



- (51) International Patent Classification: **H04L 12/721** (2013.01) **H04L 12/717** (2013.01)
- (21) International Application Number: PCT/EP2017/062851
- (22) International Filing Date: 29 May 2017 (29.05.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicants: **HUAWEI TECHNOLOGIES CO., LTD.** [CN/CN]; Huawei Administration Building Bantian Longgang District, Shenzhen, Guangdong 518129 (CN). **POLITECNICO DI MILANO** [IT/IT]; Department of Electronics, Information and Bioengineering, Piazza Leonardo da Vinci 32, 20133 Milano (IT).
- (72) Inventor; and
(71) Applicant (for US only): **PARIS, Stefano** [IT/DE]; c/o Huawei Technologies Duesseldorf GmbH Riesstr.25, 80992 Munich (DE).
- (72) Inventors: **LEGUAY, Jeremie**; c/o Huawei Technologies Duesseldorf GmbH Riesstr.25, 80992 Munich (DE). **CAPONE, Antonio**; Department of Electronics, Information and Bioengineering, Piazza Leonardo da Vinci 32, 20133 Milan (IT). **FILIPPINI, Ilario**; Department of Electronics, Information and Bioengineering, Piazza Leonardo da Vinci 32, 20133 Milan (IT). **SANVITO, Davide**; Department of Electronics, Information and Bioengineering, Piazza Leonardo da Vinci 32, 20133 Milan (IT).
- (74) Agent: **KREUZ, Georg**; Huawei Technologies Duesseldorf GmbH, Riesstr. 8, 80992 Munich (DE).
- (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,

(54) Title: APPARATUS AND METHOD FOR PROVIDING A ROBUST ROUTING CONFIGURATION




Figure 2

(57) Abstract: A system and method for providing a routing configuration is disclosed. The routing configurations are built according to monitored traffic matrices and predicted traffic matrices. The build set of routing configurations are a set of overlapping robust routing configurations. The building provides one or more routing configurations that perform better with upcoming traffic matrices. The routing configurations are computed in advance so that they can be pre-fetched before the real need for the new routing configuration commences. The pre-fetched routing configurations are easy to activate as they have been computed and fetched in advance. Because of overlapping the activating can be executed fluently.



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, 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))*

APPARATUS AND METHOD FOR PROVIDING A ROBUST ROUTING CONFIGURATION

FIELD

- 5 [0001] The present disclosure relates to the field network routing, and more particularly to providing network configurations in varying traffic conditions.

BACKGROUND

- Traffic engineering (TE) plays a crucial role in network management since it
10 permits to optimize and load balancing the utilization of network resources.
Sophisticated algorithms have been designed to implement efficiently TE.
Nowadays, Software-defined networking (SDN) controllers integrate TE methods
to continuously optimize the network. However, sudden and unpredictable changes
of the system, like traffic variations and network failures, still pose a major
15 challenge for TE.

The solutions that have been devised to cope with system changes can be broadly
classified in two main classes: dynamic TE and stable TE. While different in the
way they calculates network configurations, both techniques require the use of the
information periodically collected by a network monitoring tool.

- 20 [0002] Dynamic TE uses the information collected by the monitoring tool to
predict the next system state and compute the optimal routing configuration for the
predicted state. The measurements can be collected over a short period of time to
predict the system evolution in the near future. However, the accuracy of the
prediction highly affects the performance (e.g., traffic congestion, delay, or cost) of
25 the routing configuration. In the presence of unexpected system variations, the
routing configuration will be far from the optimal operational point of the system
due to the low accuracy of the prediction. Furthermore, Dynamic TE requires
network reconfiguration each time the system changes, thus increasing the
signaling overhead caused by the control messages sent to all network devices.
30 Furthermore, in highly dynamic systems, network reconfiguration introduces
instability that eventually affects the performance.

[0003] Dynamic TE reconfigures the network each time a new event like a traffic change or link failure occurs in the network. A monitoring system is used by Dynamic TE methods to collect information over a relatively short period of time and predict the future evolution of the system. For example, a certain number of past traffic matrices is used to predict the next traffic matrix. This prediction is used as an input of TE techniques to reconfigure the network. Sophisticated methods like linear programming can be used to compute the best of the network configuration any time the traffic changes. However, reconfiguring the network too frequently can affect its stability, since programming hardware equipment takes time. Therefore, frequent network reconfiguration can result in worse performance. For this reason, methods to reduce the network reconfiguration have been proposed. For example, a method that reconfigures the network only if the benefit captured by the objective function is larger than using the current configuration. Nonetheless, the efficacy of Dynamic TE methods depend on the accuracy of the prediction method used to estimate the evolution of the system. Indeed, prediction errors can lead to large deviations of the network performance optimized by Dynamic TE methods.

[0004] In contrast, Stable TE (e.g., oblivious routing) monitors the system evolution over a long period of time and computes the TE configuration that minimizes the worst deviation to the optimal configurations in terms of objective function (e.g., traffic congestion, delay, or cost). Stable TE keeps the network configuration stable, but it suffers low optimality/efficiency in terms of performance.

[0005] Stable TE computes a single configuration that provides bounded performance loss in case of prediction errors or unpredictable events with respect to the optimal configuration. One of the first techniques of Stable TE is oblivious routing, which randomly selects the paths to connect the source and the destination using a small subset of preselected nodes without using any knowledge of the traffic. Routing techniques that are totally oblivious to any traffic information and system change provide the highest robustness against any system evolution but the performance loss can quickly become large as the network size increases. For this

reason, a partial knowledge of the traffic can be integrated in the computation of the network configuration to improve oblivious routing. Several methods for optimizing the stable TE are known. For example, methods providing techniques to compute a network configuration with performance guarantees for any TM that belongs to a bounded domain/set obtained using the information collected by the monitoring system. In other methods different models for traffic uncertainty by expressing the maximum load can be expected over a link (pipe model) or an upper bound on the traffic emanating from a source node and directed to a destination node (hose model). Based on those models a network configuration with performance guarantees is computed. However, these uncertainty models include less information than the dynamic TM approach, thus providing larger performance losses. To further improve the optimality of the solution, the technique considering only the most likely TMs for building the bounded set and computing of the optimal configuration has been disclosed. A penalty term is added in the optimization for less probable TMs. While this technique provides good solutions when the traffic varies smoothly, it cannot guarantee a feasible configuration against unexpected or unlikely events, since they are excluded from the computation.

20 SUMMARY

[0006] A system and method for providing a routing configuration is disclosed. The routing configurations are built according to monitored traffic matrices and predicted traffic matrices. The build set of routing configurations are a set of overlapping robust routing configurations. The building provides one or more routing configurations that perform better with upcoming traffic matrices. The routing configurations are computed in advance so that they can be pre-fetched before the real need for the new routing configuration commences. The pre-fetched routing configurations are easy to activate as they have been computed and fetched in advance. Because of overlapping the activating can be executed fluently.

[0007] In an aspect an apparatus for providing a routing configuration to a switch is disclosed. The apparatus comprises a network connection, a memory configured to store a plurality of precomputed routing configurations, wherein said routing configurations are associated with overlapping clusters of traffic matrices, and a processor. The processor is configured to monitor current traffic conditions, predict, based on the monitored conditions, one or more upcoming traffic matrices, select, based on the predicted one or more upcoming traffic matrices, at least one routing configuration from the plurality of precomputed routing configurations, and transmit at least one of the selected at least one routing configuration to the at least one switch using the network connection.

[0008] A benefit of the arrangement described above is that it provides more flexibility in choosing routing configuration than the static TE, however, it requires less resources than the dynamic TE. The arrangement described above enables the possibility to compute new routing configurations as a background process so that the routing configurations are ready when they are needed. As the computation is done in background and in advance the required computing power for selecting and adapting the routing configuration is not as high as in dynamic TE implementations, wherein the computation is done on demand. A further benefit of the arrangement is that when the routing configurations are arranged to overlapping clusters of traffic matrices the required activation process can be launched flexibly.

[0009] In an implementation the processor is further configured to transmit an activation signal to at least one switch in order to activate the transmitted routing configuration. This is beneficial as the routing configuration can be transmitted to respective switches in advance and activated with one signal when the need for activation emerges.

[0010] In an implementation the processor is further configured to incorporate in the selected routing configuration activation instructions. This is beneficial when the switches comprise a logic or other means for detecting the need for activation by themselves when the conditions set in activation instructions are met. Thus, there is no need for separate activation signal in such case.

[0011] In an implementation the processor is further configured to analyze the monitored traffic conditions, and based on the analysis, the processor is further configured to predict the traffic evolution and compute the plurality of precomputed routing configurations based on analyzed traffic conditions and the predicted traffic evolution. This is beneficial as from the monitored conditions it is possible to predict the upcoming traffic evolution and compute routing configurations performing well with the predicted upcoming traffic.

[0012] In an implementation the processor is further configured to compute the plurality of precomputed routing configurations based on at least one of the following: a temporal overlap, traffic similarity, and routing similarity. This is beneficial as it helps providing routing configurations such that the transfer from a current routing configuration to the following routing configuration can be done fluently. In addition, this is further beneficial since it provides the time to decide whether to reconfigure or not.

[0013] In an implementation, when computing the plurality of precomputed routing configurations, the processor is configured to receive a set of input parameters, wherein the input parameters comprise: a set of traffic matrices received from monitoring (M); a set of traffic matrices received from prediction; a maximum number of clusters (N); a maximum number of routing configurations (K); a minimum time overlap among clusters (O); and a minimum time of utilization of a configuration (L). It is beneficial to use this set of parameters as it provides a possibility to compute in advance routing configurations in a manner that they will perform well with the predicted traffic conditions. This is also beneficial, since it permits to decide the best trade-off among network performance, network stability, robustness, and computing power.

[0014] In an implementation the processor is further configured to sort the set M according to increasing traffic matrix sampling time; and find an initial set of B routing configurations: R_1, \dots, R_B . It is beneficial to sort the traffic matrices received from monitoring and find an initial set of routing configurations in order to find the final routing configurations and speed-up the computation.

[0015] In an implementation the processor is further configured to create N clusters of traffic matrices, wherein $N < B$, $C_i=1..N$, wherein the processor is further configured to assign for each of M traffic matrices a routing configuration R_i such that an overall maximum link utilization is minimized. Choosing the parameters as above is beneficial as minimizing overall maximum link utilization is often desired in order to minimize the overall network congestion.

[0016] In an implementation the processor is configured to assign for each of M traffic matrices a routing configuration such that upcoming conditions are met:

- 10 - $\forall C_i, Tim[(C_i)] - Time[first(C_i)] \geq L$; AND
 - $\forall C_i, C_j \in adj(K_i),$
 $Tim[last(C_i)] - Time[first(C_j)] \geq O$; OR
 $Time[last(C_j)] - Time[first(C_i)] \geq O$.

15 [0017] It is beneficial to fulfill the conditions mentioned above as they permits to maintain a routing configuration for a minimum desired time (L) and to compute clusters of traffic matrices that overlap for at least a minimum time (O) in the time domain.

[0018] In an implementation the processor is further configured to select K out of N routing configurations to be assigned to the N clusters of traffic matrices. This is beneficial because sometimes the routing configurations can be reused in several clusters.

[0019] In an aspect a switch for routing data traffic is disclosed. The switch comprises at least one incoming network connection, at least one outgoing network connection, a memory configured to store a current routing configuration and at least one precomputed routing configuration, wherein routing configurations are associated with overlapping clusters of traffic matrices and a processor configured to route data packets received by the at least one incoming network connection to the at least one outgoing network connection according to the current routing configuration and activate one of the at least one precomputed routing configuration.

[0020] This is beneficial as it provides more flexibility in choosing routing configuration than the static TE, however, it requires less resources than the dynamic TE. The arrangement described above enables the possibility to compute new routing configurations as a background process so that the routing configurations are ready when they are needed. As the computation is done in background the required computing power is not as high as in dynamic implementations, wherein the computation is done on demand. A further benefit of the arrangement is that when the routing configurations are arranged to overlapping clusters of traffic matrices the required activation process can be launched flexibly.

[0021] In an implementation the processor of the switch is further configured to receive an activating signal and activate one of the at least one precomputed routing configuration as a response to the received activating signal. It is beneficial to use an activating signal as it provides the possibility to pre-fetch the routing configurations in advance.

[0022] In an aspect a method for providing routing configuration is disclosed. the method comprises monitoring current traffic conditions, predicting, based on the monitored conditions, one or more upcoming traffic matrices, selecting, based on the predicted one or more upcoming traffic matrices, at least one routing configuration from the plurality of precomputed routing configurations and transmitting at least one of the selected routing configuration to the at least one switch.

[0023] This is beneficial as it provides more flexibility in choosing routing configuration than the static TE, however, it requires less resources than the dynamic TE. The arrangement described above enables the possibility to compute new routing configurations as a background process so that the routing configurations are ready when they are needed. As the computation is done in background and in advance the required computing power for selecting and adapting the routing configuration is not as high as in dynamic TE implementations, wherein the computation is done on demand. A further benefit of the arrangement is that when the routing configurations are arranged to overlapping clusters of traffic matrices the required activation process can be launched flexibly.

[0024] In an implementation the method further comprises activating the transmitted at least one selected routing configuration in the at least one switch. It is beneficial that the selected routing configuration may be activated independently so that a good activation moment of time can be chosen after the router configuration has been transmitted to the switch.

[0025] In an implementation the method further comprises analyzing the monitored traffic conditions, predicting the traffic evolution; and computing the plurality of precomputed routing configurations based on analyzed traffic conditions and the predicted traffic evolution. This is beneficial as from the monitored conditions it is possible to predict the upcoming traffic evolution and compute routing configurations performing well with the predicted upcoming traffic.

[0026] In an aspect a computer program performing a method as described above is disclosed. It is beneficial to implement the method as a computer program as it can be configured to collect the necessary information and instruct the switches.

[0027] The aspects disclosed above are beneficial in general as they improve the routing efficiency. This provides better use of capacity of existing networks and the user experience can be improved. The arrangement above considers multiple clusters of TMs and robust routing solutions with time continuity, TM space continuity, and routing similarities. It leverages the overlap among TMs clusters to pre-fetch and anticipate network reconfigurations. As a result, it provides a good trade-off between stability and routing optimality.

DESCRIPTION OF THE DRAWINGS

[0028] The present description will be better understood from the following detailed description read in light of the accompanying drawings, wherein:

FIG. 1 illustrates an example of an apparatus for providing a routing configuration;

FIG. 2 illustrates an example of an apparatus for providing a routing configuration;

FIG. 3 is an example arrangement of overlapping traffic matrix regions;

FIG. 4 is an example method executed in an anticipatory reconfiguration module;

FIG. 5 is an example of a pre-fetching method executed in an anticipatory reconfiguration module;

FIG. 6 is an example of an activation method executed in an anticipatory reconfiguration module;

FIG. 7 is an example of a method executed in clustered robust routing (CRR) module;

FIG. 8 is an illustration of process according to the method of figure 7,

FIG. 9 is a more detailed example of a method executed in clustered robust routing (CRR) module; and

FIG. 10 is an illustration of a mathematical model.

FIG. 11 is an illustration of a mathematical model for formulating the problem of finding the minimum maximum link utilization.

DETAILED DESCRIPTION

[0029] The detailed description provided below in connection with the appended drawings is intended as a description of the embodiments and is not intended to represent the only forms in which the embodiment may be constructed or utilized. However, the same or equivalent functions and structures may be accomplished by different embodiments.

[0030] Figure 1 illustrates an example of an apparatus 100 for providing a routing configuration. The apparatus 100 comprises at least one processor 101, at least one memory 102 and a network connection 103. The apparatus may be, for example, a server, an interconnected group of servers, a personal computer, cloud computing facility, an ordinary computer or any other similar device that may be used for providing a routing configuration. The apparatus is connectable with other devices in order to receive a data communications and information on data

communications. Furthermore, the apparatus is connected to at least one switch, typically to a plurality of switches, for providing the routing configuration that it has computed. The routing configuration is provided after detecting a need for new routing configuration.

5 [0031] The apparatus for providing a routing configuration to a switch is disclosed. The apparatus 100 comprises a network connection 103, a memory 102 configured to store a plurality of precomputed routing configurations, wherein said routing configurations are associated with overlapping clusters of traffic matrices, and a processor. The processor is configured to monitor current traffic conditions,
10 predict, based on the monitored conditions, one or more upcoming traffic matrices, select, based on the predicted one or more upcoming traffic matrices, at least one routing configuration from the plurality of precomputed routing configurations, and transmit at least one of the selected at least one routing configuration to the at least one switch using the network connection.

15 [0032] Figure 2 is an a block diagram of an example comprising a network controller 200 for providing routing configurations to a plurality of switches, such as switches 203a and 203b. The block diagram of figure 2 illustrates logical operation of a network controller 200, which may be an apparatus similar to the apparatus as shown in figure 1. The network controller 200 communicates with
20 other devices, such as switches 203a – 203b and a management application 210 through application programming interfaces, including a northbound API 201 and a southbound API 202.

[0033] The operation of the network controller 200 may be divided into two domains, which are a real time operation domain 206 and a background process
25 domain 204. The background process domain comprises a Clustered Robust Routing (CRR) module that is configured to compute the plurality of precomputed routing configurations based on analyzed traffic conditions and predicted traffic evolution.

[0034] The real time operation domain 206 further comprises an anticipatory
30 reconfiguration (AR) module 207, memory storage 208 for storing traffic matrices,

routing configurations and similar. The real time operation domain 206 further comprises a monitoring and prediction module 209.

[0035] The set of robust routing solutions for overlapping traffic matrix clusters is computed by the Clustered Robust Routing (CRR) module 205. The pre-
5 fetching and activation of robust routing solutions is realized by the Anticipatory Reconfiguration module 207. Thus, the figure 2 represents a network controller 200, which control network devices (e.g., routers, switches) through a south bound interface 202. The controller 200 is configured by external applications 210 or a Network Management System (NMS) through a north bound interface 201.

10 [0036] Based on the history of traffic matrices provided by the monitoring module 209, the Clustered Robust Routing (CRR) module 205 produces a set of robust routing solutions for overlapping clusters of traffic matrices. The module can be configured via the north bound interface with the following parameters:

- 15 • Maximum number of clusters, N
- Maximum number of routing configurations, K (one configuration can be used in several clusters)
- Minimum time overlap between clusters, O, to leave enough time to pre-fetch configurations.
- 20 • Minimum time spent in each routing configuration, L, to indirectly limit the number of routing configurations.

[0037] The CRR module 205 operates in background as it implies computer intensive operations. It computes periodically or triggered on events (e.g., link
25 failures) a set of robust routing configurations associated to traffic matrix clusters. The outcome of this computation is stored in memory storage 208.

[0038] Then, in real-time, the Anticipatory Reconfiguration (AR) module 207 follows the evolution of traffic by receiving information from a monitoring system, such as the monitoring and prediction module 209. Based on an estimate of
30 actual traffic conditions and on traffic predictions, the module 209 decides to pre-fetch and activate robust routing configurations in switches. An interface between

the module 209 and the data plane devices is available to pre-fetch, delete and activate routing configurations. This interface is also used to receive routing updates which may have been decided locally by devices in case of failures or sudden traffic evolutions. This way, AR module 207 can follow the current optimality gap and takes better decisions. Robust routing configurations are stored in memory M2 of switches 211. Several of them can be stored but only one can be activated at all times. The activation may be done by transmitting an activation signal to the at least one switch in order to activate the transmitted routing configuration or by incorporating in the selected routing configuration activation instructions.

[0039] The CRR module 205 computes the plurality of precomputed routing configurations based on at least one of the following: a temporal overlap, traffic similarity, and routing similarity. Figure 3 illustrates an example of a temporal overlap and how the space of traffic matrices for the case of two flows into overlapping regions and highlights that the example provides a robust routing configuration for each cluster is divided.

[0040] This computation is realized periodically or on events by the network controller 200 as a background process. It takes as input the history of traffic matrices provided by an external monitoring system. Then, in real-time, follows the evolution of traffic and exploits the set of pre-computed clusters to decide which robust routing configurations is best. The same routing configuration may be used in several clusters. Thus the number of used routing configurations may be smaller than the number of clusters.

[0041] However, routing configurations are not usually modified at a high frequency to follow traffic variations. Flow programming in hardware equipment takes time and devices do not have enough memory to store all robust reconfigurations. To address this issue, a subset of reconfigurations inside the memory of data plane devices, such as switches 203a-203b or other similar network devices involving routing, are carefully pre-fetched. And at every time instant, activation of the best robust routing configuration is considered and decided.

[0042] The objective of the arrangement discussed in this disclosure is to find a minimum set of robust routing solutions which are computed over the history of traffic matrices and which can be pre-fetched into data plan devices. This approach greatly improves stability while optimizing routing cost. The routing configurations are rarely modified and are able to capture repeating traffic patterns to optimize cost.

[0043] The clustering of traffic matrices is guided by three objectives: 1) cover the entire traffic matrix space so that a feasible routing configuration is available for any traffic condition, 2) reduce the number of routing changes by creating a small set of robust routing configurations that can be used for a minimum duration each time it is applied, and 3) improve the routing objective function (e.g., maximum link utilization, routing cost, etc.) by computing the routing solution that minimizes objective function in each cluster.

[0044] In the following example methods will be discussed. In the example methods following notation is used. The TM at time t $y(t)$ is represented as a vector where each entry represents an origin-destination flow at time t . A cluster $C(t)$ is represented by its extreme points, namely a list of vectors. Each vector of the list represent a pseudo-matrix where each entry is an origin-destination flow of size decided by the CRR module. A TE/routing configuration at time t $R(t)$ is a matrix that associates to each origin-destination flow the set of links that are used to route the traffic. $R(t)$ has a number of rows equal to the number of links and a number of columns equal to the number of origin-destination flows. The product $R(t) * y(t)$ provides the traffic load of links at time t . The network load at time t $L(t)$ is represented as a vector where each element e represents the load of the link e . The capacity of network links is represented as vector c , where each element e represent the capacity of the link e . The element wise division $L(t) / c$ provides the network utilization at time t . The network utilization at time t $U(t)$ is represented as a vector where each element e represents the utilization of the link e . The maximum link utilization is the maximum across all links of the link utilization $U(t)$.

[0045] Figure 4 discloses an example of a method performed by an anticipatory reconfiguration module. The anticipatory reconfiguration module used

with the method illustrated in Figure 4 is assumed to be composed of two main sub-modules: the pre-fetching and the activation. Pre-fetching means that the routing configuration is retrieved and provided to the device using the routing configuration in advance and activated later. Both sub-modules are continuously executed by the controller in order to anticipate the network reconfiguration due to the traffic variation both in the near and far future.

[0046] The method receives a plurality of input parameters, step 400. The parameters received from northbound interface and monitoring system (southbound interface) include:

- 10 $y(t)$: current traffic matrix
- $y(t-1), \dots, y(t-H1)$: history of traffic matrices
- $\hat{y}(t+1)$: prediction of the next traffic matrix
- $R(t)$: active routing configuration at time

15

[0047] The input received from anticipatory reconfiguration module includes

- 20 $S(t)$: set of clusters that contain $y(t)$
- $C(t)$: cluster that corresponds to $R(t)$
- $\{[C_1, R_{j(1)}], [C_i, R_{j(i)}], \dots, [C_N, R_{j(N)}]\}$: set of pairs that maps a cluster of TMs C_i to a routing configuration $R_{j(i)}$:
- $1 \leq i \leq N \rightarrow N$ clusters of TMs
- $1 \leq j(i) \leq K \rightarrow K$ routing configurations ($K \leq N$)

25

[0048] If the current TM $y(t)$ belongs to multiple clusters ($y(t)$, step 401, is in the intersection of several clusters, then continue to step 402 (pre-fetch future TE configurations). Otherwise, go directly to step 403 (activation of TE configuration for next period of time).

30

[0049] If the current TM $y(t)$ belongs to multiple clusters, step 402, ($y(t)$ is in the intersection of several clusters) it means that the traffic in the future will vary more than the variation that can be faced by the current routing configuration. Therefore, the controller selects the best routing configurations for the most probable traffic evolution and pre-fetch in the switches these routing configurations. If there is not enough space in the memory of the switches, only a subset of routing configurations is pre-fetched.

[0050] Independently of the current status of the traffic, the controller decides which routing configuration among those stored into the switches must be activated for the next period of time, step 403. The decision is based on the prediction of the traffic evolution for the next period of time and on the robustness level provided as an input by the management app.

[0051] Figure 5 is an example of a pre-fetching method executed in an anticipatory reconfiguration module. The method of figure 5 starts with similar input parameters as the method of Figure 4.

[0052] If the current TM $y(t)$ belongs to multiple clusters, step 500, ($y(t)$ is in the intersection of several clusters), continue to the next step 501. Otherwise go to the end, step 506.

[0053] If the current TM $y(t)$ belongs to multiple clusters, step 501 ($y(t)$ is in the intersection of several clusters) the controller predicts the likelihood that the TM in the far future (e.g., H_2 period of time in the future) belongs to any cluster stored in the memory of the controller and pre-computed by CRR. The likelihood can be computed as the probability that each cluster will contain $y(t+H_2)$ in the future using for example soft-decision techniques.

[0054] Clusters are sorted in descending order of their likelihood and only the top clusters that fit in the memory of each switch (of size M_2) are selected, step 502.

[0055] The controller sends to all switches, step 503 a command to replace all existing routing configurations except the ones corresponding to the clusters $C(t)$ and $C(t+1)$, that contain the current TM and the predicted TM for the next time

period, respectively. To this end, all routing configurations except $r(C(t))$ and $r(C(t+1))$ are erased from the memory of the switch.

[0056] The routing configurations corresponding to the clusters selected in step 502 are sent to the switches, step 504. These routing configurations are stored in the tables that have been deleted in step 503. Then proceed to end of this method, step 505. The end of this method leads into the start of the activation method. An example of an activation method is shown in Figure 6.

[0057] If the current TM $y(t)$ does not belong to multiple clusters, step 600, the controller predicts the TM of the next period of time $y(t+1)$ as a function of the historical information of the TMs collected by the monitoring system. The amount of historical data (H_1) as well as other parameters of the prediction function $f_1(\dots)$ are provided as input parameters by the management app. An example of prediction function used by the controller is an ARIMA(m,n) with $H_1=(m,n)$ applied to each entry of the TM vector.

[0058] The clusters that contain the predicted TM $S(t+1)$ are computed, step 601. To this end, each entry of the predicted TM is compared to each entry (i.e., origin-destination flows) of the extreme points of each cluster to verify if the cluster contains the predicted TM.

[0059] If, step 602, the cluster $C(t)$ that contains the current TM (at time t) belongs to the set of clusters returned by step 601 (i.e., the predicted TM of the next period of time $t+1$ belongs to the same cluster of the current TM), continue to step 603. Otherwise proceed to step 604.

[0060] If, step 603, the cluster $C(t)$ that contains the current TM (at time t) belongs to the set of clusters returned by step 601, the current TE configuration $R(t)$ is used for the next period of time ($t+1$). The current TE configuration is robust against the predicted traffic change, therefore there is no need to change the configuration.

[0061] If, step 604, the cluster $C(t)$ does not belong to the set of clusters returned by step 601, the controller selects the best TE configuration among those that are stored in the memory of the switches to be activated in the next period of time $R(t+1)$. The best routing configuration $R(t+1)$ is for example the one that

minimizes the maximum link load when applied to the predicted TM $y(t+1)$. To this end, for each routing configuration R stored in the switches, the controller computes the product $L(t+1) = R * y(t+1)$ and selects the configuration R that minimizes the maximum element of all vectors $L(t+1)$.

5 [0062] The controller sends a lightweight activation command $C2$ to all switches to activate the routing configuration R selected in step 605. The command $C2$ can be as simple as an instruction that redirects all incoming traffic to the table that contains the forwarding rules corresponding to the TE configuration R .

10 [0063] Figure 7 is an example of a method executed in clustered robust routing (CRR) module. Input taken by the method may include:

M: number of Traffic Matrices

N: max number of TMs clusters

15 **K**: max number of robust configurations

O: min time overlap between clusters

L: min in time spent in each routing configuration

$y(t-M), \dots, y(t)$: Traffic Matrices (TMs)

20 [0064] First, the input is received, step 700. Then traffic Matrices (TMs) $y(t-M), \dots, y(t)$ are divided into N clusters, step 701, each cluster contains the best TMs that can be routed through a routing configuration associated to the cluster. Routing configurations to be associated are selected from a set that must be initially populated. The cluster generation must consider two constraints: 1) each cluster
25 must include a minimum number of TMs that span over a time window L ; 2) the temporal overlap among TMs of consecutive clusters must be at least O . The output of this module is a set of N clusters, each with an associated routing configuration, thus N routing configurations in total.

[0065] The N routing configurations of the step 701 are grouped in K
30 clusters, step 702, according to their reciprocal similarity in the routing domain. Such a routing clustering induces a hyper-clustering of TMs, that is, TMs

associated to clustered routing configurations are grouped into a hyper-cluster of TMs. This allows to compute K new routing configurations that must be designed to best serve the TMs in each hyper-cluster.

[0066] The performance loss obtained when applying the computed K routing configurations to the whole TM set is evaluated, step 703. If it is reduced with respect to the previous iteration and a maximum number of iterations has not been reached, the set of K new routing configurations is merged with the set of routing configurations considered for the TMs clustering. Otherwise, if the performance does not improve or a maximum number of iterations is reached, the module provides in output a mapping between N overlapping TM clusters and K routing configurations to be applied to best serve TMs of associated clusters.

[0067] A further example how the process runs is shown in figure 8, wherein the input values result into two routing configurations in a case wherein the number of traffic matrices is 17. Steps 701 and 702 of figure 7 are shown in figure 8.

[0068] Figure 9 is a more detailed example of a method executed in clustered robust routing (CRR) module. The input parameters correspond with the input parameters of figure 7.

[0069] An initial set of B ($\leq N$) routing configurations, $R_{i=1..B}$, is created, step 900. A possible procedure could be, for example, Sequential cluster M TMs in $N = \text{floor}(M/L)$ clusters; Compute a Robust Routing configuration for each cluster.

[0070] A segmentation problem is solved, step 901. It consists in finding N sequential TM clusters, $C_{i=1..N}$, by assigning each of the M TMs to one of the available routing configurations, $R_{i=1..B}$, such that

1) the time difference between the time stamp of the first and the last TM of each cluster is greater than L:

$$\forall C_i, [\text{Time}[\text{last}(C_i)] - \text{Time}[\text{first}(C_i)]] > L.$$

2) the time difference between the last TM's timestamp and the first TM's timestamp of the next cluster (or between the first TM's timestamp of the cluster and the last TM's of the previous cluster) is at least O:

$$\forall C_i, C_j \in \text{adj}(C_i),$$

$$\text{Time}[\text{last}(C_i)] - \text{Time}[\text{first}(C_j)] \geq 0$$

$$\text{OR } \text{Time}[\text{last}(C_j)] - \text{Time}[\text{first}(C_i)] \geq 0$$

5 $\text{adj}(C_i)$ is the set of clusters that adjacent to cluster C_i (e.g., adjacent in the time domain)

[0071] The goal of the segmentation problem is to assign TMs to N selected routing configurations in order to maximize the worst routing performance
10 measure. An example could be the minimization of the overall Maximum Link Utilization (MLU), the related mathematical programming model is shown in Figure 10. The MLU minimization represents in some cases a new upper bound of the utilization in every link of the network. When the bound is exceeded the information flow is typically transmitted by different transmission path.

15 [0072] For each of the N clusters of traffic matrices found in the step 901, compute, step 902, a robust routing configuration, that is, a routing configuration that maximizes the worst routing performance measure when applied to each TM in the cluster. A mathematical programming model is shown in Figure 11. The aim of this step is to tailor the routing configuration on the set of TMs in each cluster.

20 [0073] If, step 903, the maximum number of iterations is reached or the TM clustering has not changed with respect to the previous iteration, continue to the step 904. Otherwise, replace the N routing configurations in the initial set of step 900 with those computed in step 902 and go back to step 901.

[0074] Each robust routing configuration computed at step 902 can be
25 associated to one of the TM clusters computed at step 901. Represent each cluster as a point in the routing hyper-space according to the associated robust routing configuration, step 904. The dimensions of this hyper-space are given by $L \times D$, where L is the set of network links and D is the set of origin-destination pairs. The amount of traffic of given origin-destination pair routed over a given network link
30 is the value of the cluster along the corresponding dimension. This leads to consider TM clusters close if they have a similar robust routing configuration.

[0075] Then, find, step 905, K hyper-clusters in the routing hyper-space to group the N points computed at step 904. The TM clusters grouped together into a hyper-cluster are characterized by similar robust routing configurations.

[0076] Similarly to step 902, find, step 906, a robust routing configuration for each of the K hyper-clusters found at step 905. The set of TMs used to compute the robust routing configuration is the union of the TM sets associated to each cluster in the hyper-cluster.

[0077] The routing performance measure obtained when applying the K routing configurations computed at step 906 to the whole TM set $y(t-M), \dots, y(t)$ is evaluated, step 907. If it is improve with respect to the previous iteration and a maximum number of iterations has not been reached, merge, step 908, the set of K new routing configurations with the set of routing configurations considered at step 900 and go back to step 901.

[0078] Otherwise, if the routing performance measure does not improve or a maximum number of iterations is reached, the module provides in output a mapping between the N overlapping TM clusters, computed at step 901, and K routing configurations, computed at step 906, to be applied in order to best serve TMs of associated clusters.

[0079] Figure 10 is an example of a mathematical programming model that may be applied in solving the segmentation problem. Figure 10 includes a list of variables needed and the mathematical equations that may be used, for example, in step 901 of figure 9.

[0080] Figure 11 is an example of a mathematical programming model for formulating the problem of finding the minimum maximum link utilization. Figure 11 includes a list of variables needed and the mathematical equations that may be used, for example, in steps 902 and 906 of figure 9.

[0081] As explained above providing a routing configuration may be implemented as a hardware, such as the apparatus as described above, or as a method. The method may be implemented as a computer program. The computer program is then executed in a computing device.

[0082] The apparatus and method for providing a routing configuration has been described in conjunction with various embodiments herein. However, other variations to the disclosed embodiments can be understood and effected by those skilled in the art in practicing the claimed invention, from a study of the drawings, the disclosure, and the appended claims. In the claims, the word “comprising” does not exclude other elements or steps, and the indefinite article “a” or “an” does not exclude a plurality. A single processor or other unit may fulfill the functions of several items recited in the claims. The mere fact that certain measures are recited in mutually different dependent claims does not indicate that a combination of these measures cannot be used to advantage. A computer program may be stored/distributed on a suitable medium, such as an optical storage medium or a solid-state medium supplied together with or as part of other hardware, but may also be distributed in other forms, such as via the Internet or other wired or wireless telecommunication systems.

15

CLAIMS

1. An apparatus for providing a routing configuration to a switch, the apparatus comprising:

- 5 a network connection (103);
a memory (102) configured to store a plurality of precomputed routing configurations, wherein said routing configurations are associated with overlapping clusters of traffic matrices;
a processor (101) configured to:
- 10 monitor current traffic conditions;
predict, based on the monitored conditions, one or more upcoming traffic matrices;
select, based on the predicted one or more upcoming traffic matrices, at least one routing configuration from the plurality of precomputed routing configurations; and
- 15 transmit at least one of the selected at least one routing configuration to the at least one switch using the network connection.

2. The apparatus according to claim 1, wherein the processor (101) is further
20 configured to transmit an activation signal to the at least one switch in order to activate the transmitted routing configuration.

3. The apparatus according to claim 1, wherein the processor (101) is further configured to incorporate in the selected routing configuration activation
25 instructions.

4. The apparatus according to any of preceding claims 1 – 3, wherein the processor (101) is further configured to analyze the monitored traffic conditions, and based on the analysis, the processor (101) is further configured to predict the
30 traffic evolution and compute the plurality of precomputed routing configurations based on analyzed traffic conditions and the predicted traffic evolution.

5. The apparatus according to claim 4, wherein the processor (101) is further configured to compute the plurality of precomputed routing configurations based on at least one of the following: a temporal overlap, traffic similarity, and routing similarity.

5

6. The apparatus according to claim 4 or 5, wherein when computing the plurality of precomputed routing configurations the processor (101) is configured to receive a set of input parameters, wherein the input parameters comprise:

- a set of traffic matrices received from monitoring (M)
- a set of traffic matrices received from prediction;
- a maximum number of clusters (N);
- a maximum number of routing configurations (K)
- a minimum time overlap among clusters (O); and
- a minimum time of utilization of a configuration (L).

10

15

7. The apparatus according to claim 6, wherein the processor (101) is further configured to:

- sort the set M according to increasing traffic matrix sampling time; and
- find an initial set of B routing configurations: R_1, \dots, R_B .

20

8. The apparatus according to claim 6 or 7, wherein the processor (101) is further configured to create N clusters of traffic matrices, wherein $N < B$, $C_{i=1..N}$, wherein the processor (101) is further configured to assign for each of M traffic matrices a routing configuration R_i such that an overall maximum link utilization is minimized.

25

9. The apparatus according to claim 8, wherein the processor (101) is configured to assign for each of M traffic matrices a routing configuration such that upcoming conditions are met:

- $\forall C_i, Tim[(C_i)] - Time[first(C_i)] \geq L$; and
- $\forall C_i, C_j \in adj(K_i), Time[last(C_i)] - Time[first(C_j)] \geq 0$; or

30

$$Time[(C_j)] - Time[first(C_i)] \geq 0.$$

10. The apparatus according to claims 6-9, wherein the processor (101) is further configured to select K out of N routing configurations to be assigned to the
5 N clusters of traffic matrices.

11. A switch for routing data traffic comprising:
at least one incoming network connection;
at least one outgoing network connection;
10 a memory configured to store a current routing configuration and at least one precomputed routing configuration, wherein routing configurations are associated with overlapping clusters of traffic matrices;
a processor configured to:
route data packets received by the at least one incoming network connection
15 to the at least one outgoing network connection according to the current routing configuration; and
activate one of the at least one precomputed routing configuration.

12. The switch according to claim 11, wherein the processor is further
20 configured to:
receive an activating signal; and
activate one of the at least one precomputed routing configuration as a response to the received activating signal.

25 13. A method for providing a routing configuration to a switch, the method comprising:
monitoring current traffic conditions;
predicting, based on the monitored conditions, one or more upcoming traffic matrices;

selecting, based on the predicted one or more upcoming traffic matrices, at least one routing configuration from the plurality of precomputed routing configurations; and

5 transmitting at least one of the selected routing configuration to the at least one switch.

14. The method according to claim 13, wherein the method further comprises:

10 activating the transmitted at least one selected routing configuration in the at least one switch.

15. The method according to any of preceding claims 12 – 14 the method further comprising:

analyzing the monitored traffic conditions;
15 predicting the traffic evolution; and
computing the plurality of precomputed routing configurations based on analyzed traffic conditions and the predicted traffic evolution.

20 16. A computer program comprising a program code configured to perform a method according to any one of claims 12 - 15, when the computer program is executed on a computing device.




Figure 1




Figure 2




Figure 3

4/11

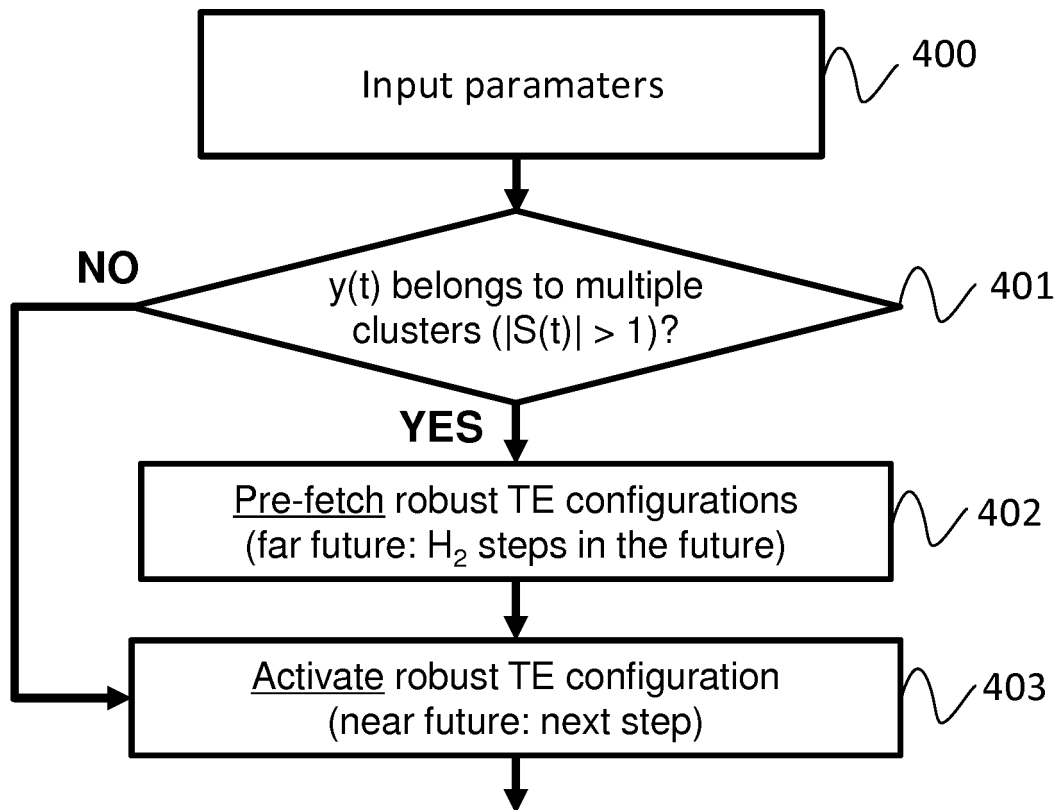



Figure 4

Figure 5






Figure 6




Figure 7




Figure 8

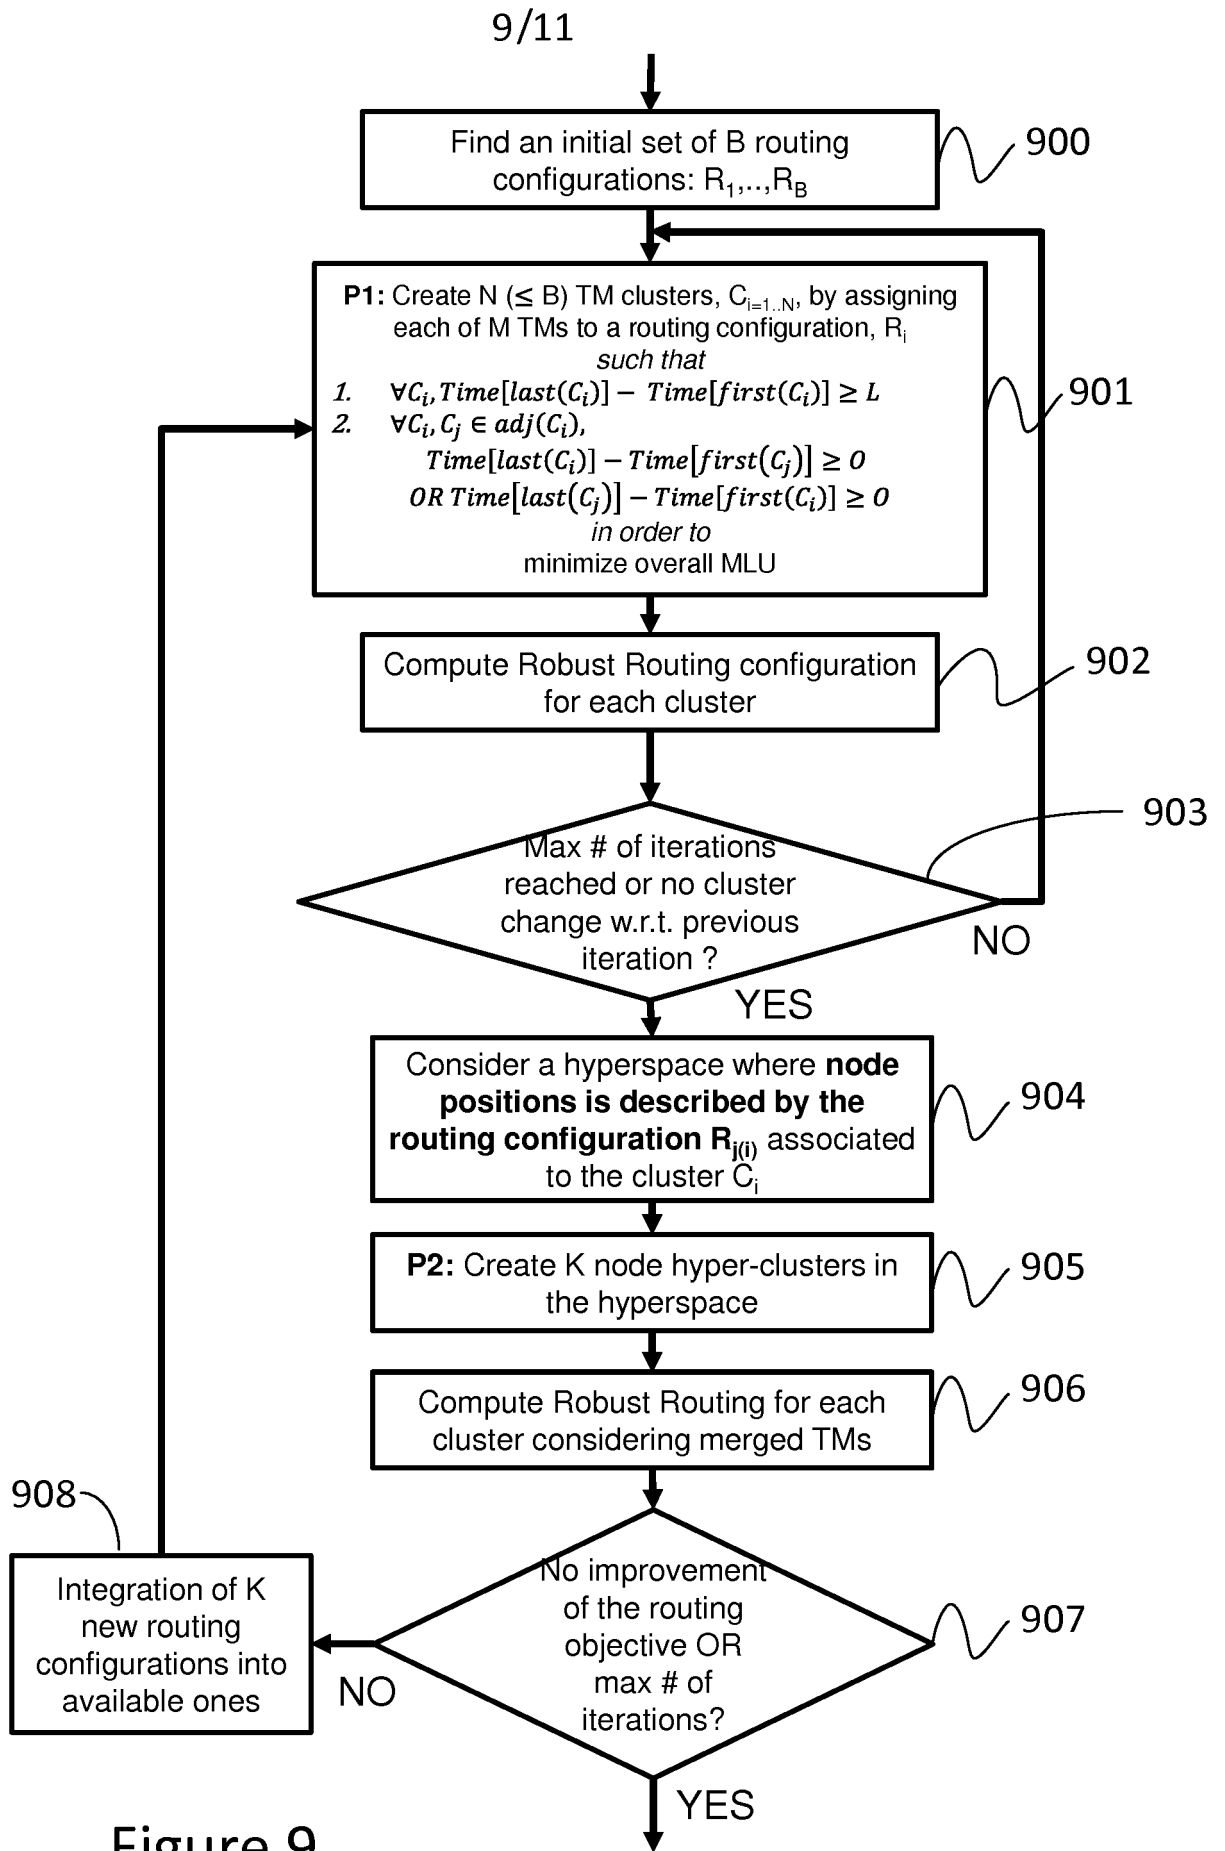


Figure 9

Input parameters

M	set of TMs
B	set of routing configurations
δ_{ij}	MLU if TM i is assigned to routing configuration j
N	Max number of clusters
O	Min time overlap among clusters
L	Min time length of (in-time spent in) a cluster

Main decision variables

x_{ij}	1 if TM i is associated to routing configuration j , 0 otherwise
z_j	1 if routing configuration j is selected

minimize γ_{max}

subject to :

$$\delta_{i+l,j} \cdot x_{i,j} \leq \gamma_{max} \quad \forall i \in M, j \in B, l \in 0, \dots, O : i+l \in I$$

$$\delta_{i-l,j} \cdot x_{i,j} \leq \gamma_{max} \quad \forall i \in M, j \in B, l \in 0, \dots, O : i-l \in I$$

$$y_{i,j} \geq x_{i+1,j} - x_{i,j} \quad \forall i \in M, j \in B$$

$$\sum_{j \in B} x_{i,j} = 1 \quad \forall i \in M$$

$$\sum_{i \in M} x_{i,j} \geq L \cdot z_j \quad \forall j \in B$$

$$\sum_{i \in M} y_{i,j} \leq z_j \quad \forall j \in B$$

$$\sum_{i \in M, j \in J} y_{i,j} = \sum_{j \in B} z_j - 1 \quad \forall j \in B$$

$$\sum_{j \in B} z_j \leq N$$

$$x_{i,j} \in \{0, 1\} \quad \forall i \in M, j \in B$$

$$y_{i,j} \in \{0, 1\} \quad \forall i \in M, j \in B$$

$$z_j \in \{0, 1\} \quad \forall j \in J$$

$$\gamma_{max} \geq 0$$

$$(N \geq 0, O \leq L)$$

Figure 10

Input parameters

M	set of TMs
D	set of origin-destination pairs
N	set of network nodes
L	set of network links
$c_{i,j}$	capacity of link (i,j)
R_m^d	traffic request of origin-destination pair d in TM m

Main decision variables

x_{ij}^d	amount of traffic of origin-destination pair d routed over link (i,j)
γ_m	MLU when TM m is considered

minimize γ_{max}

subject to :

$$\sum_{j:(i,j) \in L} x_{i,j}^d - \sum_{j:(j,i) \in L} x_{j,i}^d = \begin{cases} 1 & \text{if } i = s_d \\ -1 & \text{if } i = f_d \\ 0 & \text{otherwise} \end{cases} \quad \forall i \in N, d \in D$$

$$\sum_{d \in D} R_m^d \cdot x_{i,j}^d \leq c_{i,j} \quad \forall m \in M, (i,j) \in L$$

$$\gamma_m \geq \frac{\sum_{d \in D} R_m^d \cdot x_{i,j}^d}{c_{i,j} - \sum_{d \in D} R_m^d \cdot x_{i,j}^d} \quad \forall m \in M, (i,j) \in L$$

$$\gamma_{max} \geq \gamma_m \quad \forall m \in M$$

$$0 \leq x_{i,j}^d \leq 1 \quad \forall d \in D, (i,j) \in L$$

$$\gamma_m \geq 0 \quad \forall m \in M$$

$$\gamma_{max} \geq 0$$

Figure 11

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2017/062851

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/721 H04L12/717
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 H04W

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	US 2015/327135 A1 (LUO MIN [US] ET AL) 12 November 2015 (2015-11-12) paragraph [0003] - paragraph [0007]; figure 5 paragraph [0025] - paragraph [0030] paragraph [0047] - paragraph [0058] paragraph [0070] - paragraph [0095]	1-16
X	US 7 729 269 B1 (NUCCI ANTONIO [US] ET AL) 1 June 2010 (2010-06-01) column 1, line 62 - column 2, line 32 column 2, line 55 - column 7, line 54	1-16
A	US 2016/352648 A1 (IOVANNA PAOLA [IT] ET AL) 1 December 2016 (2016-12-01) paragraph [0035] - paragraph [0051]	1-16
	----- -/--	

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 January 2018	Date of mailing of the international search report 05/02/2018
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 Ciurel, Cristian

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2017/062851

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2010/214920 A1 (TEWANI KAMLESH [US] ET AL) 26 August 2010 (2010-08-26) claims 1-10 -----	1-16

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2017/062851

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015327135	A1	12-11-2015	NONE

US 7729269	B1	01-06-2010	NONE

US 2016352648	A1	01-12-2016	CN 106031102 A 12-10-2016
		EP 3108623 A1	28-12-2016
		US 2016352648 A1	01-12-2016
		WO 2015120912 A1	20-08-2015

US 2010214920	A1	26-08-2010	NONE
