Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2392376 A1 2001/05/25

(21) 2 392 376

(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2000/11/17
(87) Date publication PCT/PCT Publication Date: 2001/05/25
(85) Entrée phase nationale/National Entry: 2002/05/23
(86) N° demande PCT/PCT Application No.: GB 2000/004392
(87) N° publication PCT/PCT Publication No.: 2001/037605
(30) Priorité/Priority: 1999/11/19 (99309234.5) EP

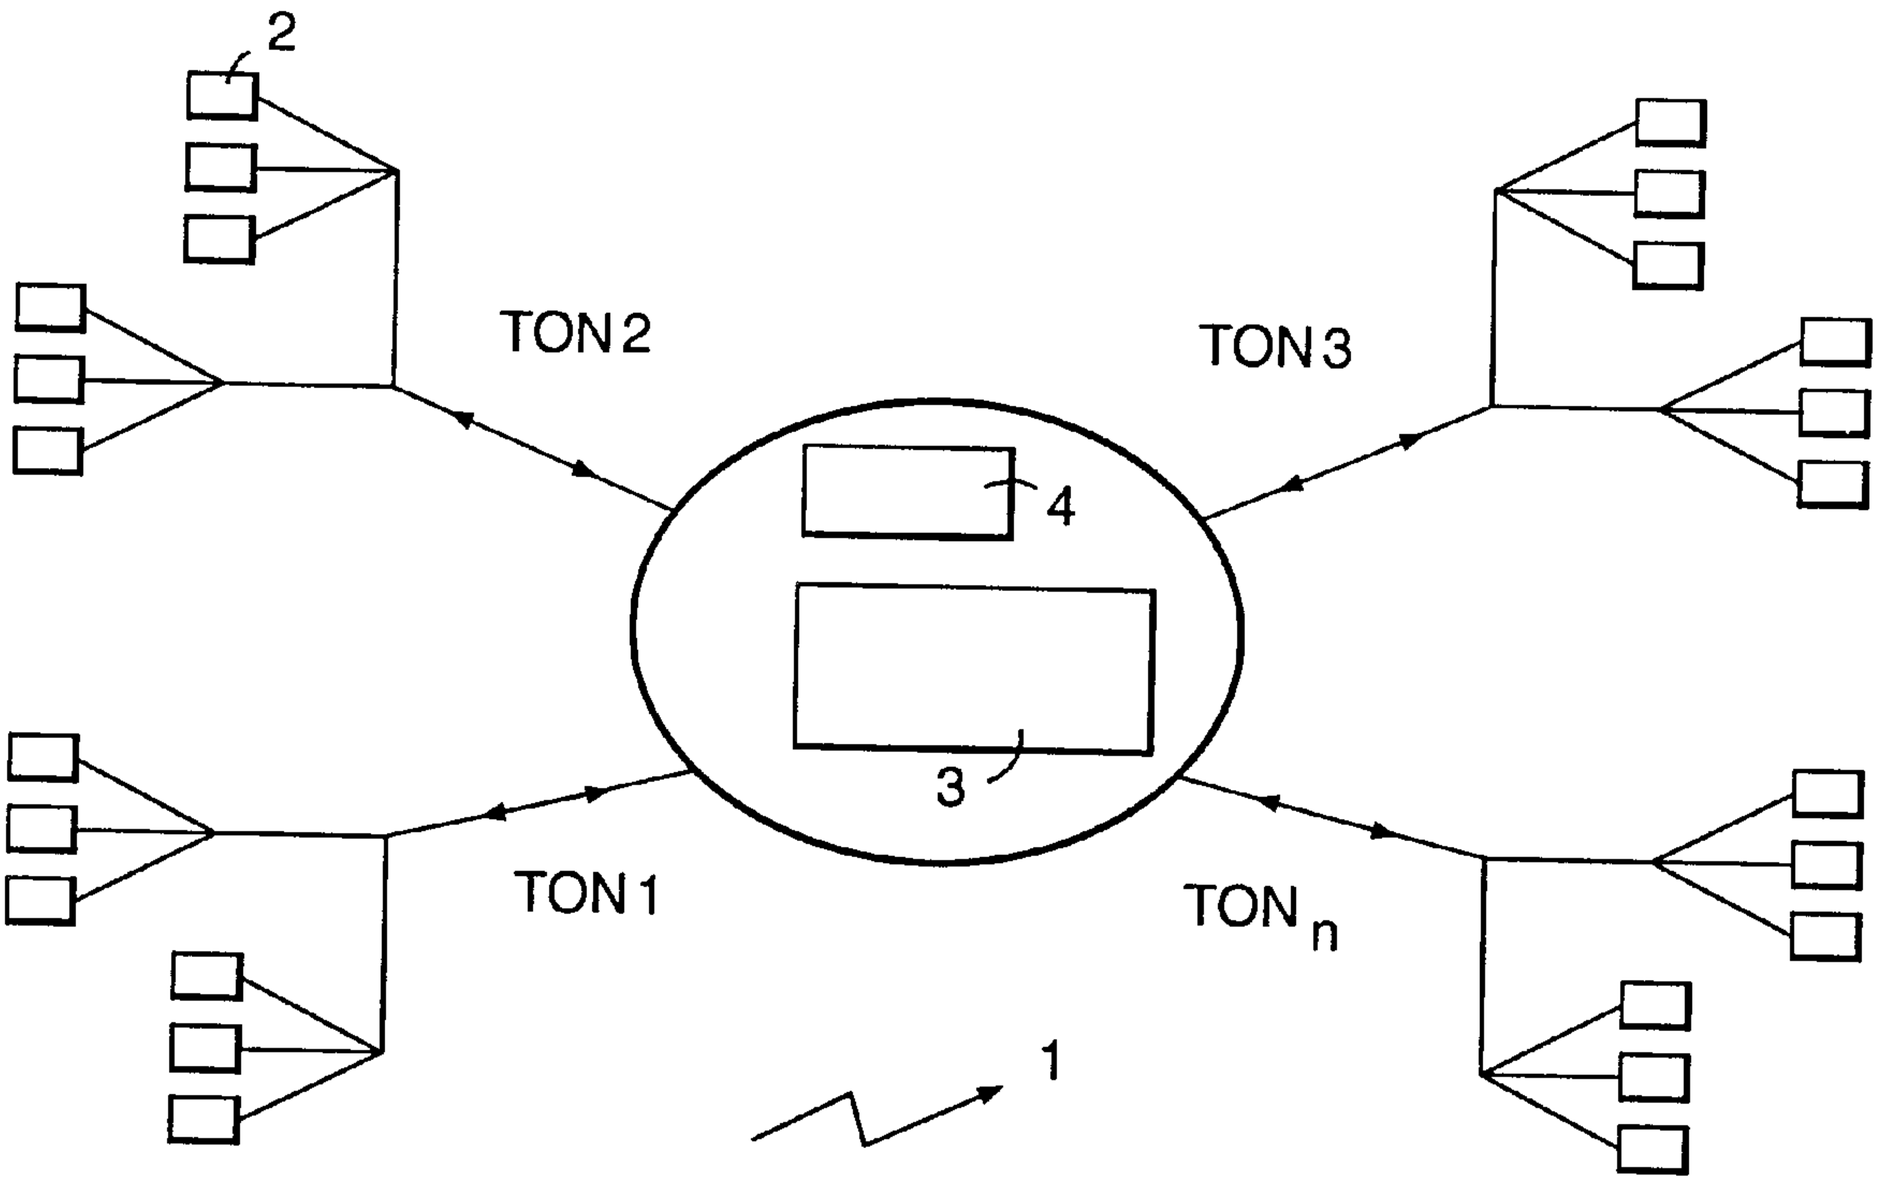
(51) Cl.Int.^{7}/Int.Cl.^{7} H04Q 11/00

(71) Demandeur/Applicant:
BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY, GB

(72) Inventeur/Inventor:
HILL, ALAN MICHAEL, GB

(74) Agent: GOWLING LAFLEUR HENDERSON LLP

(54) Titre : SYSTEME DE COMMUNICATION OPTIQUE
(54) Title: OPTICAL COMMUNICATIONS SYSTEM



(57) Abrégé/Abstract:
A passive optical network employing wavelength dependent routing uses two-phase signalling between a network controller and terminals. In a first phase, terminals are informed of the wavelength channel allocated to them for a second signalling phase, groups of terminals being allocated to each channel. In the second signalling phase, control signals indicate to the terminal the wavelength and time slots allocated to the terminal in the following data transmission phase.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
25 May 2001 (25.05.2001)

PCT

(10) International Publication Number
WO 01/37605 A1

(51) International Patent Classification[7]: H04Q 11/00

(21) International Application Number: PCT/GB00/04392

(22) International Filing Date:
17 November 2000 (17.11.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
99309234.5 19 November 1999 (19.11.1999) EP

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

(72) Inventor; and
(75) Inventor/Applicant *(for US only)*: **HILL, Alan, Michael** [GB/GB]; The Old Police House, Park Road, Grundisburgh, Woodbridge, Suffolk IP13 6TP (GB).

(74) Agent: **LIDBETTER, Timothy, Guy, Edwin**; BT Group Legal Services, Intellectual Property Department, 8th Floor, Holborn Centre, 120 Holborn, London EC1N 2TE (GB).

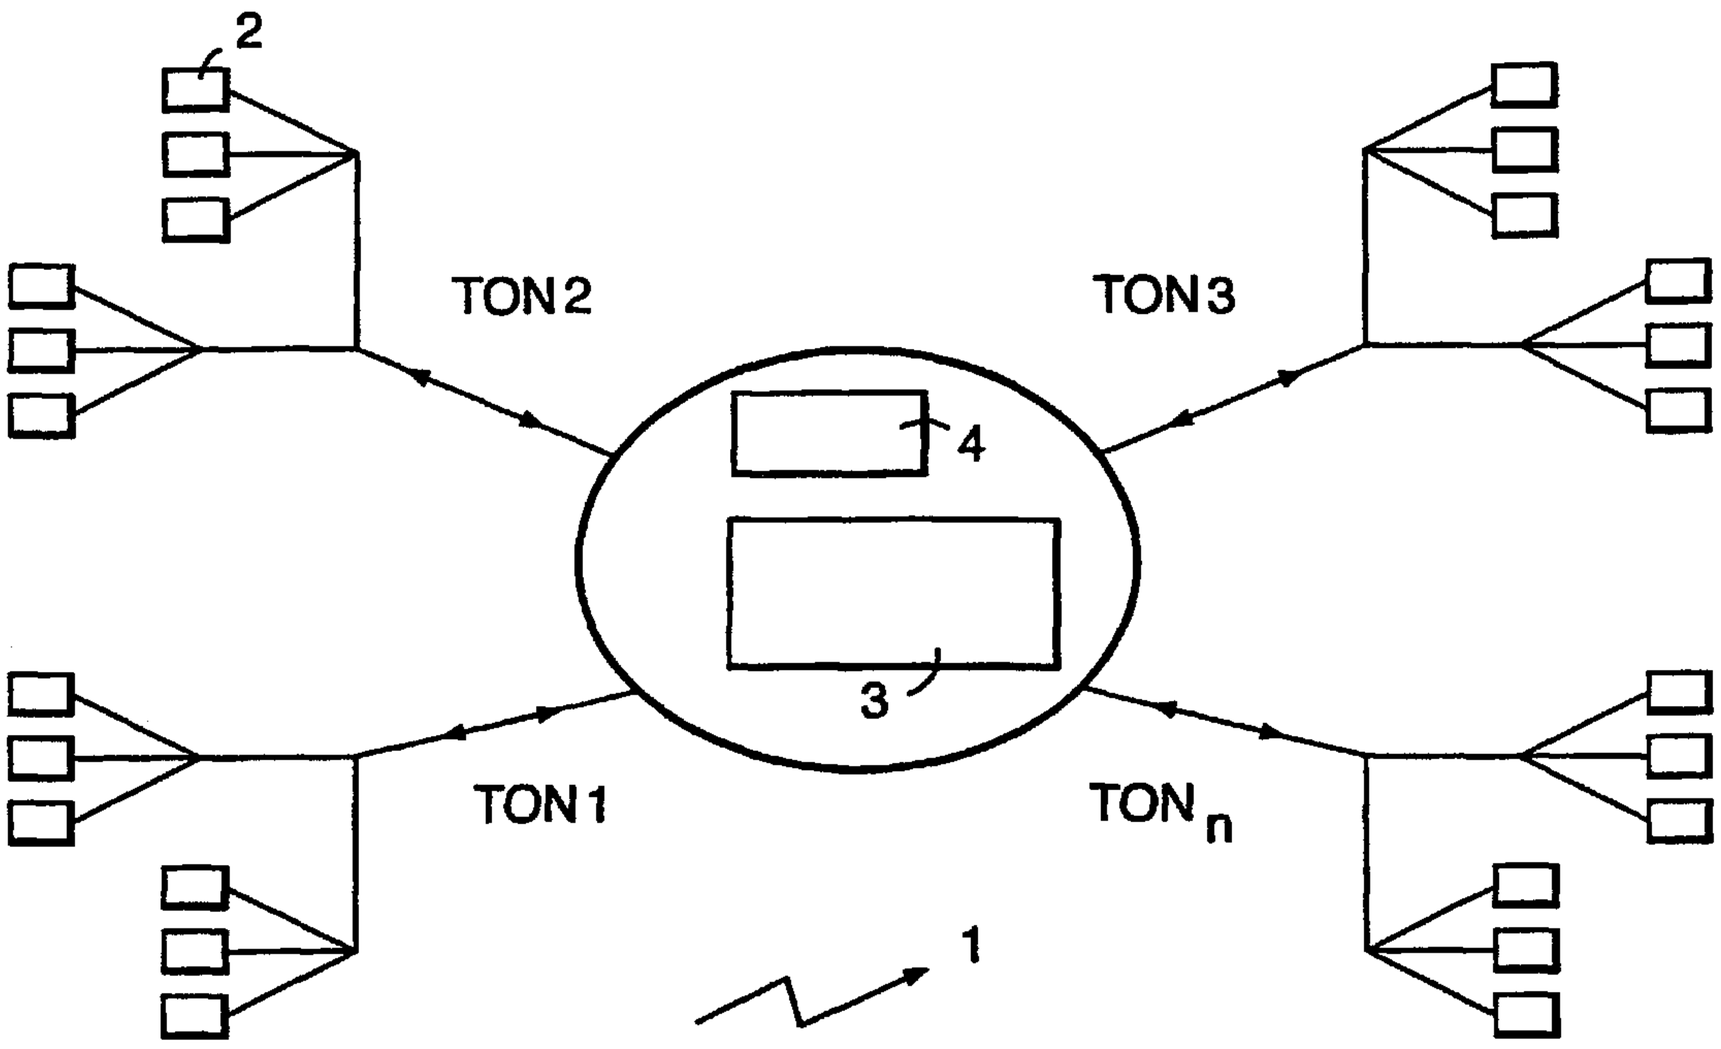
(81) Designated States *(national)*: CA, CN, JP, SG, US.

(84) Designated States *(regional)*: European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

Published:
— *With international search report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: OPTICAL COMMUNICATIONS SYSTEM



(57) Abstract: A passive optical network employing wavelength dependent routing uses two-phase signalling between a network controller and terminals. In a first phase, terminals are informed of the wavelength channel allocated to them for a second signalling phase, groups of terminals being allocated to each channel. In the second signalling phase, control signals indicate to the terminal the wavelength and time slots allocated to the terminal in the following data transmission phase.

WO 01/37605 A1

OPTICAL COMMUNICATIONS SYSTEM

The present invention relates to an optical communications system suitable for carrying broadband communications traffic.

Optical communications networks potentially offer very high capacities. However, in practice it has proved difficult to design appropriate switching technologies to realise the full potential of optical communications networks. Opto-electronic switching designs limit the bit-rate that can be handled by the network. Control and management of the switching function also provides further difficulties and there is a need to minimise software costs for call processing, network intelligence, network management and service management.

The present applicant's earlier international patent application WO95/26592, to which the reader is directed for relevant technical background, describes and claims an optical communications system in which a number of terminals are interconnected via passive optical networks (PONs) and by a wavelength-dependent router at the hub of the network. Each terminal includes means, such as a tuneable transmitter and receiver, that allows it to select one of a number of different wavelength channels for transmission and reception in different time slots. In this way, the switching function is distributed to the edges of the network, while the core of the network functions using chiefly or entirely passive routing devices. This overcomes many of the problems discussed above. However, there is a problem in that distributing the switching function to the edges of the network potentially imposes a large signalling overhead for the transmission of network control information such as bandwidth or packet requests and a resource allocation map indicating the wavelength channels and time slots allocated to each request.

According to a first aspect of the present invention, there is provided a method of operating an optical communications system comprising a wavelength-dependent optical router , a plurality of terminals, each terminal including at least one tuneable optical receiver, and at least one passive optical network interconnecting the plurality of terminals via the wavelength-dependent optical router, the method including:

(a) transmitting to each terminal in a first signalling phase, the identity of a wavelength channel assigned to the terminal, where some of the terminals are assigned a common wavelength channel;

(b) in a second signalling phase subsequent to the first signalling phase, transmitting network control signals to each terminal of the respective wavelength channel identified to the terminal in the first signalling phase; and

(c) in the data transmission phase subsequent to the first and second signalling phases, communicating data between terminals via the passive optical network and the wavelength-dependent optical router, where each terminal outputting data onto the passive optical network transmits on a respective wavelength channel selected depending on the control data received at the respective terminal in the second signalling phase.

The present invention provides a method of communicating control signalling to the edge of the network that makes highly efficient use of the network resources, thereby enabling all signalling messages transmitted from the network controller to the terminals to be received at each terminal by means of the same tuneable receiver as is used for data reception. In preferred implementations, a complete resource allocated map may be distributed to the terminals, which allows complete flexibility for the network controller to allocate time-slots and wavelengths and hence to provide faster responses to resource requests. This is achieved by using a two phase signalling process at the beginning of a data frame. The first phase, termed the meta-signalling phase, is used to assign a number of terminals to a given wavelength channel for the subsequent signalling phase. In the subsequent signalling phase control data, such as a complete wavelength/terminal allocation map, is transmitted to the terminals. The use of a two-phase process for signalling gives flexibility in the allocation of signalling wavelength channels.

According to a second aspect, the invention provides an optical communications system comprising a wavelength-dependent optical router, a plurality of terminals, each terminal including at least one tuneable optical receiver, and at least one passive optical network interconnecting the plurality of terminals via the wavelength-dependent optical router, the system being configured to operate according to a method which includes the following steps:

(a) transmitting to each terminal in a first signalling phase, the identity of a wavelength channel assigned to the terminal, where some of the terminals are assigned a common wavelength channel;

(b) in a second signalling phase subsequent to the first signalling phase, transmitting network control signals to each terminal of the respective wavelength channel identified to the terminal in the first signalling phase; and

(c) in the data transmission phase subsequent to the first and second signalling phases, communicating data between terminals via the passive optical network and the wavelength-dependent optical router, where each terminal outputting data onto the passive optical network transmits on a respective

wavelength channel and data time-slot selected depending on the control data received at the respective terminal in the second signalling phase.

The invention also encompasses a network controller or a terminal arranged to operate in accordance with the method of the first aspect, and a communications system including such a controller and terminal.

Systems embodying the present invention will now be described in further detail, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a schematic of a network embodying the invention;

Figure 2 is a timing diagram for downstream signalling;

Figure 3 is a diagram showing in further detail the network of Figure 1;

Figure 4 is a timing/wavelength diagram for upstream signalling;

Figure 5 is a diagram showing the use of optical delays at the controller receivers; and

Figure 6 is a diagram showing an alternative embodiment employing optical delays in both upstream and downstream signalling.

An optical telecommunications system comprises a plurality of transparent passive optical networks (TON1, TON2 ...TON_N). The TONs are connected in a star topology. A plurality of terminals 2 are connected to the TONs. Each terminal includes a transmit stage T arranged to select a time slot and wavelength channel for outgoing signals, and a receive stage R arranged to receive signals in a particular time slot and wavelength channel. Both the transmit stage and the receive stage are tuneable to operate at different wavelengths at different times. At the hub of the system, the TONs are connected to a wavelength-dependent router. A network controller 4 is also located at the hub of the network.

The terminals may be located at customer premises or maybe employed at an intermediate network station where customer traffic has already been aggregated and multiplexed. For example, the terminal may be located in the vicinity of a number of customer residences (a fibre to the kerb (FTTK) configuration), or in a street cabinet (fibre to the cabinet (FTTCab)) or in a local exchange (FTTExch).

The tuneable wavelength source in the transmit stage of each terminal may be, for example, a tuneable multi-section DBR (distributed Bragg reflector) commercially available from Altitune or a grating-assisted vertical coupler laser device.

Although for ease of illustration in Figure 1, a single TON is shown connected to each group of terminals, in practice a pair of such networks is used, one for downstream

signals and one for upstream signals. In operation, a terminal wishing to transmit data to another terminal located elsewhere on the network transmits to the network controller a request for one or more transmission slots. From all the requests received from the different terminals, the network controller determines an allocation of timeslots and wavelength channels. The controller returns to the terminals data indicating to each terminal the time and wavelength slot allocated to it for communication with a specified other terminal. Then, in the next transmission frame, the terminal sets its tunable laser to the appropriate wavelength value, and transmits data in the allocated timeslot. The resulting signal is received by the wavelength-routing device which routes the received signal onto the TON in which the destination terminal is located. Wavelength converters connected to dummy ports of the routers provide some active control over routing, overcoming traffic blocking problems by making more than one wavelength channel available for transmission between a given pair of TONs.

Figure 2 is a timing diagram showing the timing of transmissions in the downstream direction, that is from the hub to the terminals. The network employs two signalling phases. A first signalling phase, termed the meta-signalling phase (MP) is used for the transmission from the network controller to each terminal on the network of a signal indicating the wavelength channel to be used by a respective terminal to receive data in the second signalling phase (SP). Each terminal is assigned a specific wavelength channel and timeslot to be used during the meta-signalling phase when the terminal is initialised. In general, this allocation of timeslot and wavelength channel for the meta-signalling phase remains fixed, being changed when the terminal is brought into or out of service. Subsequent re-allocations of the wavelength channel is possible for management purposes, however on a much longer time scale than the frame duration. By contrast, the wavelength channel used for the second signalling phase (SB) can be changed for each successive frame. In general, each of the wavelength channels available on the network, is used in the second signalling phase SB to communicate data to a sub-set of the terminals. As is described in further detail below, the allocation of wavelength channels to terminals is carried out to make most effective use of the capacity available on the network.

Following the signalling phase, a data phase (DP) comprising 1000 time slots S1, S2....S1000 is used to communicate traffic between terminals.

The invention is now described in further detail with reference to an implementation in which the system carries traffic on 800 high density wavelength division multiplexed (HWDM) channels and in which 800 TONs (400 upstream and 400

downstream) are used to carry traffic from 50,000 terminals per TON pair. It should be understood that these numbers are given by way of example only, and the invention might equally be applied to larger or smaller systems, such as a fibre to the cabinet (FTTCab) system with 320,000 terminals.

Figure 3 shows some of the key elements of figure 1 in greater detail. In particular, it shows that the passive wavelength-dependent routing device 3 in this example is a NxN wavelength router connected to N upsteam PONs on one side and N downstream PONs on the other side. An appropriate device is the Stimax™ configuration router available commercially from Instruments SA. This uses a planar diffraction grating to cross-connect signals between arrays of optical fibre waveguides. The wavelength router also includes a number of dummy ports each connected to a respective wavelength converter and allowing flexible re-allocation of wavelength channels through the router. The use of dummy ports is described in further detail in our above-cited international patent application WO95/26592. The controller may comprise an appropriate computing platform, such as a UNIX workstation or a dedicated electronic processor, that controls electro-optic modulators used to modulate outgoing signalling information, and that also receives signalling input from burst mode optical receivers.

Figure 3 shows one of the tuneable transmitters connected to a respective upstream PON and one of the tuneable receivers connected to a respective downstream PON. Often, the transmitter and receiver will be co-located in a single terminal, but in some implementations, at least some terminals may have receivers only or transmitters only.

The principle employed in the network of Figure 3 is to broadcast signalling information on each of the 800 wavelength channels, but with each channel broadcasting to only a group of terminals, rather than the entire PON. The terminal groupings allocated to each wavelength channel change from frame to frame. Each wavelength is broadcast to 50,000/800=63 terminals on average. The signalling information to each terminal comprises up to 1,000 sets of time-slot connection data, since any terminal could be allocated up to 1,000 slots in a frame. However, each wavelength only needs to transmit 1,000 sets of slot data in total. Each slot may have a different source terminal for the receiver and a different destination terminal for the transmitter. The bits required for each slot are:

5 bits downstream data slot header

16 bits source/destination terminal to receive signalling data

25 bits destination/source terminal address

10 bits allocated wavelength identifier

10 bits allocated time-slot identifier

This data slot is assumed to be 9 bytes long. There are 1,000 of these data slots within a downstream reply packet. So each wavelength must transmit 2x1,000x9=18 kbytes within each frame, ie every wavelength channel must devote 18,000x8x100=14.4 Mbit/s to signalling, assuming a 10 ms frame duration. This represents only 2.3% of the capacity of the network.

Before this main signalling information can be transmitted, the 50,000 terminals on a PON are first indicated, in the meta-signalling phase, which of the 800 wavelength channels to tune to in order to receive the main signalling information, ie which group of terminals they have been allocated to. On each wavelength channel the controller transmits meta-signalling to 50,000/800=63 terminals. By always meta-signalling to the same set of terminals at each wavelength, there is no need to identify which terminal is being addressed, since they can always be addressed in sequential order. The meta-signalling capacity required on each wavelength channel is therefore very small: 63.10 bits/frame (assuming 10 bits for the wavelength identity) = 79 bytes/frame = 79.8.100= 63.2 kbit/s. The total downstream signalling bit-rate per channel is therefore 14.46 Mbit/s, which still represents only 2.3% of the network capacity.

As noted above, each wavelength channel transmits the main signalling information to 63 terminals (on average). But it may need to transmit to just 1 terminal, if that terminal is allocated 1,000 time-slots in a frame by the network controller, or, at the other extreme, to as many as 1,000 terminals, if each terminal has only 1 time-slot allocated to it. A partitioning algorithm is used to allocate groups of terminals to each wavelength channel. There are 800,000 sets of slot data to be transmitted by the 800 wavelength channels, so each wavelength can transmit 1,000 sets to its allocated group of terminals.

The allocation algorithm is as follows. Starting with the first wavelength channel, we go through each terminal in turn using the U_{tx} or U_{rx} matrices, adding up the total number of time-slots allocated to them by the Network Controller. The U_{rx} and U_{tx} matrices map respectively receiving and transmitting terminals to wavelength channels and time slots for the data transmission phase. When the total exceeds 1,000, no more terminals can be allocated to that wavelength channel. The excess number above 1,000 is carried forward into the next group for the next wavelength channel. The process starts again, until the next channel's 1,000 slots have been allocated to terminals. This continues until each wavelength has been allocated its group of terminals. With this

algorithm, it is possible for a given terminal to be in two groups (ie two different wavelength channels). It may therefore be necessary to transmit the signalling twice, so that each terminal can tune to each of the two wavelengths in turn (if it has only one tunable receiver). There is no need to transmit the identity of the second wavelength channel to tune to during meta-signalling, because it can be assumed to be the next one in order. Signalling would in this case represent 4.6% of the network capacity. Alternatively, with two tunable transmitters and receivers per terminal, signalling information would only need to be transmitted once, remaining at 2.3% of network capacity.

800 wavelength channels are needed in each PON for signalling, either for 2.3% or for 4.6% of the frame duration. The remainder of the frame is for customer data. There are several ways in which the 400.800=320,000 sources of light required can be obtained and modulated, with differing additional costs. One possibility is to provide 320,000 additional light sources. This would be very expensive, although only a fraction of the cost of e.g. 20,000,000 user terminals. Fig.3 shows several alternative approaches for the sources and their modulation.

A lower-cost approach uses the tunable lasers in the user terminals themselves. Within each PON a sub-set of 800 of the 50,000 terminals transmits preferably a continuous wave signal for the duration of the downstream signalling period. This ensures no additional cost for the light sources, but it requires 2.3% or 4.6% of the upstream capacity of the network to be lost to provide the unmodulated light for downstream signalling. The wavelength router would connect 400 different wavelengths directly into the downstream PONs, and a further 400 wavelengths into each dummy port, whence they can also be connected into the downstream PONs. These channels are then modulated with the appropriate downstream signalling information.

A slightly less cost-effective method of generating the unmodulated wavelength channels for each downstream PON employs a single wavelength reference comb using only 800 light sources, whose power is then split 400 ways. Each of the 400 multiplexes would be coupled or switched into one of the downstream PONs.

These two methods could be combined. For example, the reference comb could generate just 400 wavelengths per PON, while the remaining 400 are generated by the user terminals.

Modulation of the 320,000 CW channels could be performed in two ways. The switchless network would already demultiplex those channels that pass through the dummy ports into separate channels at the wavelength converters. For these channels,

modulation is performed by the wavelength converters 51 (preferably, to avoid additional costs), or by additional modulators. Those channels that are connected directly into the downstream PONs can be modulated by means of an SOA (semiconductor optical amplifier) array 52 sandwiched between a wavelength multiplexer and demultiplexer, one such system being deployed for each PON. A like system is used in the upstream PONs. This would demultiplex the channels, modulate each one in an SOA, then remultiplex them. If SOA array technology is employed within a dispersive optical system, then there is no need for each channel to have its own fibre termination. Only the complete SOA array would need to be aligned in the packaging process, which helps to keep costs down. At most there would be 320,000 SOA modulators required for signalling. In the rest of the network amplifier gating could be required for every 160 user terminals assuming 622 Mbit/s operation, ie (20,000,000/160)x800=100 million SOA gates (ie 5 per user terminal). The additional SOA cost is only 0.32%.

There is described above how wavelength channels are modulated with downstream signalling information. Even if the user terminal tunable transmitters themselves are used as light sources for downstream signalling, the network controller must still modulate 320,000 channels, ie 800 in each PON. An approach using optical delay lines may be employed to reduce the number of modulators required.

Each downstream signalling reply packet contains 2,000 9-byte data slots. There are 200 slots to allow for each terminal receiving on two different wavelength channels in sequence. These 18 kbytes last 231.5 μs at 622 Mbit/s. A guard-band of 2.2 μs must be added to this, making 233.7 μs packet duration. The signalling must be transmitted twice, so the total duration of the downstream signalling is 467.4 μs. Given this duration, it is possible for just 40 tunable light sources and associated modulators to be used in each PON for downstream signalling, transmitting in 10ms/467.4μs≅20 sequential slots (with frame duration of 10ms). These slots are sub-frames of the second signalling phase. The channels are generated by the tunable light sources in different blocks of 38 wavelength channels within each sub-frame. The tuneable light sources are then tuned to the next block of wavelength channels for the next sub-frame, and so on, until an entire frame has been filled. Optical delay lines are used, to cause the blocks of wavelength channels within the different time-slots to be delayed relative to each other so that they completely overlap each other in time by the end of a frame, i.e. so that the different sub-frames are aligned in the time domain. The unit of delay being 467.4 μs, the fibre delay lines need to be 93.5 km long. Each PON has 20 delay lines. The longest delay would be for 19 re-

circulations. This process is complementary to that in the upstream direction described in further detail below, where the blocks of wavelengths are spread out in time by the delay lines, rather than compressed as here.

Altogether the total number of tunable transmitters needed at the network controller for downstream signalling in this way is 400x40=16,000. This is only 0.08% of the tunable transmitter costs in all the user terminals; a negligible additional cost. No additional modulators (or SOAs) are required. There is now no need to use the user terminal transmitters to supply optical power for downstream signalling, so upstream transmission efficiency is improved.

Figure 6 shows the optical arrangements for each PON to minimise transmitters and receivers for both upstream and downstream signalling, when using 800 wavelength channels. Twelve burst-mode opto-electronic receivers 6 receive upstream optical signals via a 67-way splitter/delay network. Similarly, for the downstream PON's, 40 transmitters 7 output signals to the user terminals via a 20-way splitter/delay network.

Conventional signalling techniques allow a terminal to request packets to only a single destination terminal in each frame. This would be acceptable e.g. for telephony, if handled as TDMA circuits. Different destinations would be set up and cleared down one at a time in each frame. But, for connectionless operation, it is advantageous to provide for multiple destinations to be set up within a single frame's request packet. The techniques described here provide sufficient capacity to accommodate the necessary upstream signalling.

In each PON, (T/N)=20,000,000/400=50,000 terminals may wish to transmit up to F=1,000 destination addresses per frame, each requiring $\log_2 T=25$ bits, ie 50,000x1,000x25=1.25 Gbits. With $(1/F\tau)=100$ frames per second, this needs 125 Gbit/s upstream signalling capacity within each PON. (τ is the time-slot duration and F is the number of time-slots per frame). This would be (125/(800x0.622))x100 % = (125/497.6)x100 % = 25.1% of the entire network capacity. So it would be possible to provide complete freedom of requests in every frame. Nevertheless, if connectionless operation were to become predominant, this would be rather a high proportion of the network's capacity for upstream signalling. Interestingly, when such complete freedom is allowed to request and set up packets to 1,000 different destinations per frame, it is the upstream signalling requirement that dominates over the downstream signalling requirement, ie 25.1% vs 2.3% or 4.6% of network capacity.

More explicitly, the upstream signalling ratio (proportion of network capacity) is given by

$$\frac{\log_2 T(T/N)(F/F\tau)}{(N+d+r)B} = \frac{\log_2 T}{2\tau b} = 0.25$$

when the values for the total number of wavelength channels and router ports (N+d+r) are inserted for a "switchless" network. This ratio is independent of the number of time-slots F in the frame and the number of PONs N. It is in effect the ratio of the number of address bits per time-slot request to the average number of user data bits per time-slot. There are two ways of reducing this signalling ratio; to use longer time-slots τ or to provide greater average capacity per terminal b. Longer time-slots would cause greater difficulties for coding and other delays, and should probably be avoided. If the average terminal rate b were increased, this would represent a greater network capacity. Unless the operational bit-rate B were also increased in proportion, the numbers of wavelength channels and router ports would become very large. Let us investigate b=20 Mbit/s and B=2.5 Gbit/s. This enables the numbers of wavelength channels and router ports to remain unchanged at 800 and 400 respectively. The total network capacity becomes 800x400x2.5 Gbit/s = 800 Tbit/s, and the signalling becomes just 25/4=6.25% of the network capacity.

However, of course if the time-slot duration is not reduced when the bit-rate B is increased, the number of bits within each packet rises. The coding efficiency e.g. for telephony will therefore be reduced.

Hence by increasing the total capacity of the network, which involves increasing both the operational bit-rate B and the average bit-rate per terminal b, it is possible to reduce the upstream signalling ratio to an acceptable level, without having to increase the numbers of wavelength channels or router ports. The additional cost of doing so would be predominantly due to the reduced upstream split to the first amplifier, which would go down from around 160 at 622 Mbit/s to around 32 ways at 2.5 Gbit/s. We would also need 800/32=25 SOA gates for every terminal for amplifier gating instead of 800/160=5. The additional cost penalty due to the reduced coding efficiency, particularly for low bit-rate services like telephony, must also be taken into account.

The above analysis for allowing signalling requests to more than one destination in each frame applies to FTTH, where there are for example 20,000,000 terminals, but only 320,000 of which can transmit in a given time-slot. However, for FTTCab etc. the signalling ratio would be much lower, because there would be far fewer terminals making requests, eg 320,000 if 800 wavelengths are fully used in each of 400 PONs. For 320,000 terminals the ratio becomes

$$\frac{\log_2 T}{2\tau b} = \frac{19}{2\tau B} = 0.0015$$

Because the number of terminals is now the same as the number of wavelength channels in the network, the average terminal bit-rate b is now the same as the operational bit-rate B, and the signalling ratio is only 0.15%. Hence for FTTK, FTTCab or FTTExch it is possible for the terminals to signal requests to up to 1,000 different destinations in each frame with very small signalling inefficiency. This is advantageous because these terminals must handle the traffic of perhaps hundreds of customers, which will involve far more connection requests than from a single FTTH customer.

Approaches to upstream signalling, that is from the terminals to the hub of the network, will now be described. Use is made of all of the 800 wavelength channels for upstream signalling. This removes any need for a second, fixed-wavelength transmitter at the user terminals. Just the one tunable transmitter at each user terminal is shared between signalling and data traffic by putting the signalling with the customer data slots.

Each request packet is only 8 bytes long. At 622 Mbit/s this takes only 103 ns. However, each one of 50,000 terminals must transmit its request packet within a frame duration of 10 ms, if all terminals are to be allowed a request in every frame. Each request packet is therefore independently generated by each terminal, so the duration of each packet must be extended to accommodate the guard-time. This can be around 2.2 μs. So each mini-slot for request packets must last for 2.2+0.103=2.3 μs. Because the same transmitter must be used for data and signalling, every terminal should ideally transmit request packets simultaneously. Otherwise, the Network Controller will need to perform a more complicated algorithm to allocate data time-slots to terminals, taking into account the different slots in which each terminal is unavailable for data transmission while signalling. Since there are 800 wavelength channels available, groups of 800 terminals can signal simultaneously, followed by subsequent groups of 800, until 50,000/800=63 different groups of terminals have transmitted in 63 sequential time-slots. This takes only 63x2.3 = 145 μs, which is equivalent to only around 15 data time-slots of the 1,000 within a frame. The upstream signalling is therefore only 1.5% of the upstream network capacity. Figure 4 shows the structure of the upstream signalling packets.

Although the upstream signalling is very efficient when all 800 wavelength channels are employed for signalling, we would need 800 tunable burst-mode receivers in every PON, ie 320,000 altogether, unless some more intelligent approach to receiving the

signalling packets is employed. 320,000 is only a small proportion of the 20,000,000 user terminals, so the network cost would only be increased by 1.6% if a better approach were not employed.

At the other extreme, if we tried to use just one tunable, burst-mode receiver in each PON, we would need to translate the 800 simultaneous sets of 63 sequential request packets into one sequential stream of 50,000 packets. But this would last for 50,000x2.3 μs = 115 ms, which is 11.5 frames. This is far too long. In order to limit the duration to 1 frame, we would need to have 12 separate tunable, burst-mode receivers for signalling in each upstream PON. This would require sets of 800/12=67 wavelength channels to be translated from simultaneous transmission to sequential reception at a single receiver. In each set of 67 channels, each channel would have to be delayed relative to the others in units of 63x2.3 μs = 145 μs, with up to 67 different delays (ie 0 to 67x145μs=9.715 ms).

Figure 5 shows a possible optical arrangement for minimising the required number of tunable, burst-mode receivers for signalling in each upstream PON. At the network controller, the upstream PON fibre is split into 67 separate wavelength bands, each containing 12 different wavelength channels, by means of band-pass optical filters. Each band is delayed by units of 145 μs with respect to each other band, eg by the use of 29 km of optical fibre. To obtain different multiples of 145 μs, each band is allowed to propagate a different integer number of times around its 29 km fibre delay line, using 2x2 optical switches. The delayed wavelength bands are then split 12 ways. This 12-way split allows each of 12 tunable, burst-mode receivers to select one of the resulting 12 sets of simultaneous wavelength channels.

Altogether the total number of tunable, burst-mode receivers required for upstream signalling is just 12x400=4,800, which is a large reduction on the worst-case of 320,000, providing a large saving in the cost of receivers.

CLAIMS

1. A method of operating an optical communications system comprising a wavelength-dependent optical router, a plurality of terminals, each terminal including at least one tuneable optical receiver, and at least one passive optical network interconnecting the plurality of terminals via the wavelength-dependent optical router, the method including:

(a) transmitting to each terminal in a first signalling phase, the identity of a wavelength channel assigned to the terminal, where some of the terminals are assigned a common wavelength channel;

(b) in a second signalling phase subsequent to the first signalling phase, transmitting network control signals to each terminal of the respective wavelength channel identified to the terminal in the first signalling phase; and

(c) in the data transmission phase subsequent to the first and second signalling phases, communicating data between terminals via the passive optical network and the wavelength-dependent optical router, where each terminal outputting data onto the passive optical network transmits on a respective wavelength channel and data time-slot selected depending on the control data received at the respective terminal in the second signalling phase.

2. A method according to claim 1, including repeating steps (a), (b) and (c) for each successive data frame transmitted on the network.

3. A method according to claim 2, in which each terminal is assigned a wavelength channel for the first signalling phase on initialisation, and receives signals in the first signalling phase on the said wavelength channel in each of a plurality of successive transmission frames.

4. A method according to claim 3, in which each terminal is assigned, both a wavelength channel and a time-slot for the receipt of signals in the first signalling phase.

5. A method according to any one of the claims 2 to 4, in which, in different frames, in the first signalling phase, different allocations of wavelength channels to terminals are communicated to the terminals.

6. A method according to any one of the preceding claims, in which the data phase is divided into a plurality of time-slots, and in which in the second signalling phase, each terminal receives the identity of a time-slot and wavelength for the or each slot available to the terminal in the following transmission phase.

7. A method according to any one of the preceding claims, in which the optical communications system is carrying high-density wavelength division multiplexed (HWDM) traffic on N different wavelength channels.

8. A method according to claim 7, in which substantially all of the said N different wavelength channels are allocated to terminals for use in the second signalling phase.

9. A method according to any one of the preceding claims, further comprising passing optical signals received at one or more input ports of the wavelength-dependent optical router through a wavelength converter and returning the signal at a different wavelength to the wavelength-dependent optical router thereby routing a plurality of wavelength channels between an originating passive optical network and a destination passive optical network.

10. A method according to any one of the preceding claims, including, in step (b), communicating to a terminal control signals allocating at least two different time-slots for communication between the said terminal and a plurality of other terminals within a single data transmission frame.

11. A network controller arranged to operate by a method according to any one of the preceding claims.

12. A terminal arranged to operate by a method according to any one of the preceding claims.

13. An optical communications system including a network controller according to claim 11 and a plurality of terminals according to claim 12.

14. An optical communications system comprising a wavelength-dependent optical router, a plurality of terminals, each terminal including at least one tuneable optical receiver, and at least one passive optical network interconnecting the plurality of terminals via the wavelength-dependent optical router, the system being configured to operate according to a method which includes the following steps:

(a) transmitting to each terminal in a first signalling phase, the identity of a wavelength channel assigned to the terminal, where some of the terminals are assigned a common wavelength channel;

(b) in a second signalling phase subsequent to the first signalling phase, transmitting network control signals to each terminal of the respective wavelength channel identified to the terminal in the first signalling phase; and

(c) in the data transmission phase subsequent to the first and second signalling phases, communicating data between terminals via the passive optical network and the wavelength-dependent optical router, where each terminal outputting data onto the passive optical network transmits on a respective wavelength channel and data time-slot selected depending on the control data received at the respective terminal in the second signalling phase.

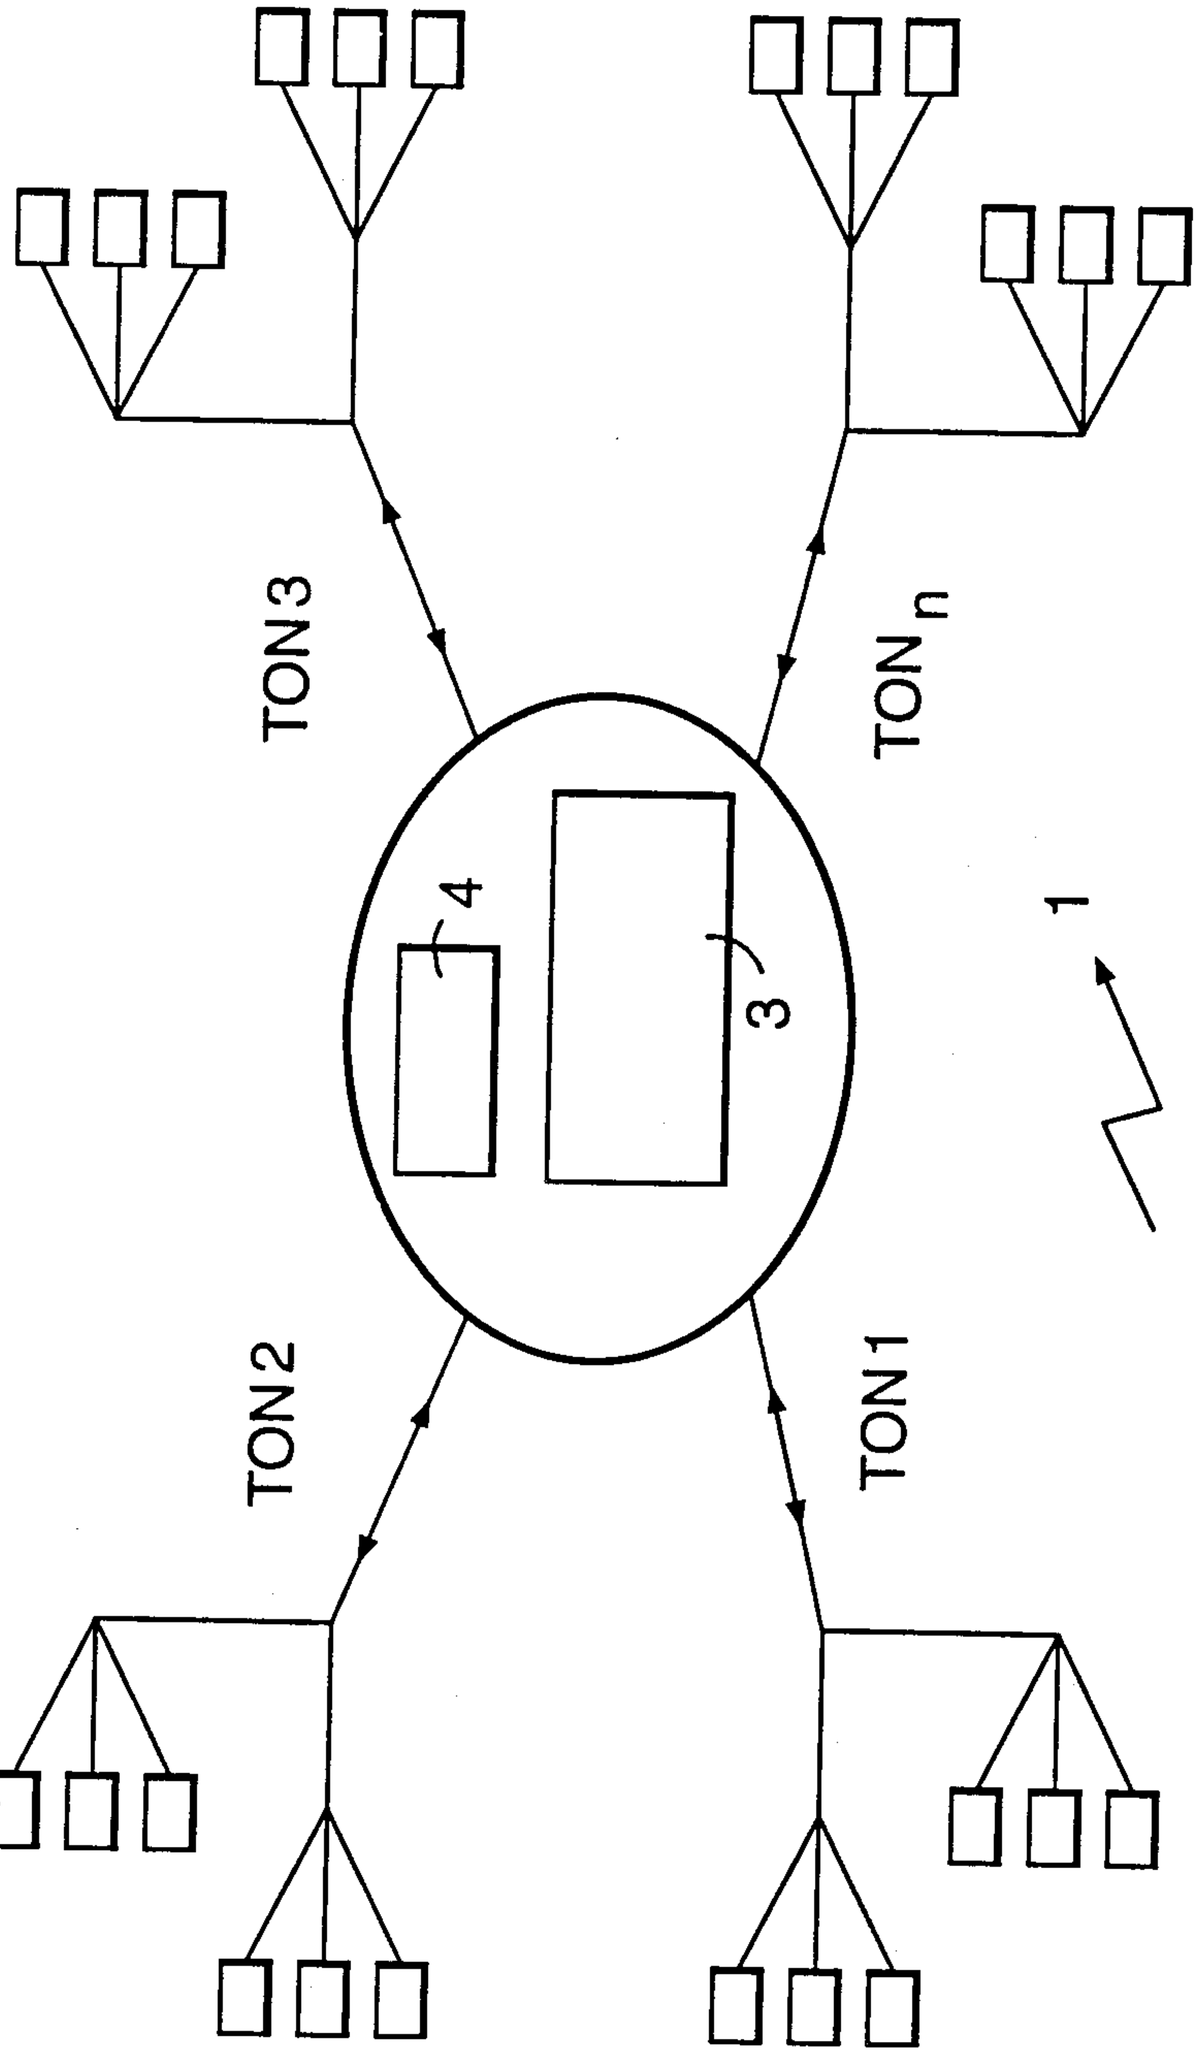


Fig.1.

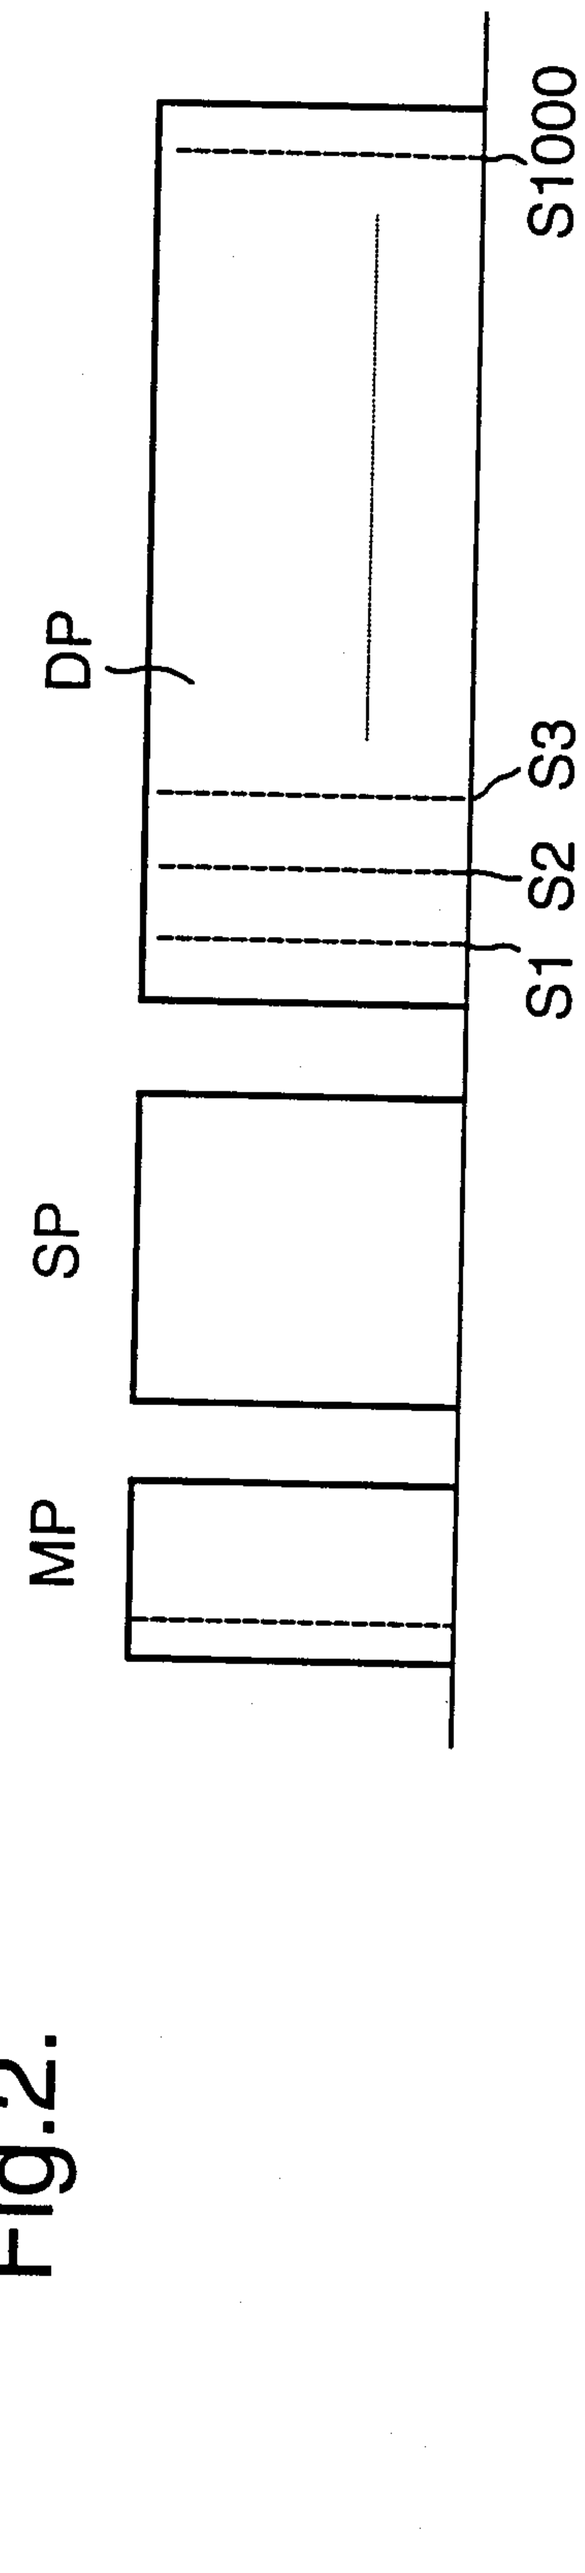


Fig.2.

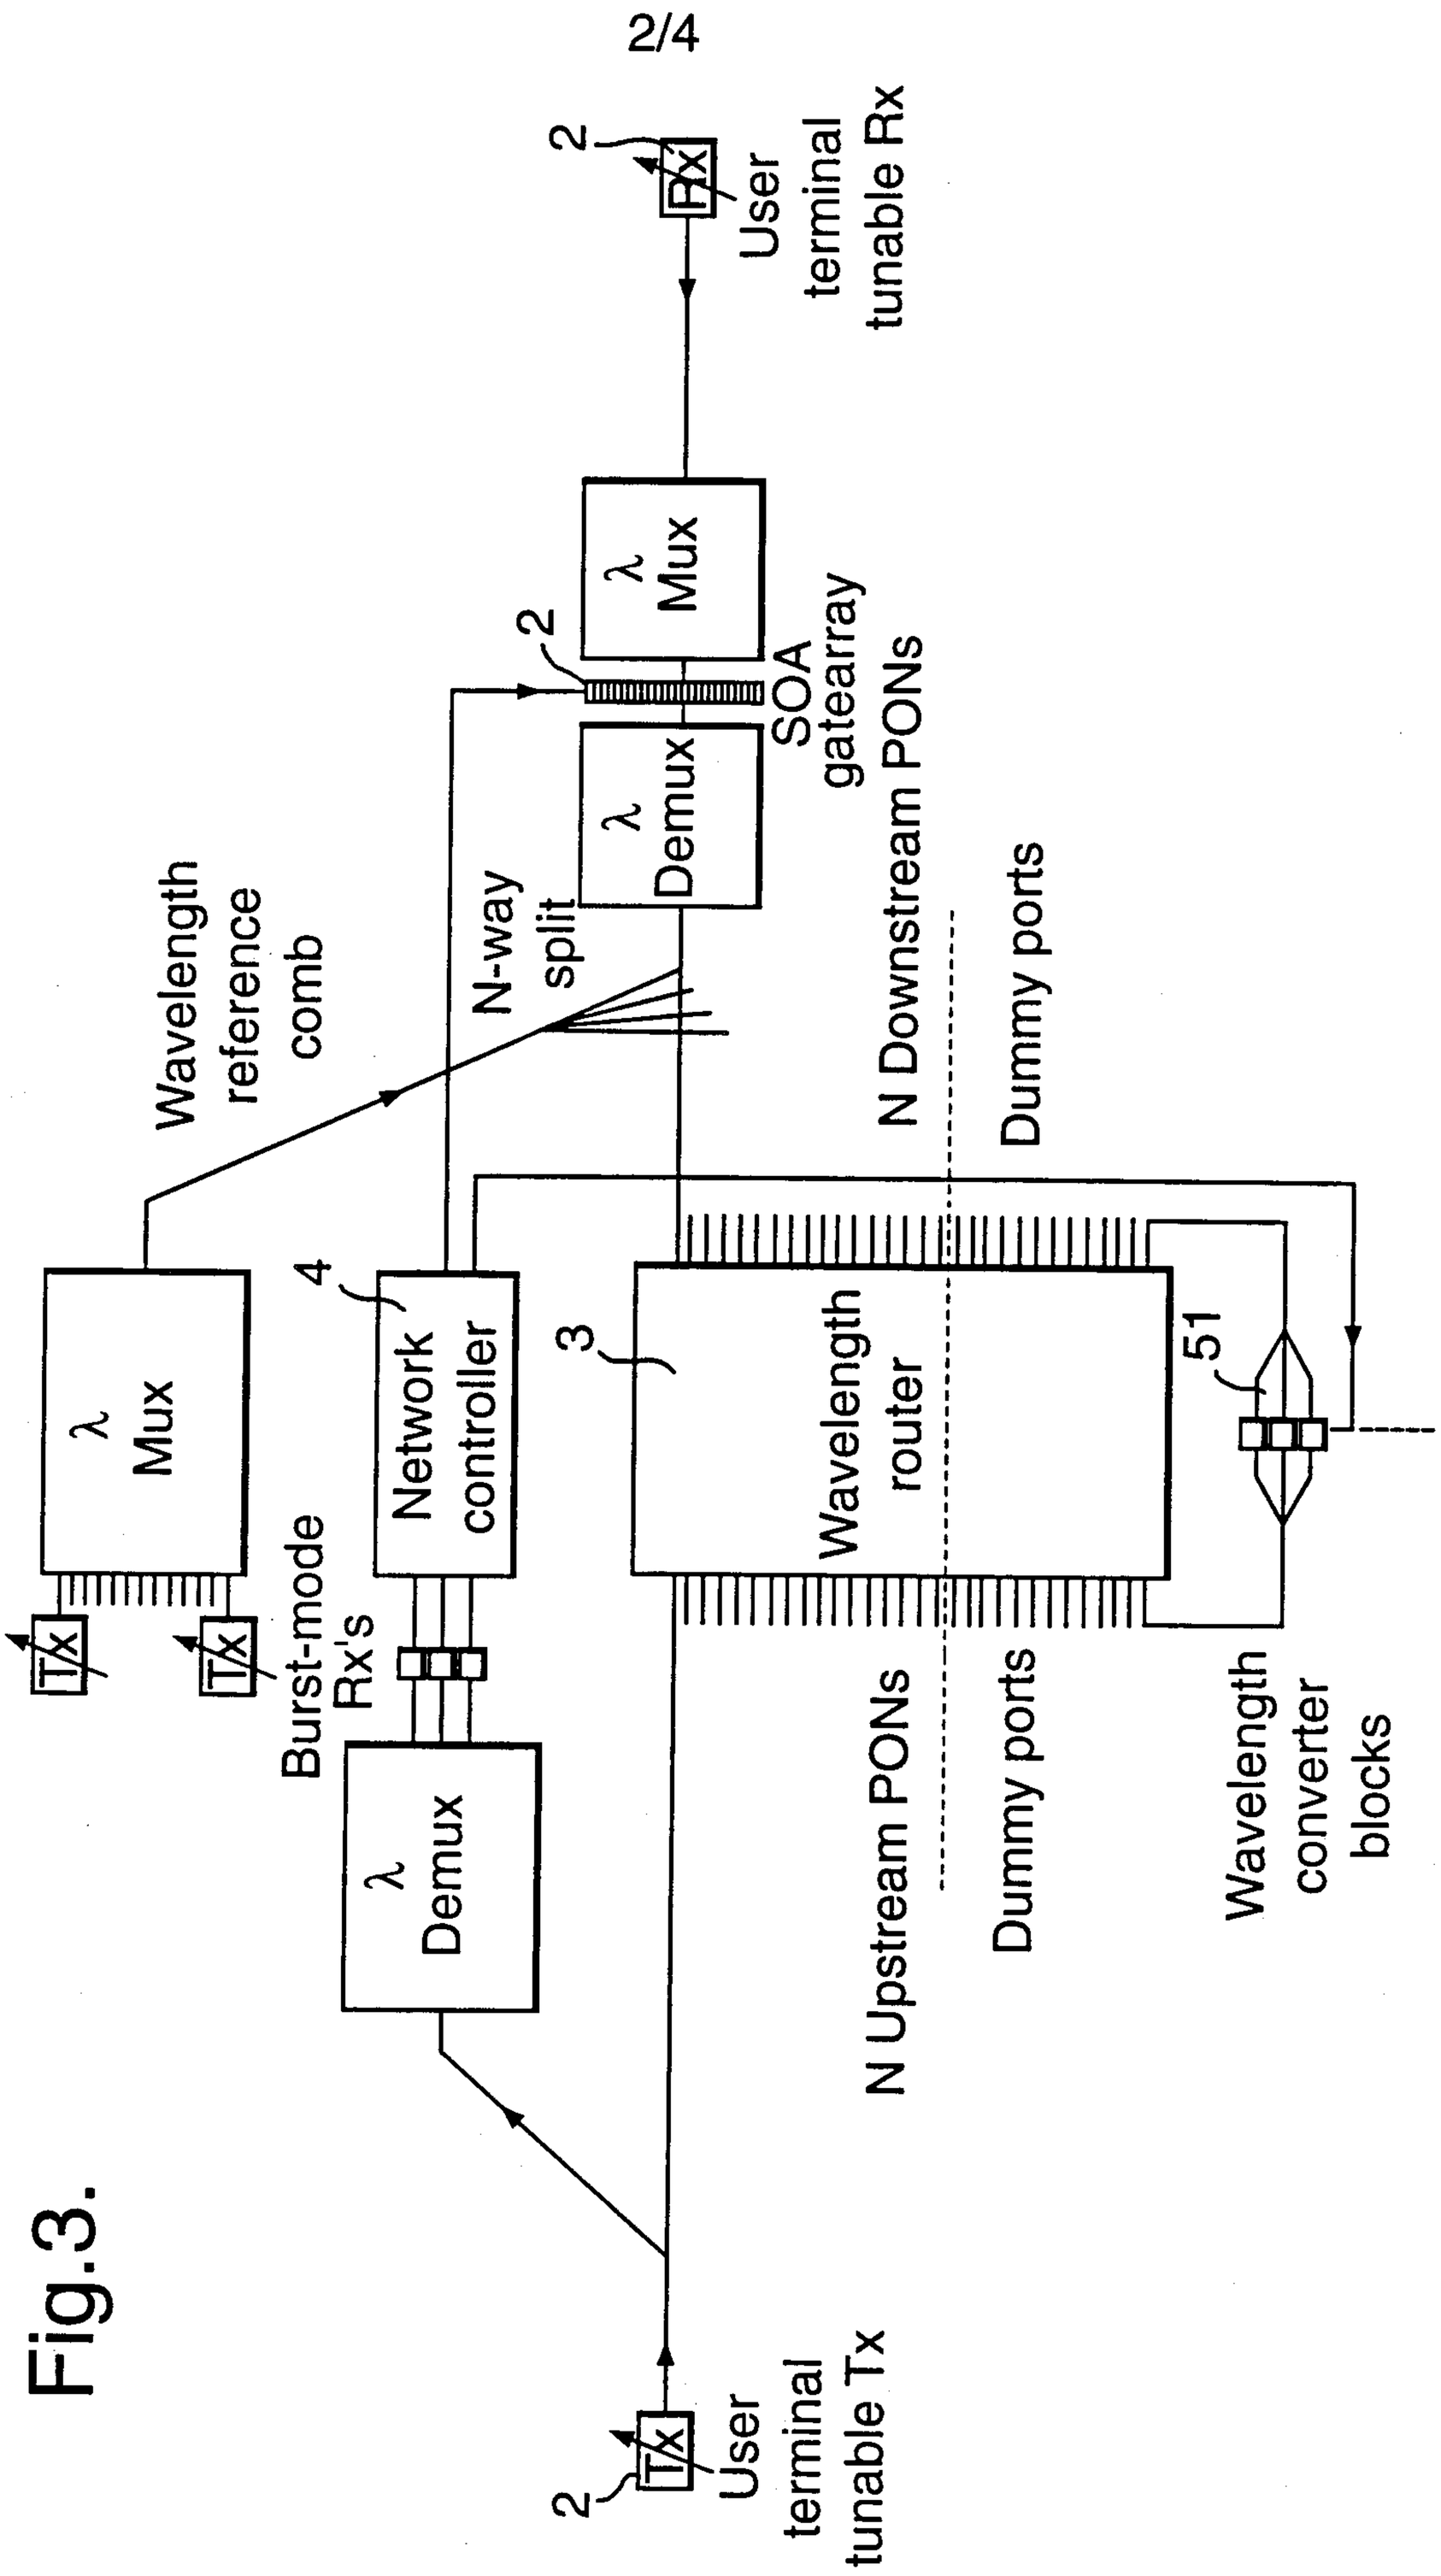

2
Rx
User terminal tunable Rx
λ Mux
2
SOA gatearray
λ Demux
Wavelength reference comb
N-way split
N Downstream PONs
Dummy ports
4
Network controller
3
Wavelength router
51
λ Mux
Tx
Tx
Burst-mode Rx's
λ Demux
N Upstream PONs
Dummy ports
Wavelength converter blocks
2
Tx
User terminal tunable Tx

Fig.3.

Fig.4.

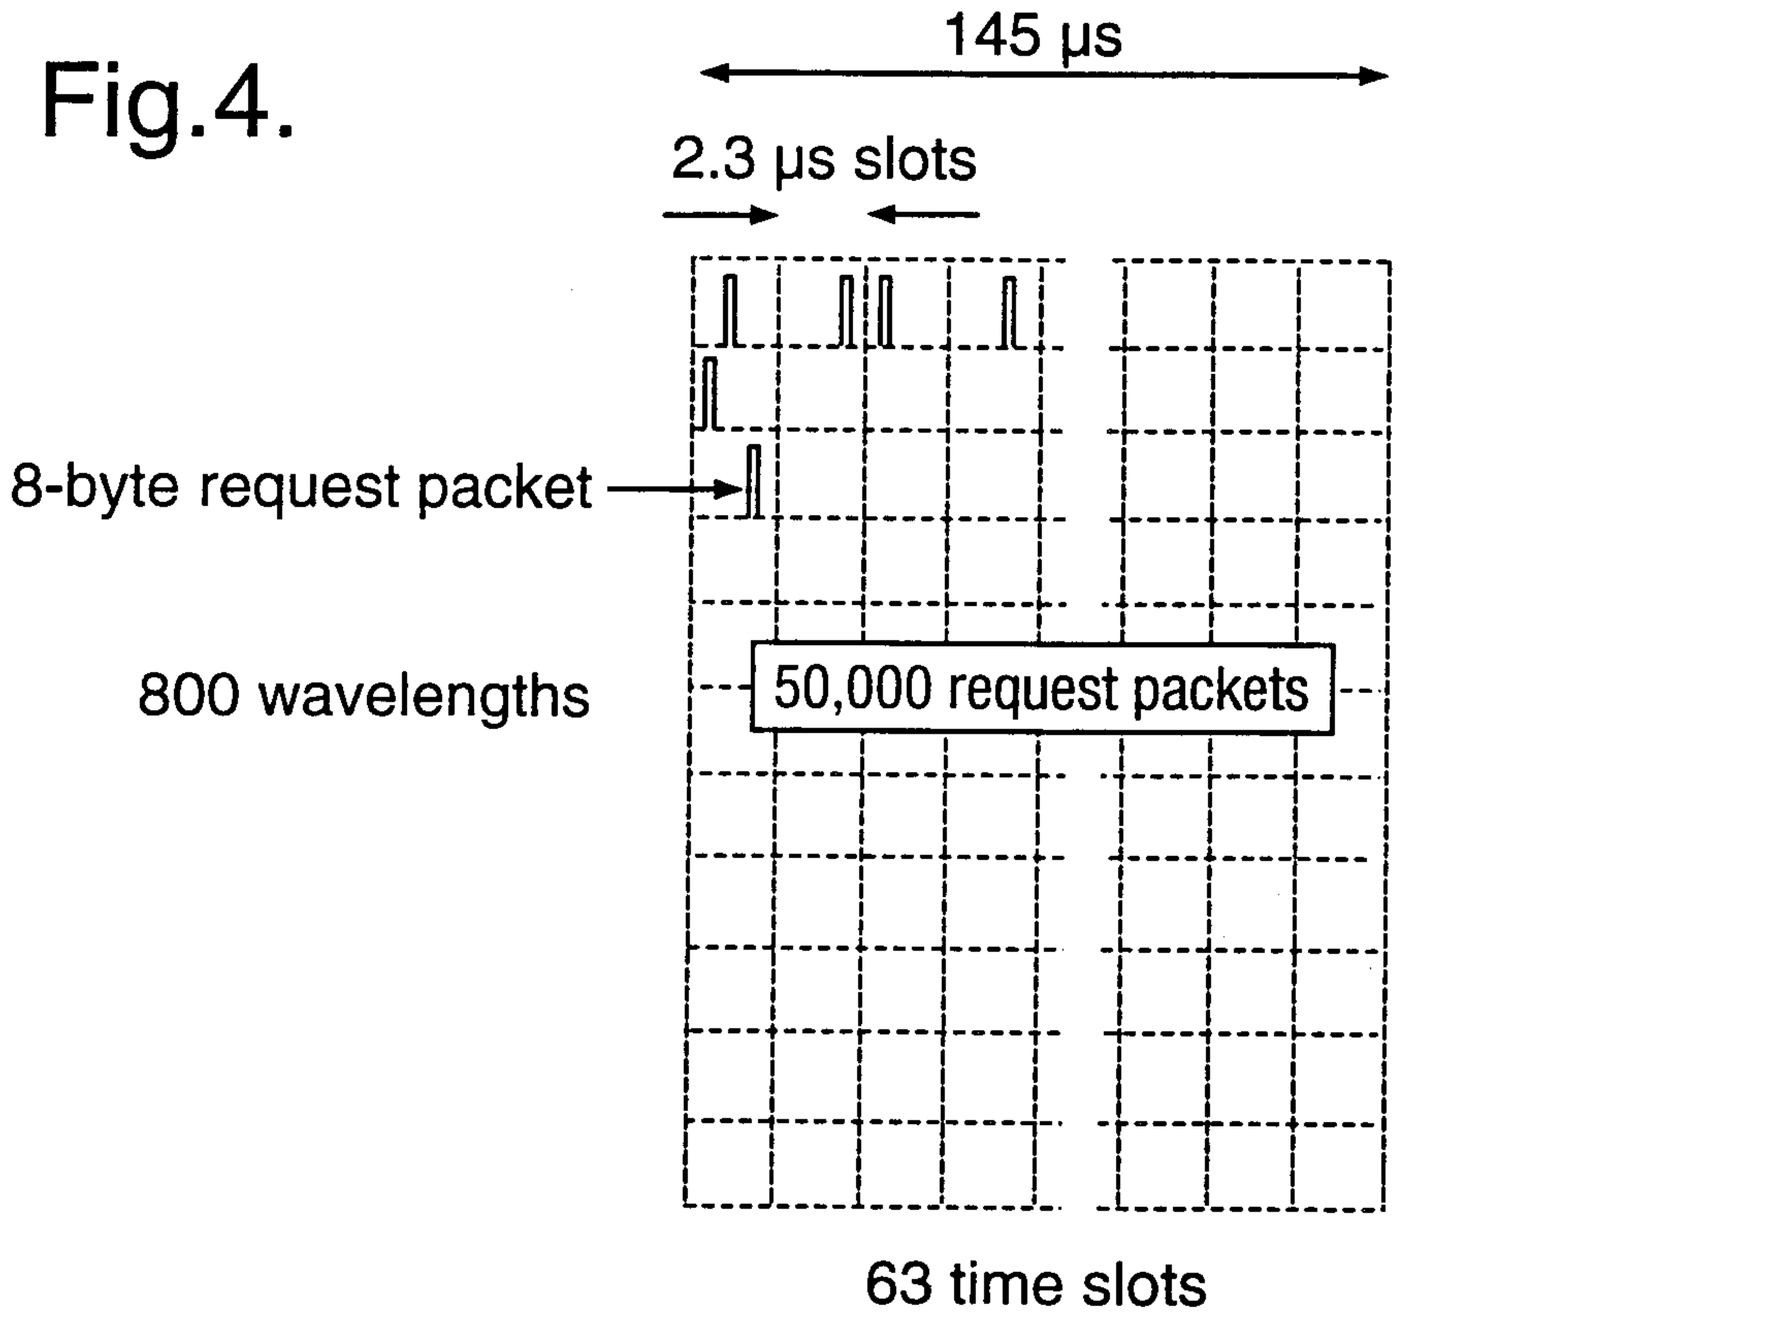
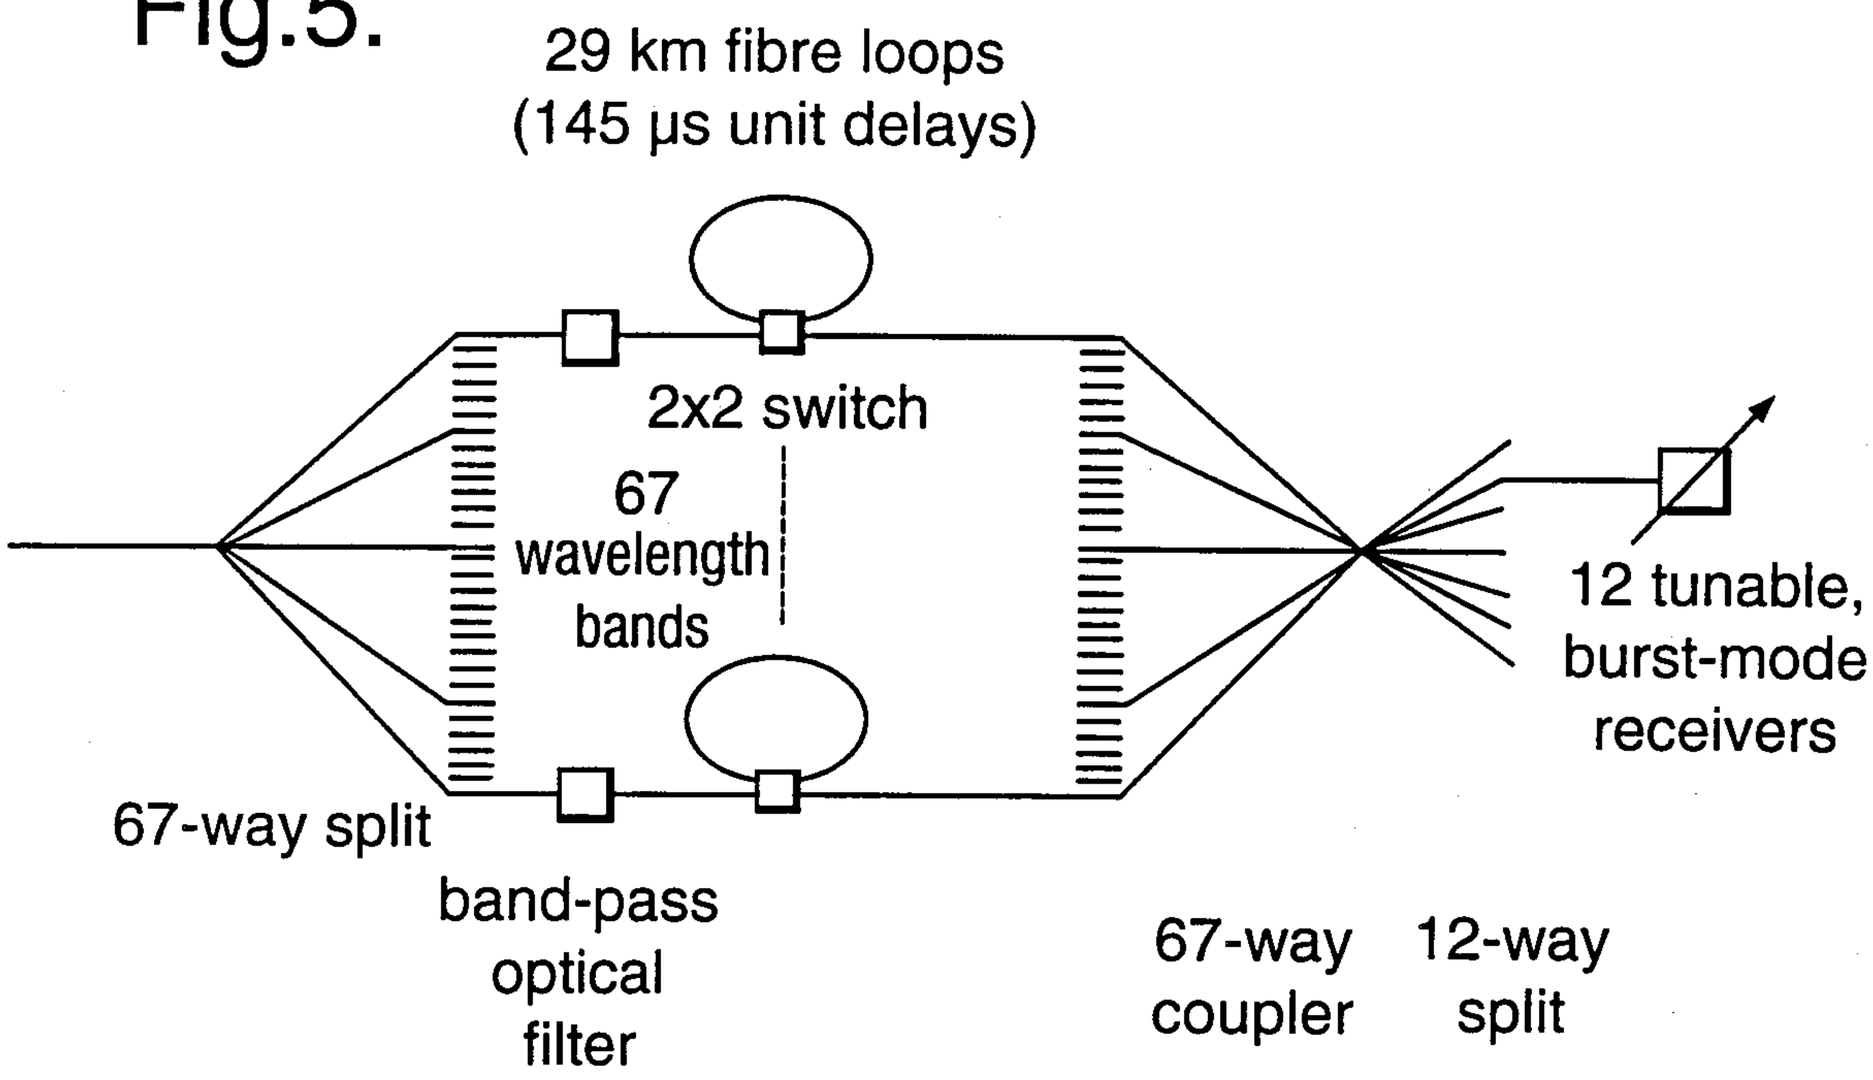

145 µs
2.3 µs slots
8-byte request packet
800 wavelengths
50,000 request packets
63 time slots

Fig.5.


29 km fibre loops
(145 µs unit delays)
2x2 switch
67
wavelength
bands
12 tunable,
burst-mode
receivers
67-way split
band-pass
optical
filter
67-way
coupler
12-way
split

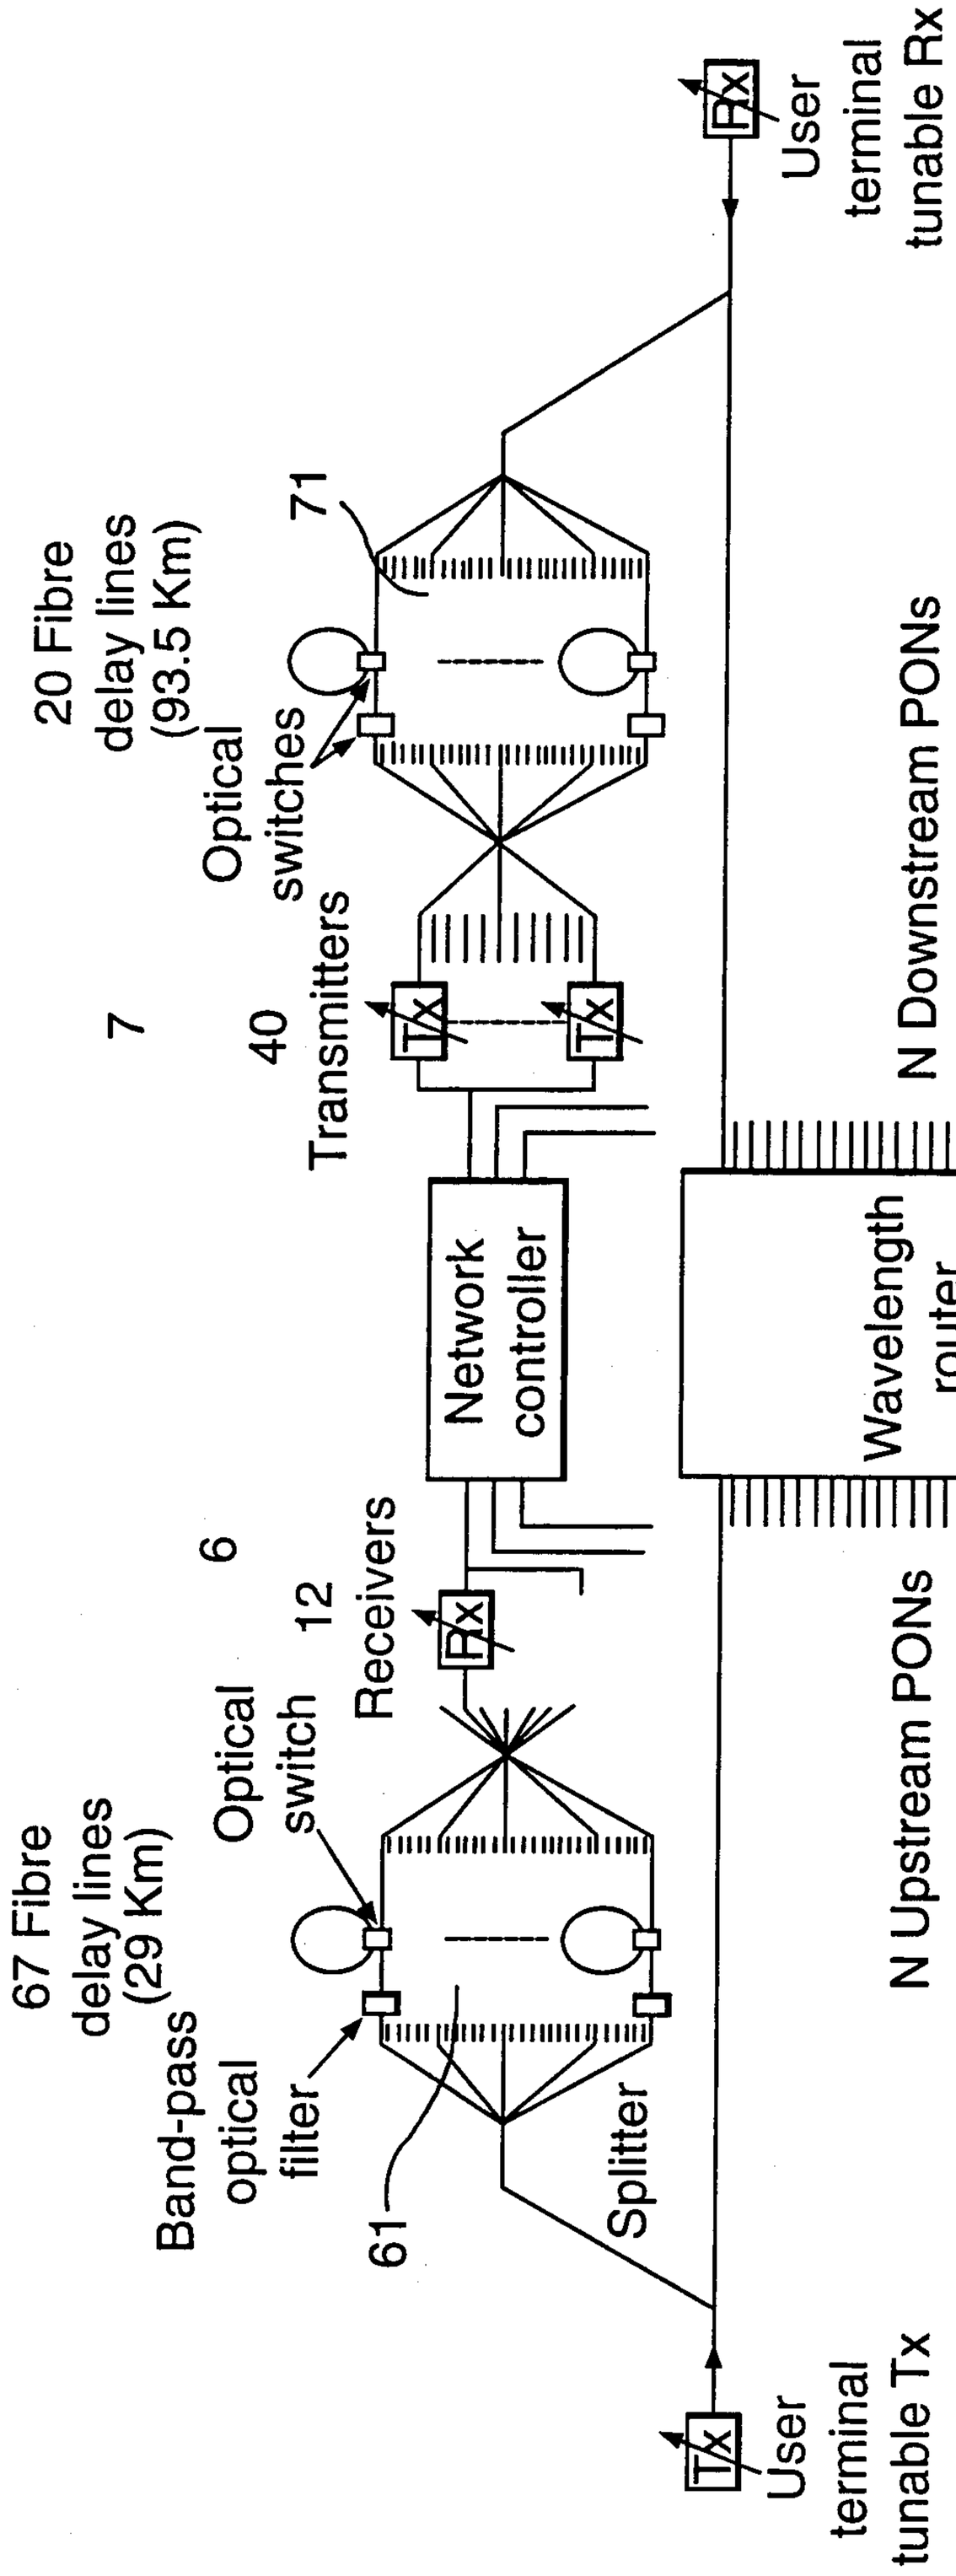
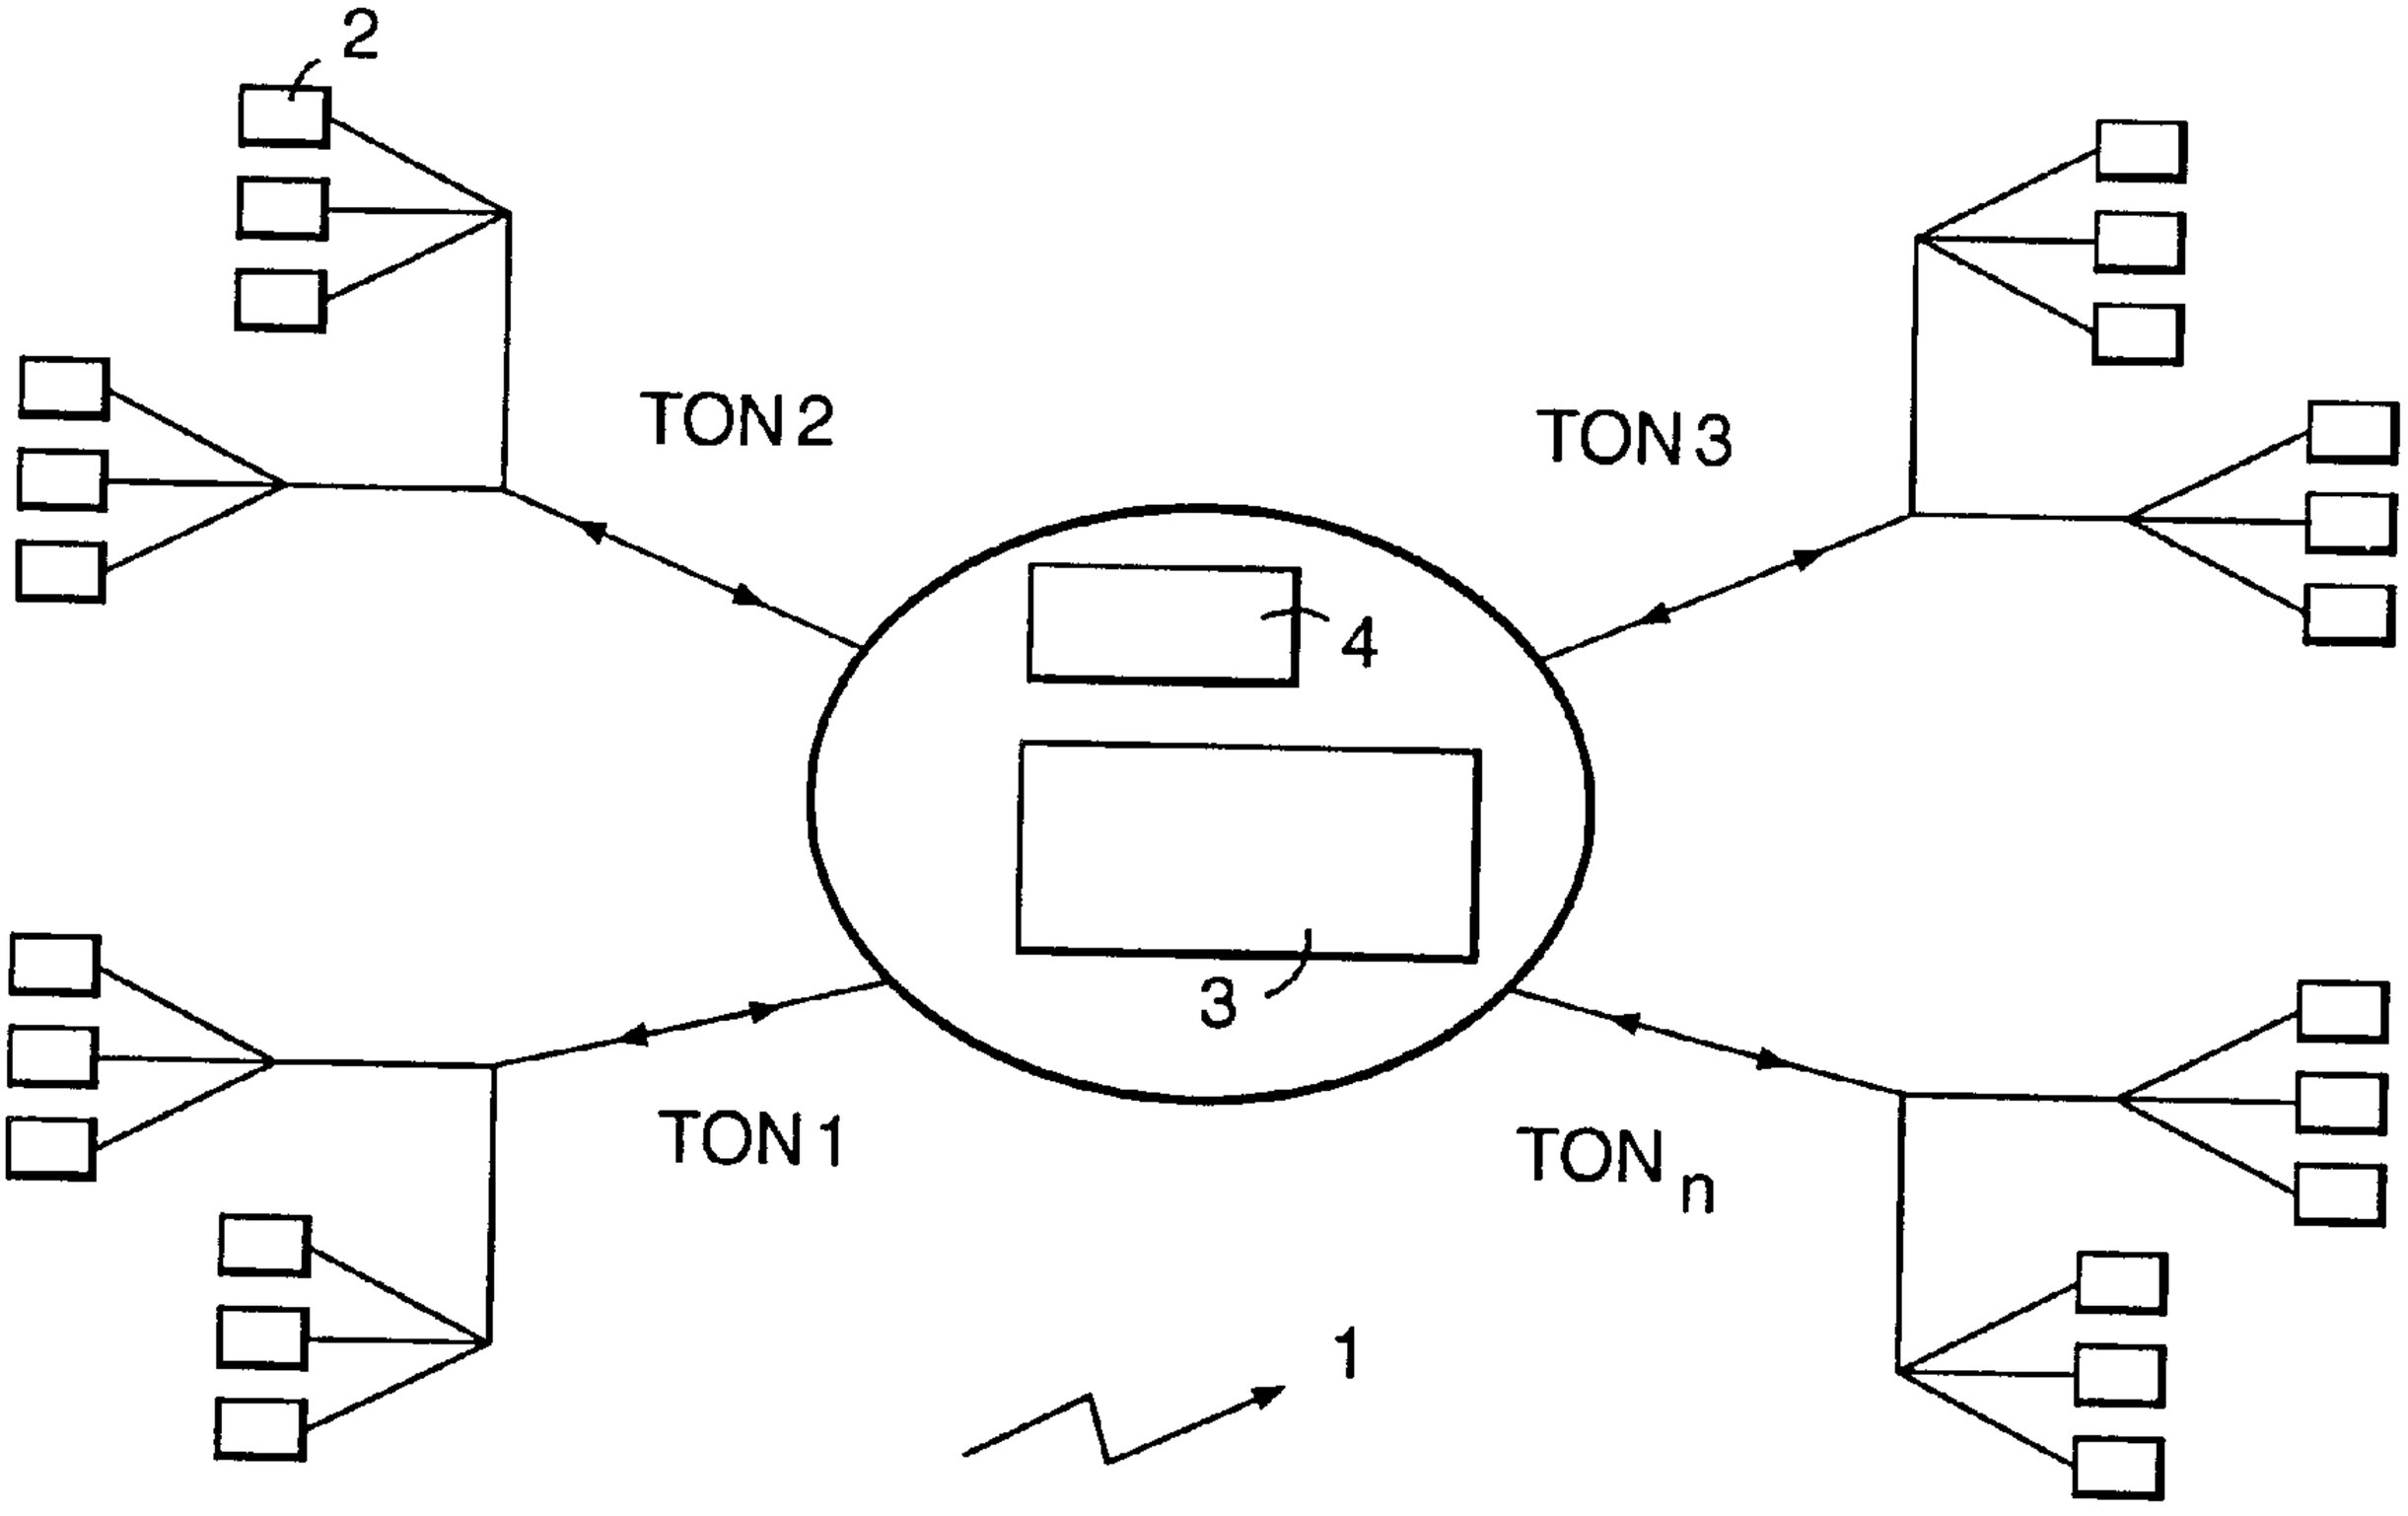

User terminal tunable Rx
Rx
71
20 Fibre delay lines (93.5 Km)
Optical switches
7
40
Transmitters
Tx
Network controller
6
Optical switch
12
Receivers
Rx
67 Fibre delay lines (29 Km)
Band-pass optical filter
61
Splitter
N Downstream PONs
Dummy ports
Wavelength router
N Upstream PONs
Dummy ports
Wavelength converter blocks
Tx
User terminal tunable Tx

Fig.6.


2
TON 2
TON 3
4
3
TON 1
TON n
1