



(51) International Patent Classification:
H04J 14/02 (2006.01)

(21) International Application Number:
PCT/EP2016/056776

(22) International Filing Date:
29 March 2016 (29.03.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15275098.0 30 March 2015 (30.03.2015) EP

(71) Applicant: **BRITISH TELECOMMUNICATIONS
PUBLIC LIMITED COMPANY** [GB/GB]; 81 Newgate
Street, London EC1A 7AJ (GB).

(72) Inventors: **WRIGHT, Paul**; PP: C5A BT Centre, London
EC1A 7AJ (GB). **LORD, Andrew**; PP: C5A BT Centre,
London EC1A 7AJ (GB). **PARKER, Michael**; 67 Temple
Road Epsom, Surrey KT19 8EY (GB).

(74) Agent: **CARDUS, Alan**; Intellectual Property Department,
Ground Floor, BT Faraday Building, 1 Knighttrider Street,
London EC4V 5BT (GB).

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

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

Published:

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

(54) Title: OPTICAL NETWORK DESIGN AND ROUTING




FIGURE 2B

(57) Abstract: A method of allocating spectrum to a signal in a wavelength division multiplex network comprising a plurality of nodes, wherein the spectrum comprises a grid divided into a plurality of slots, comprising allocating to the signal a slot immediately adjacent to a second slot occupied by a second signal, wherein the signal and the second signal are both addressed to the same destination node, and wherein the slots allocated to the signal and the second signal form a spectrum block.

OPTICAL NETWORK DESIGN AND ROUTING

This invention relates to apparatus, systems and methods relating to optical data transmission and network design and routing, particularly but not exclusively in respect
5 of the design and operation of a telecommunications network by optimising the use of optical bandwidth resource.

The way optical networks are used is undergoing significant change, driven in part by the huge growth of traffic such as multimedia services and by the increased uncertainty in
10 predicting the sources of this traffic due to the ever changing models of content providers over the Internet. Sophisticated modulation schemes for higher bandwidth 100 Gb/s services and beyond are known and come into commercial use in optical networks of large and increasing link and node numbers. A bottleneck to widespread deployment of such schemes is the "fixed" wavelength grid approach specified by the International
15 Telecommunication Union (ITU), in which the relevant optical spectrum range in the C-band is divided into fixed-sized spectrum slots. Such conventional "fixed grid" WDM (wavelength divisional multiplexed) networks work on the concept of a fixed spectrum grid typically with a spacing of typically 50GHz between channels with 80 to 100 of these channels per fibre. In these networks, an individual signal serving a demand between
20 two nodes in the network has to keep within one of these channels or slots defined by guard bands, as otherwise the signal becomes notched and degraded by the wavelength filters when being split from its neighbouring signal. As a result of this restriction, advanced modulation formats allowing up to 100Gbit/s per 50GHz channel commercially and up to 200Gbit/s experimentally, have not to now been usefully deployed in a
25 widespread manner. This is because the spectral widths of such signals are wider than can be accommodated within the 50GHz fixed grid spacing, so the potential of additional increases in transmission speed cannot be realised.

As used herein, a "slot", "wavelength" or "channel" is defined as a wavelength or a
30 spectrum of wavelengths associated with a certain signal size. A "carrier" carries a "signal" or "demand" in the known fashion. As is also known, a connection between nodes is made by assigning spectral (i.e. wavelength) slots on the optical links comprising the path between source and destination.

35 A response to the problems posed by the decade-old ITU fixed grid approach is the

flexible grid or "flexgrid", which facilitates a developing optical networking paradigm known as EON (elastic optical networking). The EON technologies allow for radically different network design and operation methodologies that can increase the amount of traffic the network can carry compared to conventional WDM networks, but need
5 different processes to make them operate effectively to get the most out of such networks. In the flexgrid approach, the optical spectrum can be divided up flexibly in dependence on requirements, and elastic optical paths (i.e. paths with variable bit rates) can be generated. This allows for operational and functional flexibility in use of both the optical spectrum and transceivers, previously unavailable in fixed grid implementations.
10 In a flexgrid, the spectrum grid is divided into much finer slot widths, typically 12.5GHz or less, compared to the 50GHz in the fixed grid approach. Significantly, adjacent channels can be joined together to form arbitrary sized slots to carry signals of a variety of widths, allowing for signals ranging in size from an individual channel to that occupying the entire optical spectrum to be carried.

15 Representations of signals carried according to the fixed and flexible grid approaches are depicted in graphs shown in Figure 1, in which graph (a) is a depiction of the fixed grid approach, in which guard bands (2) partition adjoining optical channels (4) occupied by demands or wavelengths at a particular bit rate. The guard bands serve to separate
20 a demand from other demands going to other destinations and to protect the main signal as it passes through filters in the network so as to reduce the effects of passing through the optical nodes. Graph (b) illustrates the flexgrid approach used in an EON network, in which the demands (here shown to be of various spectral widths) are not constrained within a slot of pre-defined spectral size. As illustrated by demand (4b) in graph (b) of
25 Figure 1, a high bitrate demand with a spectral width exceeding fixed grid slot sizes can be accommodated. A "superchannel" (6) (depicted in graph (c) of Figure 1) for carrying demands which are too large to be handled by a single optical channel, can similarly be accommodated in an EON network. A superchannel comprises a grouping of multiple channels and is handled as a single entity, traversing the network for demultiplexing at
30 the receiver end. Specifically, they can be produced by a bandwidth variable transponder (BVT) enabling a number of carriers. The carriers are aggregated together at the transceiver to produce an optical signal of a size which depends on the level of traffic carried by the signal. So if more traffic needs to be carried by the BVT, additional carriers can be added and conversely, if traffic levels decrease, carriers can be disabled.
35 A general description of the use of BVTs in an EON network can be found in "Elastic

Optical Networking: A New Dawn for the Optical Layer?" by O. Gerstel, M. Jinno, A. Lord, SJB Yoo (IEEE Communications Magazine, February 2012). The operational flexibility of superchannels can usefully cope with growth in traffic levels in a network over time, and significant spectral savings can be gained over the fixed grid approach.

5 BVTs can be used in both fixed grid and flexgrid systems, although they are deployed to greater effect in flexgrid networks owing to the capacity of the latter to accommodate the greater spectral widths of superchannels.

10 In a conventional WDM network, a node receives signals from one or more other nodes in the network, typically via intermediate nodes. Currently, all these signals meant for the same destination node are each considered in isolation, and known routing techniques such as routing and spectrum assignment (RSA) algorithms allocate a route and spectrum to a signal regardless of its source and destination. As a result, signals originating from different nodes but destined for the same node can be allocated

15 spectrum or blocks of spectrum potentially anywhere across the entire C-Band even on the same link. Even between the same source/destination node pair, multiple signals (caused by the need to carry increasing levels of traffic across the network) are likely to use different parts of the spectrum and possibly different routes across the network. Conversely, signals destined for different nodes can be placed on neighbouring or

20 contiguous sections on the optical spectrum along a link or route. Because no filter can be entirely precise, guard bands between demands are essential to provide a kind of "padding" between signals to be split within the spectrum for transmission to different destinations, and also to allow an operational tolerance during the splitting process. Guard bands occupy part of the spectrum alongside the demands they separate, and so

25 as might be expected, the spectral requirements for guard bands increase as the number of discrete signals or number of nodes in the network increases.

It would be desirable for more of the precious spectrum resource to be given to carrying signals or demands instead of being occupied by components such as guard bands. By

30 making better optimised use of existing optical fibre assets by serving the maximum number of customer demands possible on any given fibre link or network, the need to expensively and disruptively install additional fibre and transceivers can be prevented or at least advantageously delayed.

The applicants' co-pending applications WO2014174234 and GB1317987.4 describe a method of routing signals in a WDM network which is based on the level of "network entropy" of an identified route (representing the extent to which a spectrum comprises non-contiguous used or unused sections within it), which are incorporated by reference
5 herein. US 5596722 also considers network routing entropy in the context of satellite communication system which is not based on WDM. None of these documents address the issue of saving spectrum for more productive allocation to demands or signals in preference to other components or requirements such as guard bands.

10 According to a first aspect of the invention, there is provided a method of allocating spectrum to a signal in a wavelength division multiplex network comprising a plurality of nodes, wherein the spectrum comprises a grid divided into a plurality of slots, comprising allocating to the signal a slot immediately adjacent to a second slot occupied by a second signal unseparated by a guard band, wherein the signal and the second signal
15 are both addressed to the same destination node, and wherein the slots allocated to the signal and the second signal form a spectrum block.

In applications of the invention, signals are allocated spectrum resources on a link in a way which groups them together in dependence on their destination nodes. This is
20 achieved by placing a new demand into a slot immediately adjacent to another demand which is addressed to the same destination node as the new demand. This can result a row of slots all occupied by signals addressed to the same destination. This allows for the grouping to form a block which can be routed over one or more links through the network as a single unit, doing away with the need for guard bands between the
25 individual signals within the block. As a result, precious spectral capacity can be conserved for use by signals instead. A signal addressed to a destination node which is different from the first is treated as a different unit from the first, and separated from the first block by a guard band for purposes of passing through a filter. The signal bound for another destination can comprise a single slot or a block comprising more than one slot.
30 In a more preferred application, signals which originate and are addressed to the same source and destination nodes are allocated immediately adjacent nodes to each other.

According to a second aspect of the invention, there is provided a method for generating a spectrum allocation in a wavelength division multiplex network comprising a plurality of
35 nodes, wherein the optical spectrum comprises a grid divided into a plurality of slots, the

plurality of slots being grouped into spectrum blocks, comprising using a spectrum entropy measure indicative of the level of fragmentation of spectrum of each spectrum block in dependence on the destination node of each spectrum allocation, wherein the spectrum entropy measure is calculated based on

5

$$-\sum_{i=1}^N \frac{D_i}{D} \ln \frac{D_i}{D}$$

where N is the total number of spectrum blocks in the spectrum, D is the total number of slots, $\ln$ is the natural logarithm, and D_i is the number of slots in the i^{th} spectrum block.

10 In this aspect, a measure of the spectrum entropy of a particular block of signals all destined for the same destination, is calculated to determine how fragmented the spectrum is for that block. Ideally, signals meant for a particular destination are grouped into immediately adjacent slots, especially if they also originate from the source node, so that they are can treated as a single group without need for guard band protection
15 between individual signals. The measure can be calculated for links making up the entire route from source to destination, the entire route, or for the network as a whole. For the purpose of e.g. generating an optimal routing and spectrum assignment configuration e.g. on Day One, a genetic approach can be adopted which is based on the spectrum entropy measure of spectrum blocks.

20

According to a further aspect of the invention, there is provided a network management system module configured to generate a spectrum allocation in a wavelength division multiplex network comprising a plurality of nodes, wherein the optical spectrum comprises a grid divided into a plurality of slots, the plurality of slots being grouped into
25 spectrum blocks, comprising a processor configured to use a spectrum entropy measure obtained by methods of the invention.

According to a further aspect of the invention, there is provided an optical spectrum arrangement in a wavelength division multiplex network, wherein the spectrum
30 comprises a grid divided into a plurality of slots, comprising

- a plurality of spectrum blocks each formed from a plurality of slots, and
- a guard band separating each spectrum block from each other.

In this aspect, applications can enable an optical spectrum to be arranged into spectrum blocks which are pre-designated for signals to a specific destination, the blocks comprising slots suitable for immediate or later occupation by signals for its designated destination node. The blocks are optimally placed as far from each other as possible at the design stage or Day One, so that more signals can be added to occupy the blocks as necessary as time goes on. The skilled person would appreciate that the distance between blocks can be adjusted on the basis of e.g. foreknowledge of the traffic levels that might be expected to a particular destination.

Systems, methods and apparatus embodying the present invention will now be described by way of example only, with reference to the following drawings, wherein:

Figure 1 depicts three graphs depicting known methods of transmitting signals or demands, respectively being (a) the transmission of demands in a fixed grid system, (b) the transmission of demands in a flexgrid system, and (c) the transmission of demands in a superchannel in a flexgrid system,

Figure 2A depicts conventional routing and spectrum allocation in a network,

Figure 2B depicts routing and spectrum allocation according to the invention,

Figure 3A depicts spectrum use according to conventional methods,

Figure 3B depicts spectrum use according to the invention,

Figure 4 depicts network components arranged for the assignment of routes and spectrum for signals according to the invention, and

Figure 5 depicts a spectrum map comprising blocks of used and unused slots.

Figure 2A shows an exemplary routing and spectrum allocation according to conventional methods for an optical network consisting six nodes (10): nodes A to F. A spectrum resource (12) each comprising eight slots (14) is available on each link (18). The letters in the boxes depicting spectral slots (14) indicate the source and destination of a signal occupying each spectrum slot. So for example, the spectral resource on the link between nodes B and D is occupied by a signal originating from node B and destined for node D in slot 1 (the first from the left in the spectrum block), a signal originating from node A and destined for node D in slot 2, a signal originating from node A and destined for node F in slot 3, and so on. In the conventional approach illustrated in Figure 2, signals allocated a route and spectrum regardless of its source and destination, and hence it is possible for signals originating from different nodes to be

placed in neighbouring slots 1 and 2 on the B-D link, as is the case also for signals intended for different destinations which are placed in neighbouring spectral slots 2 and 3. Specifically, there is no grouping of destinations meaning that traffic destined for a particular node to be scattered across the spectrum.

5

This is an inefficient mode of operation for a number of reasons but one main reason is that each signal or spectrum block (of signals addressed to the same destination) requires guard bands on each side. This is illustrated in Figure 1(a) where guard bands (2) are provided for each individual signal (4). As previously noted, guard bands take up
10 spectral resource which otherwise might be available to carry signals. The more discrete signals there are on an optical link or route, the more spectrum needs to be reserved for guard bands.

In contrast, Figure 2B depicts how the same network architecture of Figure 2A can be
15 operated to carry the same traffic according to an application of the invention. The spectrum has been allocated to group together signals which share the same origination and destination points. Referring to the A-B link, it will be noted that there are two A-D signals occupying the spectrum block. By "grouping" the two A-D signals together and placing them into neighbouring slots 1 and 2 on the A-B and B-D links in the manner
20 shown in Figure 2B, they can be treated as a single unit all the way as they progress from node A until they reach their destination at node D. This is conceptually similar to how superchannels (depicted in Figure 1 (c) above), in which a number of demands are treated as a single entity for routing through the network.

25 As noted above, guard bands serve to protect the signals as they pass through the filter when being split to be directed in different directions. In the present example, the signals in slots 1 and 2 are both addressed to the same place, node D, so there is no need to split them at node B (through which they pass on the way to node D). Accordingly, there is no need to employ a guard band between slots 1 and 2, freeing up
30 the spectral resource that would otherwise have been needed. At the destination node D, the received group of signals can be safely separated from each other even in the absence of guard bands, e.g. by power splitting them and then filtering, or by using coherent transceivers which can accurately tune into any particular signal.

Figures 3A and 3B quantifies the possible savings in spectrum (hitherto taken up by guard bands) that can be realised by adoption of methods according to the invention. The figures depict a 47-slot spectral resource (12) in a network configuration comprising four nodes (A, B, C and D). In an exemplary scenario, three signals, each requiring three spectrum slots (14), is required to be received at each of the four nodes. The letters in each of the slots represents the destination node, so referring to Figure 3A for example, slots 1 to 3 represent the signal addressed to node A.

Figure 3A depicts the result in the case where no deliberate steps are taken to group signals or demands together, e.g. with the deployment of conventional, on-the-fly, allocation techniques. In this example, none of the signals destined for a particular node are located next to each other: it would be appreciated that the random nature in which signals are conventionally allocated spectrum may or may not generate such a configuration. The signal destined for node A which occupies slots 1 to 3 (first three slots on the left) is separated by a guard band ("GB") from the signal bound for node B which occupies slots 5 to 7. For simplicity, the guard band is depicted as occupying a slot's width (here, slot 4), although in general this need not be the case: e.g. if you had an underlying grid resolution of 6.25GHz, two or more adjacent slots may be needed to form the guard band. As a result, 11 guard bands (1 to 11) are needed in this example to separate 12 signals occupying a total of 36 slots, allowing them to pass through filters at intermediate node points for onward transmission to different parts of the network. All 47 slots are used, 0 slots are left and the spectrum is 100% used. As expected, the number of guard bands needed increases with the randomness of the placement of demands; so if signals destined for e.g. A, B, C, D were placed in contiguous slots in that order, this creates four spectrum blocks requiring three guard bands to separate them from each other. This aspect is central to the concept of the topological complexity entropy of a given path or link, which will be discussed in detail below.

Figure 3B shows the same signals of Figure 3A being carried but this time they are deliberately grouped together according to the invention, based on their destination node. This time, 36 slots are still required for the transmitted signals but now only 4 slots are required for the guard bands (1 to 4), i.e. one for each destination node. Seven slots are freed up, out of the total spectrum of 47 slots, meaning that the spectrum is only 85% utilised. 15% of the total spectrum resource is available to transmit further signals

meaning that the optical resource is better and more efficiently used than without the grouping being made.

Figures 3A and 3B also depict the possible profiles for the filters (20) for the four ports
5 from the Wavelength Selective Switch (WSS) that sits within an optical node and routes the traffic for the various output ports. As can be seen in Figure 3B, in the case of grouped traffic the filter profiles are much simpler with just one opening being made per destination compared with three in the non-grouped case shown in Figure 3A. This simplifies the configuration of the WSS and other devices and components within the
10 node such as amplifiers. Although frequency management is required in other communication areas (such as wireless and coaxial systems) reduction in the spectrum allocated to the guard bands is not as important as it is in optical network systems and this is due to a number of reasons. In optical systems, the usable spectrum due to attenuation and the operating window of amplification systems is fairly narrow and
15 limited to around 5THz. In other technologies this window is much greater and therefore making use of more guard bands has less impact. In the case of wireless systems, the fact that signals can come from transmitting nodes of varying distance away from the receiving node means that there can be significant differences in received power levels and therefore more advanced filtering systems are required, meaning that additional
20 guard bands are needed.

Figure 4 depicts a possible network arrangement for the management, control and operation of a network in accordance with the invention, which includes a network management system (NMS) (30) that has a current view of the network and instructs the
25 individual network elements to make changes when new customer demands are requested. The process flows between the network components used in routing and allocation management in respect of applications of the invention share much in common with that in a conventional set up based on e.g. RSA, which typically comprises the following. The process commences with a request for a connection to be made
30 between a source node and a destination node. This request is sent to the NMS, which seeks and obtains data about the network state. This information is then used to discover a suitable path based on the network state, which according to standard RSA methods would include shortest-path calculations. The identified path is then returned to the NMS. The WDM network elements or transceivers are then configured to place the
35 signal along the path identified for that purpose.

In embodiments of the invention, the conventional RSA method is substituted by a process or routine for calculating the topological complexity entropy (TCE) of the path, link, or network as the case may be. "TCE" may be understood as a measure of how
5 fragmented an optical spectrum allocation is, with respect to the destination of each demand. Generally therefore, a spectrum where all demands are grouped together in contiguous slots based on their destination has a lower TCE than a spectrum where all demands are randomly allocated. The skilled person would appreciate that as an alternative to replacement of conventional RSA techniques, the system may include
10 entropy calculations in addition to standard routing decisions, and that both the link- and path-based methods could in appropriate cases allow for deeper analyses of the network and identified routes to be carried out before a decision is taken, to obtain improved or desired outcomes, or to enable compliance with an adopted policy.

15 A preferred application of the invention uses a measure of the level of spectral entropy in an optical link or route to enable the grouping together of signals or demands addressed to the same destination node, with the aim of reduce spectrum wastage. In an embodiment, superchannels consisting of multiple signals destined for the same node are created, which require just one set of guard bands each as depicted in Figure 1(c).
20 As described above, this also has the effect of simplifying routing operations within a node, thereby reducing the topological complexity of the solution.

In one preferred application, the NMS is configured and operated to obtain a routing and spectrum assignment for each node pair in the network on Day One when the network is
25 first provisioned, such that a given spectrum allocation for signal(s) destined for a specific node is placed as far away from signals intended for other destinations as possible. This allows each signal group the greatest opportunity to "grow" as additional customer demands for the specific destination are received. So referring to Figure 3B for example, the "B" group signals could be placed several slots distant from the "A" group (using up some of the spare available spectrum), allowing for new demands
30 destined for node A to be included at a later date.

Specifically, the inputs to the process and into the NMS are the number of network nodes, fibre links and a traffic matrix; the output of the process is a routing and spectrum
35 allocation for each node pair. For this process, the route and spectrum assignment for

the demand between A and B can be different from the route and spectrum assignment from B to A. The description below is for a genetic algorithm approach, but the desired output could be obtained using any other optimisation algorithms such as simulated annealing.

5

The process starts by defining the encoding for the network configuration in terms of the routing and spectrum assignment by means of a sequence of integers. If the network consists of n nodes, there are $n*(n-1)$ node pairs. For the routing of each of these node pairs, a number of paths through the network (e.g. 32) are computed based on a standard metric such as number of hops or fibre length and the links that are used in each of these paths are stored in memory. For the spectrum assignment the available spectrum is broken into a number of slots and each node pair is assigned an allocation, the size of which is predetermined from the traffic matrix given as an input. Each spectrum assignment is identified by the base slot for its allocation. To encode the whole network configuration for a network of n nodes requires $2*n*(n-1)$ integers. For each node pair, the currently selected route out of the list of pre-computed ones is recorded along with the base spectrum allocation. These configurations are called network configuration vectors.

20 The genetic algorithm approach is initialised by randomly creating the first generation of network configurations. This is done by randomly filling in the entries in a large number (typically around 5,000) of network configuration vectors. This set of configuration vectors is known as generation zero. To generate each successive generation, the following process is followed. The first step is to check through each network configuration vector to check it is a valid configuration in that there is no signal overlap on any link in the network. If there is any overlap present, the configuration is ignored. For configurations that contain no overlaps, the network TCE of the configuration is calculated. Obtaining the TCE uses the base formula of

30

$$H = - \sum_{i=1}^N \frac{D_i}{D} \ln \frac{D_i}{D}$$

to calculate the entropy for a particular link where there are N blocks of spectrum destined for the same node, D_i is the number of slots in the current block and D is the total number of slots in the entire spectrum band. Any unused spectrum can be

considered part of adjacent blocks of used spectrum, but this is not compulsory. To calculate the TCE for the entire network, this is repeated for every link in the network and the entropies summed together.

- 5 Referring to the spectrum (12) depicted in Figure 5 as an example, this comprises a total of 16 slots (14), of which 10 slots are used (identified by the letter of the node they are destined for), and six are unused. The TCE can be calculated by first identifying four separate blocks of spectrum. The first block (B1) consists of four slots, three of which are destined for node A and one is unused. The second block (B2) consists of five slots, 10 three destined to node B and two slots unused. The third block (B3) contains one slot destined for node A again and two unused slots. This third block is considered separately from the first block (B1) as even though both blocks are destined to node A as they are not contiguous and are separated by the block (B2) of spectrum for node B. The fourth and final block (B4) contains three slots destined for node C and one slot 15 unused. As the unused slot is contained within slots destined for node C this unused slot can be considered part of the larger block.

The calculation would be:

$$H = - \sum_{i=1}^N \frac{D_i}{D} \ln \frac{D_i}{D} = - \left(\frac{4}{16} \ln \frac{4}{16} + \frac{5}{16} \ln \frac{5}{16} + \frac{3}{16} \ln \frac{3}{16} + \frac{4}{16} \ln \frac{4}{16} \right) = 1.371$$

20

- The network TCE calculation is used to quantify how optimal a solution is (in terms of how “grouped together” the signals are). A smaller value from the result of this calculation indicates that the network configuration shows a greater degree of signal 25 group based on the destination node. The process for calculating the network TCE is shown in pseudo code extract shown in Listing 1 below.

```

30 1: CalcNetworkTCEntropy (NetworkConfiguration)
2:     NetworkTCEntropy = 0
3:
4:     For Each Link in NetworkConfiguration
5:         TotalSlots = count of number of spectrum slots in Link
6:         LinkEntropy = 0

```

```

7:
8:         For Each Spectrum Block Destined for the Same Node in
LinkSpectrum
9:             BlockSlots = count of number of slots in Current
5 Block
10:             LinkEntropy = LinkEntropy + BlockSlots /
TotalSlots * ln(BlockSlots / TotalSlots)
11:             Next Block
12:
10 13:             NetworkTCEntropy = NetworkTCEntropy + Entropy
14:             Next Link
15:
16:             Return - NetworkTCEntropy
15

```

Listing 1: Pseudo code for calculating network topological complexity entropy

Once the network TCE is calculated for each of the configurations within a generation the configurations that show a better solution with an entropy below a certain value (such as above the average entropy for the generation) can be “bred” together. This means that good solutions are either spliced together at a random point in the network configuration vector (so a proportion of the new configuration comes from one candidate and the rest comes from another candidate), or a single value in the vector can be mutated to a new value chosen at random. Whether to splice or mutate can be chosen at random. The next generation, again typically containing a large number of candidates, is therefore generated taking the better traits from the previous generation. Successive generations continue to be generated until the time that the average network Shannon entropy of a number of successive generations is no longer increasing or some other threshold defined by the network operator (such as a certain percentage of the maximum network entropy) is reached.

The “best” network configuration (defined as being the one with the largest network topological complexity entropy) in the final generation is then selected as the solution to the problem and this is used as an input by the NMS uses this to configure the network nodes and for allocating future demands.

A full pseudo code extract for the process is shown in Listing 2.

```

1: TCEntropyRoutingAndSpectrumAssignment (Network)
5 2:   NumberOfPaths = 32           # Number of paths to
   calculate per node pair
3:   SpectrumSize = 8000           # Number of slots in
   spectrum
4:   GenerationSize = 5000         # Number of configurations
10 in each generation
5:   GenerationThreshold = 0.9 * MAXENT # Threshold for ideal
   situation (90% of Maximum Entropy)
6:
7:   # Determine the pre-computed routes for all node pairs in the
15 network
8:   Foreach Unique Node Pair
9:     PathTable += CalculateRoutesThroughNetwork (Source, Dest,
   NumberOfPaths)
10:   Next Unique Node Pair
20 11:
12:   # Create Generation Zero
13:   For n = 0 to GenerationSize
14:     Foreach Unique Node Pair
15:       Path = GetRandomNumber(1, NumberOfPaths)
25 16:       Spectrum = GetRandomNumber(1, SpectrumSize)
17:       ConfigurationVector(n) += <Path, Spectrum>
18:     Next Unique Node Pair
19:   Next n
20:
30 21:   GenerationEntropyAverage = 0
22:
23:   # Generate and Evaluate Generations
24:   While GenerationEntropyAverage > GenerationThreshold
25:     GenerationEntropySum = 0           # Used to
35 calculate generation average Entropy
26:     GenerationEntropyCount = 0
27:

```

```

28:          LowestEntropyConfiguration = null # Used to
determine highest Entropy solution
29:          LowestEntropyValue = 0          # in current
generation
5 30:
31:          For n = 0 to GenerationSize
32:              # Check whether current configuration vector
is valid (i.e. no overlaps)
33:              IsValid(n) =
10 ValidateConfiguration(ConfigurationVector(n))
34:
35:              If IsValid(n) is True Then
36:                  # Calculate the Network Entropy of the
current configuration
15 37:                  NetworkEntropy(n) =
CalcNetworkTCEntropy(ConfigurationVector(n))
38:
39:                  # Add to average entropy count
40:                  GenerationEntropySum =
20 GenerationEntropySum + NetworkEntropy(n)
41:                  GenerationEntropyCount =
GenerationEntropyCount + 1
42:
43:                  # Check whether this is the best
25 configuration of this generation
44:                  If NetworkEntropy < LowestEntropyValue
Then
45:                      LowestEntropyConfiguration =
ConfigurationVector(n)
30 46:                      LowestEntropyValue =
NetworkEntropy(n)
47:                  End If
48:              End If
49:          Next n
35 50:
51:          # Calculate average entropy for the Generation

```

```

52:          GenerationEntropyAverage = GenerationEntropySum /
GenerationEntropyCount
53:
54:          # Breed the next generation between those
5 configurations which are below the average network
55:          # entropy
56:          NextGeneration =
BreedNextGeneration(ConfigurationVector, IsValid, NetworkEntropy,
57:                    GenerationEntropyAverage)
10 58:
59:          # Make the next generation the current generation
60:          ConfigurationVector = NextGeneration
61:      Loop
62:
15 63:      # Return the configuration with the highest network
entropy
64:      Return LowestEntropyConfiguration

```

20 *Listing 2: Pseudo code for entire TCE network design process using genetic algorithms*

The above is centred on the TCE of each optical spectrum allocation of a signal, which as noted above is defined in terms of the destination node of the particular signal. An

25 optimal routing and spectrum assignment result following the above is a configuration in which the signals or grouped within a block so that the spectrum entropy measure within the block is minimised, while each block (being populated with signals all destined for the same node) is placed as far apart as possible from each other, allowing for new signals for that destination node to be added over time and as the network grows. The aspect of

30 placing blocks as far apart as possible, while not key to the central idea of grouping signals addressed to the same destination, is nonetheless a preferred way of designing the network especially on Day One to provide scope for growth.

The apparatus, methods and configurations described above and in the drawings are for

35 ease of description only and not meant to restrict the scope of the invention to any particular embodiment. For example, it will be apparent to the skilled person that steps

can be added or omitted from the methods and processes described herein. While the examples illustrating application of the invention are made in respect of an optical network and in particular in connection with flexgrid-based systems, it would be appreciated that other telecommunications systems as well as non-telecommunications systems can suffer from resource fragmentation as well during use, which could benefit from an analysis of entropy levels. In particular, the approach of grouping signals based on their destinations (and secondarily, on their sources) into blocks to treat them as a single entity for routing can also be applied in scenarios to select a preferred path and wavelength channel or slot more optimally than those which allocate spectrum in the current, random, manner. In particular, entropy based fragmentation RSA techniques as described herein can be adopted into existing methods with little significant component reconfiguration owing to the ease with which algorithms along the lines of the above can be loaded into existing systems. It would also be appreciated that techniques and applications of the invention can also be advantageously applied in a variety of situations, not being restricted to use only in respect of identifying network routes, to increase in network capacity by maximising spectral efficiency.

Claims:

1. A method of allocating spectrum to a signal in a wavelength division multiplex network comprising a plurality of nodes, wherein the spectrum comprises a grid divided into a plurality of slots, comprising allocating to the signal a slot immediately adjacent to a second slot occupied by a second signal unseparated by a guard band, wherein the signal and the second signal are both addressed to the same destination node, and wherein the slots allocated to the signal and the second signal form a spectrum block.
2. A method according to claim 1 wherein a third signal addressed to a second, different, destination node is allocated a third slot separated by a guard band from the spectrum block.
3. A method according to claim 1 or claim 2 wherein allocation of spectrum to the third slot results in at least one unallocated slot separating the spectrum block from the third slot.
4. A method according to any preceding claim wherein signals originating from the same source node and addressed to the same destination node are allocated slots immediately adjacent to each other.
5. A method for generating a spectrum allocation in a wavelength division multiplex network comprising a plurality of nodes, wherein the optical spectrum comprises a grid divided into a plurality of slots, the plurality of slots being grouped into spectrum blocks, comprising using a spectrum entropy measure indicative of the level of fragmentation of spectrum of each spectrum block in dependence on the destination node of each spectrum allocation, wherein the spectrum entropy measure is calculated based on

$$-\sum_{i=1}^N \frac{D_i}{D} \ln \frac{D_i}{D}$$

where N is the total number of spectrum blocks in the spectrum, D is the total number of slots, $\ln$ is the natural logarithm, and D_i is the number of slots in the i^{th} spectrum block.

6. A method according to claim 5, wherein a spectrum entropy measure is iteratively calculated for a succession of generations of signal routes in the form of spectrum allocations for each node pair comprising a source node and a destination node.

5

7. A method according to claim 6 wherein the iterative calculation of the spectrum entropy measure terminates either upon determination that a spectrum entropy measure obtained for a generation of signal routes is less than the spectrum entropy measure for a previous generation of signal routes, or upon the a spectrum entropy measure
10 obtained for a generation of signal routes meeting a predefined numeric threshold.

8. A network management system module configured to generate a spectrum allocation in a wavelength division multiplex network comprising a plurality of nodes, wherein the optical spectrum comprises a grid divided into a plurality of slots, the
15 plurality of slots being grouped into spectrum blocks, comprising a processor configured to use a spectrum entropy measure obtained by the method of claim 5.

9. An optical spectrum arrangement in a wavelength division multiplex network, wherein the spectrum comprises a grid divided into a plurality of slots, comprising
20 - a plurality of spectrum blocks each formed from a plurality of slots, and
- a guard band separating each spectrum block from each other.

10. An optical spectrum arrangement according to claim 9 wherein at least one of the slots of one of the plurality of the spectrum blocks is allocated to a signal addressed
25 to a destination node, and wherein at least one of the slots of a second one of the plurality of the spectrum blocks is allocated to a signal addressed to a second destination node.

11. An optical spectrum arrangement according to claim 9 or claim 10 wherein each
30 spectrum block comprises at least one unallocated slot.




FIGURE 1
PRIOR ART




FIGURE 4




FIGURE 5




FIGURE 2A
PRIOR ART






FIGURE 2B



INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2016/056776

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04J14/02
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04J

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

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

EPO-Internal, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GUOYING ZHANG ET AL: "Optical traffic grooming in OFDM-based elastic optical networks [Invited]", IEEE/OSA JOURNAL OF OPTICAL COMMUNICATIONS AND NETWORKING, IEEE, USA, vol. 4, no. 11, 1 November 2012 (2012-11-01), pages B17-B25, XP011472908, ISSN: 1943-0620, DOI: 10.1364/JOCN.4.000B17 abstract paragraph [II.B]; figure 2 ----- -/--	1-4,9-11



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

8 June 2016

Date of mailing of the international search report

22/06/2016

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Borsier, Celine

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2016/056776

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 2 797 247 A1 (BRITISH TELECOMM [GB]) 29 October 2014 (2014-10-29) cited in the application paragraph [0001] - paragraph [0004] paragraph [0015] paragraph [0031]; figure 6 figure 7 (a)	5-8
A	----- FAN ZHEYU ET AL: "Dynamic Multipath Routing With Traffic Grooming in OFDM-Based Elastic Optical Path Networks", JOURNAL OF LIGHTWAVE TECHNOLOGY, IEEE SERVICE CENTER, NEW YORK, NY, US, vol. 33, no. 1, 1 January 2015 (2015-01-01), pages 275-281, XP011571142, ISSN: 0733-8724, DOI: 10.1109/JLT.2014.2387312 [retrieved on 2015-01-22] paragraph [0011] -----	1-4,9-11

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2016/056776

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
EP 2797247	A1	29-10-2014	CN 105359442 A	24-02-2016
			EP 2797247 A1	29-10-2014
			EP 2989735 A1	02-03-2016
			US 2016072608 A1	10-03-2016
			WO 2014174234 A1	30-10-2014
