handover. Otherwise, second stage processing will be required.

In another embodiment, the Source BSC would generate a Source RNC to Target RNC Transparent Container IE carrying additional proprietary IEs to encode additional information about source cell in order to aid the target RNC/HNB GW to determine the target 3G cell, e.g., Geolocation of the source cell.

Case 2B: Handover of PS domain service from GERAN A/Gb mode

This case is handled similarly to case 2A. Figure 12 shows the relevant IEs of the simplest embodiment, in which the Source RNC to Target RNC Transparent Container IE carries the unique source GERAN cell ID in the Load Information Group IE and the PCI of the target 3G cell encoded in the Target Cell ID IE, which is obtained by the appropriate provisioning of the neighbour relations for the PCIs reserved for the un-coordinated 3G cell deployment.

In another embodiment, the Source BSC generates a Source RNC to Target RNC Transparent Container IE carrying additional proprietary IEs to encode additional information about the source cell, in order to aid the target RNC/HNB GW to determine the target 3G cell, e.g., Geolocation of the source cell.

Case 2C: CS/PS Handover from GERAN Iu mode to UTRAN,

Again, this case is handled similarly to Case 2A. In its simplest embodiment, the Source RNC to Target RNC Transparent Container IE carries the unique source GERAN cell ID in the Load Information Group IE and the PCI of the target 3G cell encoded in the Target Cell ID IE, which is obtained by the appropriate provisioning of the neighbour relations for the PCIs reserved for the un-coordinated 3G cell deployment.

In another embodiment, the Source BSC generates a Source RNC to Target RNC Transparent Container IE carrying additional proprietary IEs to encode additional information about the source cell, in order to aid the target RNC/HNB GW to determine the target 3G cell, e.g., Geolocation of the source cell.

Thus, step 130 of the process shown in Figure 3 determines, by the first stage processing of the information in the Source RNC to Target RNC Transparent Container IE, received in the RANAP: RELOCATION REQUEST or RNSAP: ENHANCED RELOCATION REQUEST message, whether there is just a single target 3G cell candidate.

If there is more than one candidate, then the only remaining option for the Target RNC/HNB GW to support a successful handover is to try to prepare one or more of the multiple potential 3G target cells which resulted from the first stage filtering in parallel, and to give up on pre-identifying the correct one.

Step 140 of the process shown in Figure 3 is then an optional priority ranking step, to identify the most likely targets, e.g. the set of small cells that have >90% probability of being targeted, and divide the list in multiple sets. This is particularly important when the remaining list is long, e.g. >20 targets.

The process shown in Figure 3 then passes to step 150, in which the FGW proceeds to prepare the set of target small cells for receiving the hand-in. Specifically, any cells that are identified as target small cells for a hand-in are instructed to prepare a common landing channel on which the UE can be received. Multiple target cells can be prepared in parallel, and/or hand-in can be attempted to individual or multiple cells in series if a retry is required. The preparation can be based on an assessment as to whether a cell has a high or low probability of being a target, and/or can use a trial and error process.

At this stage, the processing is heuristic in nature as there is not enough information in the Target RNC/HNB GW to identify the correct target 3G cell. The implementation of the processing is based on the fact that, if a UE fails to handover to the 3G cell it was instructed to handover to, because the Target RNC/HNB GW failed to prepare a radio link in that specific 3G cell, that failure typically does not result in the dedicated connection between the UE and the network being dropped. Instead, what will happen is that the UE will try to synchronise with the cell and will fail to achieve synchronisation. As per 3GPP standards (3GPP TS 25.331) when this happens the UE will to fall-back to the radio link(s) it had in the source cell(s). Unless the radio conditions have degraded too much, this will be a successful procedure and the UE will be able to start another handover attempt procedure. The result is that a failed handover attempt does not equate to loss of the UE-NW dedicated connection and, as such, there is room for handover failure in a heuristic procedure that tries to deliver the UE to the correct target 3G cell by trial and error.

Embodiments of the invention described here are implemented by the Target RNC/HNB GW and the underlying Target 3G cells using a parallel handover preparation procedure. This procedure allows the Target RNC/HNB GW to prepare multiple target 3G cells, which share the same PCI, to receive the UE, while generating a single handover command message to be delivered to the UE

In order to describe the procedure, we first consider a specific use case (handover to a HNB whose deployment was not coordinated with the Source RAN) which will illustrate the key general principles of the invention. We then explain how these key general principles can be re-used in a wider range of use cases.

Figure 13 illustrates a handover in a conventional situation, in which a HNB GW receives a RANAP: RELOCATION REQUEST message (message 300 in Figure 13), from which it can uniquely identity the HNB (which is registered in the HNB GW) that should be prepared to receive the UE.

In this case the HNB GW shall act according to 3GPP TS 25.467 and trigger a HNB GW-initiated UE registration procedure simultaneously with the delivery of the RANAP: RELOCATION REQUEST message 302 to the HNB. The HNB will then: prepare the Radio Link (RL) to receive the UE (step 304); start to transmit the RL bits to be used for air interface synchronisation with the incoming UE; generate the handover command message describing the radio link it has created; and place that message in the Target RNC To Source RNC Transparent Container IE and send it to the Target HNB GW via the RANAP: RELOCATION REQUEST ACK message 306.

The HNB GW then sends a RANAP: RELOCATION REQUEST ACK message 308 back in the direction of the Source RAN. When the Source RAN eventually receives a message with the Target RNC To Source RNC Transparent Container IE, it will retrieve the handover command message and send it to the UE. The UE will then attempt to synchronise with the radio link described in the handover command message (step 310).

When the HNB detects that air interface sync was achieved, it sends a RANAP: RELOCATION DETECT message 312 to the Target HNB GW, and when the UE sends the handover complete message (step 314) the HNB sends a RANAP: RELOCATION COMPLETE message 316 to the Target HNB GW, which successfully completes the handover procedures in the Target HNB GW and HNB.

Figure 14 shows the situation with which this invention is concerned, which is that, when the Target HNB GW receives a RANAP: RELOCATION REQUEST/ RNSAP: ENHANCED RELOCATION REQUEST message 330, it can use its contents to derive the Physical Cell Identity (PCI) being used by the target 3G cell but not its logical cell identity, and that the Target HNB GW is aware of more than one HNB using that Physical Cell Identity (PCI). Figure 14 therefore shows how to prepare in parallel multiple HNBs (which share the same PCI) the correct radio link configuration to receive the UE, while at the same time providing a handover command message to the Source RAN that describes a single Radio Link Configuration.

The solution is to prepare the same link configuration to receive the UE in all those HNBs, and this is logically possible if the HNBs share the same PCI. This is because a Radio Link configuration is characterised by several properties, which include its Primary Scrambling Code (PSC) and carrier frequency (UARFCN), its uplink scrambling code, the downlink Orthogonal variable spreading factor (OVSF) code(s) that it uses, etc.

The PSC and UARFCN correspond to the PCI of a 3G cell, which means that it is only possible to prepare the same RL configuration in cells that use the same PCI. However, other characteristics such as the UL scrambling code number and the DL OVSF code number can take values in the same logical ranges in all 3G cells, and so they can be kept reserved in each target 3G cell. In effect, each HNB needs only to reserve a certain number of pre-defined radio link configurations to be used as common landing channels for the incoming UEs during handover.

This means that it is possible to pre-define a set of common Radio Link configurations in HNBs operating the same PCI, and as such the HNB GW can instruct those HNBs in parallel to setup one of those common configurations and a single handover command message describing that radio link configuration can be sent by the Target HNB GW to the source RAN.

Figure 14 shows an example in which there are two potential target HNBs (Target HNB 1 and Target HNB 2). The Target HNB GW, having received the relocation request message 330 has applied filtering and detected multiple possible target small cells. It then instructs both HNBs to initialise the same reserved radio link configuration via the RANAP: RELOCATION REQUEST messages 332, 334 that it sends to each of them.

It then conveys to the source RAN a single handover command message (describing the chosen common reserved radio link configuration) in the Target RNC To Source RNC Transparent Container IE of the RANAP: RELOCATION REQUEST ACK message 336.

The source RAN sends the handover command message to the UE which will then synchronise (step 338) to the target HNB it actually measured (HNB 1) and the handover procedure for that UE will be completed as usual (i.e. as if HNB 1 had been the single target HNB). That is, when HNB 1 detects that air interface sync was achieved, it sends a RANAP: RELOCATION DETECT message 340 to the Target HNB GW, which sends it on (message 342) to the Source RAN. When the UE sends the handover complete message (step 344) HNB 1 sends a RANAP: RELOCATION COMPLETE message 346 to the Target HNB GW, which sends it on (message 348) to the Source RAN, and thus successfully completes the handover procedures in the Target HNB GW and HNB 1.

In order to release the common RL configuration in HNB 2 as soon as possible, the Target HNB GW uses the reception of the RANAP: RELOCATION DETECT message 340 from HNB 1 (sent when HNB 1 first detects the incoming UE) to trigger a resource release message 352 (RUA:DISCONNECT) towards HNB 2. Thus, if HNB 2 is again part of a set of handover target candidates, this reserved common RL configuration can be used again.

While the description of the embodiments used the HNB-HNB GW framework defined in 3GPP TS 25.467, 3GPP TS 25.468 and 3GPP TS 25.469, it is easy to see that it can be generalised to any system 1) that is composed of a Controller which appears to the PLMN as a normal RNC and is in communication with devices that operate 3G cells that appear to a UE as normal 3G cells and 2) where the Controller is aware of the Physical Cell Identity (PCI) being used by each of those 3G cells.

The effect of this embodiment is that, because a failure of a handover attempt (because the UE tries to sync to a 3G cell which was not prepared to receive it) does not typically lead to a ‘call drop’, and because it is possible to implement a parallel handover preparation procedure towards multiple 3G cells who share the same PCI, this allows the design of multiple algorithms to deliver the ability to handover a UE from a Source RAN to a 3G cell whose deployment was not coordinated with the source RAN, as long as the PCI of the target cell is known in the Target RNC/HNB GW.

One extreme possibility is to rely totally on the parallel handover preparation procedure, that is, to prepare all potential target 3G cells. This avoids any risk of a handover failure due to the fact that the intended target 3G cell was not prepared. This will be feasible in many cases, but is difficult to implement efficiently in cases where there is heavy traffic from the source RAN, and there many potential 3G cells sharing the same PCI. For example, this would be the case where a 2G cell has a large coverage area over a very dense HNB deployment. In this case, each PCI reserved for HNBs could be re-used by tens of HNBs and, if there was heavy handover traffic, this could lead to a huge number of parallel handover attempts and the need for the HNBs to reserve a significant number of radio link resources just to support these attempt procedures, which would greatly reduce their capacity for normal traffic.

In such cases, the HNB GW could adapt its strategy, for example such that, if there are, say, 30 target HNBs with the same PCI, it will first try the 10 most likely ones to be the handover targets (which it has learned from tracking handover success rates per HNB). If, unfortunately, the UE was reporting one of the other 20 cells the handover attempt will fail but the call will typically survive this failure, and so the HNB GW can then trigger a second run of targets, e.g. the next 10 most likely, etc.

At the same time, the number of reserved radio link resources for the support of the parallel handover procedure, that is the landing channels described above, could be made HNB specific. That is, if the system detects that some HNBs are very common handover targets while others are not, then a larger number of reserved radio link resources for the support of the parallel handover procedure would be configured in those high likelihood candidates, which would act as entry points into the HNB deployment. This would reserve a large capacity for support of the parallel handover preparation procedure, with the intention of then quickly triggering inter-HNB handovers to other HNBs with low inbound handover load, which would reserve only a much smaller number of resources for the support of the parallel handover preparation procedure.

Once the hand-in has been completed, the process shown in Figure 3 passes to step 160, in which it is determined whether there are steps that can be taken to learn from the process, and adapt it for further use, in order to improve the results of future hand-in attempts.

Figure 15 is a flow chart, illustrating a first method for taking account of hand-in success probability.

In step 380, a hand-in is completed, using a specific generic neighbour ID. The details of that hand-in are then stored. In particular, in this embodiment, the system stores the cell-ID of the target cell to which the hand-in was completed. In step 382, it is determined whether the number of hand-ins using that generic neighbour ID has reached a threshold. This threshold is set such that it represents the point at which the history can be used to make statistically reliable predictions about future hand-in attempts. So, for example, the threshold value might be set to 50 handovers, or might be set to a value equal to a fixed number multiplied by the number of small cells sharing that generic neighbour ID.

In step 384, statistical logic is applied, in order to determine how evenly distributed the hand-ins are, amongst the small cells that share the generic neighbour ID. In step 386, it is determined based on this logic whether the intended targets of the hand-ins are sufficiently concentrated to make it worthwhile defining a high probability set of the cells. For example, it might be decided that the targets are not sufficiently concentrated if, in a set of 10 small cells that share a generic neighbour ID, none of them is the target for more than 30% of the handovers, and/or none of them is the target for less than 3% of the handovers, and/or more than half of the cells would need to be included to reach a 90% probability of including the intended target.

If it is determined in step 386 that the targets are not sufficiently concentrated, the process passes to step 388, and no high probability set is defined. In that case, hand-in requests relating to that generic neighbour ID are handled by preparing a handover to all of the cells sharing the generic neighbour ID.

If it is determined in step 386 that the targets are sufficiently concentrated to justify the definition of a high probability set, the process passes to step 390, in which the high probability set is defined. For example, the high probability can contain the cells (or the single cell) that receives more than a configurable threshold percentage of the handovers. For example, where there are 10 cells that share a generic neighbour ID, the high probability set might contain the 2 cells that together receive more than 90% of the handovers. In that case, all of the remaining cells are placed in a low probability set. As the small cells will typically be deployed in a contiguous manner, it is quite likely that the units at the border of their coverage area will receive a high proportion of the handovers. As another example, three sets might be designed, with a high probability set containing the most common targets, a medium probability set containing the less common targets, and a low probability set containing the cells that are only very rarely targets.

The definition of the high probability and lower probability sets should be reassessed periodically. For example, the process shown in Figure 15 could be run every time that the number of hand-ins using that generic neighbour ID, since the previous assessment, has reached the threshold, e.g. 50. Alternatively, the process could build upon previous assessments, with criteria set for reassigning cells between the sets.

If hand-in requests include additional information, this can be captured, and the results of those hand-in requests can then be used to determine how to handle subsequent hand-in attempts.

Figure 16 is a flow chart illustrating a process for using this additional information. In step 420, a hand-in is completed, using a specific generic neighbour ID. The details of that hand-in are then stored. In particular, in this embodiment, the system stores the cell-ID of the target cell to which the hand-in was completed, but also stores the available additional information, such as reported UE measurements, information about the UE location, information about the timing offset between the source cell and the intended target cell, etc.

In step 422, ageing characteristics are defined for any time-dependent or variable information. For example, timing offsets between cells will typically vary over time, and so the usefulness of such information decreases over time. Step 422 deals with this by steadily reducing the weight that is given to such measurement information, until it reaches an expiry time and is completely disregarded.

In step 424, merging heuristics are applied, to determine how to capture information from subsequent successful hand-in procedures, with the rules varying by the information type. For example, in the case of timing offset information, a timing offset value previously reported by a UE can be replaced entirely by a value reported by a UE in a subsequent hand-in. In the case of UE measurements reporting extra neighbours, all neighbours over a certain power level can be added to the list associated with a particular target, and a count value associated with a previously reported neighbour that is not reported in a new hand-in can be decremented, until the count reaches zero and that other cell is eliminated from the signature associated with that target.

When the information relates to the power level of the neighbours, as reported by a UE that is requesting a hand-in to a particular target, a weighted average is formed, based on the value already stored in the signature and the value provided in the new measurement. For example, the new value might have 10 or 20% of the weight of the stored value.

When the information relates to the location of the UE requesting the hand-in, the reported location is combined with the value already stored in the signature associated with that target. For example, if the signature contains a location value which is the result of averaging the locations of the UEs that made the last 50 successful attempts to that target cell, then the new location value weighting is 2% of the stored location (so that the new stored value is the average of 51 successful attempts).

Thus, based on these rules, a signature is stored for each of the cells that share a specific generic neighbour ID. This signature contains information derived from the information provided by UEs when those UEs were attempting to hand in to the respective cell.

In step 426, heuristics are applied to be used in calculating the likelihood that a signature, provided as part of a hand-in request, matches one of the stored signatures. Thus step 426 defines the fuzzy logic (i.e. the central value and the precision) of the matching of each piece of information, and the weights for each match.

In the case of location information, the previously stored information defines an average central location, and a location range is set around that central location, so that it can be determined whether the location reported in the new hand-in request falls within that range. The range can be set based on the precision of the information provided by the UEs, and also based on the radius of the small cells. Thus, when the UE location is known as coordinates with a 500m radius imprecision, and the target small cells are deployed with a maximum radius of 200m, the location range is set to 500m + 200m = 700m. It is then determined that a new hand-in request might relate to a cell, if the location reported by the UE in the hand-in request is within 700m of the stored central location associated with the cell.

In the case of additional information relating to neighbour cells detected by the UE making the hand-in request, the rules set in step 426 can relate to the proportion of detected neighbours that must match those stored in the signature associated with a target cell, and the required similarity in the power measurements relating to those neighbours. For example, the rule might be set so that it is determined that a new hand-in request might relate to a cell, if, say, at least 3 of the neighbours reported by the UE are on the list of 5 neighbours stored in the signature associated with that cell, and if the two strongest neighbours have reported signal strengths within +-10dB of the recorded power levels in the signature.

This signature matching is done without any additional knowledge about the target femtocells on top of their identity, UARFCN and PSC. All information is captured in the FGW on the fly while the hand-in procedure happens.

As mentioned above, however, the filtering of the potential targets can be improved by using additional information.

The process as described above is generally the same, whether the handover is originating from a macrolayer cell or from a small cell. However, when the originating cell is a small cell, some or all of that process can be performed in the small cell. That is, the originating cell can apply the logic of the points above, in order to filter the potential targets, prioritise and generate one or multiple handover requests. This requires the small cell to collect information from its neighbour small cells, in order to aid selection process, and the continuous improvement and adaptation based on the hand-in signature, etc.

In the case of a small cells layer that supports anchor and drift handovers within the small cell layer itself, the handover requests are prepared over the peer-to-peer interface between the small cells. This will typically be useful in the case of small cell to small cell handovers (either in the same layer or in different layers), where the FGW does not support ambiguous hand-in.

When it is determined that such a handover is required, the originating small cell sends out multiple HO requests to various target small cells (or even ML cells), and each of these handover requests is treated independently by the chain of nodes (which may include the macrolayer core network and/or a FGW) to reach the respective target cell. Thus, this process is totally transparent to the ML CN elements and/or the FGW apart from the additional signalling loading.

As in the case where the handover requests are sent from the femtocell gateway, only one of the requests will succeed. In order to avoid the other requests being counted by the FGW (or the ML CN) as failed handovers, the originating small cell can tag the handover requests with a “multiple hand-in request” tag, so that such parallel attempts are not counted as failed handover attempts.

CLAIMS

1. A method of operation of a cellular telecommunications network, wherein the network comprises a plurality of large cells, and wherein at least one of the large cells coverage area includes a plurality of small cells, each of said small cells operating with a respective Physical Cell Identity, the method comprising:

on determining that a user equipment handover is required from a serving large cell to one of the small cells included in said serving large cell, sending a handover request from the serving cell to a node of the network, the handover request identifying the serving large cell and the Physical Cell Identity of said one of the small cells, wherein said Physical Cell Identity might not unambiguously identify said one of the small cells; and

including with said handover request additional information to allow the network node to at least reduce ambiguity in an identity of said one of the small cells from amongst a plurality of small cells included in said serving large cell operating with the same Physical Cell Identity.

2 A method as claimed in claim 1, wherein the additional information comprises information relating to a location of the user equipment within the large cell.

3. A method as claimed in claim 1, wherein the additional information comprises information relating to a timing offset between transmissions of the large cell and of said one of the small cells.

4. A method as claimed in claim 3, further comprising:

when a handover is successful, storing the information relating to the timing offset, and using said information to identify a possible handover target when receiving further handover requests.

5. A method as claimed in any preceding claim, wherein the additional information comprises information relating to neighbour cells that can be detected by the user equipment.

6. A method as claimed in one of claims 1-5, further comprising:

in the network node, comparing the additional information with corresponding information relating to the small cells included in said serving large cell operating with the Physical Cell Identity identified by the handover request.

7. A method as claimed in claim 6, wherein said corresponding information is derived from previous handovers.

8. A method as claimed in claim 7, comprising setting time parameters as to when and how said corresponding information derived from previous handovers should still be used.

9. A method as claimed in claim 6, comprising setting a degree of similarity to be met by said comparison.

10. A method as claimed in one of claims 6 to 9, comprising performing the handover request based on a result of said comparison.

11. A method as claimed in claim 1, further comprising:

including with said handover request information relating to a priority to be given to the handover.

12. A method as claimed in any preceding claim, comprising requesting a plurality of said small cells to prepare a common landing channel for the user equipment.

13. A method as claimed in claim 12, wherein the landing channel has a fixed configuration.

14. A method as claimed in claim 12, wherein the landing channel is identified on demand.

15. A method as claimed in claim 12, wherein the common landing channel is selected from multiple landing channels.

16. A method as claimed in any preceding claim, wherein each small cell can assign itself to a neighbour cell list of a large cell in which it is located.

17. A method of operation of a cellular telecommunications network, wherein the network comprises a plurality of large cells, and wherein at least one of the large cells includes a plurality of small cells, each of said small cells operating with a respective Physical Cell Identity, the method comprising:

on determining that a user equipment handover is required from a serving large cell to one of the small cells included in said serving large cell, sending a handover request from the serving cell to a node of the network, the handover request identifying the serving large cell and the Physical Cell Identity of said one of the small cells, and

in the network node:

dividing said small cells included in said serving large cell and operating with said Physical Cell Identity into at least first and second groups, wherein small cells in the first group have been targets for handovers from the serving large cell more often than small cells in the second group; and

sending a handover request from said network node to at least one small cell in the first group.

18. A method as claimed in claim 17, comprising sending parallel handover requests from said network node to each small cell in the first group.

19. A method as claimed in claim 17 or 18, further comprising:

sending a handover request from said network node to at least one small cell in the second group, only if the or each handover request from said network node to said at least one small cell in the first group is unsuccessful.

20. A method as claimed in any of claims 17 to 19, comprising requesting a plurality of said small cells in the first group to prepare a common landing channel for the user equipment.

21. A method as claimed in claim 20, wherein the landing channel has a fixed configuration.

22. A method as claimed in claim 20, wherein the landing channel is identified on demand.

23. A method as claimed in claim 20, wherein the common landing channel is selected from multiple landing channels.

24. A method of operation of a cellular telecommunications network, wherein the network comprises a plurality of large cells, and wherein at least one of the large cells includes a plurality of small cells, each of said small cells operating with a respective Physical Cell Identity, the method comprising:

on determining that a user equipment handover is required from a serving large cell to one of the small cells included in said serving large cell,

obtaining additional information and using said additional information to partially identify said one of the small cells from amongst a plurality of small cells included in said serving large cell operating with the same Physical Cell Identity, and

sending a handover request from the serving cell to a node of the network, the handover request identifying the serving large cell and the Physical Cell Identity of said one of the small cells, and the partial identification of said one of the small cells.

25. A method as claimed in claim 24, comprising requesting a plurality of said small cells in the first group to prepare a common landing channel for the user equipment.

26. A method as claimed in claim 25, wherein the landing channel has a fixed configuration.

27. A method as claimed in claim 25, wherein the landing channel is identified on demand.

28. A method as claimed in claim 25, wherein the common landing channel is selected from multiple landing channels.

29. A method of operation of a cellular telecommunications network, wherein the network comprises a plurality of large cells, and wherein at least one of the large cells includes a plurality of small cells, each of said small cells operating with a respective Physical Cell Identity, the method comprising:

dividing said small cells included in said serving large cell and operating with said Physical Cell Identity into at least first and second groups, wherein small cells in the first group have been targets for handovers from the serving large cell more often than small cells in the second group; and

on determining that a user equipment handover is required from a serving large cell to one of the small cells included in said serving large cell, sending a handover request from the serving cell to at least one small cell in the first group, the handover request identifying the serving large cell and the Physical Cell Identity of said one of the small cells.

30. A method as claimed in claim 29, comprising requesting a plurality of said small cells in the first group to prepare a common landing channel for the user equipment.

31. A method as claimed in claim 30, wherein the landing channel has a fixed configuration.

32. A method as claimed in claim 30, wherein the landing channel is identified on demand.

33. A method as claimed in claim 30, wherein the common landing channel is selected from multiple landing channels.

34. A method as claimed in any of claims 29 to 33, comprising sending a handover request from the serving cell to each small cell in the first group in parallel.

35. A method as claimed in any of claims 29 to 33, comprising sending a handover request from the serving cell to each small cell in the first group in series, until one of said handover requests is successful.

36. A method as claimed in claim 34 or 35, wherein the femto gateway generates multiple handover requests.

37. A method as claimed in claim 34 or 35, wherein the serving cell generates multiple handover requests.

38. A method of operation of a cellular telecommunications network, wherein the network comprises a plurality of large cells, and wherein at least one of the large cells includes a plurality of small cells, each of said small cells operating with a respective Physical Cell Identity, the method comprising:

dividing said small cells included in said serving large cell and operating with said Physical Cell Identity into at least first and second groups, wherein small cells in the first group have been targets for handovers more often than small cells in the second group; and

on determining that a user equipment handover is required from a serving small cell to another of the small cells included in said large cell, sending a handover request from the serving cell to at least one small cell in the first group, the handover request identifying the serving cell and the Physical Cell Identity of said one of the small cells.

39. A method as claimed in claim 38, comprising sending a handover request from the serving cell to each small cell in the first group in parallel.

40. A method as claimed in claim 38, comprising sending a handover request from the serving cell to each small cell in the first group in series, until one of said handover requests is successful.

41. A method as claimed in one of claims 38 to 40, comprising, in the serving small cell, identifying said one of the small cells by information obtained by the serving small cell.

42. A method as claimed in claim 41, comprising obtaining said information in the serving small cell from measurements made by the serving small cell.

43. A method as claimed in claim 41, comprising obtaining said information in the serving small cell from measurements made by UE devices on the serving small cell.

Application No: GB1205269.2 **Examiner:** Dan Hickery
Claims searched: 1 to 16, 24 to 28 **Date of search:** 27 July 2012

Patents Act 1977: Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X	1, 2, 6, 24	EP 2352332 A1 (PANASONIC) par.0034-0035, 0039-0041, fig.4, 6
X	1, 2, 6-10, 12-15, 24-28	EP 2343920 A1 (ALCATEL LUCENT) par.0014-0021, fig.1
X	1, 2, 5-10, 24	EP 2326122 A1 (MITSUBISHI) par.0086-0132, fig.4
X	1, 2, 24	WO 2011/133921 A1 (QUALCOMM) par.0125-0131, fig.10
X	1, 3, 4, 6, 24	WO 2011/020488 A1 (NOKIA SIEMENS) p.2 line 19-p.3 line 27
X	1, 2, 6, 24	WO 2010/149829 A1 (NOKIA) par.0062, fig.6
X	1, 2, 5-10, 24	WO 2010/033729 A2 (QUALCOMM) par.0034-0041, 0048-0050, 0070-0072, 0095-0101, fig.2, 3A, 4A, 6

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^X :

Worldwide search of patent documents classified in the following areas of the IPC

H04W

The following online and other databases have been used in the preparation of this search report

Online: EPODOC, WPI

International Classification:

Subclass	Subgroup	Valid From
H04W	0036/04	01/01/2009