



(51) International Patent Classification:

H04L 12/715 (2013.01) H04L 12/735 (2013.01)
H04L 12/707 (2013.01) H04L 12/721 (2013.01)

(21) International Application Number:

PCT/EP2016/055929

(22) International Filing Date:

18 March 2016 (18.03.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant: NEC EUROPE LTD. [DE/DE]; Kurfür-
sten-Anlage 36, 69115 Heidelberg (DE).

(72) Inventors: OECHSNER, Simon; Sandhäuser Strasse 10,
69124 Heidelberg (DE). HASSELMAYER, Peer; Kathar-
ina-Staritz-Str. 17, 60438 Frankfurt (DE).

(74) Agent: ULLRICH & NAUMANN; Schneidmühlstrasse
21, 69115 Heidelberg (DE).

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:

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,

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

(54) Title: SDN SUPPORT FOR DISJOINT MULTIPATH CONFIGURATION




Fig.

(57) Abstract: The present invention relates to SDN support for Disjoint Multipath Configuration. It allows a network controller (2) with network and traffic management functions and having knowledge of the network topology to support and control multipath connections, optimizing path selection and the number of subflows for these connections. These functions can be implemented without any changes to the client network stacks. The invention comprises, by a sender (3), initiating a multipath connection with a receiver (4) and establishing at least one subflow between the sender (3) and the receiver (4) along a first network path, and, by the network controller (2), determining the number of disjoint network paths available between the sender (3) and the receiver (4) and, based on this number, triggering the establishment of additional subflows between the sender (3) and the receiver (4) along network paths being at least partially disjoint to said first network path.



SDN SUPPORT FOR DISJOINT MULTIPATH CONFIGURATION

The present invention generally relates to a method and a system for supporting a multipath connection between two connection endpoints in a network.

5

The availability of multiple distinct network paths has given rise to the development of multipath communication, for example standardized in the form of MPTCP (Multipath Transmission Control Protocol). The basic idea behind multipath communication is that the reliability of a connection can be increased since link failures do not affect all paths, and that congestion can be avoided by distributing the traffic load intelligently among the paths.

10

However, for this to work optimally, the paths taken by the subflows of a multipath connection should be completely disjoint. This becomes clear if one imagines a jointly used link to fail, or being the bottleneck link of all paths. It is therefore of advantage to ensure that different subflows of the same multipath connection are forwarded over disjoint paths as far as possible. Since forwarding is typically not managed by the endpoints, the network entity controlling the forwarding should support multipath communication by providing this disjoint forwarding.

20

Research has been conducted on solutions to this problem. In Gregory Detal et al.: "Revisiting Flow-Based Load Balancing: Stateless Path Selection in Data Center Networks", Computer Networks, Volume 57, Issue 5, 7 April 2013, Pages 1204–1216, CFLB (Controllable per-Flow Load-Balancing) has been discussed, which implements end-host controlled next-hop selection via port number selection. The approach basically encodes the number of the outgoing interface to be taken at each hop into the source port number used for the subflow. The switches then decode this information and implement the forwarding decision. This approach necessitates global network topology knowledge at the end hosts in order to select a specific path, since otherwise different but random selection of next hop interfaces might lead again to shared bottlenecks. In addition, letting the end hosts decide on forwarding takes away control of network management from the network itself.

25

30

In Matthieu Coudron, Stefano Secci, Guy Pujolle, Patrick Raad, Pascal Gallard: "Cross-layer Cooperation to Boost Multipath TCP Performance in Cloud Networks", Cloud Networking (CloudNet), 2013, San Francisco, USA, and in Matthieu Coudron, Stefano Secci, Guido Maier, Guy Pujolle, Achille Pattavina:

5 "Boosting Cloud Communications through a Crosslayer Multipath Protocol Architecture", 2013 IEEE SDN for Future Networks and Services (SDN4FNS), Nov 2013, Trento, Italy, a cross-layer architecture to achieve the goal of disjoint forwarding is presented, which identifies the basic roles necessary (such as flow qualification and monitoring, path discovery, path computation, and path
10 enforcement). These papers show the feasibility of implementing the individual components and of the concept as a whole, but rely on active signaling between MPTCP-capable endpoints and the network controller to support disjoint forwarding.

15 Similarly, the approach proposed in Jingpu Duan, Zhi Wang, Chuan Wu: "Responsive Multipath TCP in SDN-based Datacenters", Proceedings of IEEE International Conference on Communications (IEEE ICC 2015), London, UK, June 8-12, 2015, necessitates signaling between a controller and a monitor installed on the end hosts.

20 Thus, the approaches also presuppose changes in the client network stack or alternatively in the hypervisors deployed on the physical hosts. This would necessitate, among other things, a standardized protocol for the signaling, as well as support and use of particular encapsulation techniques, creating additional
25 overhead and reducing the flexibility.

It is therefore an objective of the present invention to improve and further develop a method and a system for supporting a multipath connection between two connection endpoints in a network in such a way that the benefit of disjoint path
30 forwarding for multipath connections is achieved without needing changes to multipath capable end hosts and thus with increased flexibility.

In accordance with the invention, the aforementioned objective is accomplished by a method for supporting a multipath connection between two connection endpoints

in a network, wherein said connection endpoints include a sender and a receiver, the method comprising:

by said sender, initiating a multipath connection with said receiver and establishing at least one subflow between said sender and said receiver along a first network path,

by a network controller having knowledge of the network topology, determining the number of disjoint network paths available between said sender and said receiver and, based on said number, triggering the establishment of additional subflows between said sender and said receiver along network paths being at least partially disjoint to said first network path.

Furthermore, the above objective is accomplished by a system with multipath connection support, the system comprising

a network comprising a plurality of switches,

at least two devices connected to one or more of the plurality of switches, wherein said at least two devices include a sender and a receiver being the two connection endpoints of a multipath connection that includes at least one subflow established between said sender and said receiver along a first network path, and

a network controller being connected to one or more of the plurality of switches in the network and having knowledge of the topology of said network, wherein said network controller is configured to determine the number of disjoint network paths available between said sender and said receiver and, based on said number, to trigger the establishment of additional subflows between said sender and said receiver along network paths being at least partially disjoint to said first network path.

In accordance with embodiments of the invention it has been recognized that the above mentioned objective can be accomplished by a network controller with network and traffic management functions to support and control multipath connections, optimizing path selection and the number of subflows for these connections. The network controller, which is assumed to have knowledge of the network topology, may comprise a logic integrated with and/or executable by a processor of the network controller, wherein the logic is configured to calculate the number of available disjoint paths. Based on this number, the network controller

applies an algorithm for selecting and establishing additional forwarding paths inside the network and triggers the establishment of new subflows along these paths, thereby increasing path multiplicity. It should be noted that the network paths calculated by the network controller do not necessarily have to be disjoint completely. Depending on the network topology and current load of the network, partially disjoint paths can be beneficial to use, if a bottleneck link can be augmented or replaced by multiple other paths through the network. Up until the branching points, the paths will share certain links, but will be placed on disjoint ones for other hops in the path. Such a setup can still improve throughput as congested links can be sidestepped.

Embodiments of the present invention allow the use of multiple networking paths with the minimum number of subflows necessary, even if the endpoints do not have multiple network interfaces or multiple (IP) addresses, and even if they do not know about the existence of the different network paths. The presented method is even most useful if the clients to be connected are multipath capable, but possess only a single network interface. In this situation, multiple subflows could still be opened, but an endpoint without topology knowledge will not know what the actual number of disjoint paths available for different subflows is. Opening too many subflows incurs overhead, which can be avoided by embodiments of the present invention.

According to the present invention, the implemented functions do not require any changes to the networking stacks in the connection endpoints. In particular, any manual settings on the connection endpoints and/or on the network equipment in order to increase the number of network paths becomes dispensable due to the quasi automation of this process achieved by the present invention. In comparison to flow-based load balancing the method and system according to the present invention allow for optimal number of subflows between individual logical endpoint connections. In comparison to packet-based load-balancing the present invention avoids reordering and enables joint congestion control over subflows. In comparison to solutions where connection endpoints just open a large number of subflows until all disjoint network paths are used, or perform tests by removing

'slow' subflows and trying new ones, an implementation of the present invention significantly reduces the overhead on the connection endpoints.

5 Since embodiments of the present invention require programmatic access to networking elements, an implementation of the present invention in software-defined networks, SDN, is considered most suitable.

10 According to an embodiment of the invention the network controller may trigger the establishment of additional subflows between two endpoints of an existing connection by creating messages that contain an option of announcing additional addresses on which said connection endpoints can be reached, and injecting these messages into the network, e.g. into data streams towards the connection endpoints. In this regard the implementation of a multipath protocol that allows for the extraction of the necessary information to create protocol-conformant address
15 announcements proves to be beneficial. For instance, in case of using MPTCP (i.e. the network controller, the plurality of switches and the connection endpoints are configured to exchange MPTCP messages), the above mentioned option of announcing additional addresses could be, e.g., the add address option 'ADD_ADDR' as implemented in MPTCP.

20 According to an embodiment the network paths of the additional subflows may have forwarding rules that are based on virtual destination addresses for the two connection endpoints. The network controller may announce these virtual destination addresses in the network to be used by the connection endpoints for
25 establishing the additional subflows. From the perspective of a connection endpoint, the announced additional virtual addresses seem to belong to the remote endpoint of the connection (i.e. the connection endpoints are not aware of the announced addresses being virtual), thus implicitly informing the endpoints of the number of available paths. The connection endpoints will then in turn establish
30 new subflows between each other using said virtual addresses.

The virtual addresses may be translated within the network to simulate the necessary number of interfaces to both endpoints. To this end, i.e. for rewriting source and destination addresses contained in the headers of packets of

additional subflows, the network controller may install appropriate address translation rules in the network, for instance in those switches that are the topologically closest to the sender and to the receiver, respectively. In effect, the network thus functions as a NAT (Network Address Translation) component.

5

With respect to an efficient implementation that hinders clients/senders to establish additional subflows on their own (which, due to insufficient network topology knowledge on the side of the client/sender, will likely be suboptimal in terms of path diversity), according to an embodiment the network controller may prevent direct establishment of additional subflows between the sender and the receiver, i.e. the 'normal' establishment of subflows following the respective multipath protocol where the connection endpoints use their known addresses. For instance, such a prevention of direct subflow establishment may be achieved by blocking signaling packets related to any additional subflow handshake, e.g. MP_JOIN packets in the case of using MPTCP).

According to an alternative embodiment it may be provided that the network controller generally allows direct establishment of additional subflows between the sender and the receiver. However, in this case, in order to ensure disjoint network path configurations for these subflows, the network controller may perform monitoring of those subflows established by the sender directly. For instance, it may be provided that the network controller applies a tracking mechanism for tracking multipath protocol control packets and for retaining information about existing multipath connections and dependent subflows. In particular, the tracking mechanism may identify which subflows belong to the same connection between the two connection endpoints. Subflows identified by this tracking mechanism may then be forwarded by the network controller on paths as disjoint as possible.

Additionally or alternatively, the tracking mechanism may include informing the network controller of the initial handshake packets for the connection between the sender and the receiver (e.g., in case of using MPTCP, an MPTCP SYN with MP_CAPABLE option) and storing relevant parameters of this connection (e.g., MPTCP tokens). This tracking of MPTCP tokens in the controller, which is not part of the data end-to-end connection, will enable the controller to know dependent

30

subflows and to place dependent subflows on different network paths. Still further, the tracking mechanism may include monitoring address announcements originating from the connection endpoints (e.g., in case of using MPTCP, ADD_ADDR packets) and performing predictions of the establishment of any new subflows to such an announced address shortly after the announcement.

There are several ways how to design and further develop the teaching of the present invention in an advantageous way. To this end it is to be referred to the dependent patent claims on the one hand and to the following explanation of preferred embodiments of the invention by way of example, illustrated by the drawing on the other hand. In connection with the explanation of the preferred embodiments of the invention by the aid of the drawing, generally preferred embodiments and further developments of the teaching will be explained. In the drawing the only

Fig. is a schematic view illustrating exemplarily the establishment of a new subflow triggered by a network controller in accordance with embodiments of the present invention.

Embodiments of the present invention, as described hereinafter in connection with the figure, relate to implementations that use the MPTCP (Multipath TCP) protocol and software-defined networking (SDN) mechanisms. However, as will be appreciated by those skilled in the art, the present invention would be applicable also in the case of other protocols supporting multipath communications having similar characteristics as MPTCP, as well as with control mechanisms other than SDN that enable a controller entity to manage and control communication endpoints in the network.

According to its general implementation, MPTCP creates individual subflows for distinct address pairs created from the IP addresses of the two endpoints of a TCP connection. First, a single subflow is established, followed by setting up additional subflows. Alternative IP addresses of the server endpoint of a connection are communicated via MPTCP control messages to the client, which is then

establishing additional subflows to these addresses based on the hope that paths to the additional addresses will follow separate network paths.

In the embodiment illustrated in the figure, a data transfer network 1 having SDN capabilities is assumed. The network controller 2 is managing the network 1 and therefore has knowledge of the devices that make up that network 1. Upon the establishment of a first connection between two network endpoints, the controller 2 calculates all available disjoint network paths between these endpoints. Such calculation may be affected by the current load on the links participating in the network paths. This means that the network controller 2, while basically having identified the existence of a certain number of disjoint network paths (e.g. five), only considers a lower number of paths (e.g. three) suitable for forwarding subflows along these paths, since the remaining paths (i.e. two paths in the above example) are too heavily loaded for carrying any additional subflows.

If the controller 2 has calculated the availability of more disjoint paths between the endpoints than the number of subflows that have been set up, the controller 2 will trigger the creation of additional subflows. Preferably, the number of additionally triggered subflows is such that it exactly matches the difference between the calculated number of paths being available (and, possibly, being considered suitable by the network controller) and the number of existing subflows that have already been set up. The triggering is achieved by generating address announcements in the network 1 and sending them to the endpoint that initiated the multipath connection, which in turn typically will try and establish a new subflow to this new address. It should be noted that the announced addresses do not need to correspond to actual interfaces on the endpoints, but can be handled transparently by the network 1.

In the specific scenario illustrated in the figure, it is assumed that a transport connection exists between a first connection endpoint, which is sender 3, and a second connection endpoint, which is receiver 4. The sender 3 has a network interface with address AS, and the receiver 4 has a network interface with assigned address AR. In the figure, this existing transport connection, which is

assumed to include three network switches 5 (denoted Sw1, SwX, and Sw2), is indicated by a solid line and termed 'path P1'.

5 The network controller 2 recognizes that the multipath-capable connection between sender 3 and receiver 4 (hereinafter sometimes briefly denoted 'connection S-R') has opened less subflows than there are disjoint paths available in the network 1. For instance, in the illustrated scenario it is assumed that there is the disjoint network path including the network switches 5 denoted Sw1, SwY, and Sw2. In the figure this additional disjoint network path is indicated by a dotted line
10 and termed 'path P2'.

The network controller 2 establishes a new path in the network 1, which is yet unused by any subflow of the connection S-R. This path has forwarding rules based on virtual destination addresses for the two endpoints, VAS and VAR for
15 sender 3 and receiver 4, respectively. These virtual addresses (that distinguish different subflows without the endpoints being aware of these addresses being virtual) do not need to be assigned to any actual interface of sender 3 or receiver 4, and actually must not be assigned to any other interface.

20 In addition to establishing the forwarding based on these destination addresses, the controller 2 also installs header-rewrite rules in the network 1, specifically in switches Sw1 and Sw2. These rules rewrite packet headers in the following way: packets with source AS and destination VAR are rewritten to show source address VAS and destination address AR. On the other hand, packets with source address
25 AR and destination address VAS are rewritten to show source address VAR and destination address AS. In effect, the network 1 acts as a NAT component. Any additional information necessary to be rewritten is included in the rules as well (e.g., address IDs). It should be noted that in the embodiment illustrated in the figure rule installation is shown at exemplary network elements. As will be easily
30 appreciated by those skilled in the art, the rules may be installed at different network elements than the ones shown in the figure.

The controller 2 announces VAR to the sender 3. This triggers the establishment of a new subflow from the sender 3. From the perspective of the sender 3, this

subflow is between AS and VAR, but in effect it will be between AS and AR. However, due to the forwarding rules installed by the controller 2, it will take a path different from the path of the already established subflows between sender 3 and receiver 4, i.e. path P2 instead of path P1 of the initial subflow.

5

The embodiment described above has the advantage that connection endpoints with only one network interface that are unaware of multiple paths between themselves can be triggered to use the right number of subflows to exploit disjoint paths in the network. Furthermore, if different sub-networks are under the control

10 of different network controllers, each can increase the number of subflows based on its local observations of available paths and with a destination address selected to facilitate routing/forwarding in its own network. However, in this case, care needs to be taken to no provoke address ID collisions.

15 Although not explicitly shown in the figure, the method described above also works if the endpoints of a connection, i.e. sender 3 and/or receiver 4, have multiple interfaces configured. In this case, the controller 2 can choose between the following options:

20 According to one implementation (option 1), the controller 2 can prevent the 'normal' (i.e., following standard MPTCP, for instance) establishment of subflows between the sender 3 and the receiver 4 using their known addresses, e.g., by blocking the signaling packets (such as MP_JOIN) used for this subflow handshake.

25

According to an alternative implementation (option 2), the controller 2 can allow the 'normal' (i.e., following standard MPTCP, for instance) establishment of subflows (e.g., between all possible distinct pairs of addresses of the sender 3 and the receiver 4), with or without ensuring disjoint path forwarding for these flows.

30 Disjoint path forwarding can be ensured by letting the controller 2 track which subflows seen by its network 1 belong to the same bonded end-to-end connection, in order to forward them on paths as disjoint as possible, or according to other optimization criteria.

According to one option for this subflow tracking, the controller 2 may be informed of the initial handshake packets for the connection between the sender 3 and the receiver 4 (e.g., an MPTCP SYN with MP_CAPABLE option), and the controller 2 can store relevant parameters of the connection (e.g., MPTCP tokens). According to another option address announcements (e.g., ADD_ADDR packets in MPTCP) may be monitored in order to predict the establishment of new subflows to such an announced address shortly after the announcement. According to still another option, preconfigured and known addresses may be used on both endpoints of the connection, enabling the mapping of addresses to subflows and thus disjoint forwarding. This might be possible within a datacenter environment, but is more difficult for WAN connections.

If this disjoint path forwarding for regularly established flows is not to be supported, the controller 2 still could store the same information (i.e., tokens, IDs, keys, etc.) in order to be able to judge whether additional subflows over other paths would be beneficial for the connection S-R. Since the necessary effort is the same, the first mentioned implementation (i.e. option 1) would be preferable.

Regardless of the chosen option, the controller 2 can always trigger the controlled establishment of new subflows using the method described above in case it recognizes this as beneficial.

The embodiment of the present invention described herein includes (part of) an SDN controller 2, e.g., OpenDaylight, for the use with the standardized MPTCP. SDN allows for path establishment on the network elements, using e.g., OpenFlow. A way of implementing the monitoring of client-established subflows includes listening to MPTCP ADD_ADDR options, with which end hosts can announce a useable address to the remote endpoint, then reactively establishing the path when a MP_JOIN packet is seen in reverse direction. An alternative way includes collecting tokens exchanged between end hosts in the first handshake and storing them in a small database. These tokens are re-used when new subflows are established, allowing to reactively establish a new path for the new subflow and to announce new addresses in a valid fashion via existing subflows of a bundled connection.

The announcement of a new (virtual) address to an endpoint can be implemented by creating an OpenFlow *packet_out* message containing a valid TCP packet (such as a copy of an ACK) with the ADD_ADDR MPTCP option set and the
5 known token/keys derived from the monitored initial connection handshake.

In addition to the header addresses that need to be rewritten, in the case of MPTCP address IDs need to be assigned by the controller 2 to the virtual endpoints of the connection, and need to be replaced in the MP_JOIN options
10 together with the IP addresses and potentially port numbers.

Many modifications and other embodiments of the invention set forth herein will come to mind the one skilled in the art to which the invention pertains having the benefit of the teachings presented in the foregoing description and the associated
15 drawings. Therefore, it is to be understood that the invention is not to be limited to the specific embodiments disclosed and that modifications and other embodiments are intended to be included within the scope of the appended claims. Although specific terms are employed herein, they are used in a generic and descriptive sense only and not for purposes of limitation.

20

List of reference numbers

	1	data transfer network
	2	network controller
5	3	sender
	4	receiver
	5	network switch

10

Claims

1. Method for supporting a multipath connection between two connection endpoints in a network (1), wherein said connection endpoints include a sender (3) and a receiver (4), said method comprising:
- 5 by said sender (3), initiating a multipath connection with said receiver (4) and establishing at least one subflow between said sender (3) and said receiver (4) along a first network path,
- 10 by a network controller (2) having knowledge of the network topology, determining the number of disjoint network paths available between said sender (3) and said receiver (4) and, based on said number, triggering the establishment of additional subflows between said sender (3) and said receiver (4) along network paths being at least partially disjoint to said first network path.
- 15 2. Method according to claim 1, wherein the establishment of additional subflows between said sender (3) and said receiver (4) is triggered by said network controller (2) by creating messages that contain an option of announcing additional addresses on which said connection endpoints can be reached and injecting said messages into said network (1) towards said connection endpoints.
- 20 3. Method according to claim 1 or 2, wherein network paths of said additional subflows have forwarding rules based on virtual destination addresses for said two connection endpoints.
- 25 4. Method according to any of claims 1 to 3, wherein said network controller (2) announces virtual destination addresses in the network (1) to be used by said connection endpoints for establishing said additional subflows.
- 30 5. Method according to any of claims 1 to 4, wherein said network controller (2) installs address translation rules in the network (1) for rewriting source and destination addresses contained in the headers of packets of said additional subflows.

6. Method according to any of claims 1 to 5, wherein said network controller (2) prevents direct establishment of additional subflows between said sender (3) and said receiver (4) by blocking signaling packets related to said additional subflow handshakes.

5

7. Method according to any of claims 1 to 5, wherein said network controller (2), while allowing direct establishment of additional subflows between said sender (3) and said receiver (4), performs monitoring of subflows established by said sender (3).

10

8. Method according to any of claims 1 to 7, wherein said network controller (2) applies a tracking mechanism that identifies which subflows belong to the same connection between said two connection endpoints, and wherein said network controller (2) forwards these identified subflows on paths as disjoint as possible.

15

9. Method according to any of claims 1 to 8, wherein said tracking mechanism includes informing said network controller (2) of the initial handshake packets for the connection between said sender (3) and said receiver (4) and storing predefined parameters of the connection.

20

10. Method according to any of claims 1 to 9, wherein said tracking mechanism includes monitoring address announcements originating from said connection endpoints and performing predictions of the establishment of any new subflows to such an announced address shortly after the announcement.

25

11. System with multipath connection support, in particular for executing a method according to any of claims 1 to 10, said system comprising

a network (1) comprising a plurality of switches (5),

30

at least two devices connected to one or more of the plurality of switches (5), wherein said at least two devices include a sender (3) and a receiver (4) being the two connection endpoints of a multipath connection that includes at least one subflow established between said sender (3) and said receiver (4) along a first network path, and

a network controller (2) being connected to one or more of the plurality of switches (5) in the network (1) and having knowledge of the topology of said network (1), wherein said network controller (2) is configured

5 to determine the number of disjoint network paths available between said sender (3) and said receiver (4) and, based on said number, to trigger the establishment of additional subflows between said sender (3) and said receiver (4) along network paths being at least partially disjoint to said first network path.

10 12. System according to claim 11, wherein said sender (3) and/or said receiver (4) each comprise only one network interface.

15 13. System according to claim 11 or 12, wherein said network controller (2), said plurality of switches (5) and said connection endpoints are configured to exchange MPTCP, Multipath Transmission Control Protocol, messages among each other.

14. Network switch, in particular SDN switch (5), configured for being employed in a method and/or in a system according to any of claims 1 to 13.

20 15. Network controller, in particular SDN controller (2), configured for being employed in a method and/or in a system according to any of claims 1 to 13.




Fig.

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2016/055929

A. CLASSIFICATION OF SUBJECT MATTER INV. H04L12/715 H04L12/707 H04L12/735 H04L12/721 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) H04L		
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, 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	Savvas Zannettou ET AL: "Exploiting path diversity in datacenters using MPTCP-aware SDN", 2016 IEEE Symposium on Computers and Communication (ISCC), 30 November 2015 (2015-11-30), pages 539-546, XP055318143, DOI: 10.1109/ISCC.2016.7543794 ISBN: 978-1-5090-0679-3 Retrieved from the Internet: URL: https://arxiv.org/pdf/1511.09295v1.pdf [retrieved on 2016-11-10]	1-4,7-9, 11-15
A	figures 1,2 page 1, column 2, line 14 - line 23 page 2, column 2, line 19 - line 34 page 2, column 2, line 52 - page 3, column 1, line 7 page 3, column 1, line 25 - line 37 page 3, column 1, line 55 - column 2, line -/--	5,6,10
☒ 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 10 November 2016		Date of mailing of the international search report 28/11/2016
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 Beker, Sergio

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2016/055929

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
	6 page 3, column 2, line 8 - line 34 page 3, column 2, line 44 - line 54 page 4, column 1, line 38 - line 50 page 4, column 2, line 1 - line 37 page 4, column 2, line 40 - line 56 -----	
X	US 9 112 794 B2 (AGARWAL KANAK B [US] ET AL) 18 August 2015 (2015-08-18)	1-5,7,8, 11,12, 14,15
A	abstract figures 2,4,5 column 2, line 48 - line 50 column 2, line 67 - column 3, line 6 column 3, line 65 - column 4, line 12 column 5, line 65 - line 67 column 6, line 25 - line 50 column 6, line 61 - column 7, line 4 column 7, line 10 - line 38 column 7, line 56 - line 59 column 8, line 4 - line 27 column 10, line 30 - line 44 -----	6,9,10, 13
A	US 2014/362765 A1 (BISWAS ANUMITA [US] ET AL) 11 December 2014 (2014-12-11) the whole document	1-15
A	MATTHIEU COUDRON; STEFANO SECCI; GUY PUJOLLE; PATRICK RAAD; PASCAL GALLARD: "Cross-layer Cooperation to Boost Multipath TCP Performance in Cloud Networks", 2013 IEEE 2ND INTERNATIONAL CONFERENCE ON CLOUD NETWORKING (CLOUDNET), 11 November 2013 (2013-11-11), pages 58-66, XP032550296, DOI: 10.1109/CloudNet.2013.6710558 cited in the application the whole document -----	1-15
A	MATTHIEU COUDRON; STEFANO SECCI; GUIDO MAIER; GUY PUJOLLE; ACHILLE PATTAVINA: "Boosting Cloud Communications through a Crosslayer Multipath Protocol Architecture", 2013 IEEE SDN FOR FUTURE NETWORKS AND SERVICES (SDN4FNS), 11 November 2013 (2013-11-11), pages 1-8, XP032540939, DOI: 10.1109/SDN4FNS.2013.6702535 cited in the application the whole document -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2016/055929

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
US 9112794	B2	18-08-2015	CN	104618240 A		13-05-2015
			US	2015124812 A1		07-05-2015

US 2014362765	A1	11-12-2014	CN	104243443 A		24-12-2014
			DE 10	2014208162 A1		11-12-2014
			US	2014362765 A1		11-12-2014
