

(21) Application No: **0911274.9**
 (22) Date of Filing: **30.06.2009**

(51) INT CL:
G06N 99/00 (2010.01) **H04B 10/30** (2006.01)
H04L 9/08 (2006.01)

(71) Applicant(s):
Hewlett-Packard Development Company, L.P.
11445 Compaq Center Drive West, Houston 77070,
Texas, United States of America

(56) Documents Cited:
EP 1865657 A1 **WO 2009/155486 A1**
WO 2007/021945 A2
CHINESE PHYSICS, VOL. 13, NO. 2, 2004, Wu
Xiang-Sheng, Memory-based quantum repeater in
quantum information communication, Pages 173 - 177

National Institute of Informatics
6-7 Minami Azabu 4-chome, Minato-ku,
Tokyo 106-0047, Japan


(58) Field of Search:
 INT CL **G06N, H04B, H04L**
 Other: **ON-LINE: INSPEC, EPODOC, WPI, XPI3E**

(72) Inventor(s):
Keith Harrison
William Munro
Kae Nemoto

(continued on next page)

(54) Title of the Invention: **Quantum repeater and system and method for creating extended entanglements**
 Abstract Title: **Quantum repeater and system and method for creating extended entanglements utilising cyclic synchronised control signals at repeater relay nodes**

(57) A quantum repeater and a method for creating an end-to-end entanglement between qubits in first and second end nodes of a chain of optically-coupled repeater nodes. Local entanglements are created on an on-going basis between qubits in neighbouring pairs in the chain through interaction of the qubits with light fields transmitted between the nodes. The quantum repeaters are cyclically operated with their top-level operating cycles being synchronized. Once every top-level operating cycle, each repeater initiates a local merge operation between two qubits that are at least expected to be entangled with qubits in nodes disposed in opposite directions along the chain from the repeater.



GB 2471470 A continuation

(74) Agent and/or Address for Service:

Hewlett-Packard Limited
Bldg 3, Intellectual Property Department,
Filton Road, Stoke Gifford, BRISTOL, BS34 8QZ,
United Kingdom

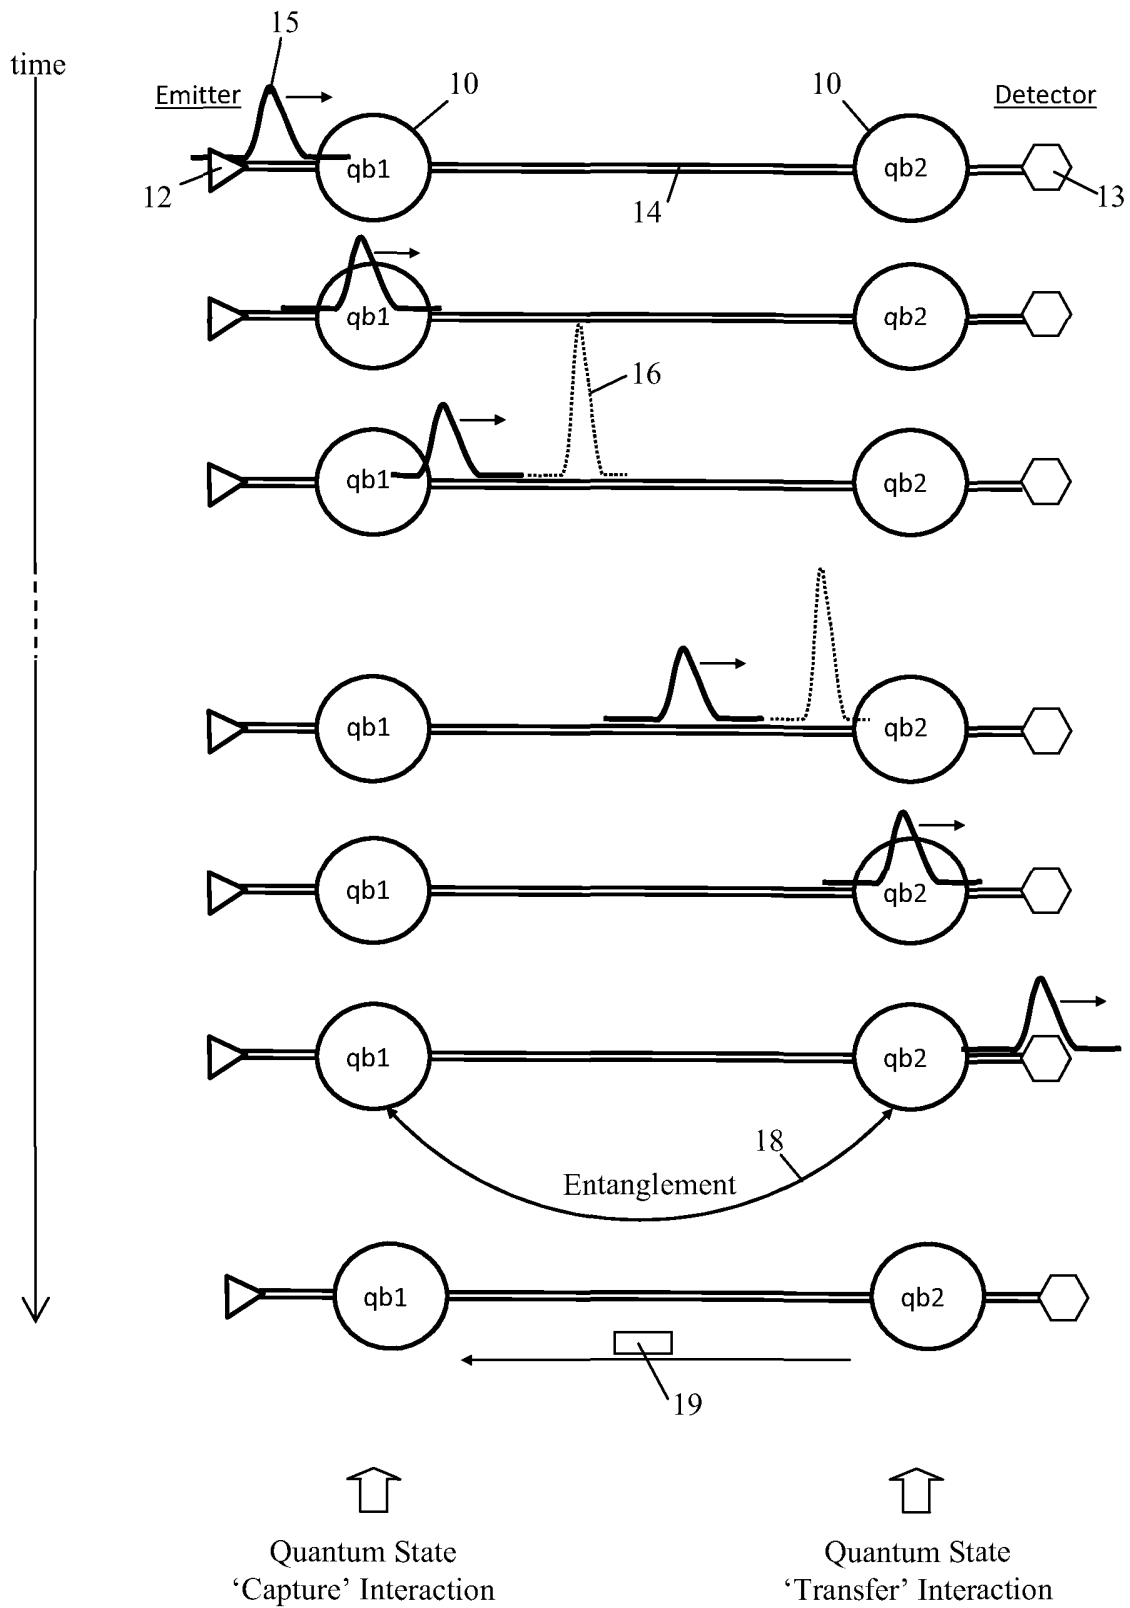


Figure 1

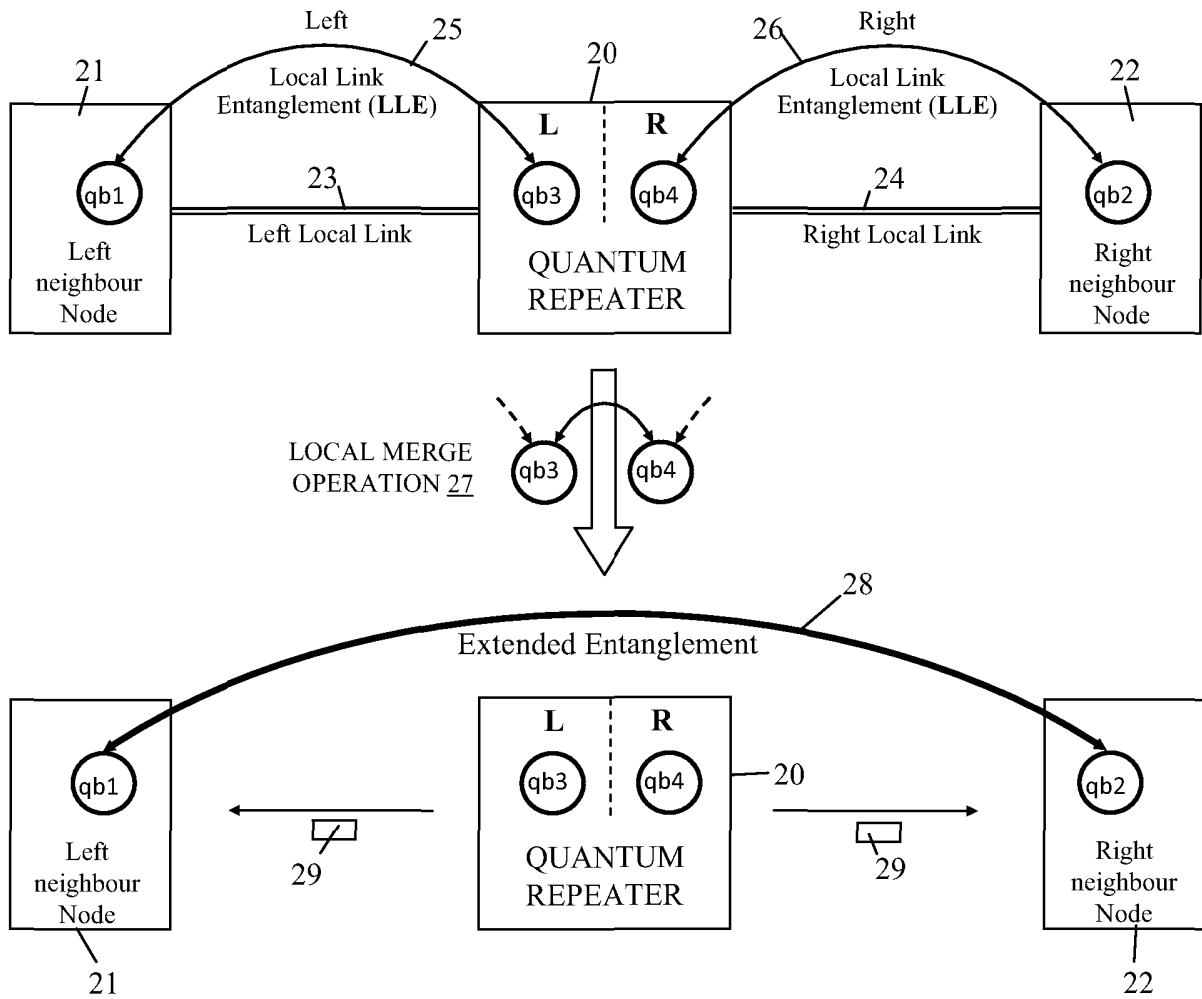


Figure 2




Figure 3




Figure 4




Figure 5

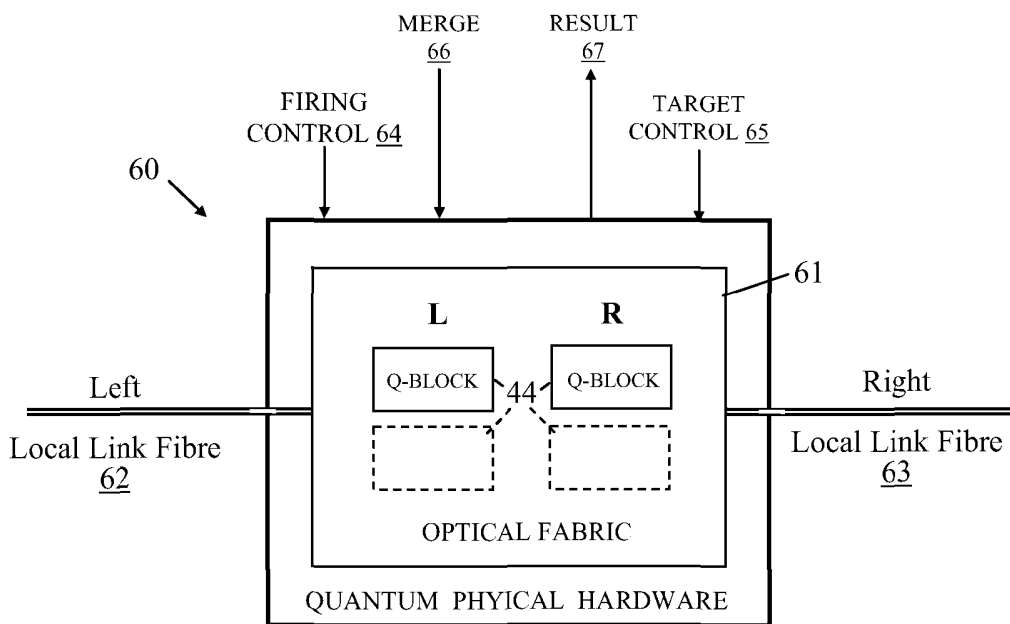


Figure 6

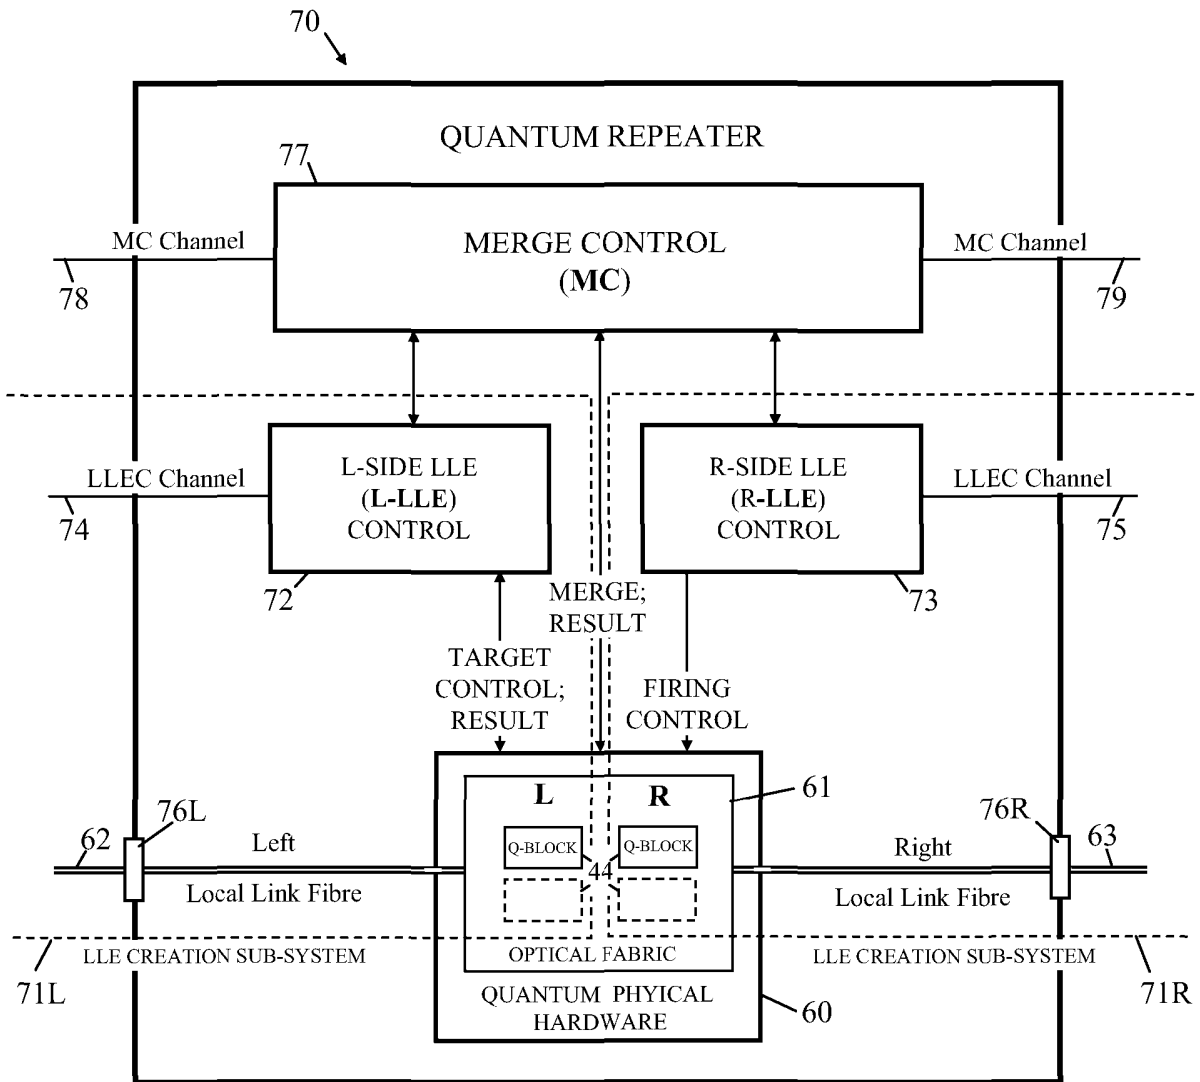
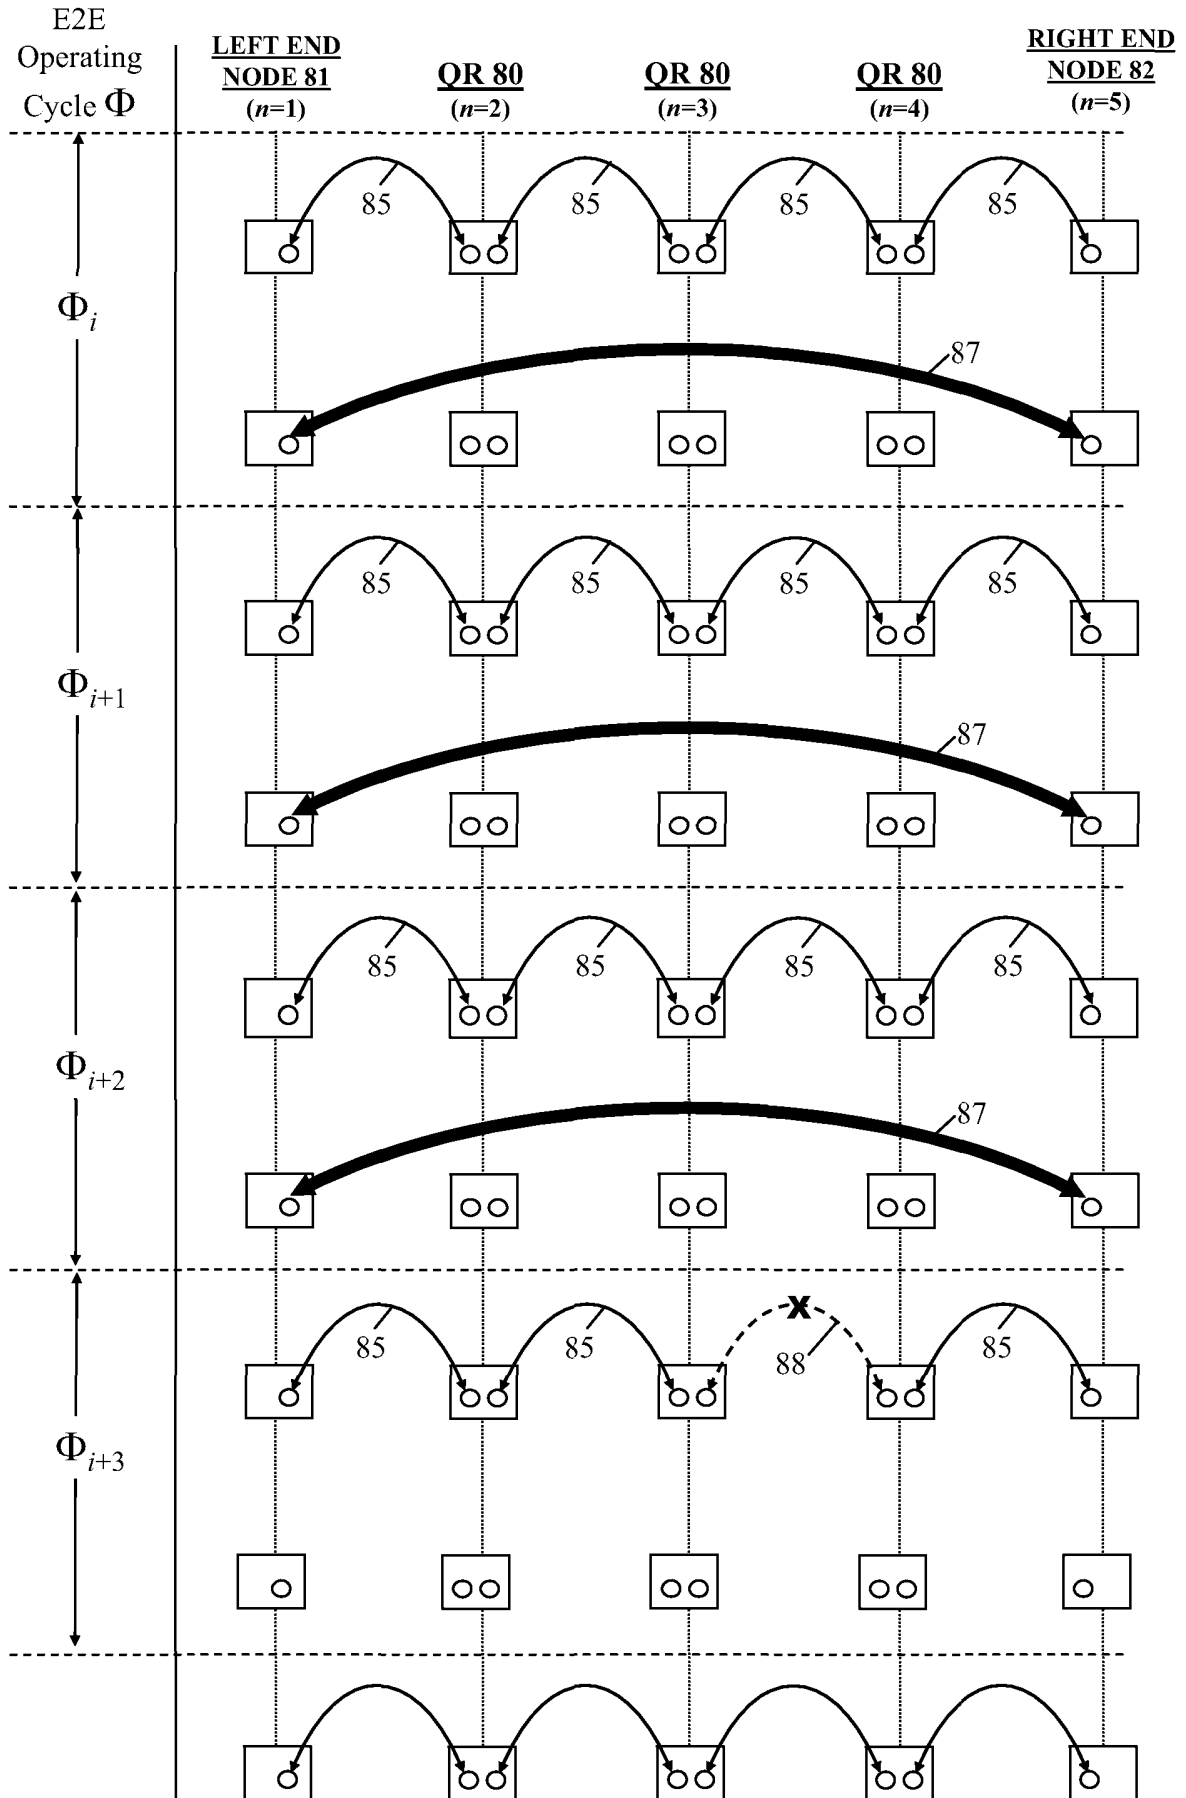


Figure 7

5/17

Figure 8



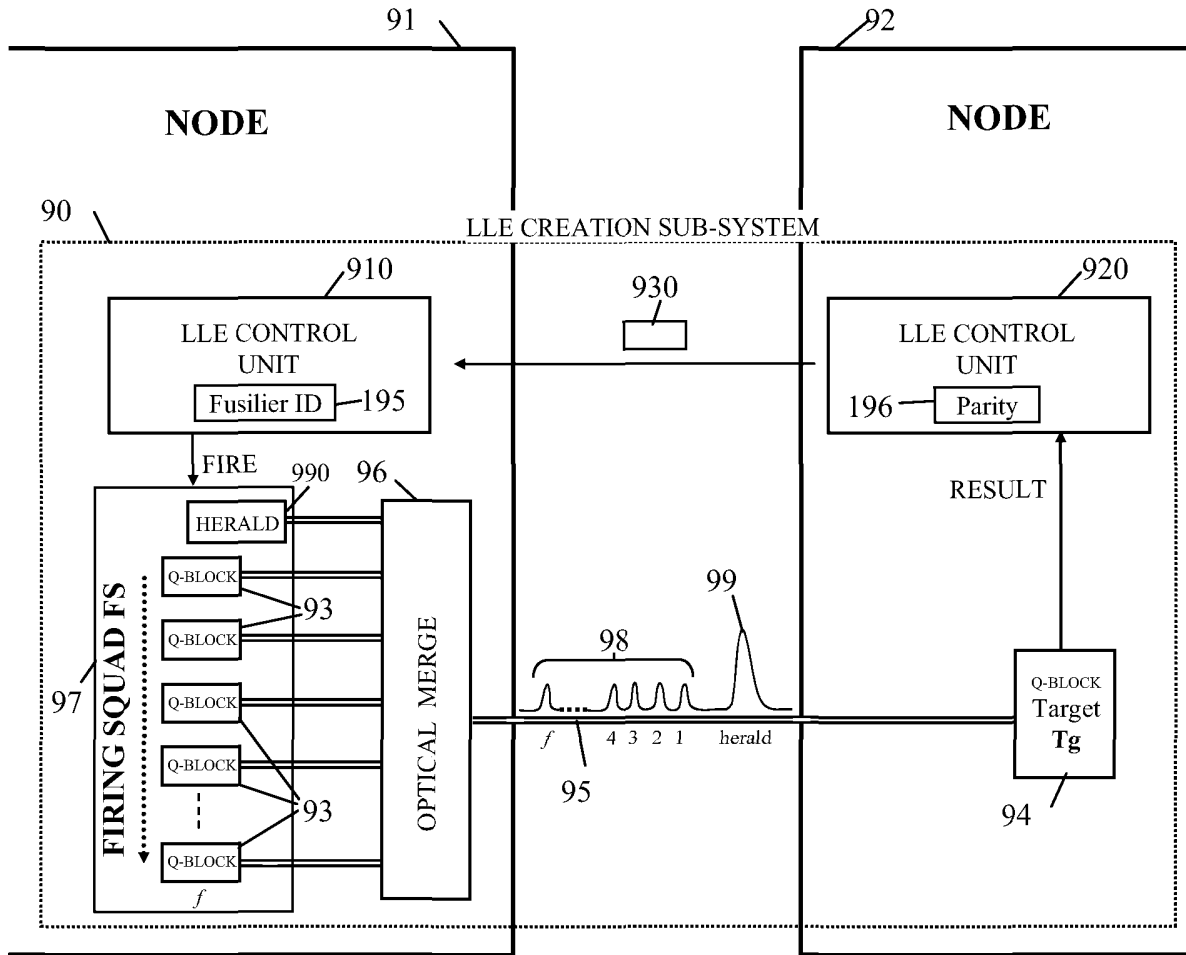


Figure 9




Figure 10

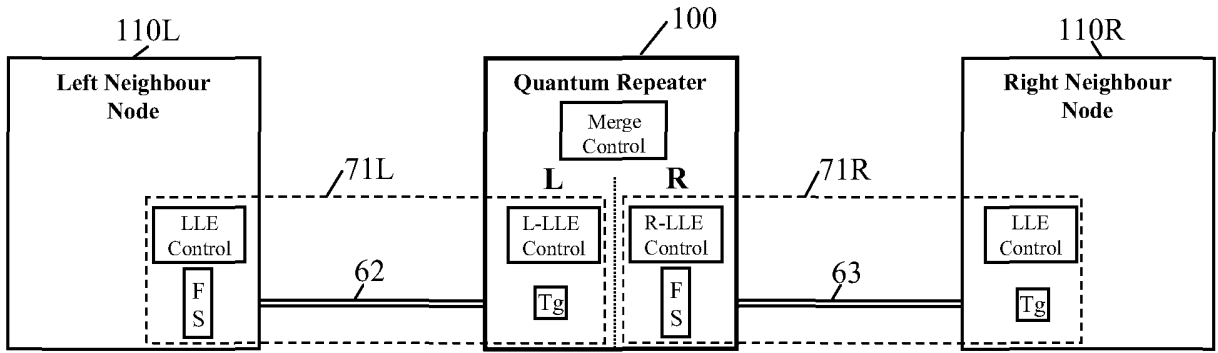


Figure 11

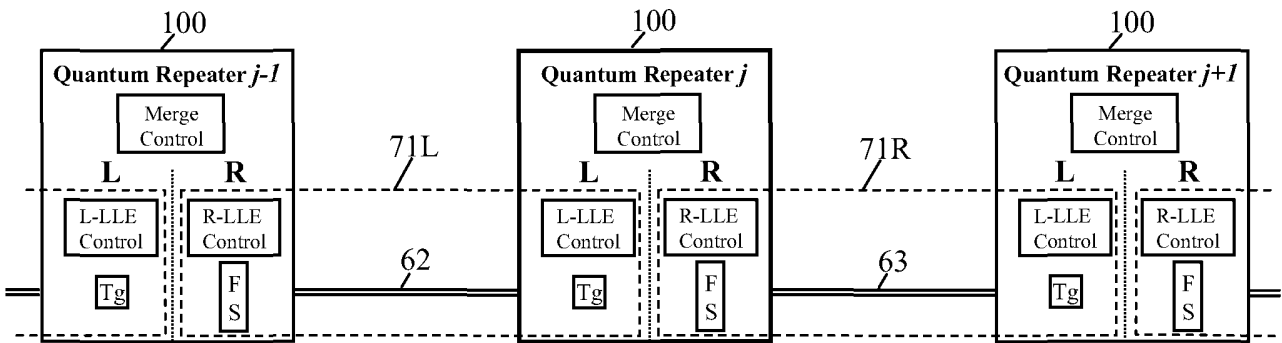


Figure 12

9/17

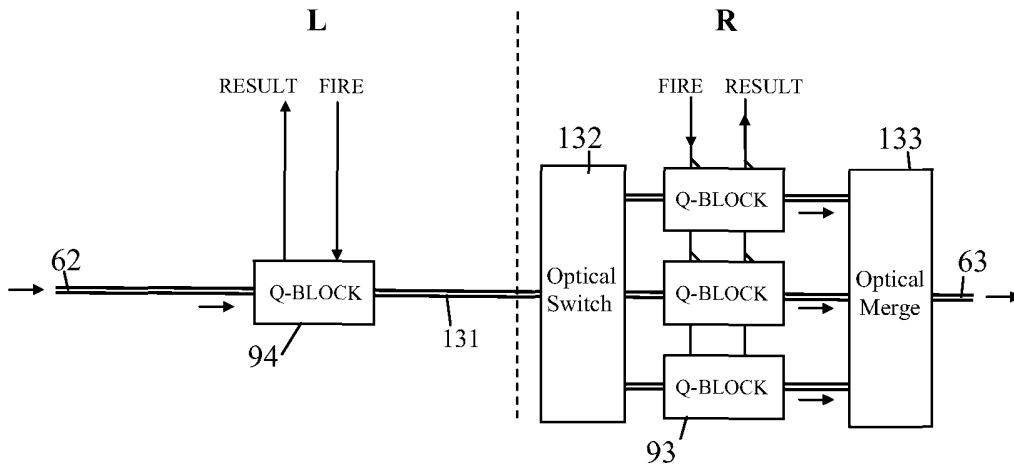


Figure 13A




Figure 13B




Figure 14




Figure 15




Figure 16




Figure 17

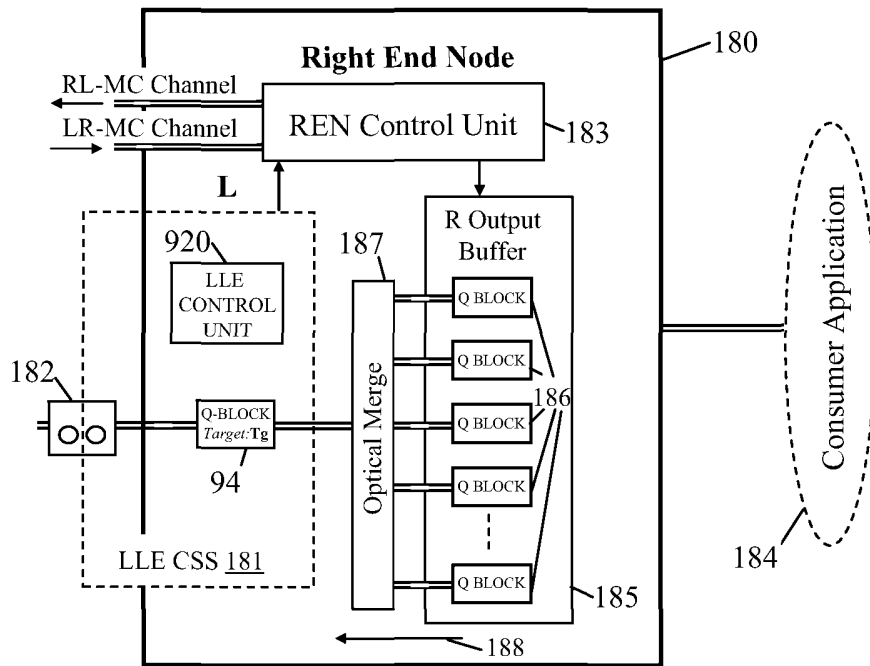


Figure 18

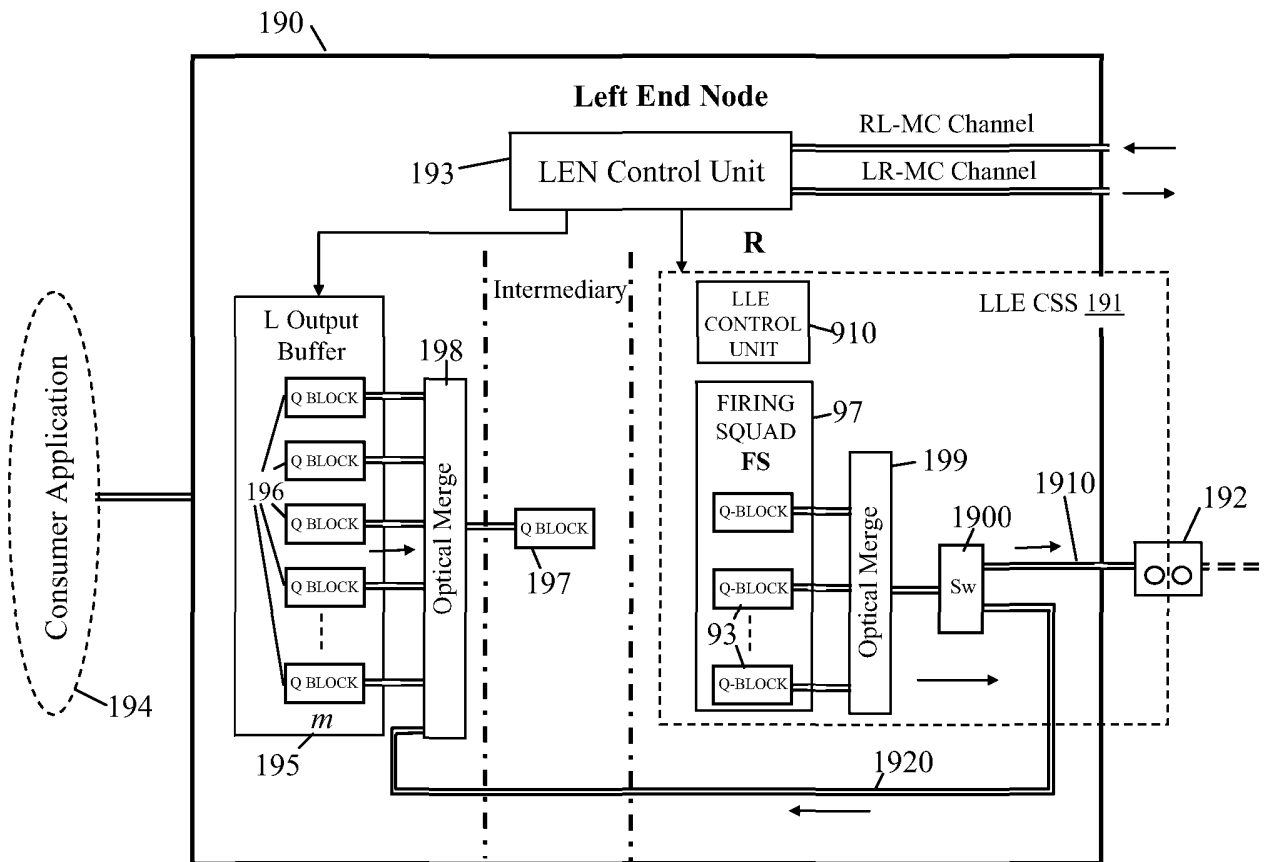


Figure 19

Figure 20

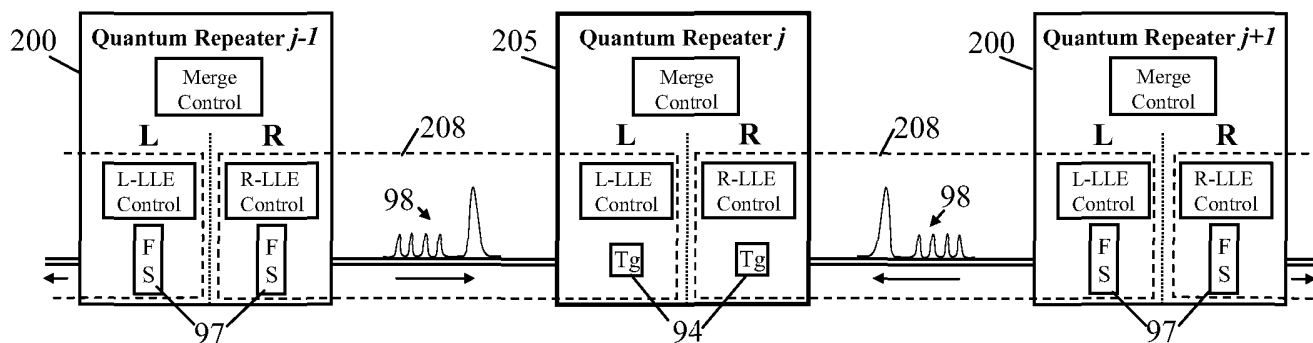


Figure 21


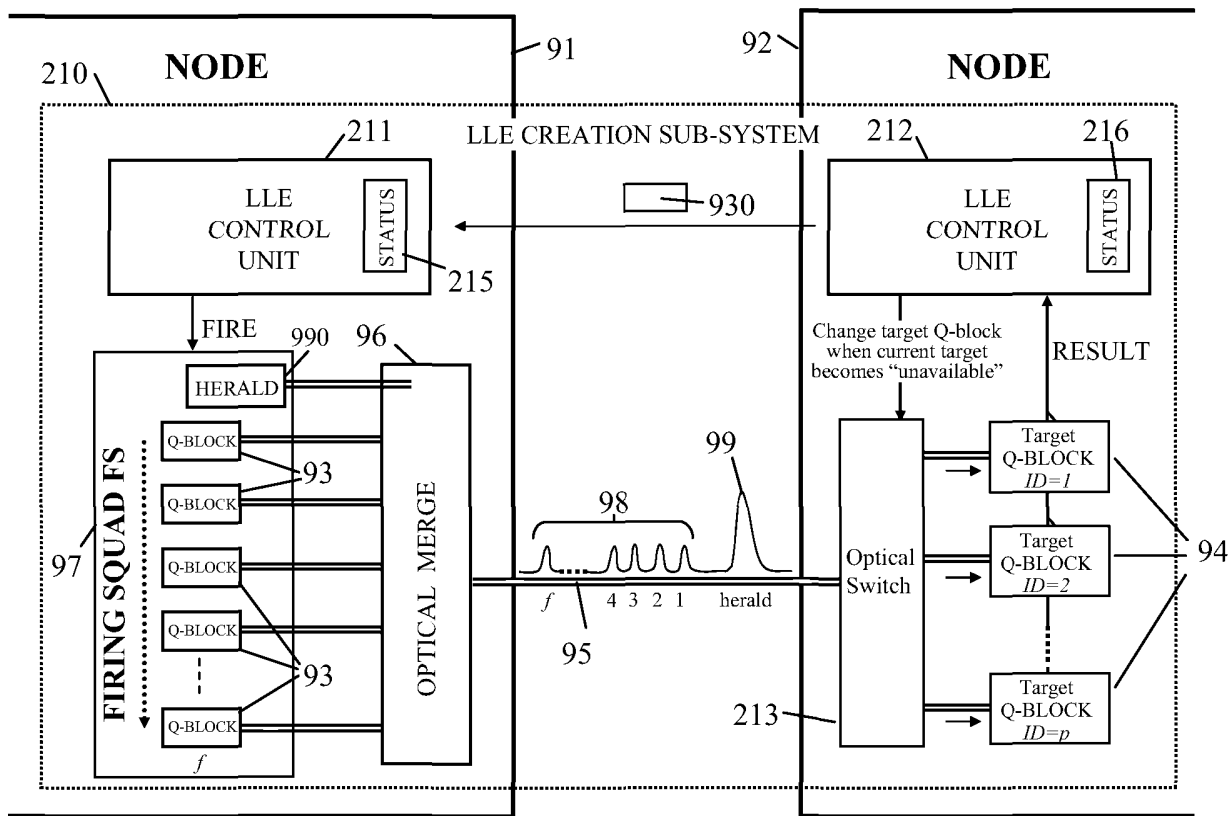


Figure 22

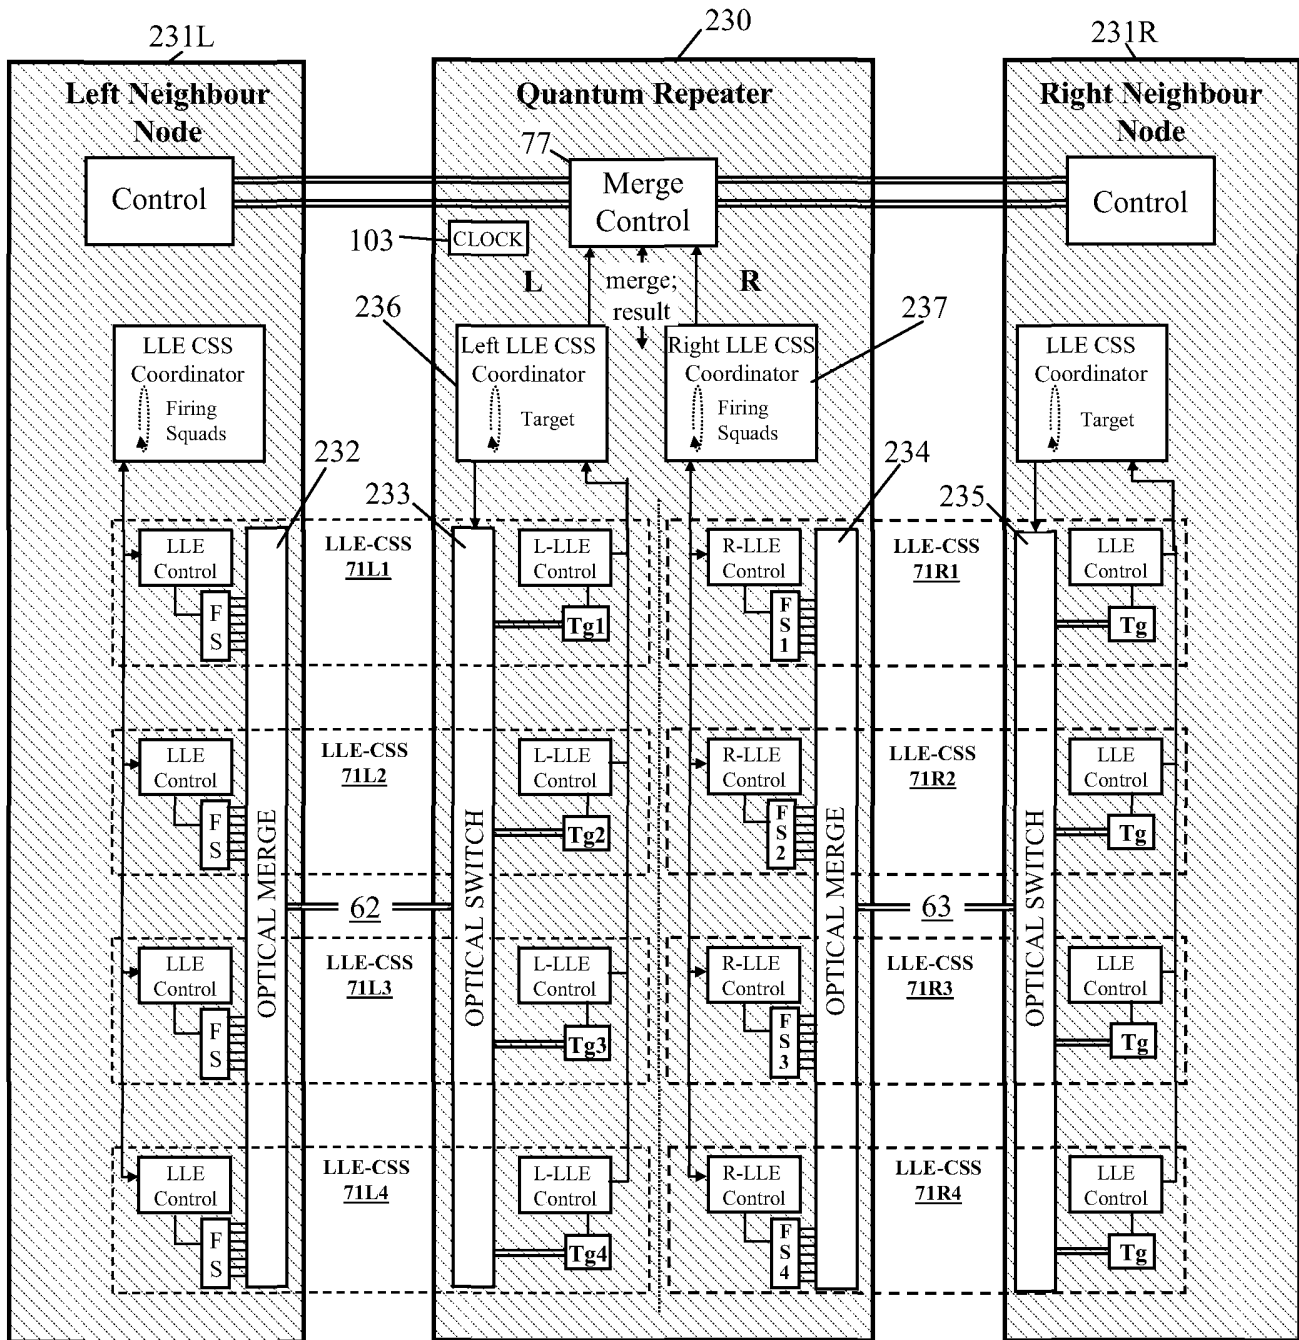


Figure 23




Figure 24




Figure 25

5

Quantum Repeater and System and Method for Creating Extended Entanglements

10

Field of the Invention

The present invention relates to quantum repeaters and to systems and methods for creating extended entanglements.

Background of the Invention

15

In quantum information systems, information is held in the "state" of a quantum system; typically this will be a two-level quantum system providing for a unit of quantum information called a quantum bit or "qubit". Unlike classical digital states which are discrete, a qubit is not restricted to discrete states but can be in a superposition of two states at any given time.

20

Any two-level quantum system can be used for a qubit and several physical implementations have been realized including ones based on the polarization states of single photons, electron spin, nuclear spin, and the coherent state of light.

25

Quantum network connections provide for the communication of quantum information between remote end points. Potential uses of such connections include the networking of quantum computers, and "quantum key distribution" (QKD) in which a quantum channel and an authenticated (but not necessarily secret) classical channel with integrity are used to create shared, secret, random classical bits. Generally, the processes used to convey the

30

quantum information over a quantum network connection provide degraded performance as the transmission distance increases thereby placing an upper limit between end points. Since in general it is not possible to copy a quantum state, the separation of endpoints

cannot be increased by employing repeaters in the classical sense.

One way of transferring quantum information between two spaced locations uses the technique known as ‘quantum teleportation’. This makes use of two entangled qubits, known as a Bell pair, situated at respective ones of the spaced locations; the term “entanglement” is also used in the present specification to refer to two entangled qubits. The creation of such a distributed Bell pair is generally mediated by photons sent over an optical channel (for example, an optical waveguide such as optical fibre). Although this process is distance limited, where a respective qubit from two separate Bell Pairs are co-located, it is possible to combine (or ‘merge’) the Bell pairs by a local quantum operation effected between the co-located qubits. This process, known as ‘entanglement swapping’, results in an entanglement between the two non co-located qubits of the Bell pairs while the co-located qubits cease to be entangled at all.

The device hosting the co-located qubits and which performs the local quantum operation to merge the Bell pairs is called a “quantum repeater”. The basic role of a quantum repeater is to create a respective Bell pair with each of two neighbouring spaced nodes and then to merge the Bell pairs. By chaining multiple quantum repeaters, an end-to-end entanglement can be created between end points separated by any distance thereby permitting the transfer of quantum information between arbitrarily-spaced end points.

It may be noted that while QKD does not directly require entangled states, the creation of long-distance Bell pairs through the use of quantum repeaters facilitates long-distance QKD. Furthermore, most other applications of distributed quantum computation will use distributed Bell pairs.

The present invention is concerned with the creation of entanglement between spaced qubits and with the form, management and interaction of quantum repeaters to facilitate the creation of entanglements between remote end points.

Summary of the Invention

According to the present invention, there is provided a quantum repeater as set out in the accompanying claim 1. The quantum repeater is usable as a intermediate node in a chain of nodes, to permit an end-to-end entanglement between qubits in end nodes of the chain of nodes

Also provided is a method of creating an end-to-end entanglement between qubits in end nodes of a chain of nodes whose intermediate nodes are quantum repeaters, the method being as set out in accompanying claim 13.

Brief Description of the Drawings

Embodiments of the invention will now be described, by way of non-limiting example, with reference to the accompanying diagrammatic drawings, in which:

- Figure 1** is a diagram depicting a known process for entangling two qubits;
- Figure 2** is a diagram depicts how a quantum repeater can be used to create an entanglement between two qubits qb1, qb2 over a distance greater than that possible using the Figure 1 process alone;
- Figure 3** is a diagram illustrating how a chain of quantum repeaters, can be used to create an extended entanglement between any arbitrarily spaced pair of nodes;
- Figure 4** is a diagram illustrating three varieties of a basic quantum physical hardware block, herein a “Q-block”, for carrying out various quantum interactions;
- Figure 5** is a diagram depicting left and right local-link entanglement creation subsystems formed between a quantum repeater and its neighbour nodes;
- Figure 6** is a generic diagram of quantum physical hardware of a quantum repeater;
- Figure 7** is a diagram depicting the general form of quantum repeater embodiments of the invention;
- Figure 8** illustrates four successive operating cycles of a process embodying the invention for creating end-to-end entanglements between end nodes of a chain of five optically-coupled nodes, the intermediate nodes of the chain

being quantum repeaters of the Figure 7 form;

Figure 9 is a diagram of a local-link entanglement creation subsystem for use in quantum repeaters of the Figure 7 form;

5 **Figure 10** is a diagram of a first quantum-repeater embodiment, this embodiment basing local-link entanglement creation on subsystems of the Figure 9 form;

Figure 11 is a diagram showing how the Figure 10 quantum repeater cooperates with neighbouring nodes to form two LLE creation subsystems;

10 **Figure 12** is a diagram showing how Figure 10 quantum repeaters can be serially optically coupled to provide LLE creation subsystems between neighbouring repeaters;

Figures 13A & 13B show respective example implementations of quantum physical hardware of the Figure 10 quantum repeater embodiment;

Figure 14 is a graphical depiction of the timing notation used herein in describing the operation of the Figure 10 quantum repeater;

15 **Figure 15** is a graphical depiction of what is happening in an LLE creation subsystem of the Figure 9 form during successive quarter operating cycle quarter periods of the Figure 10 quantum repeater;

20 **Figure 16** is a message flow diagram depicting one example messaging arrangement for accumulating the failure and parity information for a chain of repeaters of the Figure 10 form;

Figure 17 is a message flow diagram depicting a second example messaging arrangement for accumulating the failure and parity information for a chain of repeaters of the Figure 10 form;

25 **Figure 18** is a diagram of an example implementation of a right end node of a chain of nodes having intermediate nodes formed by Figure 10 quantum repeaters;

Figure 19 is a diagram of an example implementation of a left end node of a chain of nodes having intermediate nodes formed by Figure 10 quantum repeaters;

Figure 20 is a diagram showing how two complimentary varieties of a repeater based on the Figure 10 embodiment can be combined to form a repeater chain;

30 **Figure 21** is a diagram of an alternative local-link entanglement creation subsystem on which quantum repeaters of the Figure 10 form can be based;

Figure 22 is a diagram showing an example segmentation of a chain of quantum

repeater nodes;

Figure 23 is a diagram of a second quantum-repeater embodiment in which four parallel local link entanglement creation subsystems are formed with each of its neighbour nodes;

5 **Figure 24** is a diagram, similar to Figure 15, showing the staggered operation of the paralleled local link entanglement creation subsystems existing between the Figure 23 repeater and one neighbour node; and

Figure 25 is a message flow diagram depicting one example messaging arrangement for accumulating the failure and parity information for a chain of repeaters of
10 the Figure 23 form.

Best Mode of Carrying Out the Invention

15 Basic Quantum Repeater Operation

Figure 1 depicts, in general terms, a known process for entangling two qubits qb1, qb2 (referenced 10) to create a Bell pair, the Figure showing a time series of snapshots taken over the course of the entanglement process. Where, as in the present case, the qubits qb1, qb2 are separated by a distance greater than a few millimeters, the creation of a Bell pair is
20 mediated by photons, which may be sent through free space or over a waveguide such as optical fibre 14. Very generally, processes for Bell-pair creation may be divided into those that use very weak amounts of light (single photons, pairs of photons, or laser pulses of very few photons) and those that use pulses of many photons from a coherent source, such as a laser. As will be understood by persons skilled in the art, the details of the methods of
25 creating photons, performing entangling operations, and making measurements differ depending on whether very weak amounts of light or laser pulses of many photons are used; however, as the present invention can be implemented using any such approach, the following description will be couched simply in terms of a “light field” being used to create (and subsequently ‘merge’) Bell Pairs.

30

Considering Figure 1 in more detail, a light field 15 from emitter 12 is passed through the physical qubit qb1 which is in a prepared non-classical state (for example: $0,+1$); typically, the physical qubit implementation is as electron spin, the electron being set into a

predetermined state immediately prior to passage of the light field. The light field 15 and qubit qb1 interact, with the light field 15 effectively ‘capturing’ the quantum state of the qubit qb1. The light field 15 then travels down the optical fibre 14 and interacts with qubit qb2 before being measured at detector 13; if successful, this results in the ‘transfer’ of the quantum state of qubit qb1 to qubit qb2, entangling these qubits (in Figure 1, this entanglement is represented by double-headed arrow 18). The properties of the light field 15 measured by detector 13 enable a determination to be made as whether or not the entanglement operation was successful. The success or failure of the entanglement operation is then passed back to the qb1 end of the fibre 14 in a classical (non-quantum) message 19. This message can be very simple in form (the presence or absence of a single pulse) and as used herein the term “message” is to be understood to encompass both such simple forms as well as structured messages of any degree of complexity (subject to processing time constraints); in embodiments where the message 19 needs to identify a particular qubit amongst several as well as the success or failure of an entanglement operation, the message may still take the form of the presence or absence of a single pulse with the timing of the latter being used to identify the qubit concerned. Due to the need to transmit information about the success/failure of the entanglement operation back to the qb1 end of the fibre 14, the overall elapsed time from initiation of the entanglement operation to when an entanglement could be usefully available is at least the round trip propagation time along the fibre 14, even where the entanglement operation is successful.

The properties of the light field 15 measured by detector 13 also enable a determination to be made, in the case of a successful entanglement operation, as to whether the entangled states of the qb1 and qb2 are correlated or anti-correlated, this generally being referred to as the ‘parity’ of the entanglement (even and odd parity respectively corresponding to correlated and anti-correlated qubit states). It is normally important to know the parity of an entanglement when subsequently using it; as a result, either parity information must be stored or steps taken to ensure that the parity always ends up the same (for example, if an odd parity is determined, the state of qb2 can be flipped to produce an even parity whereby the parity of the entanglement between qb1 and qb2 always ends up even).

As already indicated, the qubits qb1 and qb2 are typically physically implemented as

electron spin. However, the practical lifetime of quantum information stored in this way is very short (of the order of 10^{-6} seconds cumulative) and therefore generally, immediately following the interaction of the light field 15 with qb1 and qb2, the quantum state of the qubit concerned is transferred to nuclear spin which has a much longer useful lifetime (typically of the order of a second, cumulatively). The quantum state can be later transferred back to electron spin for a subsequent light field interaction (such as to perform a merge of two entanglements).

Another practical feature worthy of note is that the physical qubits qb1 and qb2 are generally kept shuttered from light except for the passage of light field 15. To facilitate this at the qb2 end of the fibre 14 (and to trigger setting the qubit into a prepared state immediately prior to its interaction with light field 15), the light field 15 can be preceded by a 'herald' light pulse 16; this light pulse is detected at the qb2 end of the fibre 14 and used to trigger priming of the qubit qb2 and then its un-shuttering for interaction with the light field 15. Other ways of triggering these tasks are alternatively possible.

The relationship between the probability of successfully creating a Bell pair, the distance between qubits involved, and the fidelity of the created pair is complex. By way of example, for one particular implementation using a light field in the form of a laser pulse of many photons, Bell pairs are created with fidelities of 0.77 or 0.638 for 10km and 20km distances respectively between qubits, and the creation succeeds on thirty eight to forty percent of the attempts. The main point is that the entanglement creation process depicted in Figure 1 is distance limited; for simplicity, in the following a probability of success of 0.25 is assumed at a distance of 10km.

25

Figure 2 depicts, in general terms, how a quantum repeater 20 can be used to create an entanglement between two qubits qb1, qb2 which are separated by a distance greater than that over which the Figure 1 Bell Pair creation process can be employed with any reasonable probability of success.

30

The quantum repeater 20 forms one node in a chain (sequential series) of nodes that also includes left and right neighbour nodes 21, 22 to which the quantum repeater 20 is

connected by respective optical fibre links 23 and 24. In the Figure 2 example, the qubits qb1, qb2 that are to be entangled are located in nodes 21 and 22 respectively. The optical fibre 23 connecting the quantum repeater 20 to its left neighbour node 21 is referred to herein as the 'left local link' and the optical fibre 24 connecting the quantum repeater 20 to its right neighbour node 22 is referred to as the 'right local link'. The terms "left" and "right" as used throughout the present specification are simply to be understood as convenient labels for distinguishing opposite senses (directions along; ends of ; and the like) of the chain of nodes that includes the quantum repeater 20.

- 10 The quantum repeater 20 effectively comprises left and right portions or sides (labeled "L" and "R" in Figure 2) each comprising a respective qubit qb3, qb4.

In operation, the Figure 1 Bell-Pair creation process is applied over the left local link 23 to set up an entanglement 25 between qubit qb1 of node 21 and qb3 of quantum repeater 20; 15 similarly, the Figure 1 Bell-Pair creation process is applied over the right local link 24 to set up an entanglement 26 between qubit qb2 of node 22 and qb4 of quantum repeater 20. These directly created, node-to-node, entanglements are hereinafter referred as "local link entanglements" (abbreviated to "LLE"), the LLE 25 set up between the quantum repeater 20 and its left neighbour node 21 being labeled the 'left' LLE and the LLE 26 set up 20 between the quantum repeater 20 and its right neighbour node 22 being labeled the 'right' LLE.

The direction of travel (left-to-right or right-to-left) of the light field used to set up each LLE is not critical whereby the disposition of the associated emitters and detectors can be 25 set as desired. For example, the light fields involved in creating both LLE 25, 26 could both be sent out from the quantum repeater 20 meaning that the emitters are disposed in the quantum repeater 20 and the detectors in the left and right nodes 21, 22. However, to facilitate chaining of quantum repeaters of the same form, it is convenient if the light fields all travel in the same direction along the chain of nodes; for example, the light fields can 30 be arranged all to travel from left to right in which case the left side L of the quantum repeater 20 will include the detector for creating the left LLE 25 and the right side R will include the emitter for creating the right LLE 26. For simplicity, and unless otherwise

stated, a left-to-right direction of travel of the light field between the nodes will be assumed hereinafter in the described exemplary embodiments; the accompanying Claims are not, however, to be interpreted as restricted to any particular direction of travel of the light field, or to the direction of travel being the same across different links, unless so
5 stated or implicitly required.

The order in which the left and right LLEs 25 and 26 are created is not relevant (indeed they could be created simultaneously); all that is required is that both entanglements exist in a usable condition at a common point in time. At such a time, the left and right LLEs 25,
10 26 are “merged” by a quantum operation 27 carried out locally in the quantum repeater 20 on qubits qb3 and qb4. This local operation (referred to hereinafter as a “merge” operation) involves a first process akin to that of Figure 1 process with Capture and Transfer interactions effected by passing a light field successively through the two qubits qb3 and qb4, or vice versa, and then measuring the light field. This first process, if successful,
15 results in the qubits qb1 and qb2 becoming entangled; a second measurement process effectively removes the qubits qb3 and qb4 from the entangled whole leaving an ‘extended’ entanglement (medium thick arc 28 in Figure 2) between the qubits qb1, qb2. The qubits qb3 and qb4 finish up neither entangled with each other nor with the qubits qb1, qb2. Because the merge operation 27 is a local operation between two co-located
20 qubits, the probability of success is very high.

The measurements made as part of the merge operation 27 provide both an indication of the success or otherwise of the merge, and an indication of the “generalized parity” of the merge operation (that is, both a qubit parity value and a conjugate qubit parity value of the
25 merged LLEs). For example, the first merge-operation process may measure a qubit value parity and the second merge-operation process, the conjugate qubit parity value. In this case, the second process can be effected either as a single measurement on a light field passed through both qubits qb3 and qb4 (in which case the light field has a different value to that used in the first process e.g. 0,+1 as opposed to 0,-1), or as individual
30 measurements, subsequently combined, made individually on qb3, and qb4. “Generalized parity” requires two classical bits to represent it. In certain applications of the resultant extended entanglement (such as QKD), knowledge of the conjugate qubit parity value

information may not be required and so would not need to be passed on from the quantum repeater. Hereinafter, reference to “parity” is to be understood to mean “generalized parity” but with the understanding that in appropriate cases, the conjugate qubit parity value information can be omitted.

5

The parity of the extended entanglement 28 will be a combination of the parities of the LLEs 25, 26 and the parity of the merge operation; for the LLEs, the conjugate qubit parity value information is effectively even parity. Where qubit parity value information and conjugate qubit parity value information are each represented by binary values ‘0’ and ‘1’
10 for even and odd parity respectively, the qubit parity value information and conjugate qubit parity value information of the extended entanglement are respective XOR (Exclusive OR) combinations of the corresponding component parities.

Information about the success or otherwise of the merge operation is passed in classical
15 messages 29 from the quantum repeater 20 to its left and right neighbour nodes 21, 22 as otherwise these nodes do not know whether the qubits qb1, qb2 are entangled; alternatively since the failure probability of a merge operation is normally very low, success can be assumed and no success/failure message sent – in this case, it will be up to applications consuming the extended entanglement to detect and compensate for merge
20 failure leading to absence of entanglement.

As the parity of the extended entanglement will normally need to be known to make use of the entangled qubits, parity information is also passed on from the quantum repeater 20 in one of the messages 29 to one or other of the nodes 21, 22. Where the parity of the local
25 link entanglements has been standardized (by qubit state flipping as required), only the merge parity information needs to be passed on by the quantum repeater and either node 21 or 22 can make use of this information. However, where LLE parity information has simply been stored, then the quantum repeater needs to pass on whatever parity information it possesses; for example, where the parities of the left and right LLEs 25, 26
30 are respectively known by the quantum repeater 20 and the node 22, the quantum repeater 20 needs to pass on to node 22 both the parity information on LLE 25 and the merge parity information, typically after combining the two. Node 22 can now determine the parity of

the extended entanglement by combining the parity information it receives from the quantum repeater 20 with the parity information it already knows about LLE 26.

From the foregoing, it can be seen that although the merge operation itself is very rapid (of the order of 10^{-9} seconds), there is a delay corresponding to the message propagation time to the furthest one of the nodes 21, 22 before the extended entanglement 28 is usefully available to these nodes.

It will be appreciated that where the quantum states of qubits qb3, qb4 have been transferred from electron spin to nuclear spin immediately following the creation of the LLEs 25, 26 respectively, these states need to be transferred back to electron spin before the merge operation 27 is effected.

By chaining together multiple quantum repeaters, it is possible to create an extended entanglement between any arbitrarily spaced pair of nodes. Figure 3 illustrates this for a chain of N nodes comprising left and right end nodes 31 and 32 respectively, and a series of $(N-2)$ quantum repeaters 30 (each labeled "QR" and diagrammatically depicted for simplicity as a rectangle with two circles that represent Land R qubits). The nodes 30-32 are interconnected into a chain by optical fibre (not shown) and are numbered from left to right - the number n of each node is given beneath each node and node number " j " represents an arbitrary QR node 30 along the chain. The node number of a QR node can be used as a suffix to identify the node; thus "QR $_j$ " is a reference to the quantum repeater node numbered j . This node representation, numbering and identification is used generally throughout the present specification.

25

In Figure 3, three existing entanglements 33, 34, and 35 are shown between qubits in respective node pairings; for convenience, when referring at a high level to entanglements along a chain of nodes, a particular entanglement will herein be identified by reference to the pair of nodes holding the qubits between which the entanglement exists, this reference taking the form of a two-element node-number tuple. Thus, entanglement 35, which is a local link entanglement LLE between qubits in the neighbouring nodes numbered $(N-1)$ and N , is identifiable by the node number tuple $\{(N-1), N\}$. Entanglements 33 and 34

30

(shown by medium thick arcs in Figure 3) are extended entanglements existing between qubits in the node pairings $\{1, j\}$ and $\{j, (N-1)\}$ respectively, these entanglements having been created by the merging of LLEs. To create an end-to-end (abbreviated herein to “E2E”) entanglement between qubits in the left and right end nodes 31, 32 (see thick arc 5 39 in Figure 3), entanglements 33 and 34 can first be merged by QR_j with the resultant extended entanglement then being merged with LLE 35 by $QR_{(N-1)}$; alternatively, entanglements 34 and 35 can first be merged by $QR_{(N-1)}$ with the resultant extended entanglement then being merged with entanglement 33 by QR_j .

10 Representation of Low Level Quantum Physical Hardware

The particular form of physical implementation of a qubit and the details of the methods of performing entanglement and merge operations (for example, whether very weak amounts of light or laser pulses of many photons are used) are not of direct relevance to the present invention and accordingly will not be further described herein, it being understood that 15 appropriate implementations will be known to persons skilled in the art. Instead, the physical hardware for implementing the quantum operations (the “quantum physical hardware”) will be represented in terms of a basic block, herein called a “Q-block”, and an associated optical fabric.

20 Figure 4 depicts three varieties of Q-block, respectively referenced 40, 42 and 44.

Q-block variety 40 represents the physical hardware needed to carry out the “Capture” interaction of Figure 1, that is, the controlled sending of a light field through a qubit in a prepared state. This variety of Q-block – herein called “a Capture Q-block ” (abbreviated 25 in the drawings to “Q-block (C)”) – comprises a physical implementation of a qubit 10 and a light-field emitter 12, together with appropriate optical plumbing, functionality for putting the qubit in a prepared state and for shuttering it (for example, using an electro-optical shutter) except when a light field is to be admitted, functionality (where appropriate for the qubit implementation concerned) for transferring the qubit state between electron 30 spin and nuclear spin (and vice versa) as needed, and control functionality for coordinating the operation of the Capture Q-block to send a light field through its qubit (and on out of the Q-block) upon receipt of a “Fire” signal 41.

Q-block variety 42 represents the physical hardware needed to carry out the “Transfer” interaction of Figure 1, that is, the passing of a received light field through a qubit in a prepared state followed by measurement of the light field. This variety of Q-block – herein called “a Transfer Q-block ” (abbreviated in the drawings to “Q-block (T)”) – comprises a physical implementation of a qubit 10 and a light-field detector 13, together with appropriate optical plumbing, functionality (responsive, for example to a herald light pulse 16) for putting the qubit in a prepared state and for shuttering it except when a light field is to be admitted, functionality (where appropriate for the qubit implementation concerned) for transferring the qubit state between electron spin and nuclear spin (and vice versa) as needed, and control functionality for coordinating the operation of the Transfer Q-block and for outputting the measurement results 43.

Q-block variety 44 is a universal form of Q-block that incorporates the functionality of both of the Capture and Transfer Q-block varieties 40 and 42 and so can be used for both Capture and Transfer operations. For convenience, this Q-Block variety is referred to herein simply as a “Q-block” without any qualifying letter and unless some specific point is being made about the use of a Capture or Transfer Q-block 40, 42, this is the variety of Q-block that will be generally be referred to even though it may not in fact be necessary for the Q-block to include both Capture and Transfer interaction functionality in the context concerned – persons skilled in the art will have no difficulty in recognizing such cases and in discerning whether Capture or Transfer interaction functionality is required by the Q-block in its context. One reason not to be more specific about whether a Q-block is of a Capture or Transfer variety is that often either variety could be used provided that a cooperating Q-block is of the other variety (the direction of travel of light fields between them not being critical).

It may be noted that where there are multiple Q-blocks in a node, the opportunity exists to share certain components between Q-blocks (for example, where there are multiple Q-blocks with Capture interaction functionality, a common light-field emitter may be used for all such Q-blocks. Persons skilled in the art will appreciate when such component sharing is possible.

An entanglement operation will thus involve a Q-block with Capture interaction functionality (either a Transfer Q-block 40 or a universal Q-block 44) optically coupled to a Q-block with Transfer interaction functionality (either a Transfer Q-block 42 or a universal Q-block 44), the entanglement operation being initiated by a Fire signal 41 sent to the Q-block with Capture interaction functionality and the success/failure of the operation being indicated in the result signal 43 output by the Q-block with Transfer interaction functionality.

Where a merge operation is to be effected, this will also involve a Q-block with Capture interaction functionality and a Q-block with Transfer interaction functionality. The latter will include additional elements associated with the second merge-operation process mentioned above; indeed, depending on how this second process is effected, the Q-block with Capture interaction functionality may also include additional elements. Hereinafter it is to be understood that where a Q-block is involved in a merge operation, it includes the required additional elements.

Figure 5 depicts a quantum repeater 50 optically coupled to left and right neighbour nodes 51, 53 by local link fibres 52, 54 respectively. Quantum repeater 50 includes respective Q-blocks 44 on its left side L and right side R. The L side Q-block 44 of the quantum repeater 50 is optically coupled by the local link fibre 52 to a Q-block 44 of the left neighbour node 51 thereby to form a sub-system 55 for creating local link entanglements (LLEs) between the quantum repeater 50 and its left node 51; such a node-spanning sub-system, along with the associated control logic, is herein called a “LLE creation subsystem”. A similar LLE creation subsystem 56 is formed by the R side Q-block 44 of the quantum repeater 50, the local link fibre 54, and a Q-block 44 in the right neighbour node 53. Because the Q-blocks 44 depicted in Figure 5 are of the universal variety, the direction of travel along the local link fibres of light fields involved in entanglement creation is not tied down; thus for example, in respect of LLE creation, the Q-blocks of the quantum repeater 50 could both be operated as of the same variety (Capture or Transfer) in which case the light fields in the two LLE creation subsystems pass in opposite directions along

their respective local link fibres 52, 54 (either both away from the quantum repeater 50 or both towards the quantum repeater 50).

5 A merge operation will also involve a Q-block with Capture interaction functionality (either a Transfer Q-block 40 or a universal Q-block 44) optically coupled to a Q-block with Transfer interaction functionality (either a Transfer Q-block 42 or a universal Q-block 44); the Q-blocks involved will both be located in the same quantum repeater. The result signal 43 output by the Q-block with Transfer interaction functionality not only provides a success/failure indication for the merge operation, but also a parity indication.

10

Generally, within a quantum repeater, a controllable optical fabric will need to be provided to appropriately guide light fields to/from the Q blocks of the repeater depending on its current operational requirements, for example, whether LLE creation is being effected in cooperation with Q-blocks of neighbour nodes or whether two local Q blocks are involved
15 in a merge operation. Also, where one or both of the L side and R side of the quantum repeater comprises multiple Q-blocks, an optical fabric will generally be needed to enable the local link fibres to be shared between the Q-blocks.

In general terms, therefore, the quantum physical hardware of a node, that is, the physical
20 elements that implement and support qubits and their interaction through light fields, comprises not only one or more Q-blocks but also an optical fabric in which the Q-block(s) are effectively embedded. By way of example, Figure 6 depicts such a representation for a quantum repeater node; thus, quantum physical hardware 60 is shown as comprising an optical fabric 61 for guiding light fields to/from the Q-blocks 44 and the Q-blocks 44 are
25 depicted as existing within the optical fabric 61 with the local link fibres 62, 63 coupling directly to the optical fabric. One L-side and one R-side Q-block are shown in solid outline and possible further L-side and R-side Q-blocks are indicated by respective dashed-outline Q-blocks.

30 As employed herein, any instance of the above-described quantum physical hardware representation (such as the instance shown in Figure 6 in respect of a quantum repeater), is intended to embrace all possible implementations of the quantum physical hardware

concerned, appropriate for the number and type of Q-blocks involved and their intended roles.

The quantum physical hardware block 60 is arranged to receive Firing Control and Target
 5 Control signals 64, 65 for controlling LLE creation operations, to receive “Merge” signals
 67 for controlling merge operations, and to output “Result” signals 66 indicative of the
 outcome of these operations. The signals 64-67 may be parameterized to indicate particular
 Q-blocks. Target Control signals are not needed in some quantum repeater embodiments as
 will become apparent hereinafter.

10

In one implementation, the Firing Control signals 64 comprise both:

- set-up signals for appropriately configuring the optical fabric 61 (if not already so
 configured) to optically couple one or more Q-block(s) with Capture interaction
 functionality to one of the local link fibres, and
- 15 - the previously-mentioned “Fire” signal(s) for triggering light-field generation by
 one or more of the Q-block(s) with Capture interaction functionality;

and the Target Control signals 65 comprise:

- set-up signals for appropriately configuring the optical fabric 61 (if not already so
 configured) to optically couple a Q-block with Transfer interaction functionality
 20 to one of the local link fibres.

Furthermore, in this implementation, the Merge signals 66 comprise both:

- set-up signals for appropriately configuring the optical fabric 61 (if not already so
 configured) to effect a merge operation involving a L-side and R-side Q-block of
 the repeater, and
- 25 - first and second “Fire” signals for triggering the first and second merge-operation
 processes respectively.

The optical fabric 61 may have a default configuration. For example, where the quantum
 physical hardware 60 only includes one L-side and one R-side Q-block, the optical fabric
 30 61 may be arranged to default to an LLE creation configuration optically the Q-blocks to
 respective ones of the local link fibres. In this case, the merge signals 66 are arranged to
 only temporarily optically couple the two Q-blocks to each other for the time needed to

carry out a merge operation. In cases such as this, the Target Control signals 65 can be dispensed with entirely and the Firing Control signals 64 simply comprise Fire signals sent to the appropriate Q-block.

5

General Form of Quantum Repeater Embodiments

Figure 7 depicts the general form of the quantum repeater embodiments to be described hereinafter.

10

More particularly, quantum repeater 70 comprises quantum physical hardware 60 of the form described above with respect to Figure 6 and including one or more L-side and R-side Q-blocks 44, and optical fabric 61 coupled to left and right local link fibres 62, 63 via respective optical interfaces 76L, 76R. As already indicated, for convenience and without
 15 limitation, the light fields involved in LLE creation will be taken (unless otherwise stated) as travelling from left to right along the local link fibres between nodes, whereby the R-side Q-block(s) of the Figure 7 repeater 70 act as Capture Q-block(s) during LLE creation (forming a right-side LLE creation subsystem 71R with L-side Q-block(s) in a right neighbour node, not shown), and the L-side Q-block(s) of the repeater 70 act as Transfer
 20 Q-block(s) during LLE creation (forming a left-side LLE creation subsystem 71R with R-side Q-block(s) in a left neighbour node, not shown).

An R-side LLE (“R-LLE”) control unit 73 is responsible for generating the Firing Control signals that select (where appropriate) and trigger firing of the R-side Q-block(s) in respect
 25 of LLE creation. An L-side LLE (“L-LLE”) control unit 72 is responsible for generating, where appropriate, the Target Control signals for selecting the L-side Q-block(s) to participate in LLE creation; the L-LLE control unit 72 is also arranged to receive the Result signals from the quantum physical hardware 60 indicative of the success/failure of the LLE creation operations involving the L-side Q-blocks.

30

It will thus be appreciated that initiation of right-side LLE creation is effectively under the control of the R-LLE control unit 73 of the repeater 70 (as unit 73 is responsible for generating the Fire signal for the R-side Q-block involved in creating the right-side LLE);

initiation of left-side LLE creation is, however, effectively under the control of the R-LLE control unit in the left neighbour node.

5 An LLE control (“LLEC”) classical communication channel 74 inter-communicates the L-LLE control unit 72 with the R-LLEC unit of the left neighbour node (that is, the R-LLE control unit associated with the same LLE creation subsystem 71L as the L-LLE control unit 72); the L-LLEC unit 72 uses the LLEC channel 74 to pass LLE creation success/failure messages (message 15 in Figure 1) to the R-LLE control unit of the left neighbour node.

10

An LLE control (“LLEC”) classical communication channel 75 inter-communicates the R-LLE control unit 73 with the L-LLE control unit of the right neighbour node (that is, the L-LLE control unit associated with the same LLE creation subsystem 71R as the R-LLE control unit 73); the R-LLE control unit 73 receives LLE creation success/failure messages
15 (message 15 in Figure 1) over the LLEC channel 75 from the L-LLE control unit of the right neighbour node.

Messages on the LLEC channels 74, 75 are referred to herein as ‘LLEC’ messages.

20 It will be appreciated that where the light fields involved in LLE creation are arranged to travel from right to left along the local link fibres between nodes (rather than from left to right), the roles of the L-side and R-side LLE control units 72, 73 are reversed.

25 A merge control (“MC”) unit 77 is responsible for generating the Merge signals that select, where appropriate, local Q-blocks to be merged, and trigger their merging. The MC unit 77 is also arranged to receive from the quantum physical hardware 60, the Result signals indicative of the success/failure and parity of a merge operation.

30 A merge control (“MC”) classical communication channel 78, 79 inter-communicates the MC unit 77 with corresponding units of its left and right neighbour nodes to enable the passing of parity information and, if needed, success/failure information concerning merge operations.. Messages on the MC channels 78, 79 are referred to herein as ‘MC’ messages.

The LLEC communication channel 74, 75 and the MC communication channel 78, 79 can be provided over any suitable high-speed communication connections (such as radio) but are preferably carried as optical signals over optical fibres. More particularly, the LLEC communication channel 74, 75 and the MC communication channel 78, 79 can be carried over respective dedicated optical fibres or multiplexed onto the same fibre (which could be the fibre used for the local links optically coupling Q-blocks in neighbouring nodes – for example, the MC communication channel can be implemented as intensity modulations of the herald signal 79, particularly where only parity information is being sent on this channel). More generally, the LLEC and MC communication channels can be combined into a single duplex classical communications channel.

In the embodiments described hereinafter, the LLEC communication channel 74, 75 is carried by the local link fibres and the MC communication channel 78, 79 is split into a left-to-right (“LR”) channel and a right-to-left (“RL”) channel, these two MC channels being carried by respective optical fibres distinct from that used for the local links. Messages on the LR MC channel and the RL MC channel are respectively referred to herein as ‘LR-MC’ messages and ‘RL-MC’ messages. It will be appreciated that this arrangement of channels and fibres is merely exemplary and other arrangements could alternatively be used.

It may be noted that the end nodes linked by a chain of quantum repeaters will each contain functionality for inter-working with the facing side (L or R) of the neighbouring quantum repeater. Thus, the left end node will include functionality similar to that of the R-side of a quantum repeater thereby enabling the left end node to inter-work with the L-side of the neighbouring repeater, and the right end node will include functionality similar to that of the L-side of a quantum repeater to enable the right end node to inter-work with the R-side of the neighbouring repeater.

With regard to entanglement parity, in the embodiments described below, rather than the parity of entanglements being standardized by qubit state flipping, at each quantum repeater LLE parity information is stored and subsequently combined with merge parity

information for passing on along cumulatively to an end node thereby to enable the latter to determine the parity of end-to-end entanglements.

In the following description of the quantum repeater embodiments, the same reference numerals are used for the main repeater components as are used in the generic diagram of Figure 7, it being understood that the specific implementations of these components will generally differ.

10 “Synchronized” Quantum Repeater Embodiments

The first and second quantum repeater embodiments, illustrated in Figures 10 and 23 respectively, both operate on a “Synchronized” basis to build an end-to-end (E2E) entanglement between qubits in left and right end nodes of a chain of nodes whose intermediate nodes are quantum repeaters. Building an E2E entanglement on the “Synchronized” basis involves all repeaters in a chain of quantum repeaters being time synchronized (for example, by timing taken from a GPS constellation or by phase lock looping clocks in the repeaters by signals sent between them). Also required are reliable LLE creation subsystems between repeaters (in this context, ‘reliable’ means LLEs can be created with a high probability of being available in successive operating periods).

20 During each successive synchronized operating cycle of the quantum repeaters, each repeater creates a right-side LLE (or brings into play a previously-created right-side LLE) whereby LLEs are operatively available between all repeaters, and then all the repeaters carry out a merge operation – provided all the merges are successful, the effect is to create an E2E entanglement.

Figure 8, which uses the same notation as Figure 3, depicts four successive operating cycles Φ for a chain of five optically-coupled nodes comprising left and right end nodes 81, 82 and three quantum repeaters 80 (QR₂, QR₃, QR₄); the optical fibres coupling the nodes are omitted for clarity. The four operating cycles are labelled Φ_i to Φ_{i+3} respectively. For the first three cycles $\Phi_i, \Phi_{i+1}, \Phi_{i+2}$, LLEs are successfully created between each pair of neighbour nodes (arched arrows 85), these LLEs then being merged to form an E2E

entanglement (thick arched arrow 87). However, for the fourth cycle Φ_{i+3} , although LLEs are successfully created between the first - second, second – third, and fourth - fifth node pairs, there is a failure to create an LLE between the third and fourth nodes (crossed-through dashed arrow 88) whereby the merging of LLEs in operating cycle Φ_{i+3} fails to produce an E2E entanglement.

The operating cycles Φ referred to above are strictly operating cycles of the node-chain as a whole (referred to below as ‘E2E operating cycles’) with each cycle encompassing the merging of LLEs established between every pair of neighbour nodes to create one E2E entanglement. Each E2E operating cycle also corresponds time-wise to a top-level operating cycle of each repeater, each such top-level repeater operating cycle generally encompassing one merge operation. Where there exists only a single LLE creation subsystem between each pair of neighbour nodes, the operating cycle of each such LLE creation subsystem (the ‘LLE creation subsystem operating cycle’) will also correspond time-wise to the E2E operating cycle (and the top-level repeater operating cycle) with LLEs being created between pairs of neighbour nodes at the same rate as E2E entanglements.

However, as will be described hereinafter in relation to the second quantum repeater embodiment, it is possible to provide multiple LLE creation subsystems between each pair of neighbour nodes to increase the rate of LLE creation between nodes and thereby enable E2E entanglements to be created at a correspondingly increased rate effectively shortening the E2E operating cycle (and the top-level repeater operating cycle). In this case, the LLE creation subsystem operating cycle will typically be longer than (and usually an integer multiple of) the E2E operating cycle (and the top-level repeater operating cycle). Quantum repeaters arranged to operate in this manner carry out LLE creation according to a timing determined by the LLE creation subsystem operating cycle but effect merge operations according to a timing determined by the E2E operating cycle. The symbol Φ is always used herein to refer to the E2E operating cycle.

30

It may be noted that the *overall* (as opposed to top-level) operating cycle of a quantum repeater, that is, a cycle of operation starting and ending in substantially the same state for

all components of the repeater, will typically correspond time-wise to the LLE creation subsystem operating cycle and may encompass multiple E2E operating cycles; however, the following description will generally only be couched in terms of E2E and LLE creation subsystem operating cycles without specific reference to the overall operating cycle of a
 5 quantum repeater.

Before describing the first and second quantum repeater embodiments in detail, consideration will be given to a form of LLE creation subsystem suitable for use in quantum repeaters operating on the “synchronized” basis, that is, one that can create LLEs
 10 with a high probability of success in successive operating periods. Of course, with long enough operating periods for multiple firings and/or favourable operating conditions (such as a short distance between nodes), even a simple LLE creation subsystem such as depicted in Figure 5 (or multiple paralleled subsystems of that form) can create LLEs with high probability and would be usable in quantum repeaters operating on the
 15 “synchronized” basis in such an environment. However, for multi-kilometre inter-node distances and operating periods of the order of 10^{-6} s, the simple LLE creation subsystem depicted in Figure 5 is unlikely to be adequate whereas the LLE creation subsystem now to be described with reference to Figure 9 offers much higher reliability.

20 “Firing Squad” LLE Creation Subsystem

Figure 9 depicts a “firing squad” form of LLE creation subsystem 90 formed between two nodes 91 and 92 that are optically coupled by local link fibre 95.

The node 91 comprises an LLE control unit 910, and quantum physical hardware formed
 25 by f Q-blocks 93 (with respective IDs 1 to f) that have Capture interaction functionality, and an optical merge unit 96. The Q-blocks 93 (herein “fusilier” Q-blocks) collectively form a “firing squad” 97. The node 92 comprises an LLE control unit 920, and quantum physical hardware formed by a single Q-block 94 with Transfer interaction functionality. The fusilier Q-blocks 93 of the firing squad 97 of node 91 are optically coupled through
 30 the optical merge unit 96 and the local link optical fibre 95 to the single target Q-block 94 of node 92. Thus, as can be seen, all the Q-blocks 93 of the firing squad 97 are aimed to fire at the same target Q-block 94.

When the LLE control unit 910 of node 91 outputs a Fire signal to its quantum physical hardware to trigger an LLE creation attempt, the fusilier Q-blocks 93 of the firing squad 97 are sequentially fired and the emitted light fields pass through the merge unit 96 and onto
5 the fibre 95 as a light-field train 98. It may be noted that there will be an orderly known relationship between the fusilier Q-block IDs and the order in which the light fields appear in the train. Rather than each light field being preceded by its own herald, a single herald 99 preferably precedes the light-field train 98 to warn the target Q-block 94 of the imminent arrival of the train 98, this herald 99 being generated by emitter 990 in response
10 to the Fire signal and in advance of the firing of the fusilier Q-blocks 93.

As each light field arrives in sequence at the target Q-block 94 of node 92, the shutter of the target Q-block is briefly opened to allow the light field to pass through the qubit of the target Q-block to potentially interact with the qubit, the light field thereafter being
15 measured to determine whether an entanglement has been created. If no entanglement has been created, the qubit of target Q-block 94 is reset and the shutter is opened again at a timing appropriate to let through the next light field of the train 98. However, if an entanglement has been created by passage of a light field of train 98, the shutter of the target Q-block is kept shut and no more light fields from the train 98 are allowed to
20 interact with the qubit of target Q-block 94. The measurement-result dependent control of the Q-block shutter is logically part of the LLE control unit 920 associated with the target Q-block 94 though, in practice, this control may be best performed by low-level control elements integrated with the quantum physical hardware.

25 It will be appreciated that the spacing of the light fields in the train 98 should be such as to allow sufficient time for a determination to be made as to whether or not a light field has successfully entangled the target qubit, for the target qubit to be reset, and for the Q-block shutter to be opened, before the next light field arrives.

30 In fact, rather than using an explicit shutter to prevent disruptive interaction with the target qubit of light fields subsequent to the one responsible for entangling the target qubit, it is possible to achieve the same effect by transferring the qubit state from electron spin to

nuclear spin immediately following entanglement whereby the passage of subsequent light fields does not disturb the captured entangled state (the target qubit having been stabilized against light-field interaction). It may still be appropriate to provide a shutter to exclude extraneous light input prior to entanglement but as the qubit is not set into its prepared state until the herald is detected, such a shutter can generally be omitted.

The LLE control unit 920 is also responsible for identifying which light field of the train successfully entangled the target qubit of Q-block 94 and thereby permit identification of the fusilier Q-block 93 (and thus the qubit) entangled with the target Q-block qubit (as already noted, there is a known relationship between the fusilier Q-block IDs and the order in which the light fields appear in the train). For example, the light fields admitted to the target Q-block may simply be counted and this number passed back by the LLE control unit 920 to the node 91 in a 'success' form of a message 930, the LLE control unit 910 of node 91 performing any needed conversion of this number to the ID number of the successful fusilier Q-block 93 before storing the latter in a register 195 for later reference (alternatively, the fusilier ID may be passed on immediately). Of course, if none of the light fields of train 98 is successful in creating an entanglement, a 'fail' form of message 930 is returned and a corresponding indication stored in register 195.

With regard to the parity information contained in the measurement result in respect of the successful entanglement of the target qubit, this parity information is passed to the control unit 920 which may either store it for later use (for example in a register 196) or pass it on, for example to node 91 in the message 930.

Rather than sequentially firing the fusilier Q-blocks 93 of node 91 to produce the train of light fields 98, an equivalent result can be achieved by firing them all together but using different lengths of fibre to connect each fusilier Q-block to the optical merge unit 96, thereby introducing different delays and creating the light-field train 98.

The number of fusilier Q-blocks 93 in the firing squad 97 is preferably chosen to give a very high probability of successfully entangling target Q-block 94 at each firing of the firing squad, for example 99% or greater. More particularly, if the probability of

successfully creating an entanglement with a single firing of a single fusilier Q-block is s , then the probability of success for a firing squad of f fusilier Q-blocks will be:

$$\text{Firing squad success probability} = 1 - (1-s)^f$$

whereby for $s = 0.25$, 16 fusilier Q-blocks will give a 99% success rate and 32 fusilier Q-blocks a 99.99% success rate. Typically one would start with a desired probability $P_{success}$ of successfully entangling the target qubit with a single firing (i.e. a single light-field train) and then determine the required number f of fusilier qubits according to the inequality:

$$P_{success} \leq 1 - (1-s)^f$$

- 10 The time interval between adjacent light fields in the train 98 is advantageously kept as small as possible consistent with giving enough time for the earlier light field to be measured, the target qubit reset and its shutter opened before the later light field arrives. By way of example, the light fields are spaced by 1 – 10 nanoseconds.
- 15 It will be appreciated that with the Figure 9 form of LLE creation sub-system 90, because there is only one target Q-block 94, the firing squad 97 cannot in practice be re-triggered until the whole sub-system is freed up by the most recently created entanglement being consumed or timing out (or otherwise ceasing to be of use). The minimum time between triggering of the firing squad 97 is thus the round trip time between the nodes (that is, the
- 20 minimum time for the light train 98 to reach node 92 and for message 930 to be returned to node 91) plus a time for consuming the entanglement (for example, in a merge operation).

First “Synchronized” Quantum Repeater Embodiment (Figure 10)

The first “Synchronized” quantum repeater embodiment will now be described with reference to Figure 10, it being understood that the quantum repeater operates in the context of being an intermediate node in a chain of N nodes (such as depicted in Figure 8 for $N=5$) between the left and right end nodes of which E2E entanglements are to be created.

- 30 The general form of the Figure 10 quantum repeater corresponds to that shown in Figure 7, and comprises: quantum physical hardware 60; left and right local link fibres 62, 63, interfacing via optical interfaces 76L, 76R; L-side and R-side LLE control units 72, 73 and

merge control unit 77.

The quantum physical hardware 60 (depicted in the generalized manner explained with respect to Figure 6) comprises:

- 5 a L-side (left-side) target Q-block 94 that forms part of a left LLE creation subsystem 71L;
- multiple R-side fusilier Q-blocks 93 that form the firing squad 97 of a right LLE creation subsystem 71R; and
- an optical fabric 61 coupled to left and right local link fibres 62, 63.

10

The left and right LLE creation subsystems 71L, 71R are substantially of the form illustrated in Figure 9 for LLE creation subsystem 90. As graphically depicted in Figure 11, the left LLE creation subsystem 71L comprises:

- (a) in repeater 100, the above-mentioned L-side elements of the quantum physical hardware 60 (in particular, the target Q-block 94, depicted in Figure 11 by a box with the letters 'Tg' inside), and the left LLE (L-LLE) control unit 72 with parity register 196;
- 15 (b) the left local link fibre 62; and
- (c) in a left neighbour node 110L, a firing squad of fusilier Q-blocks 93 (depicted in Figure 11 by a box with the letters 'FS' inside) and its associated optical fabric and LLE control unit.
- 20

The right LLE creation subsystem 71R comprises:

- (a) in repeater 100, the above-mentioned R-side elements of the quantum physical hardware 60 (in particular, the firing squad 97 depicted in Figure 11 as box 'FS'), and the right LLE (R-LLE) control unit 73 with fusilier ID register 195;
- 25 (b) the right local link fibre 63; and
- (c) in a right neighbour node 110R, a target Q-block (box 'Tg') and its associated optical fabric and LLE control unit.

30 With this arrangement of complementary firing squad and target portions of a Figure 9 LLE creation subsystem 90, multiple quantum repeaters 100 can be optically coupled in series such as to form one LLE creation subsystem between every pairing of neighbouring

repeaters as is illustrated in Figure 12 for quantum repeaters $j-1$, j , $j+1$ (the quantum repeater j forming an LLE creation subsystem 71L with its left neighbour repeater $j-1$ and an LLE creation subsystem 71R with its right neighbour repeater $j+1$).

- 5 As only a single LLE creation subsystem exists between each pair of neighbour nodes, the E2E operating cycle Φ and the top-level repeater operating cycle both correspond time-wise to the LLE creation subsystem operating cycle. Therefore, for convenience in describing the Figure 10 embodiment, the term ‘operating cycle’ is used without qualification unless such qualification (e.g. ‘E2E’ or ‘LLE creation subsystem’) is helpful
 10 to understanding; furthermore, although the symbol Φ was specifically designated above as referring to the E2E operating cycle, for the Figure 10 embodiment the symbol Φ can also be understood as an indicator of LLE creation subsystem operating cycle.

The optical fabric 61 of the quantum repeater 100, as well as coupling the L-side and R-side Q-blocks 94, 93 to the left and right local link fibres 62, 63 respectively for LLE
 15 creation, also provides for the selective optical coupling of the L-side target Q-block 94 to a selected one of the R-side fusilier Q-blocks 93 for the purpose of effecting a local merge operation on the qubits of these Q-blocks.

- 20 During LLE creation, the quantum physical hardware 60 receives firing control signals from the R-LLE control unit 73 for controlling the R-side elements (in particular, the triggering of the firing squad 97), and outputs result signals (success/failure; parity; fusilier-identifying information) from the L-side target Q-block 94 to the L-LLE control unit 72. For a local merge operation, the quantum physical hardware 60 receives merge
 25 control signals from a merge control unit 77 (these signals selecting the fusilier Q-block 93 that is to participate in the merge, and triggering the merge itself), and outputs back to the unit 77 results signal (success/failure; parity) regarding the outcome of the merge operation.

- 30 Figures 13A and 13B illustrated two possible implementations of the optical fabric 61 depending on the nature of the Q-blocks 93 and 94.

The Figure 13A optical fabric implementation is applicable to the case where the fusilier and target Q-blocks 93, 94 are universal Q-blocks 44 (c.f. Figure 4). In this case, the left local link fibre 62 interfaces directly with the optical input of the target universal Q-block 94, and the optical output of this universal Q-block is optically coupled to an intermediate optical fibre 131. An active optical switch 132 interfaces the intermediate fibre 131 with the inputs of the fusilier universal Q-blocks 93 and a passive optical merge unit 133 puts the outputs of the fusilier Q-blocks 93 onto the right local link fibre 63. During LLE creation operation, the target Q-block 94 is set up for Transfer interaction and light fields coming in over the left link fibre 62 are fed to the target Q-block; the fusilier Q-blocks 93 are set up for Capture interaction and the optical merge unit 133 couples the fusilier Q-blocks 93 to the right local link fibre 63. For a merge operation, the target Q-block 94 is set up for Capture interaction and the fusilier Q-block involved in the merge is set up for Transfer interaction (the fusilier Q-block concerned will have been indicated in the merge set-up signals fed to the quantum physical hardware 60); the optical switch 132 is also set by the merge set up signals to optically couple the target Q block 94 to the fusilier Q-blocks 93 involved in the merge.

The Figure 13B optical fabric implementation is applicable to the case where the target Q-block 94 is a Transfer Q-block 42 (c.f. Figure 4) and the fusilier Q-blocks 93 are Capture Q-blocks 40. In this case, a passive optical merge unit 135 puts the outputs of the fusilier Capture Q-blocks 94 onto a single fibre which is then switched by an active optical switch 136 either to the right local link fibre 63 or to a loop-back optical fibre 137. A passive optical merge unit 134 fronts the target Transfer Q-block 93, the optical merge unit 134 being coupled on its input side to the left local link fibre 62 and the loop-back optical fibre 137. For an LLE creation operation, the optical switch 135 is set to feed the light fields output by the fusilier Capture Q-blocks 93 to the right local link fibre 63. For a merge operation, the optical switch 135 is set to feed the light field output by a selected one of the fusilier Capture Q-blocks 93 to the loop-back fibre 137 (the Q-block concerned will have been indicated in the merge set-up signals fed to the quantum physical hardware 60).

30

Returning to a consideration of Figure 10, the left LLE control unit 72 associated with the L-side target Q-block 94 of LLE creation subsystem 71L, communicates with the firing-

squad-associated LLE control unit of the same LLE creation subsystem (this control unit being in the left neighbour node) via left LLEC channel 74. In the present example embodiment, the left LLEC channel 74 is imposed on the left local link fibre 62 via optical interface 76L and used to pass LLE creation “success / failure” messages (the message 930 of Figure 9 with fusilier ID being included as appropriate) from the L-LLE control unit 72 to the LLE control unit of the left neighbour node.

Similarly, the right LLE control unit 73 associated with the R-side fusilier Q-blocks 93 of the firing squad 97 of LLE creation subsystem 71R, communicates with the target-associated LLE control unit of the same LLE creation subsystem (this control unit being in the right neighbour node) via right LLEC channel 75. The right LLEC channel 75 is imposed on the right local link fibre 63 via optical interface 76R and used to pass, to the R-LLE control unit 73, LLE creation “success / failure” messages (the message 930 of Figure 9 with fusilier ID as appropriate) from the LLE control unit of the right neighbour node.

Merge control is effected by merge control (MC) unit 77 which, as well as interfacing with the quantum physical hardware to initiate a merge operation and receive back result signals, is arranged to exchange various signals with the L-LLE control unit 72 and R-LLE control unit 73 and to communicate with the merge control units of other nodes by messages sent over a left-to-right merge control (LR-MC) channel carried on left and right optical fibres 78LR and 79LR respectively, and a right-to-left merge control (RL-MC) channel carried on left and right optical fibres 78RL and 79RL respectively. Transmit (Tx) and receive (Rx) interfaces 101, 102 couple the fibres 78, 79 to the MC unit 77.

As will be more fully explained below, the LR-MC channel carries cumulative failure and parity messages each concerning a respective E2E operating cycle Φ , and the RL-MC channel carries cumulative failure messages each concerning a respective E2E operating cycle. As well as controlling the execution of a local merge operation, the merge control unit 77 is responsible for contributing its own parity and failure information to the appropriate cumulative messages on the LR-MC and RL-MC channels.

To permit the synchronized operation of a chain of geographically separated quantum

repeaters of the Figure 10 form, the repeater 100 includes a clock 103 that can be synchronized with similar clocks provided in other such repeaters. One way to provide synchronized clocks in a set of quantum repeaters 100 would be to transmit a common timing signal to all repeaters from a source a known distance from each repeater thereby enabling each repeater to make its own adjustments to the timing signal to take account of propagation delays; in fact, the GPS system provides a very convenient synchronization system of this type and the clock 103 could be based on such a system. One drawback to using GPS is that its signals are not readily available underwater and so could not be used by quantum repeaters placed in submarine fibre optic cables.

10

In the present embodiment, therefore, a different approach is taken where the repeater clocks are synchronized based on a synchronization signal cyclically transmitted from one end node of the chain including the quantum repeaters. Assuming an initial calibration has been carried out to determine the propagation delay between each pair of neighbour nodes, a node receiving the synchronization signal from its neighbour can determine, on the basis that the signal was transmitted by its neighbour at a synchronization-cycle start at the originating end node, the point in time corresponding to the next synchronization-cycle start; the node is then arranged to transmit on the synchronization signal at that point in time to its neighbour node away from the originating end node. The overall result, after a number of synchronization cycles sufficient to allow the synchronization signal to have propagated fully down the chain of nodes, is that every node transmits a synchronization signal at each synchronization-cycle start time at the originating end node. As each node can determine the synchronization-cycle start times, these are used as common timing points for the internal operation of the nodes. Preferably, this approach to synchronization is implemented by synchronizing operation of the internal clock 103 of each node to the synchronization start times using the received synchronization signals and a phase-locked loop arrangement that takes account of the propagation delay of the received synchronization signal from the sending neighbour node; internal operation of each node is then effected relative to timing signals output by the clock 103 whereby if, for any reason, a synchronization signal should be missed, internal operation of the node is unaffected.

30

As will be more fully described below, in the present embodiment each operating cycle Φ

starts with the firing squad 97 of every quantum repeater 100 being triggered to fire a light-field train 98 (see Figure 9) preceded by a herald 99, towards the neighbour node containing the target Q-block 94 of the LLE creation subsystem concerned. It is therefore convenient (though not essential) to use the herald 99 as the synchronization signal for effecting clock synchronization; thus, when the herald 99 is received at the target end of the LLE creation subsystem concerned, it is extracted and passed to the internal clock 103 of the node to keep the clock in synchronization with the clocks 103 of the other nodes (for example, using the phase-locked loop arrangement mentioned above). In the present embodiment, the synchronization cycle period therefore corresponds to the LLE creation subsystem operating cycle period, though this need not be the case. Furthermore, the period of clock 103 will be taken to be the same as the LLE creation subsystem operating cycle period, though again this need not be the case (for example, the clock could be run with a period of one quarter of the LLE creation subsystem operating cycle period).

The operation of the LLE control units 72, 72 (and therefore of the quantum physical hardware 60) and of the merge control unit 77 are timed relative to timing signals fed to these units from clock 103. In particular, the clock 103 is arranged to output a cycle-start timing signal at the start of each operating period, this cycle-start signal triggering the generation of a Fire signal by the R-LLE control unit 73 to trigger the firing squad 97 and thus the sending of the herald 99 followed by the light-field train 98.

For the Figure 10 embodiment, the operating cycle period is necessarily at least as long as the longest light-speed round-trip time between pairs of neighbour nodes – this is because during an operating cycle, an LLE must be created between each pair of neighbour nodes by sending a light-field train between the nodes, and the identity of the fusilier Q-block involved in the resulting LLE must be returned to the node that sent the train. The operating cycle period in fact needs to be greater than the longest light-speed round-trip time by an amount sufficient to accommodate certain other operating-cycle operations such as the merging of entanglements and to allow for a degree of synchronization error between clocks in different nodes. For ease of explanation of the operation of the Figure 10 embodiment, it is assumed that the operating cycle period is four times the longest inter-node light-speed single trip time τ (also referred to as the inter-node “transit time”);

as already indicated, shorter operating cycle periods are possible for the Figure 10 embodiment down to just over twice the longest inter-node transit time.

Figure 14 graphically depicts the timing notation used below in describing the operation of the Figure 10 quantum repeater. The i th operating cycle is designated Φ_i and extends between times t_i and t_{i+1} . For convenience of explanation, each operating cycle is divided into four equal quarters each of duration P (and since in this example the operating cycle period is set to be four times the longest inter-node transit time τ , the quarter periods of the operating cycle have a duration P that corresponds to the longest inter-node transit time τ). Individual operating-cycle quarters are referenced by the operating cycle designation followed by the quarter number 1, 2, 3 or 4 in curly brackets – thus the second quarter of operating cycle Φ_i is referenced as $\Phi_i\{2\}$.

Operation of the Figure 10 quantum repeater over the course of an operating cycle Φ_i will now be described with reference to Figure 15 which depicts what is happening in an example LLE creation subsystem, LLE CSS 150, formed by and between two neighbour quantum repeaters 100 in a repeater chain, midway through each quarter of the operating cycle. The LLE creation subsystem, LLE CSS, 150 comprises firing squad (FS) 97 and target Q-block (Tg) 94 in respective quantum repeaters 100 that are spaced by a distance corresponding to the longest transit time τ in the node chain (this time here being the same as the clock quarter period P).

Considering operation in each operating-cycle quarter in turn:

25 *First operating-cycle quarter $\Phi_i\{1\}$*

At the start of operating cycle Φ_i the firing squad (FS) 97 of LLE CSS 150 is triggered whereby midway through the first operating-cycle quarter $\Phi_i\{1\}$, the herald 99 emitted by the firing squad is halfway toward the target Q-block (Tg) 94 followed by the light-field train 98. It should be noted that Figure 15 exaggerates the spatial distribution of the herald and light-field train relative to what would actually be the case for typical inter-light-field time intervals and typical inter-node distances. In any event, it is here assumed that the firing squad 97 will have fully

sent out the herald and light-field train within the duration P of the first operating-cycle quarter $\Phi_i\{1\}$.

Second operating-cycle quarter $\Phi_i\{2\}$

5 The light-field train begins to arrive at the target Q-block (Tg) at the start of the second operating-cycle quarter $\Phi_i\{2\}$. By midway through this quarter all or most of the train 98 will have arrived node holding the target Q-block and an LLE will probably already have been created (indicated by dashed arrowed arc 151 in Figure 15).

10

Third operating-cycle quarter $\Phi_i\{3\}$

 By midway through the third operating-cycle quarter $\Phi_i\{3\}$, the light-field train has all arrived at the node holding the target Q-block (Tg) and an LLE will, with high probability (dependent on the factors discussed above in relation to Figure 9), have been created (indicated by solid arrowed arc 152 in Figure 15). The parity information measured at the target Q-block will have been stored in register 196 of the L-LLE control unit associated with the target Q-block, and a success message 930 (including a successful fusilier indicator) will be on its way back to the R-LLE control unit 73 associated with the firing squad (FS). By the end of the third quarter
15 at the latest, this control unit 73 will have stored the successful fusilier ID in its register 195 (see Figure 10). If LLE creation was unsuccessful, a fail message 930 would have been returned and a corresponding indication stored in register 195.
20

Fourth operating-cycle quarter $\Phi_i\{4\}$

25 Midway through the fourth operating-cycle quarter $\Phi_i\{4\}$, each and every quantum repeater of the repeater chain, including the two neighbour quantum repeaters 100 forming LLE CSS 150, effect a merge operation between its L-side qubit and the successful R-side fusilier qubit. More particularly and with reference to Figure 10, in each repeater 100 a merge firing functional block 104 operating in
30 dependence on cycle timing signals fed it from clock 103, retrieves the identity of the successful R-side fusilier Q-block from register 195 and incorporates this ID into the merge control signals fed to the quantum physical hardware 60 thereby to set up

and trigger a merge between the L-side qubit and the identified R-side fusilier qubit of the repeater. In Figure 15, the merges effected by the two quantum repeaters 10 forming LLE CSS 150 are indicated by the small block arrows labelled with 'M'.

5 Assuming LLEs were successfully created between all pairs of neighbour nodes and all merge operations are successful, the result is the creation of an E2E entanglement 153, regardless of the order in which the less extensive entanglements are merged by the individual merge operations effected by the quantum repeaters.

10 If any LLE has not been successfully created, this is preferably detected by the merge control unit 77 on reading the register 195 causing the merge control unit to skip the merge operation and proceed as if a failure indication had been returned in the merge results.

15 Although in the present example, the merge operation is initiated midway through the fourth operating-cycle quarter $\Phi_i\{4\}$ on the expectation that left and right LLEs will be present, it is in fact possible for the repeater to initiate the merge at any point in a time slot of the operating cycle when known L-side and R-side qubits are known to be (through measurement results), or at least can be expected to be (due to the elapsed time being sufficient for reliable LLE creation to have been effected), entangled leftwards and rightwards respectively - in the present example, this time slot corresponds to the fourth operating-cycle quarter $\Phi_i\{4\}$.

20 In each repeater, the merge result measurements are passed to the merge control unit 77 and stored. More particularly, the merge parity information is first combined by an exclusive OR operation with the L-side LLE parity information retrieved from the register 196 before being stored to FIFO parity queue 106. Merge success/failure information is stored to a FIFO failure queue 105 (hereinafter this success/failure information is taken to be, by way of example, a binary value with '1' representing 'failure'; this information is accordingly simply referred to as 'failure' information).
25 The FIFO queues 105 and 106 are thus advanced just after the midpoint of the fourth quarter of each operating cycle as failure and parity information relevant to that cycle is added whereby at the start of each operating cycle the most recent queue
30 entries relate to the immediately preceding operating cycle.

Although the foregoing describes how an E2E entanglement is created in the course of an

operating cycle, this is not the end of the story because at least cumulative parity information collected from all the quantum repeaters needs to be passed to at least one end node. In the present embodiment, cumulative parity information is passed to the right end node in LR-MC messages and cumulative failure information is passed to both end nodes in LR-MC and RL-MC messages to indicate whether an E2E entanglement has been successfully formed.

In general terms, at the start of an operating cycle Φ_{i+1} , the left end node sends out a LR-MC message to accumulate failure and parity information from each quantum repeater for the E2E entanglement just created in the preceding operating cycle Φ_i ; similarly, the right end node sends out a RL-MC message to accumulate failure information from each quantum repeater for the E2E entanglement just created in the operating cycle Φ_i . As each of these messages traverses each quantum repeater, the latter combines the cumulative information in the message with its own contribution from the relevant queue 105, 106 for the operating cycle Φ_i . More particularly, in the example implementation of the Figure 10 embodiment, local failure information from the failure queue 105 is combined into the cumulative failure information in the LR-MC and RL-MC messages through respective OR functions 109, 107 whereby any local failure sets the cumulative failure value to '1', that is, to 'failure'; local parity information is combined into the cumulative parity information in the LR-MC message through XOR function 108 (operating independently on the two parity bits) whereby any uneven parity will toggle the corresponding cumulative parity bit value. Of course, it is necessary to carefully control extraction of information from the queues 105, 106 to ensure that the extracted information corresponds to the same operating cycle as the cumulative information with which the extracted information is to be combined. An alternative approach would be for each repeater simply to store parity and failure information in memory along with an operating cycle indicator which could then be matched up to an operating cycle indicator accompanying the cumulative information carried by the LR-MC and RL-MC messages.

The LR-MC message for operating cycle Φ_i reaches the right end node at the same time as the RL-MC message for the same cycle reaches the left end node whereby both end nodes come to know at the same time the operative status (success/failure) of the E2E

entanglement created in the operating cycle Φ_i ; the right end node will also simultaneously know the parity information relevant to use of the E2E entanglement.

Figure 16 depicts a message flow diagram in respect of one example messaging arrangement for accumulating the failure and parity information in the LR-MC and RL-MC messages as they pass along a chain of five nodes (such as depicted in Figure 8) between the left and right end nodes of which E2E entanglements are to be created. For simplicity, the nodes are all spaced by the same distance corresponding to a transit time of τ (operation is not affected by shorter inter-node spacing other than by the earlier arrival of messages than depicted). Figure 16 covers the time period t_i to t_{i+5} (that is, the duration of five operating cycles Φ_i to Φ_{i+4}). The failure and parity information being accumulated is represented by boxes whose fill is specific to the operating cycle to which the information relates; at least partial information for operating cycles Φ_{i-3} to Φ_{i+3} is depicted. LR-MC message flows are shown in solid lines with the LR-MC message flow accumulating failure and parity information for operating cycle Φ_i being shown by the thick solid line 160. RL-MC message flows are shown in dashed lines with the RL-MC message flow accumulating failure information for operating cycle Φ_i being shown by the thick dashed line 161.

In the example messaging arrangement illustrated in Figure 16, each intermediate node (which will be a quantum repeater node) sends out one LR-MC cumulative failure and parity information message and one RL-MC cumulative failure information message at the start of each operating cycle, having combined its own local contribution with the cumulative information in LR-MC and RL-MC messages received during the preceding operating cycle (and temporarily stored in the receive interfaces 102 - see Figure 10). It should be noted that the LR-MC failure and parity message and RL-MC failure message being prepared and sent out at the same time by a repeater node will generally not relate to the same operating cycle as each other, the exception being where the messages are being processed by a middle node of a chain (such as node $n=3$ in Figure 16).

30

For the LR-MC message being sent out in the c^{th} operating cycle Φ_c by quantum repeater node j , the cumulative failure and parity information will relate to the $(c-j)^{\text{th}}$ operating

cycle, the relevant local failure and parity information being extracted from the queues 105, 106 by tapping the appropriate queue slots (as depicted in Figure 10, it being assumed that this information is extracted after the entry of the merge result information for the ($c-1$) cycle).

5

For the RL-MC message being sent out in the c^{th} operating cycle Φ_c by quantum repeater node j in a chain of N nodes, the cumulative failure information will relate to the ($c-1-(N-j)$)th operating cycle, the relevant local failure information being extracted from the queue 105 by tapping the appropriate queue slot (as depicted in Figure 10, again on the
10 assumption that this information is extracted after the entry of the merge result information for the ($c-1$) cycle).

Considering more specifically the flow of the LR-MC message accumulating failure and parity information for the operating cycle Φ_i (the thick solid line 160 in Figure 16), at the
15 start of the first quarter of operating cycle Φ_{i+1} , the left end node ($n=1$) sends a LR-MC message to its neighbour quantum repeater node ($n=2$) with one cumulative failure bit and two cumulative parity bits. As the left end node does not itself effect a merge operation, the cumulative failure bit has a value '0' unless no LLE was created between the left end node and its neighbour quantum repeater node (in which case the failure bit has a value
20 '1'). The cumulative parity bits are both set to indicate even parity. The LR-MC message sent by the left end node is received by its neighbour quantum repeater node ($n=2$) at the start of the second quarter of operating cycle Φ_{i+1} and the cumulative failure and parity bits are temporarily stored in receive interface 102 (as indicated by the vertical portion of line 160 at node $n=2$ during the second, third and fourth quarters of operating cycle Φ_{i+1}).

25 At the start of the first quarter of the operating cycle Φ_{i+2} , quantum repeater node ($n=2$) combines its local failure and parity information for operating cycle Φ_i with the cumulative failure and parity information temporarily stored in interface 102 and sends on the new cumulative failure and parity information in a LR-MC message to its right neighbour node ($n=3$) – and so on rightwards along the chain from node to node until the right end node is
30 reached in operating cycle Φ_{i+4} . The right end node completes the process by adding its own local parity information in respect of the left LLE set up for operating cycle Φ_i .

As regards the flow of the RL-MC message accumulating failure information for the operating cycle Φ_i (the thick dashed line 161 in Figure 16), at the start of the first quarter of operating cycle Φ_{i+1} , the right end node ($n=5$) sends a RL-MC message to its neighbour quantum repeater node ($n=4$) with one cumulative failure bit. This cumulative failure bit has a value '0' since the right end node does not itself effect a merge operation and any failure to create an LLE between the right end node and its neighbour quantum repeater node is picked up by the latter. The RL-MC message sent by the right end node is received by its neighbour quantum repeater node ($n=4$) at the start of the second quarter of operating cycle Φ_{i+1} and the cumulative failure bit is temporarily stored in receive interface 102 (as indicated by the vertical portion of line 161 at node $n=4$ during the second, third and fourth quarters of operating cycle Φ_{i+1} . At the start of the first quarter of the operating cycle Φ_{i+2} , quantum repeater node ($n=4$) combines its local failure information for operating cycle Φ_i with the cumulative failure information temporarily stored in interface 102 and sends on the new cumulative failure information in a RL-MC message to its left neighbour node ($n=3$) – and so on leftwards along the chain from node to node until the left end node is reached in operating cycle Φ_{i+4} . The left end node completes the process by adding its own local failure information in respect of the right LLE that should have been set up for operating cycle Φ_i .

Figure 17 depicts (in a manner similar to Figure 16 for the first example messaging arrangement) a message flow diagram in respect of a second example messaging arrangement for accumulating the failure and parity information in the LR-MC and RL-MC messages passing along a chain of five nodes all spaced by the same distance corresponding to a transit time of τ . Figure 17 covers the time period t_i to t_{i+3} (that is, the duration of three operating cycles Φ_i to Φ_{i+2}). The LR-MC message flow accumulating failure and parity information for operating cycle Φ_i is shown by the thick solid line 170 and the RL-MC message flow accumulating failure information for operating cycle Φ_i is shown by the thick dashed line 171.

The second example messaging arrangement speeds the transmission of the cumulative failure and parity information by cutting out most of the message holding time at the nodes that is present in the first example messaging arrangement (represented by the vertical

portions of the message flow line 160, 161 in Figure 16). This is done by having each node transmit on the cumulative failure and parity information at the start of the next cycle quarter following its receipt rather than waiting to the start of the next full cycle. With the inter-node transit time corresponding to the longest transit time τ and thus to the cycle quarter period P , the effect is a substantially continuous onward transmission of the cumulative failure and parity information (it being appreciated that in practice the cycle quarter period P would need to be slightly greater than τ to allow for processing at each node, in particular the incorporation of the local failure and parity information).

10 For the LR-MC message being sent out in the c^{th} operating cycle Φ_c by quantum repeater node j , the cumulative failure and parity information will relate to the $(c-1-\text{floor}[(j-1)/4])^{\text{th}}$ operating cycle and will be transmitted at the start of cycle quarter $(1+[(j-1) \bmod 4])$. The relevant local failure and parity information is extracted from the queues 105, 106 by tapping the appropriate queue slots (different to those depicted in Figure 10 but readily
15 determined by persons skilled in the art).

For the RL-MC message being sent out in the c^{th} operating cycle Φ_c by quantum repeater node j in a chain of N nodes, the cumulative failure information will relate to the $(c-1-\text{floor}[(N-j)/4])^{\text{th}}$ operating cycle and will be transmitted at the start of cycle quarter $(1+[(N-j) \bmod 4])$. The relevant local failure information is extracted from the queue 105 by
20 tapping the appropriate queue slot (again, different to that depicted in Figure 10 but readily determined by persons skilled in the art).

It may be noted that although in the Figure 17 example the end-to-end message
25 transmission time is one full operating cycle, this is merely an artefact of there being five nodes in the example chain.

With regard to the left and right end nodes between which the E2E entanglements are
30 created, these nodes are not themselves quantum repeaters though, of course, they comprise functionality for completing the LLE creation subsystems involving their respective neighbour quantum repeaters, and functionality for sending/receiving the LR-

MC cumulative failure and parity message and RL-MC cumulative failure messages as described above. In the present example, where the firing squads 97 fire left to right along the node chain, the left end node also provides the master clock and sends out synchronization signals in the form of the heralds 99 sent by its firing squad.

5

The left and right end nodes also serve a further function, namely to free up at the end of each operating cycle the entangled end-node LLE creation subsystem qubits between which an E2E has just been formed. This is done by providing each end node with an output buffer comprising multiple Q-blocks and shifting each newly created E2E entanglement across into qubits of the buffers pending their consumption by consumer applications associated with the end nodes. Of course, such buffering may not be required where the consumer applications are arranged to consume E2E entanglements as they become available at the end of each operating cycle and can tolerate the loss of such entanglements if not timely consumed.

10

Figures 18 and 19 depict example implementations of right and left end node respectively.

The right end node 180 shown in Figure 18 comprises:

- a target Q-block 94 and associated LLE control unit 920 of an LLE creation subsystem 181 formed with left neighbour quantum repeater node 182;
- a high-level right end node (REN) control unit 183 that includes a clock (not separately shown) synchronized by the herald signals of the LLE creation subsystem 181; the control unit interfaces with the RL-MC and LR-MC channel fibres and sends out a RL-MC cumulative failure message at the start of each cycle (as explained above) and receives LR-MC cumulative failure and parity messages;
- an output buffer 185 comprising multiple Q-blocks 186 into a selected one of which the end of an entanglement rooted in target Q-block 94 can be shifted (this is done under the control of REN control unit 183 at the end of the relevant operating cycle).

15

The right end node 180 also interfaces with a local E2E entanglement consumer application 184 (shown dashed).

Figure 18 depicts a particular optical fabric implementation that uses an optical merge unit 187 to couple the buffer Q-blocks 186 to the target Q-block 94. The buffer Q-blocks 186 have Capture interaction functionality and the selection, under the control of REN control unit 183, of a particular buffer Q-block is effected by selecting only its light-field emitter for energisation. The light field from the selected buffer Q-block 186 travels right to left (arrow 188) and is coupled by the optical merge unit 187 to the target Q-block 94 (this Q-block 94 already possesses the required Transfer interaction capability). As this operation involving Capture and Transfer interactions is carried out over a short distance, the probability of success is high.

The REN control unit 183 is responsible for keeping track of which buffer Q-blocks 186 are currently entangled and also of the operating cycle in which the E2E entanglement concerned was created (whereby to enable cumulative parity information received in LR-MC messages to be associated with the relevant buffer Q-block 186).

The left end node 190 shown in Figure 19 comprises:

- a firing squad 97 with fusilier Q-blocks 93, and associated LLE control unit 910 of an LLE creation subsystem 191 formed with right neighbour quantum repeater node 192;
- a high-level left end node (LEN) control unit 193 that includes the master clock (not separately shown); the control unit interfaces with the RL-MC and LR-MC channel fibres and sends out a LR-MC cumulative failure and parity message at the start of each cycle (as explained above) and receives RL-MC cumulative failure messages;
- an output buffer 195 comprising m Q-blocks 196 into a selected one of which the end of an entanglement rooted in a fusilier Q-block 93 can be shifted (this is done under the control of LEN control unit 193 at the end of the relevant operating cycle).

The right end node 190 also interfaces with a local E2E entanglement consumer application 194 (shown dashed).

Figure 19 depicts a particular optical fabric implementation for coupling a selected one of the fusilier Q-blocks 93 to a particular buffer Q-block 196. The depicted optical fabric implementation avoids the use of an $f \times m$ optical switch that would otherwise be required to interface the f fusilier Q-blocks 93 with the m Q-blocks of the output buffer 195, this
 5 being achieved through the provision of an intermediary Q-block 197..

More particularly, in the Figure 19 implementation, the f fusilier Q-blocks 93 are optically coupled through an optical merge unit 199 and local link fibre 1910 to the repeater node chain. The fusilier and buffer Q-blocks 93 and 196 all have Capture interaction
 10 functionality whereas the intermediary Q-block 197 has Transfer interaction capability. A 1×2 optical switch 1900 enables the output of the optical merge unit 199 to be switched between the local link fibre 1910 and a loopback fibre 1920 that feeds an input of an optical merge unit 198; the outputs of the buffer Q-blocks are also coupled as inputs to the optical merge unit 198. The output of the optical merge unit 198 is coupled to the
 15 intermediary Q-block 197. This arrangement permits any selectively-fired one of the fusilier Q-blocks 93 or any selectively-fired one of the output-buffer Q-blocks 196 to be coupled to the intermediary Q-block 197. As a result, the end of an E2E entanglement anchored in one of the fusilier Q-blocks 93 can be shifted across to the intermediary Q-block 197 and from there into a selected one of the output-buffer Q-blocks 196.

20

The LEN control unit 193 is responsible for controlling the selection of fusilier Q-block and buffer Q-block involved in the transfer of an E2E entanglement into the buffer 195, and for keeping track of which buffer Q-blocks 196 are currently entangled.

25 It will be appreciated that different optical fabric implementations are possible for the left and right end nodes to those illustrated in Figures 18 and 19; for example, to reverse the light-field direction of travel 188 in the right end node, an active optical switch could be used to optically couple the target Q-block 94 to a selected buffer Q-block 186 (in this case, the target Q-block 94 would need Capture interaction capability and the buffer Q-
 30 blocks 186 would need Transfer interaction capability).

It will further be appreciated that associated with the operation of moving of an E2E

entanglement into a buffer Q-block, will be one or more parity measurements. If a measured parity is even, no further action is needed as the parity of the E2E entanglement unchanged; however, if a measured parity is odd, then to keep the E2E entanglement the same, the buffer qubit concerned is flipped.

5

Various modifications, additional to those already alluded to above, can be made to the Figure 10 quantum repeater embodiment. For example:

Right-to-left LLE creation. As already indicated, the terms “left” and “right” are simply convenient labels for relative directions along the node chain. The Figure 10 embodiment could equally as well been described in terms of the light-field trains 98 passing from right to left in the LLE creation subsystems (in which case, for LLE creation, the repeater L-side comprises fusilier Q-blocks and the repeater R-side is a target Q-blocks).

Passing LLE Parity Information to Firing-Squad End of LLE Creation Subsystem. Rather than LLE parity information being held in register 196 of the LLE control unit 920 at the target end of each LLE creation subsystem, this parity information could be passed in message 930 to the LLE control unit 910 at the firing-squad end the LLE creation subsystems for storage in register 195. After the merge operation in the same cycle, this parity information would them be XORed with the merge parity information for storage in the parity queue 106.

Passing Cumulative Parity to Left End Node. Cumulative parity information could be passed to the left end *node* in the RL-MC messages rather than to the right end node in LR-MC messages.

Complimentary Repeater Varieties. A hybrid form of quantum repeater, with two complimentary *varieties*, is possible in which the direction of travel of the light-field train 98 during LLE creation, is opposite for the left and right sides of the repeater. Thus, as depicted in Figure 20, in one variety 200 of this hybrid repeater, light-field trains 98 are generated by the left and right side firing squads 97 of the repeater variety 200 and after passage through L and R fusilier Q-blocks respectively, are sent out over left and right local link fibres to the left and right neighbour nodes; in the other variety 205 of this hybrid repeater, light-field trains 98 are received by the left and right sides of the repeater variety 205 over left and right local link fibres

30

respectively from the left and right neighbour nodes, are passed through L and R target Q-blocks 94 respectively, and are then measured. It will be appreciated that in a chain of quantum repeaters of the foregoing hybrid form, it is necessary to alternate the two varieties of repeater 200, 205 in order to create LLE creation subsystems.

Modifications can also be made with a view to increasing the rate of successful E2E entanglement creation. Several such modifications are identified below (it being understood that these modifications can be used alone or in combination to increase the rate of E2E entanglement creation):

Shortening the operating cycle. In the implementation described above, the length of the operating cycle Φ is four times the duration of the longest inter-node transit time τ ; however, as already noted there is scope to substantially reduce the operating cycle duration to between two and three times τ .

Enhancing LLE creation success rate. An example modification of this nature is described below with reference to Figure 21.

Parallel operation of node chain segments. An example modification of this nature is described below with reference to Figure 22;

Parallel LLE Creation. An example modification of this nature is described below in relation to the second “Synchronized” quantum repeater embodiment and with reference to Figures 23 to 26.

Enhancing LLE creation success rate (Fig. 21)

Figure 21 shows a modified form of the Figure 9 LLE creation subsystem 90 in which more than one target Q-block 94 is provided. More particularly, in the Figure 21 LLE creation subsystem 210 the basic arrangement of the quantum physical hardware (firing squad 97 and optical merge unit 96) in node 91 is the same as for the Figure 9 subsystem; the LLE control unit 211 of the Figure 21 subsystem does, however, differ in certain respects from the control unit 910 of Figure 9 as will be explained below. The main difference between the Figure 9 and Figure 21 subsystems, is to be found in node 92 where the quantum physical hardware now comprises multiple (p in total) target Q-blocks 94 with respective IDs 1 to p , and an optical switch 213 for directing light fields received over

the local link 95 to a selected one of the target Q-blocks 94. The optical switch 213 is controlled by LLE control unit 212 of node 92 such that all the incoming light fields are directed by the optical switch 213 to the same target Q-block 94 until a successful entanglement is created whereupon the optical switch 213 is switched to pass the incoming
5 light fields to a new, available (un-entangled), target Q-block 94. The optical switch thus effectively performs the role of shuttering an entangled target qubit from subsequent light fields and thereby preventing interaction of these light fields with that qubit. Each successful entanglement is reported to the node 91 in a 'success' message 930 which may now also include (in addition to information permitting identification of the involved
10 fusilier Q-block 93 and possibly parity information) the ID of the target Q-block 94 concerned.

Of course, the control unit 212 must keep track of the availability status of each of the target Q-blocks 94 since the control unit 212 is tasked with ensuring that the optical switch
15 213 only passes the incoming light fields to a target Q-block with an un-entangled qubit. This availability status can be readily tracked by the control unit 212 using a status register 216 arranged to store a respective entry for each target Q-block 94. Each register entry not only records the availability of the corresponding target Q-block but is also used to record,
20 of a fusilier Q-block), related parity information unless this is passed back to node 91 instead.

Operating node 92 in this way ensures an efficient use of the light fields fired by the firing squad 97 as they are all used to attempt entanglement creation.

25

The control unit 211 of node 91 also includes a status register 215, this register being arranged to store a respective entry for each fusilier Q-block 93. Each register entry records the availability of the corresponding fusilier Q-block 93; a fusilier Q-block is 'unavailable' between when its qubit is entangled with the qubit of a target Q-block 94 (as
30 indicated by a message 930) and when the entanglement concerned is used up. (All fusilier Q-blocks 94 are, of course, effectively 'unavailable' for the round trip time between when the firing squad is triggered and a message is received back from node 92

since it is not known whether any particular fusilier Q-block is, or is about to become, involved in an entanglement; however, such ‘unavailability’ may be ignored since whether any particular fusilier Q-block has become entangled will be known before the next firing of the firing squad 97. Each entry of register 215 is also used to record, in the case where
 5 the corresponding Q-block 93 is unavailable because its qubit is entangled, and parity information where such information has been provided in the related message 930.

Where multiple LLEs are created by a single triggering of the firing squad 97, the one or more LLEs created over and above the one to be used in the merge operation to be effected
 10 in the same operating cycle, can be put to a number of uses. Thus, one, some or all of these excess LLEs can be kept in reserve (‘banked’) in a queue and so immediately available to become the LLE to be merged in a following operating cycle should the LLE creation subsystem 210 fail to create any LLE in that cycle. This, of course, requires the relevant Q-blocks 93, 94 to be kept unavailable for participation in LLE creation which can be readily
 15 achieved through reference to the status registers 215, 216. Also, the nodes sharing banked LLEs must use them in the same order (for example, the order in which they are reported in messages 930) otherwise a disjunction could occur in the line of merged LLEs intended to make up an E2E entanglement.

20 Excess LLEs can also be used in the process known as ‘purification’. Purification raises the fidelity of an entanglement by combining two entanglements, via local quantum operations and classical communication, into one higher-fidelity pair.

It should be noted that ‘banked’ LLEs have a limited lifetime even where qubit state has
 25 been transferred without delay from electron spin to nuclear spin; accordingly a track should be kept of the remaining lifetime of the qubits involved in banked LLEs with LLEs that include an expiring qubit being discarded.

Parallel operation of node chain segments (Fig. 22)

30 By splitting the chain of nodes into multiple segments each with its own pair of left and right end nodes, creation of extended entanglements can be effected in parallel (over respective segments); these E2E segment entanglements can then be merged to create the

final E2E entanglement.

One particular example arrangement of such segmentation is depicted in Figure 22. In this case, the ultimate end nodes between which it is desired to create an E2E entanglement are
5 left end node 221 and right end node 222. The chain of nodes (end nodes 221, 222 and intermediate repeater nodes) is divided into a first segment 223 and a second segment 224.

The end nodes of the first segment 223 are the left end node 221 and a sub-node 225 of a segment-spanning node 229, this sub-node 225 serving as a right end node for the first
10 segment. The end nodes of the second segment 224 are right end node 222 and a sub-node 226 of the segment-spanning node 229, this sub-node 226 serving as a left end node for the second segment.

The segment-spanning node 229 includes the master clock. The firing squads of the first
15 segment 223 fire their light-field trains right to left (as indicated by arrow 227) whereas the firing squads of the second segment 224 fire their light-field trains left to right (as indicated by arrow 228); in this way, the synchronization signals formed by the light-field train heralds are propagated outwards from the segment spanning node 229.

20 The first and second segments 223, 224 create E2E segment entanglements in parallel time-wise in coordinated segment operating cycles. In due course, in each operating cycle the segment spanning node 229 will receive cumulative failure messages in respect of both segments indicating the success or failure of E2E segment entanglement creation. Assuming both E2E segment entanglements have been successfully created, the segment-
25 spanning node 229 now merges these E2E segment entanglements to generate the desired E2E entanglement between nodes 221 and 222. The segment-spanning node 229 thus not only possesses end node functionality but also merge functionality.

30 Second “Synchronized” Quantum Repeater Embodiment (Figure 23)

The second “Synchronized” quantum repeater embodiment 230 is shown in Figure 23, it being understood that the quantum repeater 230 operates in the context of being an

intermediate node in a chain of nodes (such as depicted in Figure 8) between the left and right end nodes of which E2E entanglements are to be created.

The second quantum repeater embodiment 230 of Figure 23 is similar in form and operation to the first quantum repeater embodiment 100 of Figure 10 and the same references are used in Figure 23 as in Figure 10 for elements which are substantially the same. The main differences between the first and second embodiments is that in the second quantum repeater embodiment 230 the E2E operating cycle period is reduced by a factor of four, that is, to the previous quarter-cycle time P while the LLE creation subsystem operating cycle time remains at $4P$. This is achieved by having four LLE creation subsystems in parallel between each pair of neighbour nodes (rather than just one) with the operation of the four paralleled LLE creation subsystems being staggered relative to each other by P – as a result, LLEs are created between each pair of neighbour nodes substantially at a rate of one per time period P , thereby enabling a merging of LLEs to form an E2E entanglement once every P .

More particularly, and as depicted in Figure 23, quantum repeater 230 comprises:

- synchronized clock 103 (with clock cycle period P , for example);
- merge control unit 77, now arranged to trigger a merge operation every time period P ;
- the target Q-blocks $Tg1 - Tg4$ and associated L-LLE control units of four leftward LLE creation subsystems 71L1 to 71L4 (the firing squads FS of these LLE creation subsystems being provided by the left neighbour node 231L);
- a left LLE CCS coordinator 236 for coordinating the staggered cyclic operation of the four leftward LLE creation subsystems 71L1 to 71L4, and for interfacing the L-LLE control units of these LLE creation subsystems with the merge control unit 77;
- the firing squads FS1 – FS4 and associated R-LLE control units of four rightward LLE creation subsystems 71R1 to 71R4 (the target Q-blocks Tg of these LLE creation subsystems being provided by the right neighbour node 231R); and
- a right LLE CCS coordinator 237 for coordinating the staggered cyclic operation of the four rightward LLE creation subsystems 71R1 to 71R4, and for interfacing the R-LLE control units of these LLE creation subsystems with the merge control unit 77.

The four leftward LLE creation subsystems 71L1 to 71L4 share the same left local link fibre 62. The light trains from the different firing squads FS in the left neighbour node 231L are combined onto the fibre 62 via optical merge unit 232, these light trains being
 5 subsequently directed to the correct target Q-block Tg1 to Tg4 in the repeater 230 by optical switch 233 under the control of left LLECCS coordinator 236.

Similarly, the four rightward LLE creation subsystems 71R1 to 71R4 share the same right local link fibre 63. The staggering of the firing of the firing squads FS1 to FS 4 in the
 10 repeater 230 is controlled by the right LLE CCS coordinator 237 and the resultant light trains are combined onto the fibre 63 via optical merge unit 234; these light trains are subsequently directed to the correct target Q-block in the right neighbour node 231R by optical switch 235.

15 By way of example, Figure 24 shows, for each of the four rightward LLE creation subsystems 71R1 to 71R4, a respective depiction (similar that of Figure 15) of what is happening midway through each quarter period P of the LLE creation subsystem operating cycle. As already explained above, each such quarter period P now corresponds to a E2E operating cycle Φ and accordingly each quarter period P is identified by a corresponding
 20 E2E operating cycle Φ_i to Φ_{i+3} (see top row of Figure 24). In Figure 24, the firing squads FS1 to FS4 of LLE creation subsystems 71R1 to 71R4 are shown as being fired in successive E2E operating cycles Φ_i to Φ_{i+3} . Merge operations involving the qubits entangled through operation of the LLE creation subsystems 71R1 to 71R4, are shown as occurring in E2E operating cycles to Φ_{i+3} , Φ_i , Φ_{i+1} , and Φ_{i+2} respectively, these merge
 25 operations each being part of the distributed merging that creates an E2E entanglement in each E2E operating cycle (c.f. Figure 8).

With regard to the merge operations effected by the merge control unit 77, in the Figure 23 repeater 230, the merge control unit 77 interfaces with the LLE CSS coordinators 236 and
 30 237 rather than with the LLE control units of the LLE creation subsystems as was the case for the Figure 10 embodiment. The right LLE CSS coordinator 237 is responsible for providing to the merge control unit 77 the identity of the rightward LLE creation

subsystem and the fusilier qubit of that subsystem next to be merged, and the left LLE CSS coordinator 236 is responsible for providing to the merge control unit 77 parity information on the target qubit to be involved in the same merge from one of the leftward LLE creation subsystems (this being the leftward LLE creation subsystem whose cycle of operation is in step with the rightward LLE creation subsystem concerned, it being assumed that the merge control unit knows or can be made aware of the appropriate in-step association between the leftward and rightward LLE creation subsystems).

Figure 25 depicts for the Figure 23 repeater embodiment (in a manner similar to Figures 16 and 17 for the Figure 10 repeater embodiment) a message flow diagram in respect of an example messaging arrangement for accumulating failure and parity information in the LR-MC and RL-MC messages passing along a chain of five nodes all spaced by the same distance corresponding to a transit time of τ . It will be recalled that, in the present example case, the value of the time period P is set to be equal to this transit time of τ whereby, for the Figure 23 repeater embodiment, the transit time between nodes is equal to the cycle time of the E2E operating cycle Φ . Figure 25 covers the time period t_i to t_{i+5} (that is, the duration of three operating cycles Φ_i to Φ_{i+4}). The LR-MC message flow accumulating failure and parity information for operating cycle Φ_i is shown by the thick solid line 250 and the RL-MC message flow accumulating failure information for operating cycle Φ_i is shown by the thick dashed line 251.

As can be seen, each node transmits on cumulative failure and parity information at the start of the next E2E cycle following its receipt and the effect is a substantially continuous onward transmission of the cumulative failure and parity information (it being appreciated that in practice the period P would need to be slightly greater than τ to allow for processing at each node, in particular the incorporation of the local failure and parity information). This is similar to the message flow diagram shown in Figure 17 in respect of the second example messaging arrangement for the Figure 10 embodiment. However, because only one cumulative failure and parity message is sent per direction per E2E operating cycle, the relationship between current E2E cycle, node position, and the cumulative failure and parity information being sent out at the start of the current E2E cycle, is as in the first example messaging arrangement for the Figure 10 embodiment, that is:

- for the LR-MC message being sent out in the c^{th} E2E operating cycle Φ_c by quantum repeater node j , the cumulative failure and parity information will relate to the $(c-j)^{\text{th}}$ E2E operating cycle;
- for the RL-MC message being sent out in the c^{th} E2E operating cycle Φ_c by quantum
5 repeater node j in a chain of N nodes, the cumulative failure information will relate to the $(c-1-(N-j))^{\text{th}}$ E2E operating cycle.

10

It will be appreciated that many of the modifications discussed above in relation to the Figure 10 quantum repeater embodiment are equally applicable to the Figure 23 quantum repeater embodiment.

15 In addition, further modifications can be made to both the Figure 10 and Figure 23 quantum repeater embodiments. For example, the possibility described above with reference to Figure 21 of ‘banking’ LLEs makes it possible to decouple the LLE creation subsystem operating cycle from the E2E / repeater-top-level operating cycle provided the average on-going (though not necessarily continuous) LLE creation rate is substantially at
20 least as great as the frequency of the E2E operating cycle (thus, if a Figure 21 LLE creation subsystem formed between two repeaters results in an average creation rate of 1.5 LLEs per triggering of its firing squad, the LLE creation subsystem period can be longer than the E2E cycle period).

25 As already mentioned, the reliable creation of LLEs is not limited to the use of the “firing squad” form of LLE creation subsystem, other LLE creation subsystems being appropriate in certain cases. By way of example, high reliability can be achieved by using in parallel multiple LLE creation subsystems of the Figure 5 form though this would require multiple local link fibres.

30

With regard to the implementation of the LLE control units 72, 73 and the merge control unit 77 of both the Figure 10 and Figure 23 embodiments, as well as the coordinators 236, 237 of the Figure 23 embodiment, it will be appreciated that typically the described

functionality will be provided by a program-controlled processor or corresponding dedicated hardware. Furthermore, the functionality of the LLE control units, the coordinators where present, and the merge control unit may in practice be integrated, particularly in cases where the LLE control unit functionality is minimal. Of course, the
5 division of control functionality is to a degree arbitrary; however, LLE control functionality merits separation into the LLE control units because in certain repeater embodiments LLE creation is free-running, that is, decoupled from the timing of higher level operations such as merge control. Overlying the LLE control functionality is the control functionality associated with merge control and operating-cycle synchronisation –
10 this latter control functionality effectively provides top level control of the repeater and can be considered as being provided by a top-level control arrangement (in the described embodiments this is formed by the merge control unit and the clock).

Although in the foregoing description light fields have generally been described as being
15 sent over optical fibres both between nodes and between components of the quantum physical hardware of a repeater, it will be appreciated that light fields can be sent over any suitable optical channel whether guided (as with an optical waveguide) or unguided (straight line) and whether through free space or a physical medium. Thus, for example, the optical fabric of the quantum physical hardware of a repeater may comprise silicon
20 channels interfacing with a qubit provided by a nitrogen atom in a diamond lattice located within an optical cavity.

As already indicated, persons skilled in the art will understand how the Q-blocks can be physically implemented and relevant example implementation details can be found in the
25 following papers, herein incorporated by reference:

- “Fault-tolerant quantum repeaters with minimal physical resources, and implementations based on single photon emitters” L. Childress, J. M. Taylor, A. S. Sørensen, and M. D. Lukin; *Physics Review A* 72, 052330 (2005).
- “Fault-Tolerant Quantum Communication Based on Solid-State Photon Emitters”
30 L. Childress, J. M. Taylor, A. S. Sørensen, and M. D. Lukin *Physical Review Letters* 96, 070504 (2006).
- “Hybrid quantum repeater based on dispersive CQED interactions between matter

qubits and bright coherent light” T D Ladd, P van Loock, K Nemoto, W J Munro, and Y Yamamoto; *New Journal of Physics* 8 (2006) 184, Published 8 September 2006.

- “Hybrid Quantum Repeater Using Bright Coherent Light” P. van Loock, T. D. Ladd, K. Sanaka, F. Yamaguchi, Kae Nemoto, W. J. Munro, and Y. Yamamoto; *Physical Review Letters* 96, 240501 (2006).
- “Distributed Quantum Computation Based-on Small Quantum Registers” Liang Jiang, Jacob M. Taylor, Anders S. Sørensen, Mikhail D. Lukin; *Physics. Review. A* 76, 062323 (2007).

CLAIMS

1. A quantum repeater (100; 230) optically couplable to left and right neighbour nodes
5 through local-link optical channels (62, 63); the repeater comprising:
quantum physical hardware (60) providing left-side and right-side repeater portions (L,
R) respectively arranged to support left-side and right-side qubits for entanglement
with qubits in the left and right neighbour nodes respectively by light fields
transmitted over the local-link channels (62, 63) thereby to form respective local
10 link entanglements, herein “LLE”s; the quantum physical hardware (60) being
operable to merge two entanglements respectively involving a left-side and a right-
side qubit, by a local merge operation on these qubits;
left and right LLE control units (72, 73) for controlling the quantum physical hardware
(60) to effect on-going creation of left and right LLEs in cooperation with the left
15 and right neighbour nodes; and
a top-level control arrangement (77, 103) operative to control top-level cyclic
operation of the repeater in synchronism with its neighbour nodes, the top-level
control arrangement (77, 103) being arranged once every top-level operating cycle
to initiate a said local merge operation involving left-side and right-side qubits
20 when the latter are at least expected to be entangled leftwards and rightwards
respectively.
2. A quantum repeater (100, 230) according to claim 1, wherein the LLE control units (72,
73) are arranged to effect cyclic creation of said left and right LLEs in synchronism with
25 the top-level operating cycle.
3. A quantum repeater (100; 240) according to claim 1, wherein the left-side repeater
portion (L) and the right-side repeater portion (R) are complimentary in form; one of these
repeater portions (L, R) being operative to generate a light field, pass it through its qubit,
30 and then send the light field out over a local link channel; and the other repeater portion
(R, L) being operative to receive a light field over a local link channel, pass it through its
qubit and then measure the light field.

4. A quantum repeater (100; 230) according to claim 3, wherein:

one of the left-side and right-side repeater portions (L, R) comprises a plurality of fusilier Q-blocks (93) each arranged to support a fusilier qubit and to pass a light field through that qubit, and an optical fabric (61) for orderly coupling light fields that have passed through fusilier qubits, onto the corresponding local link channel (63); a corresponding one of the LLE control units (73) being arranged to control this repeater portion to cause the coordinated passing of respective light fields through the fusilier qubits whereby to produce an outgoing train (228) of closely-spaced light fields on the local link channel (63); and

the other of the left-side and right-side repeater portions (R, L) comprises a target Q-block (94) arranged to support a target qubit and to measure a light field passed through that qubit whereby to determine whether the target qubit has been successfully entangled, and an optical fabric (61) for coupling the corresponding local link channel (62) with the target Q-block (94) to enable light fields of an incoming train (228) of light fields received over the local link channel (62) from a neighbour node to pass through the target qubit and be measured; a corresponding one of the LLE control units (72) being arranged to control this repeater portion to allow a first light field of the train (228) to pass through and potentially interact with the target qubit and thereafter only to allow a next light field through and potentially interact with the target qubit upon the target Q-block (94) indicating that the preceding light field was unsuccessful in entangling the target qubit, this LLE control unit being responsive to the target Q-block (94) indicating that the target qubit has been successfully entangled to pass, to the neighbour node originating the train, information identifying the light field of the train which successfully entangled the target qubit whereby to permit identification of the fusilier qubit entangled with the target qubit.

5. A quantum repeater (100; 230) according to claim 4, wherein the number f of fusilier Q-blocks is such as to satisfy the inequality:

$$P_{success} \leq 1 - (1-s)^f$$

where:

s is the probability of successfully creating an entanglement with a single light field

for a predetermined operating environment; and

$P_{success}$ is a desired probability of successfully entangling the target qubit with a single light-field train, $P_{success}$ being selected to be at least 99%.

- 5 **6.** A quantum repeater (100; 230) according to claim 4, wherein the LLE control unit (73) associated with the repeater portion (R) including the fusilier Q blocks (93) is arranged to trigger the generation of a said outgoing train (228) of light fields once every top-level operating cycle, and wherein the top-level control arrangement (77) is arranged to trigger a merge operation in each top-level operating cycle at a timing sufficient to allow
- 10 information identifying, from amongst the light fields of the outgoing train (228) generated in that cycle, the light field which has successfully entangled a neighbour-node target qubit to be returned to the repeater from the neighbour node concerned.
7. A quantum repeater (100; 230) according to claim 1, wherein for each top-level
- 15 operating cycle, the top-level control arrangement (77) is arranged to store parity information based on:
- merge parity information in respect of the local merge operation carried out in that cycle; and
- parity information in respect of an LLE involving a said qubit subject of the local
- 20 merge operation effected in the cycle.
- 8.** A quantum repeater (100; 230) according to claim 7, wherein the top-level control arrangement (77) is arranged once per top-level cycle to receive from one neighbour node cumulative parity information relating to a predetermined previous top-level operating
- 25 cycle, to combine its stored parity information relating to that cycle with the received cumulative parity information to form updated cumulative parity information, and to send on the updated cumulative parity information to its other neighbour node.
- 9.** A system, comprising a chain of nodes, for creating an end-to-end entanglement (87)
- 30 between working qubits in left and right opposite end nodes (81, 82) of the chain, intermediate nodes (80) of the chain being formed by quantum repeaters (100; 230) according to any one of claims 1 to 8 with each quantum repeater being linked to its

neighbour nodes by local link optical channels (62, 63); the top-level control arrangements (103) of the repeaters being arranged to cause the synchronization of their top-level operating cycles whereby in each cycle the combined effect of the merge operations is to create a said end-to-end entanglement (87).

5

10. A system according to claim 9, wherein each end node (180, 190) includes an output buffer (185, 195) arranged to provide a qubit into which the end of an end-to-end entanglement that is anchored in a working qubit of the end node, can be transferred in order to free up that working qubit.

10

11. A system according to claim 9, wherein the top-level control arrangement of each repeater includes a synchronization arrangement (103), these synchronization arrangements being coordinated through a signal passed along said chain from one end node.

15

12. A system according to claim 9, wherein the top-level control arrangement of each repeater includes a synchronization arrangement (103), these synchronization arrangements being coordinated through signals independently received at each repeater from a common external source.

20

13. A method of creating an end-to-end entanglement (87) between qubits in first and second end nodes (81, 82) of a chain of optically-coupled nodes whose intermediate nodes (80) are quantum repeaters (100; 230), the method comprising:

creating, on an on-going basis, local link entanglements (85), herein "LLE"s, between qubits in each pair of neighbour nodes in said chain, the LLEs (85) being created through interaction of the qubits with light fields transmitted between the nodes; and

effecting synchronized top-level cyclic operation of the quantum repeaters (100; 230) with each repeater initiating, once every top-level operating cycle, a local merge operation between two qubits that are at least expected to be entangled with qubits in nodes disposed in opposite directions along the chain from the repeater.

30

14. A method according to claim 13, wherein LLEs (85) are cyclically created in synchronism with the top-level cyclic operation of the quantum repeaters (100; 230).

15. A method according to claim 14, wherein LLEs are created between at least one pair
5 of neighbour nodes (91, 92) coupled by an optical channel (95), by:

passing respective light fields through a plurality of fusilier qubits (93) in one node
(91) of said pair and into the optical channel (95), the generation and organization
of the light fields being such as to result in a train (98) of closely-spaced light
fields being transmitted along the optical channel (95);

10 receiving, at the second node (92) of said pair, light fields of said train (78) over the
optical channel (95) and while a target qubit (94) remains un-entangled, allowing
each light field to pass in turn through, and potentially interact with, the target
qubit (94), each light field thereafter being measured to determine whether the
target qubit has been entangled,

15 upon successful entanglement of the target qubit (94), inhibiting interaction of further
light fields of the train (98) with the target qubit and identifying which light field
successfully entangled the target qubit whereby to permit identification of the
fusilier qubit (93) entangled with the target qubit (94).

20 16. A method according to claim 13, wherein for each top-level operating cycle, each
quantum repeater (100; 230) stores parity information based on:

merge parity information in respect of the local merge operation carried out in that
cycle; and


25 parity information in respect of an LLE involving a said qubit subject of the local
merge operation effected in the cycle.

17. A method according to claim 16, wherein once per top-level cycle each quantum
repeater (100; 230):

30 receives from one neighbour node cumulative parity information relating to a
predetermined previous top-level operating cycle,

combines its stored parity information relating to that cycle with the received
cumulative parity information to form updated cumulative parity information, and

sends on the updated cumulative parity information to its other neighbour node.



Application No: GB0911274.9

Examiner: Ben James

Claims searched: 1-17

Date of search: 7 September 2010

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
A	-	CHINESE PHYSICS, VOL. 13, NO. 2, 2004, Wu Xiang-Sheng, Memory-based quantum repeater in quantum information communication, Pages 173 - 177
A	-	EP1865657 A1 (TOSHIBA KK) Whole document, see particularly paragraphs 25-65.
A	-	WO2007/021945 A2 (HEWLETT PACKARD DEVELOPMENT CO) Whole document.
A	-	WO2009/155486 A1 (TELCORDIA TECH INC) Whole document

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

G06N; H04B; H04L

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

ON-LINE: INSPEC, EPODOC, WPI, XPI3E

International Classification:

Subclass	Subgroup	Valid From
G06N	0099/00	01/01/2010
H04B	0010/30	01/01/2006
H04L	0009/08	01/01/2006