



(51) International Patent Classification:

H04L 12/715 (2013.01) H04L 12/803 (2013.01)
H04L 12/729 (2013.01)

(21) International Application Number:

PCT/US2017/027190

(22) International Filing Date:

12 April 2017 (12.04.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/321,730 12 April 2016 (12.04.2016) US
15/485,089 11 April 2017 (11.04.2017) US

(71) Applicant: NICIRA, INC. [US/US]; 3401 Hillview Avenue, Palo Alto, California 94304 (US).

(72) Inventors: HIRA, Mukesh; 3401 Hillview Avenue, Palo Alto, California 94304 (US). KATTA, Naga; 737 Post Street, Apt 721, San Francisco, California 94109 (US). KESLASSY, Isaac; 3401 Hillview Avenue, Palo Alto, California 94304 (US). GHAG, Aditi; 3401 Hillview Avenue, Palo Alto, California 94304 (US).


(74) Agents: HEYMAN, Leonard E. et al.; 3401 Hillview Avenue, Palo Alto, California 94304 (US).

(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, DJ, 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, KH, KN, KP, KR, KW, 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,

[Continued on next page]

(54) Title: VIRTUAL TUNNEL ENDPOINTS FOR CONGESTION-AWARE LOAD BALANCING



Legend: [] : Second packet with inner header information (I) ; [] : Second packet encapsulated with outer header information (O) that includes SP1

Congestion state information at VTEP-A 190

Outer source port number (source_PN) 192	Path (path_ID) 194	Congestion flag (congestion_flag) 196	Path weight (weight) 198
SP1	P1 (via S1)	false	w1
SP2	P2 (via S2)	true	w2

(57) Abstract: Example methods are provided for a source virtual tunnel endpoint (VTEP) to perform congestion-aware load balancing in a data center network. The method may comprise the source VTEP learning congestion state information associated with multiple paths provided by respective multiple intermediate switches connecting the source VTEP with a destination VTEP. The method may also comprise the source VTEP receiving second packets that are sent by a source endpoint and destined for a destination endpoint; and selecting a particular path from multiple paths based on the congestion state information. The method may further comprise the source VTEP generating encapsulated second packets by encapsulating each of the second packets with header information that includes a set of tuples associated with the particular path; and sending the encapsulated second packets to the destination endpoint.

WO 2017/180731 A1

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))*

VIRTUAL TUNNEL ENDPOINTS FOR CONGESTION-AWARE LOAD
BALANCING

MUKESH HIRA, NAGA KATTA, ISAAC KESLASSY, ADITI GHAG

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims the benefit of U.S. Provisional Application No. 62/321,730, filed April 12, 2016, and U.S. Non-Provisional Application No. 15/485,089, filed April 11, 2017, which are incorporated by reference herein in their entirety.

BACKGROUND

[0002] Unless otherwise indicated herein, the approaches described in this section are not admitted to be prior art by inclusion in this section.

[0003] Data center networks generally employ multi-rooted topologies that are characterized by a large degree of multipathing. For example, physical servers are connected with each other using a number of switches that provide alternative paths for packet forwarding. When a physical server has data to send to another physical server, one of the paths may be selected to transmit the data as a flow of packets. In practice, traffic may not be evenly distributed across the different paths, which may cause over-utilization of one path and under-utilization of another. Load balancing is important to spread the traffic as evenly as possible to reduce congestion and improve network performance.

BRIEF DESCRIPTION OF DRAWINGS

[0004] Fig. 1 is a schematic diagram illustrating example data center network in which congestion-aware load balancing is performed;

[0005] Fig. 2 is a flowchart of an example process for a source virtual tunnel endpoint (VTEP) to perform congestion-aware load balancing in a data center network;

[0006] Fig. 3 is a flowchart of a first example process for a source VTEP to learn congestion state information in a data center network;

[0007] Fig. 4A is a flowchart of a second example process for a source VTEP to learn congestion state information in a data center network;

[0008] Fig. 4B is a schematic diagram illustrating example congestion state information learned according to the example in Fig. 4A; and

[0009] Fig. 5 is a flowchart of an example process for a source VTEP to perform data packet processing in a data center network.

DETAILED DESCRIPTION

[0010] In the following detailed description, reference is made to the accompanying drawings, which form a part hereof. In the drawings, similar symbols typically identify similar components, unless context dictates otherwise. The illustrative embodiments described in the detailed description, drawings, and claims are not meant to be limiting. Other embodiments may be utilized, and other changes may be made, without departing from the spirit or scope of the subject matter presented here. It will be readily understood that the aspects of the present disclosure, as generally described herein, and illustrated in the drawings, can be arranged, substituted, combined, and designed in a wide variety of different configurations, all of which are explicitly contemplated herein.

[0011] The challenges of load balancing in data center networks will be described in more detail with reference to Fig. 1, which is a schematic diagram illustrating example data center network 100 in which congestion-aware load balancing is performed. It should be understood that data center network 100 may include additional and/or alternative components than that shown, depending on the desired implementation.

[0012] In the example in Fig. 1, data center network 100 includes first endpoint 102 (see "EP-A") and second endpoint 104 (see "EP-B") that are connected via multiple paths provided by virtual tunnel endpoints (VTEPs), such as "VTEP-A" 110 and "VTEP-B" 120. "VTEP-A" 110 is connected to "VTEP-B" 120 via multiple paths

provided by multiple intermediate switches, such as “A1” 130, “S1” 140, “S2” 150 and “A2” 160. When forwarding data packets from source “EP-A” 102 to destination “EP-B” 104, the data packets may travel over one of the following: a first path via “A1” 130, “S1” 140 and “A2” 160, and a second path via “A1” 130, “S2” 150 and “A2” 160.

[0013] In practice, the term “virtual tunnel endpoints” (e.g., “VTEP-A” 110 and “VTEP-B” 120) may refer to any suitable network elements configured to provide packet forwarding services, load balancing services, gateway services, etc., to endpoints (e.g., “EP-A” 102 and “EP-B” 104). VTEP 110/120 may be implemented by one or more physical or virtual entities. For example, VTEP 110/120 may be implemented by a hypervisor (e.g., a virtual switch of the hypervisor) supported by physical computing device (e.g., edge device, physical server, etc.). VTEP 110/120 and its associated endpoint 102/104 may reside on the same physical computing device, or on different computing devices. For example, “EP-A” 102 may be a virtual machine and “VTEP-A” 110 a virtual switch supported by the same physical server. In another example, “EP-A” 102 may be a virtual machine supported by a first physical server, and “VTEP-A” 110 a virtual switch supported by a second physical server or a physical top-of-rack (ToR) switch connected to the first physical server.

[0014] The term “endpoint” (e.g., “EP-A” 102 and “EP-B 104) may refer generally an originating node (“source endpoint”) or terminating node (“destination endpoint”) of a bi-directional inter-process communication flow. In practice, an endpoint may be a physical computing device (e.g., physical server, physical host), virtualized computing instance supported by a physical computing device, etc. A virtualized computing instance may represent a workload, virtual machine, addressable data compute node, isolated user space instance, etc. In practice, any suitable technology may be used to provide isolated user space instances, including but not limited to hardware virtualization. Other virtualized computing instances may include containers (e.g., running on top of a host operating system without the need for a hypervisor or separate operating system such as Docker, etc.; or implemented as an operating system level virtualization), virtual private servers, etc. The virtual machines may also be complete computational environments, containing virtual equivalents of the hardware and software components of a physical computing

system. The term “hypervisor” may refer generally to a software layer or component that supports the execution of multiple virtualized computing instances, including system-level software that supports namespace containers such as Docker, etc.

[0015] The term “switch” (e.g., “A1” 130, “S1” 140, “S2” 150 and “A2” 160) may refer generally to any suitable network element configured to receive and forward packets, which may layer-3 router, layer-2 switch, gateway, bridge, etc. Depending on the network topology, a switch may be a ToR switch, aggregate switch, spine switch, etc. Although two alternative paths are shown in Fig. 1 for simplicity, the number of paths depends on the number of inter-connected switches and the topology of data center network 100, such as a multi-rooted topology (e.g., leaf-spine topology, fat-tree topology, etc.). Further, there may be additional switches connecting “VTEP-A” 110 to “VTEP-B” 120 than that shown in Fig. 1.

[0016] The term “layer-2” generally refers to a data link layer (e.g., Media Access Control (MAC) or Ethernet layer), “layer-3” to a network layer (e.g., Internet Protocol (IP) layer), and “layer-4” to a transport layer (e.g., Transmission Control Protocol (TCP) layer) in the Open System Interconnection (OSI) model, although the concepts described herein may be applicable to other networking models. The term “packet” may refer generally to a group of bits that can be transported together, and may be in another form, such as “frame”, “message”, “segment”, etc.

[0017] To provide connectivity between “VTEP-A” 110 and “VTEP-B” 120, a “tunnel” (not shown for simplicity) may be established between the VTEPs using any suitable protocol (e.g., Generic Network Virtualization Encapsulation (GENEVE), Stateless Transport Tunneling (STT) or Virtual eXtension Local Area Network (VXLAN)). The term “tunnel” may generally refer to an end-to-end, bi-directional communication path between a pair of VTEPs. In this case, before forwarding data packets (see 170 in Fig. 1) from “EP-A” 102, “VTEP-A” 110 performs encapsulation to generate encapsulated packets (see 172 in Fig. 1).

[0018] In more detail, each data packet 170 includes “inner header information” (labelled “I” in Fig. 1) and application data as payload. The inner header information may include a layer-2 header, a layer-3 header, a layer-4 header, etc. After encapsulation, each encapsulated packet 172 includes outer header information

(labelled “O” in Fig. 1), and data packet 170 as payload. The “outer header information” (also known as “outer tunnel header”) may include an outer layer-2 header, an outer layer-3 header, an outer layer-4 header, etc. The encapsulation is performed such that the fabric overlay (e.g., formed by 130-160) only needs to perform packet forwarding between a pair of VTEPs based on the outer tunnel header.

[0019] In practice, traffic load may be unevenly spread among different paths in data center network 100, which may cause congestion and performance degradation. Conventionally, equal cost multipath routing (ECMP) is commonly used as a data plane load balancing mechanism to spread traffic uniformly across multiple paths with equal costs (e.g., equal number of hops). ECMP switches use a simple, hash-based load balancing scheme to assign each new traffic flow to one of the available paths at random. ECMP is usually implemented in custom silicon (e.g., application-specific integrated circuit (ASIC)), which lacks flexibility to update the load balancing scheme. Further, ECMP is congestion-agnostic and does not protect against oversubscription of paths that causes performance degradation.

[0020] For example in Fig. 1, links connecting different pairs of switches have different congestion levels, as indicated using queue occupancy levels (see 180-186 in square brackets) for packets travelling from “EP-A” 102 to “EP-B” 104. Along the first path via “S1” 140, the queue occupancy levels are 40% (see 180) and 50% (see 182). Along the second path via “S2” 150, the queue occupancy levels are 30% (see 184) and 80% (see 186). Since ECMP does not consider the different congestion levels, it is possible that long-running flows are assigned to the second path via “S2” 150, which is suffering congestion with a queue occupancy level of 80% (see 186).

[0021] Conventionally, control-plane load balancing mechanisms have also been used to address the shortcomings of ECMP. In this case, instead of selecting paths at random, a central controller is deployed in data center network 100 to collect statistics from, and push forwarding rules to, “A1” 130, “S1” 140, “S2” 150 and “A2” 160 to implement control-plane load balancing. However, since a central controller is required, control-plane mechanisms are relatively slow due to high control loop latency and incapable of handling highly volatile traffic.

[0022] Conventionally, host-based approaches have also been used to address the shortcomings of ECMP. For example, a modified version of transmission control protocol (TCP) called multipath TCP (MPTCP) may be used to establish multiple subflows between endpoints to split traffic over different paths. However, host-based approaches usually require changes to all the endpoints, such as modifying the TCP/IP stack of the “EP-A” 102 and “EP-B” 104 in the case of MPTCP. Such changes are usually challenging (and impossible in some cases), especially when “EP-A” 102 and “EP-B” 104 are running different operating systems, or controlled by different entities.

[0023] Congestion-aware load balancing

[0024] According to examples of the present disclosure, a congestion-aware load balancing approach may be implemented by “VTEP-A” 110 in a manner that is completely oblivious to associated “EP-A” 102. Unlike the conventional approaches discussed above, examples of the present disclosure may be implemented without necessitating modification to “EP-A” 102 to implement MPTCP, or modification to intermediate switches 130-160 to implement a new load balancing scheme. Further, unlike control plane load balancing mechanisms, it is not necessary to deploy a central controller to perform congestion monitoring and push forwarding rules intermediate switches 130-160.

[0025] In more detail, Fig. 2 is a flowchart of example process 200 for source VTEP 110 to perform congestion-aware load balancing in data center network 100. Example process 200 may include one or more operations, functions, or actions illustrated by one or more blocks, such as 205 to 240. The various blocks may be combined into fewer blocks, divided into additional blocks, and/or eliminated depending on the desired implementation.

[0026] In the following, “VTEP-A” 110 will be used as an example source VTEP; “VTEP-B” 120 as an example destination VTEP; “S1” 140 and “S2” 150 as example intermediate switches; “EP-A” 102 as “source endpoint”; and “EP-B” 104 as “destination endpoint.” Although queue occupancy level is used as one example to indicate congestion in data center network 100 in Fig. 1, it should be understood that

any other suitable indicator of congestion may be used, such as link utilization level, round trip time (RTT), etc.

[0027] At 205 in Fig. 2, “VTEP-A” 110 learns congestion state information (see 190 in Fig. 1) associated with multiple paths provided by intermediate switches 130-160 that connect “VTEP-A” 110 with “VTEP-B” 120. The congestion state information may be learned based on first packets representing destination-to-source feedback information (see 188 in Fig. 1) from “VTEP-B” 120. At 210 in Fig. 2, “VTEP-A” 110 receives second (data) packets 170 that are sent by “EP-A” 102 and destined for “EP-B” 104. For example, second packets 170 may include application data from an application running on “EP-A” 102 to another application running on “EP-B” 104. Each second packet 170 generally includes inner header information (labelled “I” in Fig. 1) associated with the inter-process communication between “EP-A” 102 and “EP-B” 104.

[0028] At 220 in Fig. 2, “VTEP-A” 110 selects a particular path (also referred to as “selected path”) from the multiple paths. For example in Fig. 1, based on congestion state information associated with the multiple paths, “VTEP-A” 110 may select the first path via “A1” 130, “S1” 140 and “A2” 160. At 230 in Fig. 2, “VTEP-A” 110 generates encapsulated second packets 172 by encapsulating each of second packets 170 with (outer) header information that includes a set of tuples associated with the path selected at 220. At 240 in Fig. 2, “VTEP-A” 110 sends encapsulated second packets 172 to destination “EP-B” 104 such that encapsulated second packets 172 are forwarded via the selected path based on the set of tuples.

[0029] As will be explained further using Fig. 3, “VTEP-A” 110 may rely on congestion state information (see 190 in Fig. 1), which associates different outer source port numbers (see *source_PN* 192) with respective paths (see *path_ID* 194) and flags indicating congestion (see *congestion_flag* 196). For example, prior to performing packet forwarding, “VTEP-A” 110 may perform path learning to learn the mapping between *source_PN* 192 and *path_ID* 194. For the first path via “S1” 140, *source_PN* = SP1 and *path_ID* = P1. For the second path via “S2” 150, *source_PN* = SP2 and *path_ID* = P2.

[0030] As will be explained further using Fig. 3, “VTEP-A” 110 may learn the mapping between different pairs of *source_PN* 192 and *congestion_flag* 196. In one example, “VTEP-A” 110 may rely on existing capabilities intermediate switches 130-160, such as Explicit Congestion Notification (ECN) that facilitates end-to-end notification of congestion state information in data center network 100, etc. In this case, instead of dropping packets, intermediate switches 130-160 perform packet marking as a form of congestion notification to inform “VTEP-B” 120 of a present or pending congestion associated with a particular path. “VTEP-B” 120 subsequently reports any congestion notification to “VTEP-A” 110 (see 188 in Fig. 1). Besides ECN, any other suitable approach may be used to learn the congestion state information. For example, as will be described further using Fig. 4A and Fig. 4B, “VTEP-A” 110 may measure an RTT between “VTEP-A” 110 and “VTEP-B” 120.

[0031] In the example in Fig. 1, each encapsulated second packet 172 includes outer header information that includes a set of tuples associated with the selected path. In particular, each encapsulated second packet 172 includes outer header information (labelled “O”) such as an outer layer-2 header, an outer layer-3 header and an outer layer-4 header. The set of tuples may include a source port number, destination port number, source IP address, destination IP address and protocol. For the first path via “S1” 140 (i.e., *path_ID* = P1), the outer layer-4 header may include a source port number having the value of *source_PN* = SP1. The source port number is set such encapsulated packets 172 are forward via “A1” 130, “S1” 140 and “A2” 160 along the selected path. At destination “VTEP-B” 120, decapsulation is performed to remove outer header information, data packets 174 sent to “EP-B” 104.

[0032] Using example process 200, “VTEP-A” 110 may distribute virtual network traffic over different paths in data center network 100, taking into account congestion state information 190 associated with the different paths. “VTEP-A” 110 may select a different outer source port number every time it wants to send encapsulated second packets 172 on a different path. Since intermediate switches 130-160 connecting “VTEP-A” 110 and “VTEP-B” 120 will perform load balancing based on the outer header information, the outer source port number in the outer layer-4 header may be used as an entropy to exploit multiple paths (e.g., equal-cost paths) in data center network 100.

[0033] Congestion state information

[0034] Fig. 3 is a flowchart of first example process 300 for source VTEP 110 to learn congestion state information 190 in data center network 100 according to a first example. Example process 300 may include one or more operations, functions, or actions illustrated by one or more blocks, such as 310 to 365. The various blocks may be combined into fewer blocks, divided into additional blocks, and/or eliminated depending on the desired implementation. Using the example in Fig. 1, example process 300 may be performed by “VTEP-A” 110.

[0035] At 310 in Fig. 3, “VTEP-A” 110 learns the mapping or association between multiple pairs of *source_PN* 192 and *path_ID* 194. For example in Fig. 1, different values of *source_PN* 192 may lead to different, non-overlapping ECMP paths in data center network 100. The mapping represents *a priori* knowledge of different paths from which “VTEP-A” 110 may choose for packet forwarding. When ECMP hashing is applied to a set of tuples that includes a particular value of *source_PN* 192, the result is the associated path identified by *path_ID* 194.

[0036] In practice, “VTEP-A” 110 may implement a background daemon (e.g., modeled after Paris traceroute) to send periodic probe packets to all other VTEPs in data center network 100 to collect “traceroute” style information about all interface IPs encountered on each path. “VTEP-A” 110 may rotate the outer source port number in the outer header information of each probe packet to collect the path trace for each port number. “VTEP-A” 110 then updates congestion state information 190 to add or update *source_PN* 192 every time the corresponding probe pathtrace differs from the collected traces so far.

[0037] Blocks 315 to 360 in Fig. 3 relate to “VTEP-A” 110 learning congestion state information 190 that associates a path identified by *path_ID* 194 and *congestion_flag* 196 indicating congestion along that path. In particular, at 315 and 320 in Fig. 3, “VTEP-A” 110 encapsulates packets with an outer tunnel header that includes a particular value of *source_PN* 192 and sends encapsulated packets.

[0038] In the example in Fig. 3, “VTEP-A” 110 may rely on a congestion notification capability of existing switches, such as ECN marking to indicate a

present or pending congestion, etc. Detailed information of ECN may be found in the Internet Engineering Task Force (IETF) Request for Comments number 3168 and entitled "The Addition of Explicit Congestion Notification (ECN) to IP," which is incorporated herein in its entirety by reference. Although ECN is described as an example, it should be understood that any other suitable packet marking approach may be used.

[0039] At 325, 330, 335 and 340 in Fig. 3, in response to receiving encapsulated packets from "VTEP-A" 110, a switch (e.g., "S1" 140, "S2" 150) may perform ECN marking to flag congestion at that switch before forwarding the encapsulated packets. In this case, the switch is known as an ECN-enabled switch that is capable of modifying header information (e.g., in a reserved field of a TCP header) of the encapsulated packets as a form of congestion notification.

[0040] For example in Fig. 1, when the queue occupancy level (e.g., 80%) at "S2" 150 exceeds a predetermined threshold (e.g., $T_Q = 60\%$), "S2" 150 marks packets received from "VTEP-A" 110 to notify "VTEP-B" 120 of the congestion, thereby indicating that the second path is a congested path. On the other hand, it is not necessary for "S1" 140 to perform any packet marking because associated queue occupancy level (e.g., 50%) is below the threshold, thereby indicating that the first path is an uncongested path. .

[0041] At 345 and 350 in Fig. 3, destination "VTEP-B" 120 receives an encapsulated packet and reports the mapping between outer source port number in the outer header information and any congestion notification (e.g., ECN marking) to source "VTEP-A" 110 (see also 188 in Fig. 1). At 355 and 360 in Fig. 3, "VTEP-A" 110 updates the mapping between *source_PN* 192 and associated *congestion_flag* 196 (i.e., flag information). For example in Fig. 1, "VTEP-A" 110 determines that *congestion_flag* = false for *source_PN* = SP1 for the first path via "S1" 140, and *congestion_flag* = true for *source_PN* = SP1 for the second path via "S2" 150.

[0042] The *congestion_flag* may be used in a weighting algorithm that influences the likelihood of the selection of its associated path. At the beginning of the load balancing process, "VTEP-A" 110 may start off with equal weighting for all equal-cost paths discovered by the path learning function at 310 in Fig. 3. Subsequently, at 365

in Fig. 3, the path weight (see *weight* 198 in Fig. 1) is adjusted based on *congestion_flag*. For example, in response to determination that *congestion_flag* = false (i.e., clear), the *weight* = w_1 associated with the first path via “S1” 140 is increased to increase its likelihood for selection. On the other hand, in response to determination that *congestion_flag* = true (i.e., set), the *weight* = w_2 associated with the second path via “S2” 150 is decreased to reduce its likelihood for selection.

[0043] Another metric that may be used to indicate congestion is RTT that may be measured and actively kept track of for each path. An example is described using Fig. 4A, which is a flowchart of second example process 400 for source VTEP 110 to learn congestion state information in data center network 100. Example process 400 may include one or more operations, functions, or actions illustrated by one or more blocks, such as 410 to 455. The various blocks may be combined into fewer blocks, divided into additional blocks, and/or eliminated depending on the desired implementation.

[0044] At 410, 415 and 420 in Fig. 4A, “VTEP-A” 110 periodically sends probe packets on each path via associated switch 140/150, where each probe packet identifies a *source_PN* and a transmit (*Tx*) timestamp. At 425, 430 and 435 in Fig. 4A, “VTEP-B” 120 replies to each probe packet by sending an acknowledgement (ACK) packet that also identifies the *source_PN* and *Tx* timestamp via switch 140/150. At 440 and 445 in Fig. 4A, in response to receiving an ACK packet at a particular receive time, “VTEP-A” 110 may determine the RTT based on difference between the receive time and the *Tx* timestamp. At 450 and 455 in Fig. 4A, “VTEP-A” 110 associates the RTT with the *source_PN*, and adjusts a weight associated with the path.

[0045] Fig. 4B is a schematic diagram illustrating example congestion state information 460 learned using the example in Fig. 4A. For example, the first path via “S1” 140 is associated with *source_PN* = SP1 (see 462), *path_ID* = P1 (see 464), *RTT* = R1 (see 466) and *weight* = w_1 (see 468). The second path via “S2” 150 is associated with *source_PN* = SP2, *path_ID* = P2, *RTT* = R2 and *weight* = w_2 . In the example in Fig. 1 where the second path has a higher congestion level than the first path, R2 should be greater than R1. In this case, w_2 is adjusted to be smaller than

w1 to reduce the likelihood of selecting the second path. In practice, the closer that the timestamping and acknowledgement are performed, the more accurately the RTT reflects actual network latency for the path as it does not include latency introduced by the software stack at the transmitter and the receiver.

[0046] Data packet processing

[0047] According to examples of the present disclosure, load balancing may be performed at the granularity of flowlets to avoid or ameliorate packet reordering issues associated with transport layer protocols such as TCP. This may be achieved by splitting a flow of packets into multiple smaller groups called “flowlets.” As used herein, the term “flowlet” may refer generally to a group or burst of packets within a flow.

[0048] Fig. 5 is a flowchart of example process 500 for virtual tunnel endpoint 110 to perform data packet processing in data center network 100. Example process 500 may include one or more operations, functions, or actions illustrated by one or more blocks, such as 510 to 560. The various blocks may be combined into fewer blocks, divided into additional blocks, and/or eliminated depending on the desired implementation.

[0049] At 510 and 515, in response to receive data packet 170 from “EP-A” 102, “VTEP-A” 110 determines whether data packet 170 belongs to a new flowlet or a current flowlet. For example, a new flowlet may be detected whenever a time interval between the arrivals of two consecutive packets within the same flow (i.e., inter-packet gap) exceeds a predetermined threshold (e.g., $T_{flowlet}$ seconds; see 515). All subsequent packets that do not exceed the threshold are considered to be part of the same flowlet. The two scenarios are discussed below.

[0050] (a) When a new flowlet is detected (i.e., inter-packet gap $> T_{flowlet}$ or first packet of the flow is detected), “VTEP-A” 110 assigns a new *flowlet_ID* (e.g., “F1”) to data packet 170 at 520 in Fig. 5. At 525 in Fig. 5, a path identified by *path_ID* (e.g., “P1”) is selected for the new flowlet (see 525) based on congestion state information 190/460 at that time. At 530 in Fig. 5, *source_PN* associated with the selected path is determined, an association between *flowlet_ID* and *source_PN*

(e.g., "SP1") associated with the selected path is stored. "VTEP-A" 110 also stores *flowlet_time* to record the time at which the most recent data packet 170 of the flowlet is received.

[0051] (b) When data packet 170 of an existing flowlet is detected (i.e., inter-packet gap = current time - *flowlet_time* $\leq T_{flowlet}$), "VTEP-A" 110 retrieves the current *flowlet_ID* (e.g., "F1") and associated *source_PN* (e.g., "SP1") at 550 and 555 in Fig. 5. Similarly, at 560 in Fig. 5, "VTEP-A" 110 also stores *flowlet_time* to record the time at which the most recent data packet 170 of the flowlet is received.

[0052] In both cases (a) and (b), example process 500 continues at 535 and 540 in Fig. 5, where "VTEP-A" 110 encapsulates data packet 170 with an outer tunnel header that is configured to include an outer source port number having the value of *source_PN* (e.g., "SP1") associated with *flowlet_ID* (e.g., "F1") and sends encapsulated data packet 172. The outer tunnel header also includes a source IP address associated with "VTEP-A" 110, and a destination IP address associated with "VTEP-B" 120. At the destination, the outer tunnel header is removed by "VTEP-B" 120 before the packet is forwarded to "EP-B" 104.

[0053] It should be understood that the above may be repeated for subsequent data packets 170 that belong to another flowlet, such as *flowlet_ID* = "F2" when the inter-packet gap is exceeded at 515 in Fig. 4. In this case, congested state information 190/460 may have changed, and the second path via "S2" 150 (i.e., *path_ID* = P2) may now be selected instead of the first path via "S1" 140. In this case, associated *source_PN* = SP2 is included in the outer header information such that encapsulated data packets 172 are forwarded by "A1" 130, "S2" 150 and "A2" 160 along the second path to "VTEP-B" 120.

[0054] Using examples of the present disclosure, "VTEP-A" 110 remembers the *flowlet_ID* of a new flowlet, and reuses the associated *source_PN* for packets of the same flowlet. In practice, threshold $T_{flowlet}$ may be set based on the estimated RTT in the network (e.g., in the order of RTT). A large inter-packet gap between flowlets ensures that when multiple flowlets (that are technically part of the same flow) take different paths, the resulting packet reordering is minimal. Further, "VTEP-A" 110 may assign the outer source port number with minimal congestion to each new

flowlet to ensure that they take the selected path. This in turn keeps the queues on all the active paths substantially low at all times, thereby leading to better throughput and latency for workloads.

[0055] Examples of the present disclosure are implemented in a manner that is oblivious to end-user guest virtual machines (e.g., acting as “EP-A” 102 and “EP-B” 104) by configuring “VTEP-A” 110 to split a network flow into multiple uncongested paths in a manner completely oblivious to the guest virtual machines. In one example, both source virtual switch (e.g., acting as “VTEP-A” 110) and destination virtual switch (e.g., acting as “VTEP-B 120) may be used. In this case, a flow may be divided into flowlets arbitrarily, and not contingent upon an idle gap between flowlets.

[0056] For example, each TCP segmentation offload (TSO) segment may be treated as a flowlet. If flowlets arrive out of order at the destination virtual switch because of the different paths taken, the destination virtual switch may reorder the flowlets before delivering them to the destination endpoint (e.g., guest virtual machine). This allows the destination virtual switch to hide the out-of-order arrival from the destination protocol stack (e.g., TCP/IP stack) of the destination endpoint, preventing the destination protocol stack from slowing down the source protocol stack at the source endpoint due to the out-of-order delivery.

[0057] Examples of the present disclosure may be deployed in the virtual switch of a hypervisor leading to intelligent path selection from the first point of entry of traffic. For example, there could be multiple ECMP paths from the virtual switch onwards, each path using a different physical NIC. In this case, examples of the present disclosure may be implemented using a host that includes a processor, and a non-transitory computer-readable storage medium storing a set of instructions. In response to execution by the processor, the set of instructions cause the processor to implement a source VTEP (e.g., “VTEP-A” 110 at a virtual switch supported by the host) to perform congestion-aware load balancing in data center network 100 according to the examples in Fig. 1 to Fig. 5.

[0058] In one example, the source VTEP implemented by the host may learn, based on first packets from the destination VTEP, congestion state information

associated with multiple paths provided by respective multiple intermediate switches connecting the source VTEP with the destination VTEP. Also, in response to receiving second packets that are sent by a source endpoint and destined for a destination endpoint associated with the destination VTEP, the source VTEP may select a particular path from multiple paths based on the congestion state information, generate encapsulated second packets by encapsulating each of the second packets with header information that includes a set of tuples associated with the particular path, and send the encapsulated second packets to the destination endpoint such that the encapsulated second packets are forwarded via the particular path based on the set of tuples.

[0059] In scenarios where all physical NICs connect to the same layer-3 next-hop and path diversity starts beyond the first-hop switch (e.g., ToR switch), examples of the present disclosure may be implemented in the NIC driver/hardware or on the first-hop switch at faster speeds than in the virtual switch software. Compared to conventional approaches that require advanced switch architectures, examples of the present disclosure may be performed in the edge hypervisor (e.g., entirely in software) and scale to any number of hops between sources and destinations. In this case, examples a (first-hop) switch may be used to implement a source VTEP (e.g., "VTEP-A" 110) to perform congestion-aware load balancing in data center network 100 according to the examples in Fig. 1 to Fig. 5. The switch may include any suitable switch logic, such as hardware logic (e.g., hardware circuitry), programmable logic, a combination thereof, etc.

[0060] In one example, the switch may include first port(s), second port(s) and a switch logic. The first port(s) may be used to receive first packets from a destination VTEP connected to the source VTEP via multiple paths provided by respective multiple intermediate switches connecting the source VTEP with the destination VTEP. The second port may used to receive second packets that are sent by the source endpoint and destined for the destination endpoint associated with the destination VTEP. The switch logic may be configured to learn congestion state information associated with the multiple paths based on the first packets. In response to receiving the second packets. switch logic may be configured to select a particular path from multiple paths based on the congestion state information,

generate encapsulated second packets by encapsulating each of the second packets with header information that includes a set of tuples associated with the particular path, and send the encapsulated second packets to the destination endpoint such that the encapsulated second packets are forwarded via the particular path based on the set of tuples.

[0061] The techniques introduced above can be implemented in special-purpose hardwired circuitry, in software and/or firmware in conjunction with programmable circuitry, or in a combination thereof. Special-purpose hardwired circuitry may be in the form of, for example, one or more application-specific integrated circuits (ASICs), programmable logic devices (PLDs), field-programmable gate arrays (FPGAs), programmable switch architectures, and others. The term 'processor' is to be interpreted broadly to include a processing unit, ASIC, logic unit, or programmable gate array, etc.

[0062] The foregoing detailed description has set forth various embodiments of the devices and/or processes via the use of block diagrams, flowcharts, and/or examples. Insofar as such block diagrams, flowcharts, and/or examples contain one or more functions and/or operations, it will be understood by those within the art that each function and/or operation within such block diagrams, flowcharts, or examples can be implemented, individually and/or collectively, by a wide range of hardware, software, firmware, or any combination thereof.

[0063] Those skilled in the art will recognize that some aspects of the embodiments disclosed herein, in whole or in part, can be equivalently implemented in integrated circuits, as one or more computer programs running on one or more computers (e.g., as one or more programs running on one or more computing systems), as one or more programs running on one or more processors (e.g., as one or more programs running on one or more microprocessors), as firmware, or as virtually any combination thereof, and that designing the circuitry and/or writing the code for the software and or firmware would be well within the skill of one of skill in the art in light of this disclosure.

[0064] Software and/or to implement the techniques introduced here may be stored on a non-transitory computer-readable storage medium and may be executed

by one or more general-purpose or special-purpose programmable microprocessors. A “computer-readable storage medium”, as the term is used herein, includes any mechanism that provides (i.e., stores and/or transmits) information in a form accessible by a machine (e.g., a computer, network device, personal digital assistant (PDA), mobile device, manufacturing tool, any device with a set of one or more processors, etc.). A computer-readable storage medium may include recordable/non recordable media (e.g., read-only memory (ROM), random access memory (RAM), magnetic disk or optical storage media, flash memory devices, etc.).

[0065] The drawings are only illustrations of an example, where the elements or procedure shown in the drawings are not necessarily essential for implementing the present disclosure. Those skilled in the art will understand that elements in the examples can be arranged in the device in the examples as described, or can be alternatively located in one or more devices different from that in the examples. The elements in the examples described can be combined into one module or further divided into a plurality of sub-elements.

WE CLAIM:

1. A method for a source virtual tunnel endpoint (VTEP) to perform congestion-aware load balancing in a data center network that includes the source VTEP, a destination VTEP, a source endpoint, a destination endpoint and multiple intermediate switches, the method comprising:

learning, by the source VTEP based on first packets from the destination VTEP, congestion state information associated with multiple paths provided by respective multiple intermediate switches connecting the source VTEP with the destination VTEP;

receiving, by the source VTEP, second packets that are sent by the source endpoint and destined for the destination endpoint associated with the destination VTEP;

selecting, by the source VTEP, a particular path from multiple paths based on the congestion state information;

generating, by the source VTEP, encapsulated second packets by encapsulating each of the second packets with header information that includes a set of tuples associated with the particular path; and

sending, by the source VTEP, the encapsulated second packets to the destination endpoint such that the encapsulated second packets are forwarded via the particular path based on the set of tuples.

2. The method of claim 1, wherein learning the congestion state information comprises:

determining a congestion flag associated with the particular path based on a particular first packet from the destination VTEP, wherein the particular first packet includes a congestion notification from at least one of the multiple intermediate switches to indicate whether the particular path is congested.

3. The method of claim 1, wherein learning the congestion state information comprises:

determining a round trip time associated with the particular path based on a receive time of a particular first packet from the destination VTEP, and a transmit time of a probe packet sent by the source VTEP to trigger the particular first packet.

4. The method of claim 1, wherein learning the congestion state information comprises:

based on a particular first packet from the destination VTEP, determining a weight associated with the particular path to increase or decrease its likelihood for selection.

5. The method of claim 1, wherein learning the congestion state information comprises:

learning mapping between multiple outer source port numbers and the respective multiple paths.

6. The method of claim 5, wherein encapsulating each of the second packets comprises:

determining, from the multiple outer source port numbers, a particular outer source port number associated with the particular path; and

configuring the set of tuples in the header information in each of the second packets to include the particular outer source port number.

7. The method of claim 6, wherein the method further comprises:

assigning each of the second packets, being a current second packet, to a flowlet based on an inter-packet gap between the current second packet and a previous second packet destined for the destination endpoint; and

storing flowlet information that associates the flowlet with the particular outer source port number associated with the particular path.

8. A non-transitory computer-readable storage medium that includes a set of instructions which, in response to execution by a processor of a host, cause the processor to implement a source virtual tunnel endpoint (VTEP) to perform a method of congestion-aware load balancing in a data center network that includes the source VTEP, a destination VTEP, a source endpoint, a destination endpoint and multiple intermediate switches, wherein the host supports the source VTEP, and the method comprises:

learning, by the source VTEP based on first packets from the destination VTEP, congestion state information associated with multiple paths provided by respective multiple intermediate switches connecting the source VTEP with the destination VTEP;

receiving, by the source VTEP, second packets that are sent by the source endpoint and destined for the destination endpoint associated with the destination VTEP;

selecting, by the source VTEP, a particular path from multiple paths based on the congestion state information;

generating, by the source VTEP, encapsulated second packets by encapsulating each of the second packets with header information that includes a set of tuples associated with the particular path; and

sending, by the source VTEP, the encapsulated second packets to the destination endpoint such that the encapsulated second packets are forwarded via the particular path based on the set of tuples.

9. The non-transitory computer-readable storage medium of claim 8, wherein learning the congestion state information comprises:

determining a congestion flag associated with the particular path based on a particular first packet from the destination VTEP, wherein the particular first packet includes a congestion notification from at least one of the multiple intermediate switches to indicate whether the particular path is congested.

10. The non-transitory computer-readable storage medium of claim 8, wherein learning the congestion state information comprises:

determining a round trip time associated with the particular path based on a receive time of a particular first packet from the destination VTEP, and a transmit time of a probe packet sent by the source VTEP to trigger the particular first packet.

11. The non-transitory computer-readable storage medium of claim 8, wherein learning the congestion state information comprises:

based on a particular first packet from the destination VTEP, determining a weight associated with the particular path to increase or decrease its likelihood for selection.

12. The non-transitory computer-readable storage medium of claim 8, wherein learning the congestion state information comprises:

learning mapping between multiple outer source port numbers and the respective multiple paths.

13. The non-transitory computer-readable storage medium of claim 12, wherein encapsulating each of the second packets comprises:

determining, from the multiple outer source port numbers, a particular outer source port number associated with the particular path; and

configuring the set of tuples in the header information in each of the second packets to include the particular outer source port number.

14. The non-transitory computer-readable storage medium of claim 13, wherein the method further comprises:

assigning each of the second packets, being a current second packet, to a flowlet based on an inter-packet gap between the current second packet and a previous second packet destined for the destination endpoint; and

storing flowlet information that associates the flowlet with the particular outer source port number associated with the particular path.

15. A switch configured to implement a source virtual tunnel endpoint (VTEP) to perform congestion-aware load balancing in a data center network, wherein the switch comprises:

one or more first ports to receive first packets from a destination VTEP connected to the source VTEP via multiple paths provided by respective multiple intermediate switches connecting the source VTEP with the destination VTEP;

a second port to receive second packets that are sent by the source endpoint and destined for the destination endpoint associated with the destination VTEP; and

a switch logic configured to:

learn congestion state information associated with the multiple paths based on the first packets, and in response to receiving the second packets, select a particular path from multiple paths based on the congestion state information;

generate encapsulated second packets by encapsulating each of the second packets with header information that includes a set of tuples associated with the particular path; and

send the encapsulated second packets to the destination endpoint such that the encapsulated second packets are forwarded via the particular path based on the set of tuples.

16. The switch of claim 11, wherein the switch logic is configured to learn the congestion state information by determining a congestion flag associated with the particular path based on a particular first packet from the destination VTEP, wherein the particular first packet includes a congestion notification from at least one of the multiple intermediate switches to indicate whether the particular path is congested.

17. The switch of claim 11, wherein the switch logic is configured to learn the congestion state information by determining a round trip time associated with the particular path based on a receive time of a particular first packet from the destination VTEP, and a transmit time of a probe packet sent by the source VTEP to trigger the particular first packet.


18. The switch of claim 11, wherein the switch logic is configured to learn the congestion state information by determining a weight associated with the particular path to increase or decrease its likelihood for selection based on a particular first packet from the destination VTEP.

19. The switch of claim 11, wherein the switch logic is configured to learn the congestion state information by learning mapping between multiple outer source port numbers and the respective multiple paths.

20. The switch of claim 19, wherein the switch logic is configured to learn the congestion state information by determining, from the multiple outer source port numbers, a particular outer source port number associated with the particular path; and configuring the set of tuples in the header information in each of the second packets to include the particular outer source port number.

21. The switch of claim 20, wherein the switch logic is further configured to:
assign each of the second packets, being a current second packet, to a flowlet based on an inter-packet gap between the current second packet and a previous second packet destined for the destination endpoint; and
store flowlet information that associates the flowlet with the particular outer source port number associated with the particular path.

1/5



: Second packet with inner header information (I)

: Second packet encapsulated with outer header information (O) that includes SP1

Congestion state information at VTEP-A ¹⁹⁰

Outer source port number (source_PN) <u>192</u>	Path (path_ID) <u>194</u>	Congestion flag (congestion_flag) <u>196</u>	Path weight (weight) <u>198</u>
SP1	P1 (via S1)	false	w1
SP2	P2 (via S2)	true	w2

Fig. 1

2/5

200




Fig. 2




Fig. 3




Fig. 4A

Congestion state information at VTEP-A 460

Outer source port number (source_PN) 462	Path (path_ID) 464	Round trip time (RTT) 466	Path weight (weight) 468
SP1	P1 (via S1)	R1	w1
SP2	P2 (via S2)	R2	w2

Fig. 4B




Fig. 5

INTERNATIONAL SEARCH REPORT

International application No PCT/US2017/027190
--

A. CLASSIFICATION OF SUBJECT MATTER
 INV. H04L12/715 H04L12/729 H04L12/803
 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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	MOHAMMAD ALIZADEH ET AL: "CONGA", SIGCOMM, ACM, 2 PENN PLAZA, SUITE 701 NEW YORK NY 10121-0701 USA, 17 August 2014 (2014-08-17), pages 503-514, XP058053862, DOI: 10.1145/2619239.2626316 ISBN: 978-1-4503-2836-4 the whole document -----	1-21

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

29 May 2017

Date of mailing of the international search report

09/06/2017

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

Raible, Markus