

(51) International Patent Classification:
H04L 12/28 (2006.01)(21) International Application Number:
PCT/US2010/051698(22) International Filing Date:
6 October 2010 (06.10.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/250,013 9 October 2009 (09.10.2009) US
12/723,697 15 March 2010 (15.03.2010) US(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).(72) Inventors: **PADHYE, Jitendra D.**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **KANDULA, Srikanth**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **BAHL, Paramvir**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (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, BR, BW, BY, BZ,

CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: FLYWAYS IN DATA CENTERS




FIG. 1

(57) Abstract: Described is a technology by which additional network communications capacity is provided to an oversubscribed base network where needed, through the use of dynamically provisioned communications links referred to as flyways. A controller detects a need for additional network communications capacity between two network machines, e.g., between two racks of servers with top-of-rack switches. The controller configures flyway mechanisms (e.g., one per rack) to carry at least some of the network traffic between the machines of the racks and thereby provide the additional network communications capacity. The flyway mechanisms may be based on any wireless or wired technologies, including 60GHz technology, optical links, 802.11n or wired commodity switches.

FLYWAYS IN DATA CENTERS

BACKGROUND

[0001] Large network data centers provide economies of scale, large resource pools, simplified IT management and the ability to run large data mining jobs. Containing the network cost is an important consideration when building large data centers. Networking costs are one of the major expenses; as is known, the cost associated with providing line speed communications bandwidth between an arbitrary pair of servers in a server cluster generally grows super-linearly to the size of the server cluster.

[0002] Production data center networks use high-bandwidth links and high-end network switches to provide the needed capacity, but they are still over-subscribed (lacking capacity at times) and thus suffer from sporadic performance problems. Oversubscription is generally the result of a combination of technology limitations, the topology of these networks (e.g., tree-like) that requires expensive “big-iron” switches, and pressure on network managers to keep costs low. Other network topologies have similar issues.

SUMMARY

[0003] This Summary is provided to introduce a selection of representative concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used in any way that would limit the scope of the claimed subject matter.

[0004] Briefly, various aspects of the subject matter described herein are directed towards a technology by which additional network communications capacity is provided to an oversubscribed base network through the use of communications links referred to as flyways. In general, each flyway comprises an additional network communication path between network machines that may be dynamically set up as needed, and torn down as desired.

[0005] In one implementation, each rack of servers in a network has an associated flyway mechanism that can communicate with other flyway mechanisms. The flyway mechanisms may be based on wireless or wired technologies. A controller dynamically provisions a flyway upon detecting a need for additional network communications capacity between the two racks by having the flyway mechanisms communicate with one another. Once provisioned, the additional communications link between machines is used to carry at least some of the network traffic between those machines.

[0006] Other advantages may become apparent from the following detailed description when taken in conjunction with the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The present invention is illustrated by way of example and not limited in the
5 accompanying figures in which like reference numerals indicate similar elements and in which:

[0008] FIGURE 1 is a block diagram showing an example data center incorporating flyway mechanisms by which flyways may be established.

[0009] FIG. 2 is a representation of flyways set up between network machines.

10 [0010] FIG. 3 is a representation of a partial data center showing server racks and the range of an example 60GHz device mounted atop a rack near the center of a set of racks.

[0011] FIG. 4 is a flow diagram showing example steps for provisioning a flyway between a pair of network machines.

DETAILED DESCRIPTION

15 [0012] Various aspects of the technology described herein are generally directed towards a network design that retains the benefits of existing networks, yet substantially improves their performance. To this end, there is described a hybrid architecture in which a base network (including whether tree-like or mesh-like) is provisioned for the average case (and is thus oversubscribed), but handles hotspots through the inclusion of extra links that
20 are added on an on-demand basis to provide additional capacity where and when it is needed in the network. As used herein, such links are referred to as flyways. Note that future network designs (e.g., Clos networks) also can be enhanced by the technology described herein. As will be understood, the enhanced design allows networks to maintain their performance as they scale-up in server counts, without significant increase in the
25 quantity and the cost of networking equipment.

[0013] It should be understood that any of the examples described herein are non-limiting examples. As such, the present invention is not limited to any particular embodiments, aspects, concepts, structures, functionalities or examples described herein. Rather, any of the embodiments, aspects, concepts, structures, functionalities or examples
30 described herein are non-limiting, and the present invention may be used in various ways that provide benefits and advantages in computing and computer networks in general.

[0014] FIG. 1 shows a production network based upon a tree-like topology. A plurality of racks 102₁-102_n each have servers, which communicate through a top of rack switch 104₁-104_n. A typical network has twenty to forty servers per rack, with increasingly

powerful links and switches going up the tree. Note that flyways are not limited to tree-like topologies, but can be used in any topology, including clos networks and other forms of mesh topologies, and FatTree topologies.

[0015] As represented in FIG. 1, each top of rack switch 104₁-104_n is coupled to one
5 another through one or more aggregation switches 106₁-106_k. In this way, each server may communicate with any other server, including a server in a different rack. Note that in this example, a higher-level aggregation switch 108 couples the rack-level aggregation switches 106₁-106_k, and there may be one or more additional levels of aggregation switch couplings.

10 [0016] Application demands generally can be met by an oversubscribed network, but occasionally the network does not have sufficient capacity to handle "hotspots." The technology described herein provides the additional capacity through the use of flyways to handle extra data traffic as needed.

[0017] As represented in FIGS. 1 and 2, flyways (the curved arrows in FIG. 2) are
15 implemented via flyway mechanisms 110₁-110_n that are controlled by a flyway controller 112. The flyways may be dynamically set up by the flyway controller 112 on an as-needed basis, and taken down when not needed (or needed elsewhere). FIG. 2 shows how one flyway 220 may be used to link racks 102₁ and 102_n and their respective top-of-rack switches 104₁ and 104_n, while another flyway 222 links racks 102₂ and 102_m and their
20 respective top-of-rack switches 104₂ and 104_n. Note that a rack / top-of-rack switch may have more than one flyway at any time, as represented by the flyway 221 between racks 102₁ and 102₂. While a single flyway mechanism is shown per rack, it can be appreciated that there may be more than one flyway mechanism per rack (or multiple devices in a single flyway mechanism), possibly using different communications technologies (e.g.,
25 wireless and optical).

[0018] Analysis of traces from data center networks shows that, at any time, only a few top-of-rack switches are "hot," that is, they are sending and/or receiving a large volume of traffic. Moreover, when hot, top-of-rack switches typically exchange much of their data with only a few other top-of-rack switches. This translates into skewed bottlenecks, in
30 which just a few of the top-of-rack switches lag behind the rest and hold back the entire network. The flyways described herein provide extra capacity to these few top-of-rack switches and thus significantly improve overall performance. Indeed, only a few flyways, with relatively low bandwidth, significantly improve the performance of an oversubscribed data center network.

[0019] Flyways may be added to a network at a relatively small additional cost. This may be accomplished by the use of wireless links (e.g., 60GHz, optical links and/or 802.11n) and/or the use of commodity switches to add capacity in a randomized manner. In general, any flyway mechanism can link to any other flyway mechanism, as long as they meet coupling requirements (e.g., within range for wireless, has line-of-sight for optical and so on).

[0020] Thus, the flyways may be implemented in various ways, including via wireless links that are set up on demand between the flyway mechanisms (e.g., suitable wireless devices), and/or commodity switches that interconnect subsets of the top-of-rack switches. As described below, 60 GHz wireless technology is one implementation for creating the flyways, as it supports short range (1-10 meters), high-bandwidth (1Gbps) wireless links. Further, the high capacity and limited interference range of 60GHz provides benefits.

[0021] The performance of a flyway-enhanced oversubscribed network may approach or even equal to that of a non-oversubscribed network. One way to achieve the most benefit is to place flyways at appropriate locations. Note that network traffic demands are generally predictable / determinable at short time scales, allowing the provisioning of flyways to keep up with changing demand. As described herein, the central flyway controller 112 gathers demand data, adapts the flyways in a dynamic manner, and switches paths to route traffic.

[0022] Wireless flyways are able to form links on demand, and thus may be used to distribute the available capacity to whichever top-of-rack switch pairs need it as determined by the central flyway controller 112. Wired flyways provide equivalent benefits; when inexpensive switches are connected to subsets of the top-of-rack switches, the limited backplane bandwidth at these switches can be divided among whichever of the many top-of-rack switch pairs that are connected need it. Wired flyways are somewhat more restrictive in that only if a top-of-rack switch pair happen to be connected via one of the flyway switches, can they directly benefit from a flyway (although multi-hop flyways are one alternative that allows indirect benefits, as described below). However, wired flyways are more likely to keep up with wired speeds (for example, as NICs go up to 10Gbps and links go to 40Gbps). In any event, flyways provide benefits relative to the alternative of spreading the same bandwidth across all racks, as much of it will go unused on links that do not need it. Further, flyways can be deployed today on existing topologies of production data centers.

[0023] To add flyways to provide extra capacity between hot top-of-rack switches and some other top-of-rack switches with which they exchange traffic, pairs of flyway mechanisms that provide the most speedup may be selected. Each flyway also needs sufficient capacity. Establishing a flyway between a rack with a top-of-rack switch that is
5 the most congested and another rack with which the switch exchanges the most data is straightforward. However subsequent choices are less clear, e.g., whether to establish another flyway at the same rack or elsewhere needs to be considered.

[0024] In general, spreading flyways too thinly, or concentrating them at a few top-of-rack switches does not work particularly well. For example, placing one flyway each
10 between the top fifty top-of-rack switches and their largest correspondent does not sufficiently reduce the completion time of the hot top-of-rack switch. Conversely, placing flyways between the top five top-of-rack switches and each of their ten largest correspondents eliminates congestion at the top five, but results in the sixth top-of-rack switch ending up as the bottleneck. In general, achieving a proper balance between
15 helping more top-of-rack switches and reducing enough congestion at every one of the hot top-of-rack switches obtains the most speedup. One suitable algorithm for flyway placement is described below.

[0025] With respect to how much capacity each flyway needs, consider an example of adding flyways between the top ten top-of-rack switches and each of the five other top-of-rack switches with which they exchange the most data (i.e., a total of fifty flyways). Most
20 flyways need less than ten percent the top-of-rack switch's uplink bandwidth to be useful; one reason is that while the top-of-rack switch's uplink carries traffic to all of the other top-of-rack switches, a flyway has to only carry traffic to one other top-of-rack switch.

[0026] The usefulness of flyways generally results from application characteristics that
25 cause sparse demand matrices. Though flyways may not provide benefits for all demands, a large set of practical applications benefit from flyways. For example, in data centers that support web services, the request traffic is load balanced across servers, each of which in turn assembles a response page by possibly asking a few other servers to generate parts of the response (e.g., advertisements). The reduce part of map-reduce in data mining jobs is
30 possibly a worst case scenario, with each reducer pulling data from all the mappers; the job is bottlenecked until all the reducers complete. Even then, it is rare to have so many mappers and reducers that all top-of-rack switches are congested simultaneously.

[0027] Turning to 60 GHz wireless communications and the like, millimeter wavelength wireless communications has rapidly advanced, and is a suitable technology for

constructing flyways in a data center. The 60 GHz frequency band is a 7 GHz wide band (57-64 GHz) that was set aside by the FCC in 2001 for use by unlicensed devices. The channel bandwidth specified by ISO and IEEE is 2.160 GHz. The substantially larger bandwidth in the 60 GHz band facilitates higher capacity links. For example, a simple
5 encoding scheme that achieves as little as 1 bps/Hz makes it possible to construct a link with a nominal bandwidth of 7 Gbps. As better modulation and coding techniques are implemented, 60 GHz devices will likely become spectrally efficient and seven orthogonal 1 Gbps channels will be common, and 1 Gbps link speeds are adequate for constructing flyways. Note that 4-5 Gbps link speeds are already available using multiple channels in
10 the 60GHz band.

[0028] In addition to capacity, other characteristics of 60 GHz communications make it particularly attractive for instantiating flyways. For one, 60 GHz devices have a relatively short range (5 to 10 meters). In wireless communications path, loss is directly proportional to the square of the frequency and so 60 GHz signals attenuate more quickly with distance
15 when compared with 2.4 and 5 GHz Wi-Fi signals. For another, highly directional 60 GHz communications are easy to achieve. Because signal wavelength is inversely proportional to frequency, 60 GHz signals have very short wavelengths (5 mm) and this allows RF chip designers to incorporate antenna arrays into the transceiver chips. The ability to integrate many antennas and transceivers with proper phasing into a single chip
20 allows beam-forming and hence directional communication. Additionally, phased-array antennas boost signal quality leading to a cleaner signal with low bit error rates at short distances. For example, a one-square inch antenna provides a signal boost of 25dBi at 60 GHz. A side benefit of the integration of the antennas on the chip is that this avoids the need for wires to carry signals to and from the chip, reducing the cost of packaging by up
25 to two orders of magnitude.

[0029] The low range and high capacity directional links are beneficial in data centers where server racks are typically packed tightly as physical space is a premium resource, and mobility and a power draw of approximately ten watts per server rack is a non-issue. Directionality allows the network designers to increase the overall spectrum efficiency
30 through spatial reuse. For example, two sets of communications between four top-of-rack switches can occur simultaneously because of directionality and range.

[0030] In a typical data center server, racks are about 24 inches wide and are arranged in a dense, regular pattern as shown in FIG. 3. In FIG. 3, each box represents a 24x48 inch rack, for example, arranged in rows of ten. The circle represents the 10m range of a 60

GHz device mounted atop a rack in the center, which contains about 70 other racks. Using directional antennas, small (2-3 cubic inches) 60 GHz network cards mounted on top of the racks can instantiate multiple non-interfering gigabit flyways between several racks on an as needed basis. Electronically steerable phased-array antennas with beam widths of
5 about 60 degrees can be steered with millisecond latency.

[0031] Note that directionality comes with consequences. For example, discovery is no longer straightforward and the popular medium access control protocols that work well for omni-directional communications have to be modified. Notwithstanding, this is a well-understood problem in the research community and many solutions exist. The IEEE and
10 ISO 60 GHz working groups are actively working on incorporating support for directional communications and spatial reuse. Although their primary scenario is in-home multimedia communications, the solutions they are discussing can be modified to apply to data center networks. Specifically, because of the way the network is configured and the scale of the network, the problem may be solved because there exists a back-channel
15 connection, often wired, between all discoverable top-of-rack switches that can be turned into a control channel. The control channel can then be used to discover top-of-rack switches, align 60 GHz beams, allocate channels for interference mitigation, and arbitrate collision-free medium access.

[0032] Another issue with 60 GHz is that at these high-frequencies, signals are absorbed
20 by objects and non-line-of-sight (NLOS) communications is not possible. However, in data centers, devices may be mounted on top of the racks and out of the way of human operators, whereby line-of-sight is not interrupted. Moreover, recent advances in technology have started to overcome this limitation.

[0033] As noted above, the large amount of spectrum combined with improving spectral
25 efficiency allows many flyways to be constructed and operated independently (with a 10 bits/Hz spectral efficiency, up to 70 flyways each operating at 1Gbps can be constructed). This combined with the properties of low interference due to directionality, and low range due to path loss allows significant frequency reuse over thousands of server racks present in the data center.

30 [0034] Moreover, unlike other high bandwidth wireless technologies, the 60 GHz band is available world-wide. Thus, chipsets and devices are usable all over the world and providing economies of scale and pushing the cost down.

[0035] Turning to implementing a network with flyways, to form a flyway between a pair of top-of-rack switches, one or more devices placed atop the corresponding racks

create a wireless link. The choice of technology affects the available bandwidth, the number of channels available for spatial re-use, interference patterns and the range of the flyway. The antenna technology dictates the time needed to setup and tear down a flyway. Note that adding a few wireless devices per top-of-rack switch increases cost only

5 marginally.

[0036] Wired flyways may be constructed by using additional switches, such as of the same make as contemporary top-of-rack switches that interconnect random subsets of the top-of-rack switches, e.g., commodity switches may interconnect twenty top-of-rack switches with 1Gbps links each. To keep links short, the flyway switches may connect
10 racks that are near to one another in the datacenter. When deploying wired flyways, spectrum allocation or interference is not an issue. However, their random construction constrains wired flyways; e.g., top-of-rack switch pairs that exchange a lot of traffic and can benefit from surplus capacity may end up without a wired flyway.

[0037] In any event, either wired or wireless flyway construction is an improvement
15 compared to dividing the same amount of bandwidth uniformly across all pairs of top-of-rack switches. Rather than spread bandwidth uniformly and have much of it wasted, as happens when the demand matrix is sparse, flyways provide a way to use the spare bandwidth to target the parts of the demand matrix that can benefit the most from surplus capacity.

[0038] The central flyway controller 112 (FIG. 1) gathers estimates of demands between
20 pairs of the rack switches. For example, information can be gathered from lightweight instrumentation at the end servers themselves, or by polling SNMP counters at the switches. Using these estimates, the controller 112 may periodically or otherwise run a placement algorithm (described below) to place the available flyways.

[0039] The topology of a flyway-based network is thus dynamic, and requires multipath
25 routing. Solutions to similar problems are known. The controller 112 determines how much of the traffic between a pair of top-of-rack switches go along the base network, or instead take a flyway from the sending top-of-rack switch to the receiving top-of-rack switch, if one exists. The top-of-rack switch splits traffic as per this ratio by assigning
30 different flows onto different MPLS (Multiprotocol Label Switching) label switched paths. Note that only a few flyways, if any, are available at each top-of-rack switch. Hence, the number of LSPs (label switched paths) required at each switch is small and the problem of splitting traffic across the base and flyways that are one hop long is significantly simpler than other traffic splitting techniques.

[0040] The problem of creating optimal flyways can be considered an optimization problem. Given D_{ij} demand between top-of-rack switches i, j and if C_l is the capacity of link l , an optimal routing is the one that minimizes the maximum completion time:

$$\begin{aligned} & \text{such that} \quad \min \max_{ij} \frac{D_{ij}}{r_{ij}} \quad (1) \\ & \sum_{l \in \text{incoming}} r_{ij}^l - \sum_{l \in \text{outgoing}} r_{ij}^l = \begin{cases} D_{ij} & \text{at ToR } j \\ -D_{ij} & \text{at ToR } i \\ 0 & \text{at all other ToRs} \end{cases} \\ & \sum_{ij} r_{ij}^l \leq C^l \quad \forall \text{ links } l \end{aligned}$$

where r_{ij} is the rate achieved for the top-of-rack switch pair i, j and r_{ij}^l is the portion of that pair's traffic on link l .

[0041] Computing an optimal / near optimal flyway placement involves suitably changing the topology and re-solving the above optimization problem. For example, all possible flyways may be added, with the constraint that no more than a certain number can be simultaneously active, or that none of the flyways can have a capacity larger than a certain amount. Not all the variants of the above optimization problem, induced due to technological constraints, may be tractable. Instead, greedy procedures may be used that add one flyway at a time by solving the above optimization problem to find the flyway that reduces completion times the most. Other procedures may be better approximations.

[0042] An example is generally represented in the example flow diagram of FIG. 4, beginning at step 402 which represents gathering the network traffic data for the pairs being considered. In general, based on the traffic data, the most optimal flyway that is needed (e.g., the switches are hot but a flyway is not already provisioned) between two selected pairs of machines (top-of-rack switches) is selected at step 404.

[0043] Note that a flyway also may be used to provide connectivity to racks whose primary connections to the rest of the racks have failed. Such a failure-compensating flyway (or flyways) may be considered the most optimal flyway that is needed at step 404, for example, or may be separately considered, such as before selecting others based upon capacity.

[0044] Step 406 represents evaluating whether that most optimal flyway can be provisioned; for example, if the total number of flyways is a constraint, there may not be another flyway available at this time. However, if there is not a flyway available, yet the

most optimal non-provisioned flyway is more needed than another flyway that is currently provisioned, step 408 may tear down that other flyway (or flyways, if a multi-hop scenario is in use as described below) so that a flyway becomes available. For example, the least needed flyway may be torn down.

5 [0045] Note however that it may not be appropriate to tear down a flyway. For example, it may be that all of the already provisioned flyways are handling even more traffic and thus are more needed than a flyway for the selected pair at step 404, in which case no action is taken (e.g., the dashed line from step 408).

[0046] Step 410 provisions the most optimal flyway (if appropriate). Once provisioned,
10 the flyway may be used to handle at least some of the network traffic between the selected pair, as represented by step 412. The process then repeats. Note that FIG. 4 is only one example; as an alternative, up to the top-N most optimal flyways that are not already provisioned may be added between each polling of the pairs at step 402.

[0047] With flyways, the physical topology of a network needs to be considered, as well
15 as constraints due to the technology. Wireless flyways are constrained by range, while wired flyways, which are constructed by inter-connecting random subsets of top-of-rack switches, can only provide surplus capacity between these random subsets. 60 GHz flyways span a distance of ten meters and use the datacenter layout (e.g., FIG. 3) from a production data-center. For wired flyways, 24 port, 1Gbps switches may be used, with
20 some of the ports used for the flyway channel. Although such constraints impact the benefit of flyways, the gains are still significant.

[0048] Note that many more wired flyways need to be added to obtain the same benefits of wireless flyways. For example, when fifty 24-port switches are added, 1,200 duplex links are added to the network. Wireless flyways provide equivalent performance with
25 only 50 half-duplex links. Thus is because wireless flyways are added in a targeted manner: they help speed up those pairs of top-of-rack switches that need additional capacity. Wired flyways are added at random, and benefit pairs only if they happen to be among the ones chosen.

[0049] Multi-hop flyways are feasible. Consider that a machine in FIG. 3 within the
30 circle wants to communicate with a machine outside of the circle. However, a second hop from a machine inside the circle that is close enough to another machine can be used to accomplish the communication.

[0050] Further, the controller 112 (FIG. 1) can consider whether two machines (e.g., A and Z) are outside of single hop range. If so, a flyway may be constructed between two

other machines (e.g., C and M) that are close but communicating via the top-of-rack switches, with their communication switched to the flyway. The bandwidth savings provided by the newly constructed flyway for C and M then may be used to free up the top-of-rack switches to allow the machines A and Z to communicate using the top-of-rack switches.

[0051] While the invention is susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific forms disclosed, but on the contrary, the intention is to cover all modifications, alternative constructions, and equivalents falling within the spirit and scope of the invention.

CLAIMS

1. In a computer networking environment, a system comprising, a set of flyway mechanisms, including one flyway mechanism coupled to a rack of servers, the flyway mechanism communicating with another flyway mechanism to provide additional network communications capacity to an oversubscribed base network.
2. The system of claim 1 further comprising a controller that detects a need for additional network communications capacity, wherein the flyway mechanisms are dynamically provisioned to communicate with one another by the controller.
3. The system of claim 1 wherein the oversubscribed base network is configured in a tree-like topology or in a mesh topology.
4. The system of claim 1 wherein the flyway mechanisms are based at least in part on a physical layout of the network.
5. The system of claim 1 wherein the flyway mechanism communicates with the other flyway mechanism over a physical wire, or wherein the flyway mechanism communicates with the other flyway mechanism via wireless technology.
6. The system of claim 1 wherein the flyway mechanism communicates with the other flyway mechanism via wireless technology, and wherein the wireless technology comprises 60 GHz frequency band technology, optical links or 802.11-based Wi-Fi technology, or any combination of 60 GHz frequency band technology, optical links or 802.11-based Wi-Fi technology.
7. One or more computer-readable media having computer-executable instructions, which when executed perform steps, comprising, determining network traffic (402), and based upon the network traffic, provisioning a communications link between machines in the network to carry at least some of the network traffic between those machines.
8. The one or more computer-readable media of claim 7 having further computer-executable instructions comprising, tearing down the communications link.
9. The one or more computer-readable media of claim 7 having further computer-executable instructions comprising, provisioning another communications link between machines, including using at least one intermediary machine for multi-hop communication.
10. The one or more computer-readable media of claim 7 wherein determining the network traffic comprises communicating with server machines or polling switches, or both communicating with server machines and polling switches.

11. The one or more computer-readable media of claim 7 wherein provisioning the communications link comprises running a placement algorithm to select the machines between which the communications link is provisioned.

12. In a computing environment, a method performed on at least one processor, comprising:

determining network traffic between pairs of network machines;

selecting, based on the network traffic, a selected pair of machines that does not already have a flyway provisioned between that selected pair of machines; and

provisioning a flyway between that selected pair of machines.

13. The method of claim 12 further comprising, tearing down a flyway between another pair of machines.

14. The method of claim 12 wherein selecting the selected pair of machines comprises running a placement algorithm to compute an optimal routing.

15. The method of claim 12 further comprising, provisioning at least two flyways to enable multi-hop network communication between two machines that are unable to communicate over a single flyway.

1/4




FIG. 1






FIG. 2

3/4

**FIG. 3**

4/4

**FIG. 4**