



(19) **United States**

(12) **Patent Application Publication**
Senarath et al.

(10) **Pub. No.: US 2016/0359725 A1**

(43) **Pub. Date: Dec. 8, 2016**

(54) **METHOD AND SYSTEM TO REPRESENT THE IMPACT OF LOAD VARIATION ON SERVICE OUTAGE OVER MULTIPLE LINKS**

H04L 12/721 (2006.01)
H04L 12/707 (2006.01)
H04W 40/16 (2006.01)

(71) Applicant: **Huawei Technologies Co., Ltd.**,
Shenzhen (CN)

(52) **U.S. Cl.**
CPC *H04L 45/125* (2013.01); *H04L 45/24*
(2013.01); *H04W 40/16* (2013.01); *H04L*
45/38 (2013.01); *H04L 43/16* (2013.01);
H04L 43/0888 (2013.01)

(72) Inventors: **Nimal Gamini Senarath**, Ottawa (CA);
Aaron James Callard, Ottawa (CA);
Ho Ting Cheng, Stittsville (CA)

(21) Appl. No.: **15/243,734**

(57) **ABSTRACT**

(22) Filed: **Aug. 22, 2016**

Increased resource utilization efficiency can be improved by modeling path costs during admission and path-selection. Specifically, path costs for candidate paths are modeled based on load characteristics (e.g., current load, load variation, etc.) of links in the candidate paths. Path costs can represent any quantifiable cost or liability associated with transporting a service flow over the corresponding path. For example, path costs can correspond to a probability that at least one link in the path will experience an outage when transporting the service flow, a price charged by a network operator (NTO) for transporting the traffic flow over the candidate path, or a total network cost for transporting the flow over a candidate path. The candidate path having the lowest path cost is selected to transport a service flow.


Related U.S. Application Data

(63) Continuation of application No. 14/203,276, filed on Mar. 10, 2014, now Pat. No. 9,426,075.

(60) Provisional application No. 61/778,104, filed on Mar. 12, 2013.

Publication Classification

(51) **Int. Cl.**
H04L 12/729 (2006.01)
H04L 12/26 (2006.01)






FIG. 1

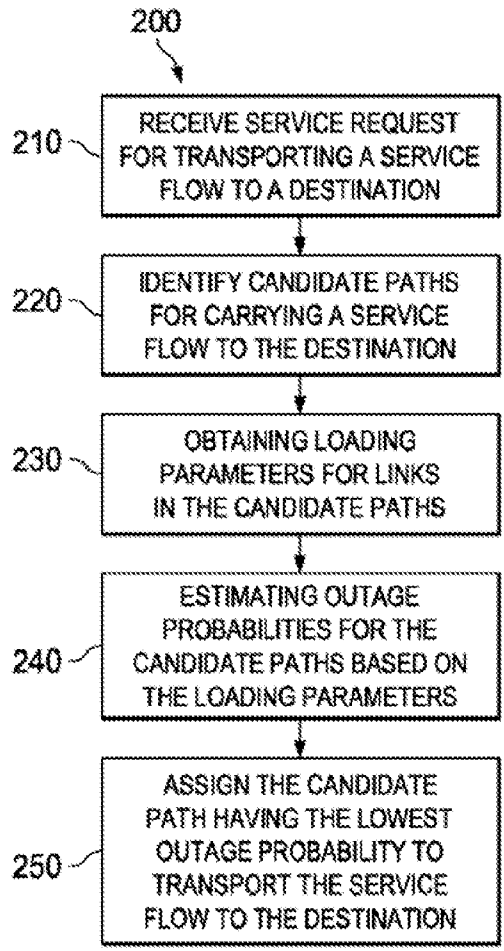


FIG. 2




FIG. 3

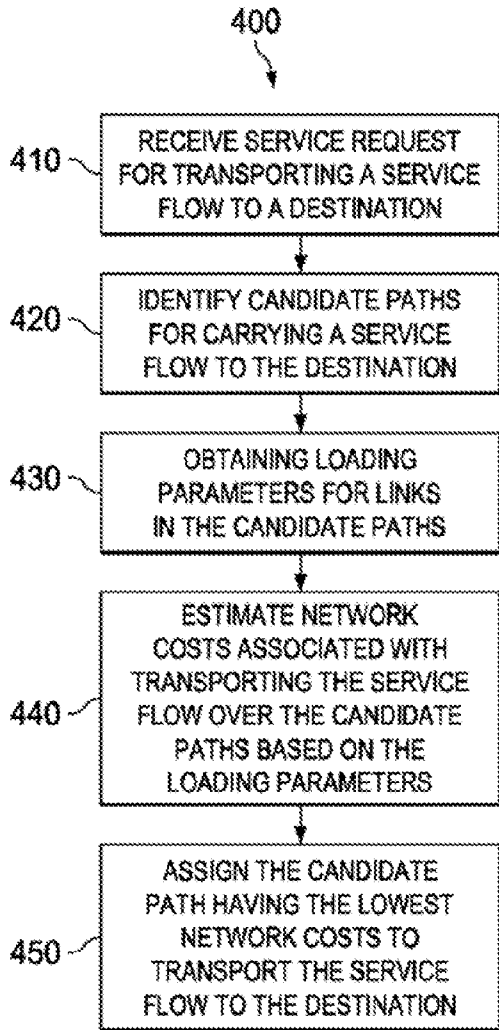


FIG. 4




FIG. 6

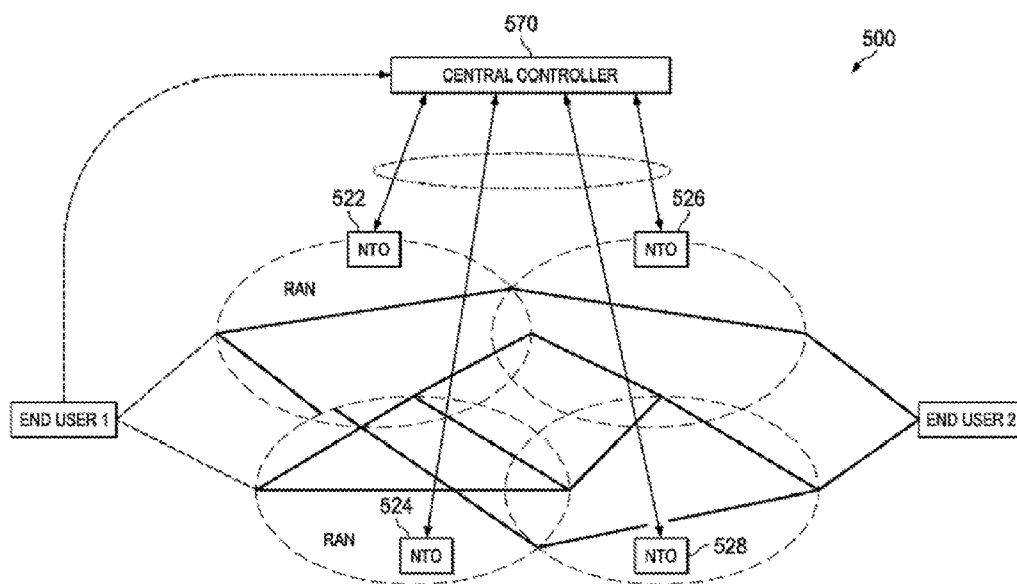


FIG. 5

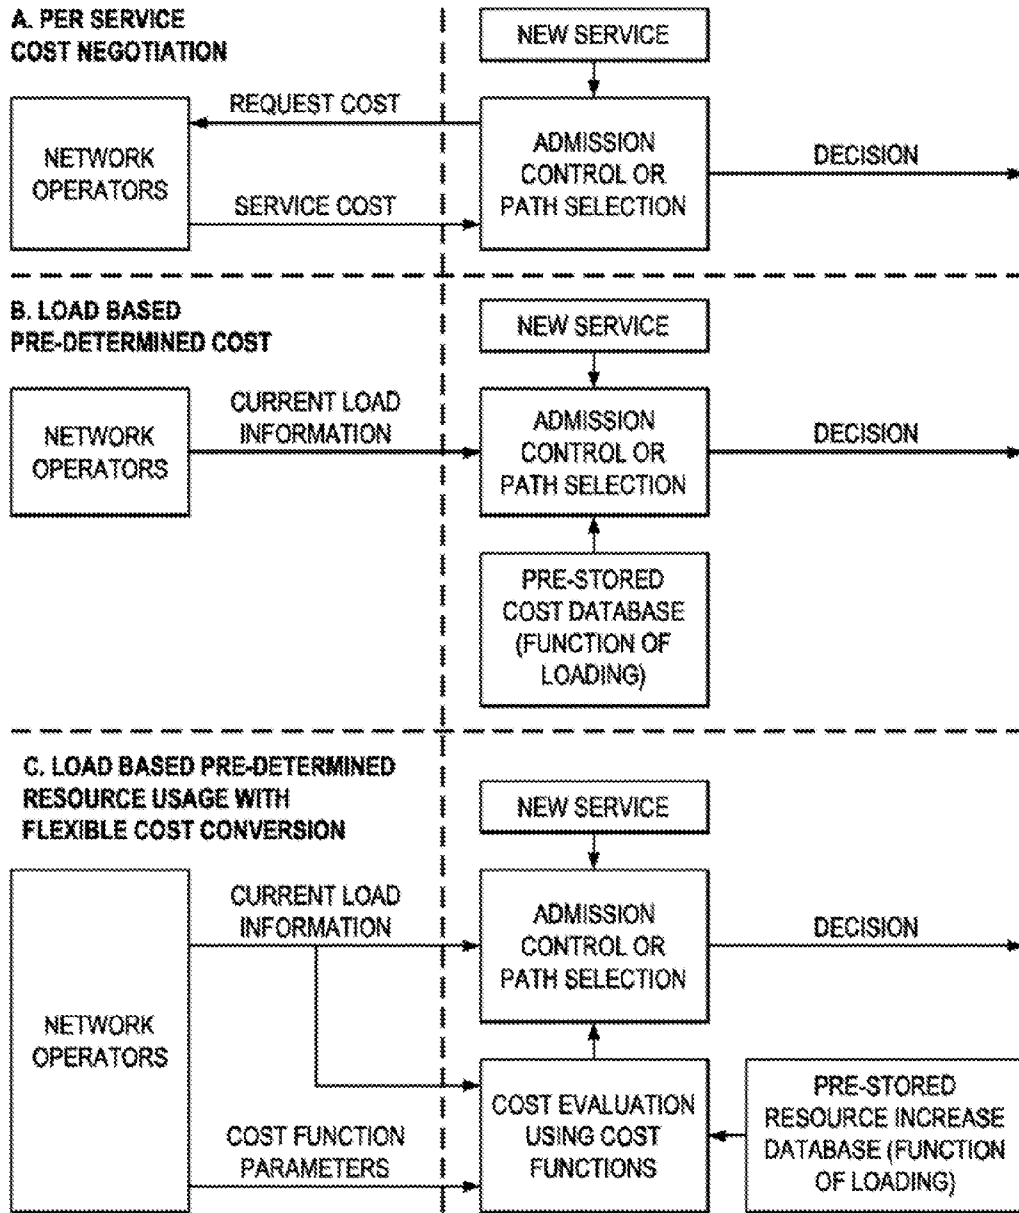


FIG. 7

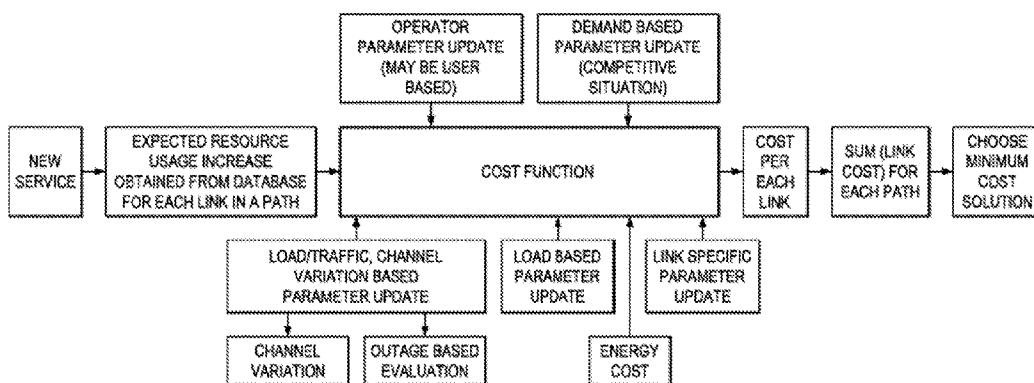


FIG. 8




FIG. 9




FIG. 10




FIG. 11

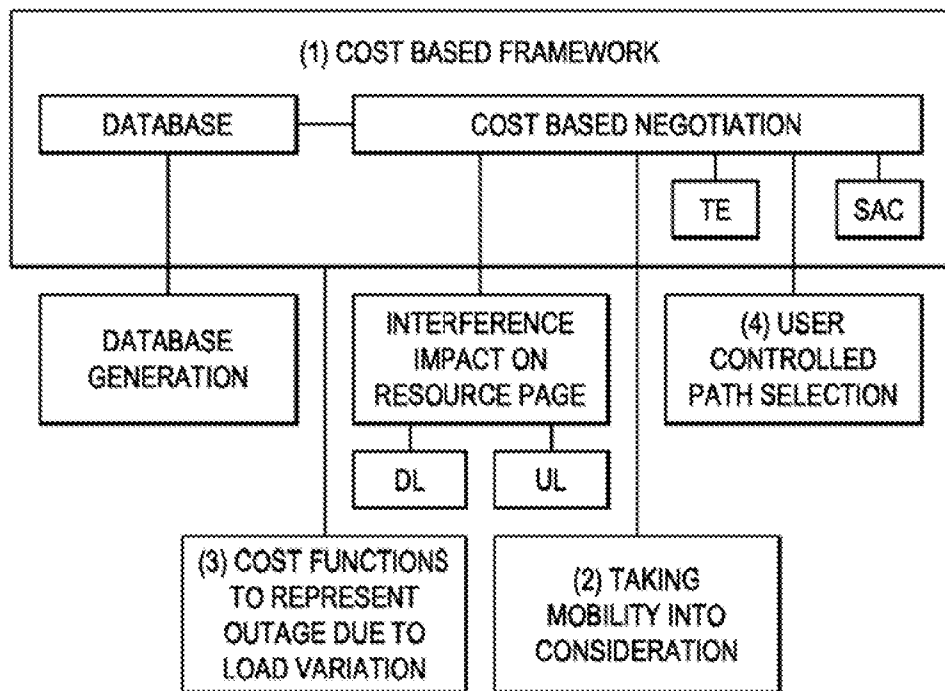


FIG. 12




FIG. 13




FIG. 14

**METHOD AND SYSTEM TO REPRESENT
THE IMPACT OF LOAD VARIATION ON
SERVICE OUTAGE OVER MULTIPLE LINKS**

[0001] This patent application is a continuation of U.S. Non-Provisional patent application Ser. No. 14/203,276 filed on Mar. 10, 2014, which claims priority to U.S. Provisional Application No. 61/778,104, filed on Mar. 12, 2013, both of which are hereby incorporated by reference herein as if reproduced in its entireties.

TECHNICAL FIELD

[0002] The present invention relates to a method and system for network load monitoring, and, in particular embodiments, to techniques for representing the impact of load variation on service outage over multiple links.

BACKGROUND

[0003] Network operators are tasked with equitably distributing finite shared resources (e.g., bandwidth, etc.) amongst multiple users in a manner that satisfies the users' collective quality of service (QoS) requirements. Conventional techniques allocate network resources in an ad hoc manner (e.g., on a case-by-case basis), which satisfies QoS requirements at the expense of overall resource utilization efficiency. For example, in wireless environments, spectrum bandwidth may be allocated to satisfy an individual service request without considering how interference resulting from increased traffic load will reduce spectral efficiency over nearby interferences. Accordingly, mechanisms and techniques for more efficiently allocating resources in a network are needed in order to satisfy ever increasing demands of next generation networks.

SUMMARY OF THE INVENTION

[0004] Technical advantages are generally achieved, by embodiments of this disclosure which describe techniques for representing the impact of load variation on service outage over multiple links.

[0005] In accordance with an embodiment, a method for resource provisioning is provided. In this example, the method includes identifying one or more candidate paths for carrying a service flow to a destination. The one or more candidate paths include at least a first path comprising a first set of links connected in series. The method further includes obtaining load characteristics associated with the first set of links, determining a first cost for carrying the service flow over the first path in accordance with the load characteristics associated with the first set of links, and prompting establishment of the first path when the first cost is less than a threshold. An apparatus for performing this method is also provided.

[0006] In accordance with another embodiment, another method for resource provisioning is provided. In this example, the method includes identifying a first path for carrying a service flow to a destination. The first path includes a first set of links managed by one or more network operators. The method further includes dynamically obtaining cost function parameters for links in the first set of links from the one or more network operators, computing a network cost for transporting the service flow over the first path in accordance with cost function parameters; and prompting transportation of the service flow over the first

path when the network cost is below a threshold. An apparatus for performing this method is also provided.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] For a more complete understanding of the present invention, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawing, in which:

[0008] FIG. 1 illustrates a diagram of an embodiment network;

[0009] FIG. 2 illustrates a flowchart of an embodiment method for resource provisioning;

[0010] FIG. 3 illustrates a diagram of another embodiment network;

[0011] FIG. 4 illustrates a flowchart of an embodiment method for resource provisioning;

[0012] FIG. 5 illustrates a diagram of yet another embodiment network;

[0013] FIG. 6 illustrates a flowchart of an embodiment method for resource provisioning;

[0014] FIG. 7 illustrates a diagram of an embodiment cost based model;

[0015] FIG. 8 illustrates a diagram of an embodiment cost function;

[0016] FIG. 9 illustrates a graph depicting load variation;

[0017] FIG. 10 illustrates a diagram of another embodiment cost function;

[0018] FIG. 11 illustrates a diagram of yet another embodiment cost function;

[0019] FIG. 12 illustrates a diagram of embodiment cost based submissions;

[0020] FIG. 13 illustrates a diagram of an embodiment network device; and

[0021] FIG. 14 illustrates a computing platform that may be used for implementing, for example, the devices and methods described herein, in accordance with an embodiment.

[0022] Corresponding numerals and symbols in the different figures generally refer to corresponding parts unless otherwise indicated. The figures are drawn to clearly illustrate the relevant aspects of the embodiments and are not necessarily drawn to scale.

DETAILED DESCRIPTION OF ILLUSTRATIVE
EMBODIMENTS

[0023] The making and using of the presently disclosed embodiments are discussed in detail below. It should be appreciated, however, that the present invention provides many applicable inventive concepts that can be embodied in a wide variety of specific contexts. The specific embodiments discussed are merely illustrative of specific ways to make and use the invention, and do not limit the scope of the invention.

[0024] Aspects of this disclosure provide a cost-based model for resource allocation that models link/path costs using load characteristics of a network. In one example, a controller identifies one or more candidate paths that are capable of transporting a service flow from a source to a destination. The controller estimates a path cost for transporting the service flow over each candidate path based on dynamically reported load characteristics, e.g., a current load on each link, a load variation on each link, etc. Path cost may represent any quantifiable cost or liability associated

with transporting the service flow over the corresponding path. In one embodiment, the path cost corresponds to a probability that at least one link in the path will experience an outage when transporting the service flow. In another embodiment, the path cost corresponds to price charged by a network operator (NTO) for transporting the traffic flow over the candidate path. In yet another embodiment, the path cost corresponds to a total network cost for transporting the service flow over the candidate path. The total network cost may include various direct cost components and indirect cost components. The direct cost component(s) correspond to costs borne by links in the candidate path, while the indirect cost component(s) correspond to costs borne by other links/interfaces (e.g., links excluded from the candidate path) as a result of transporting the service flow over the candidate path. For example, in the context of wireless networks, the total network cost may include a direct component corresponding to the bandwidth needed to transport the flow over the candidate interface, as well as an indirect cost component corresponding to interference experienced on neighboring radio interfaces as a result of transporting the flow over the candidate interface. Cost functions used for estimating the path costs may be developed by analyzing historical network data (e.g., interference, throughput, loading, etc.) to obtain correlations between costs and loading on the various links in the network. Network costs may also include energy costs for activating otherwise de-activated links along a candidate path. These and others aspects are described in greater detail below.

[0025] Embodiments of this disclosure provide techniques for estimating path costs based on dynamically reported load parameters, e.g., current load level, load variation, etc. These techniques may be used to increase provisioning efficiency during, inter alia, admission and path selection. Path costs may correspond to outage probabilities (or alternatively, to success probabilities) for transporting a service flow over candidate paths. FIG. 1 illustrates a network 100 in which outage probabilities are estimated for a first path 110 and a second path 120 extending from a source to a destination. The path 110 includes links 111, 112, and 113, where link 111 extends from node 101 to node 102 and is associated with a current load (x_1) and a load variation (σx_1), link 112 extends from node 102 to node 103 and is associated with a current load (x_2) and a load variation (σx_2), and link 113 extends from node 103 to node 104 and is associated with a current load (x_3) and a load variation (σx_3). The path 120 includes links 121, 122, and 123, where link 121 extends from node 105 to node 106 and is associated with a current load (y_1) and a load variation (σy_1), link 122 extends from node 106 to node 107 and is associated with a current load (y_2) and a load variation (σy_2), and link 123 extends from node 107 to node 108 and is associated with a current load (y_3) and a load variation (σy_3).

[0026] The current load values (x_n , y_n) correspond to an instantaneous load on the respective link, while the load variations (σx_n , σy_n) correspond to a load variation on the link. Load variations may correspond to a function (e.g., distribution, etc.) representing the average load fluctuation (e.g., median, mean, etc.) on the link over an interval, and may correspond to the relative stability of loading on the link. By way of example, links having high load variations may experience relatively higher load fluctuations than links having low load variations. The load parameters may be reported dynamically to a network operator, where they can

be used to project outage probabilities for the links 111-113 and 121-123 of the paths 110 and 120 (respectively). For example, if it is assumed that a load parameter (L_n) is a random variable having a mean (u_i) equal to the current load value, and a variance (σ^2) equal to the load variation squared, then the outage probability (α_i) for a given link can be expressed as $\alpha_i = P(X_i > T)$, where T is the maximum load on the link. Accordingly, the outage probabilities for each link in the path can be summed (directly or using a linear function) to determine the total cost of the path. As a result, the cost function can be modified to determine a cost increase as an increased outage probability as a result of transporting the traffic flow over the path.

[0027] Notably, if it is assumed that a single link failure will lead to a total path failure, then the probability of success (e.g., the inverse of the outage probability) over multiple serially links can be expressed as follows: $P_s = \pi(\Psi_i) = \pi(1 - P_{oi})$, where P_{oi} is the probability of outage on a given link (i), Ψ_i is the probability of success (i.e., no outage) on the given link (i). Moreover, if the load variation is assumed to be a normal distribution with a mean (L_i) and standard deviation (σ_i), then the probability of outage can be expressed as follows: $P_{oi} = 0.5 - 0.5 \cdot \text{erf}((1 - L_i) / (\sigma_i \cdot \sqrt{2}))$. Additionally, cost can be taken as being inversely proportional to the probability of success, where $\text{cost} = \log(1/P_s) = -\log(P_s) = -\log(\pi(1 - P_{oi})) = \sum -\log((1 - P_{oi}))$. If it is assumed that the link cost function is $C(L_i, \sigma_i) = -\log((1 - P_{oi}))$, then the total cost can be expressed as the sum of the link costs, which may be expressed as $\text{cost} = \sum C(L_i, \sigma_i)$.

[0028] Aspects of this disclosure provide methods for predicting outage probabilities during path selection and/or user admission. FIG. 2 illustrates an embodiment method 200 for predicting outage probability during admission/path-selection, as might be performed by a controller. As shown, the method 200 begins with step 210, where the controller receives a service request requesting resources for transporting a service flow to a destination. Thereafter, the method 200 proceeds to step 220, where the controller identifies candidate paths for carrying the service flow to the destination. The candidate paths may each include a series of links capable of transporting the service flow from a source to a destination, and may include wireline interfaces, wireless interfaces, or combinations thereof. Subsequently, the method 200 proceeds to step 230, where the controller obtains loading parameters for links in the candidate paths. The loading parameters may include any characteristic or value corresponding to loading on a link. In one example, the loading parameters include a current load and a load variation. Next, the method 200 proceeds to step 240, where the controller estimates outage probabilities for the candidate paths based on the loading parameters. An outage probability may correspond to a likelihood that at least one link in a corresponding path will fail (e.g., links' maximum capacity exceeded, etc.) if the traffic flow is transported over the path. Thereafter, the method 200 proceeds to step 250, where the controller assigns the candidate path having the lowest outage probability to transport the service flow. In some embodiments, the service request may be rejected altogether if all outage probabilities exceed a threshold.

[0029] Path costs may also correspond to network costs related to transporting a service flow over candidate paths. The network costs may have direct and indirect cost components. By way of example, network costs in a wireless network may include a direct component corresponding to

resources of the candidate interface needed to transport the flow, as well as an indirect cost component corresponding to interference on neighboring interferences. FIG. 3 illustrates a wireless network 300 in which direct and indirect costs are shown for a candidate link. As shown, the network 300 includes access points 310, 320 and users 301, 302. The user 302 is connected to the access point 320 via an interface 322, and the user 301 has submitted a service request requesting transportation of a service flow. Either of the candidate links 311, 321 are capable of transporting the service flow to the user 302. Network costs for transporting the traffic flow over the candidate interface 311 may include a direct cost component corresponding to network resources needed to transport the service flow over the candidate interface 311, as well as an indirect cost component corresponding to a reduction in capacity on the link 322 due to interference 312 resulting from transportation of the service flow over the candidate interface 311. There may be similar direct and indirect costs for the candidate interface 321, and the candidate interface 311, 321 having the lower total cost (e.g., weighted sum of direct and indirect cost components) may be selected to transport the service flow to the user 302.

[0030] In some embodiments, indirect cost components may include interference cost components that vary based on a loading of the neighboring links, as neighboring links having higher traffic loading may experience comparatively greater interference costs. For example, adding a new service flow to the candidate interface 311 may include an indirect cost component (e.g., corresponding to a reduction in capacity on the link 322) that varies based on a loading of the link 322. Moreover, interference costs may also depend on the location of impacted receiver(s) (e.g., receivers receiving traffic over the neighboring link) and the average amount of traffic communicated to each impacted receiver over the corresponding neighboring link. For example, adding a new service flow to the candidate interface 311 may include an indirect cost component (e.g., a reduction in capacity on the link 322) that depends on a relative location of the user 302 as well as an amount of traffic communicated to the user 302 via the link 322. These cost factors can be integrated into the link cost formation in various ways. In one example, the controller assigned to the neighbor links can dynamically update the interference cost component values as the receiver location (e.g., loading distribution) and/or the traffic loading associated with the neighboring links (e.g., loading) varies. In another example, interference cost components associated with a new flow can be modeled as a function of load and/or load-distribution on the neighbor links, e.g. the cost function of a link is evaluated as a function of its own loading/load-distribution and the loading/load-distributions of neighboring links.

[0031] The following example demonstrates how cost components can be modeled for two neighboring wireless links. Let L1 and L2 are the loading of each link and cost functions (as a function of individual node load or 'self load') are $f_1(L1)$ and $f_2(L2)$ respectively. When L1 is increased by $\Delta L1$, the interference to the second link increases which results in increased resource usage in the second link. This means the loading of the second link is increased by $\Delta L2$ which in turn increases the cost to the second link. It can measure this increased cost by its cost function, $f_2(\cdot)$. This can then be informed to the first link to adjust its final cost function value at $L1+\Delta L1$, $F1(L1+\Delta L1, L2)$ ='cost due to self load'+ 'cost to neighbour'. This may be

repeated for various values of L1 and L2 and $\Delta L1$ values. Once the function $F1(\cdot)$ is obtained, the cost of a link can be obtained taking the impact to the neighbor as a function of self load and the neighbor load, and the neighbor's load can be updated dynamically. The above example considers the neighboring link's load as a single entity. However, this can be obtained for different neighbor load distributions if multiple receivers are involved and the modifications could be done in a similar manner by repeating the above described steps for different neighbor load distributions. In one example, path cost may be computed in accordance with the following formula: $cost = \sum_{i=1}^n C(L_i, \sigma_i, L_{i1}, \sigma_{i1}, L_{i2}, \sigma_{i2}, \dots, L_{im}, \sigma_{im})$, where $C(L_i, \sigma_i)$ is the cost function for the path, L_i is a loading parameter on an i^{th} link in the path, σ_i is the load variation on the i^{th} link, L_{ij} is a loading parameter of the j^{th} neighbor of the i^{th} link, σ_{ij} is the load variation of the j^{th} neighbor of the i^{th} link, n is the number of links in the path, and m is the number of neighbors for the i^{th} link. In some embodiments, a loading parameter corresponds to a mean or average load. In other embodiments, loading parameters correspond to different loading characteristics, e.g., instantaneous load, median load, etc.

[0032] Aspects of this disclosure provide methods for predicting network costs during path selection and/or user admission. FIG. 4 illustrates an embodiment method 400 for predicting network costs during admission/path-selection, as might be performed by a controller. As shown, the method 400 begins with step 410, where the controller receives a service request requesting resources for transporting a service flow to a destination. Thereafter, the method 400 proceeds to step 420, where the controller identifies candidate paths for carrying the service flow to the destination. Subsequently, the method 400 proceeds to step 430, where the controller obtains loading parameters for links in the candidate paths. Next, the method 400 proceeds to step 440, where the controller estimates network costs (e.g., direct, indirect, or otherwise) associated with transporting the service flow over the each of the candidate paths based on the loading parameters. In some embodiments, estimating the network costs may utilize cost functions stored in a resource cost database. The cost functions may be formed by analyzing historical network information to identify correlations between costs (e.g., interference, reductions in spectral efficiency, etc.) and network loading (e.g., throughput, etc.). Thereafter, the method 400 proceeds to step 450, where the controller assigns the candidate path having the lowest outage probability to transport the service flow. In some embodiments, the service request may be rejected altogether if all outage probabilities exceed a threshold.

[0033] Network costs can also correspond to a price paid to use or reserve a network resource. More specifically, next generation networks may provision network resources using a marketplace architecture in which virtual or physical resources are offered for sale at prices that vary with supply and demand. For example, pricing for wireless spectrum bandwidth (virtual or otherwise) may be adjusted based on resource availability (or on average spectral-efficiency-per-resource-unit). Accordingly, the price for each additional resource unit may increase as network loading increases, e.g., as resource availability decreases. In some embodiments, resource pricing may be negotiated between the user and the network operator, or by an intermediary, e.g., a

telephone network service provider, etc. In other embodiments, resource pricing may be set according to a function/formula.

[0034] FIG. 5 illustrates an embodiment network 500 for provision network resources using a marketplace architecture. As shown, the network 500 includes a central controller 570 configured to purchase resources from network operators (NTOs) 522-528. The NTOs 522-528 may operate different types of networks, or different domains within the same network. As an example, the NTOs 522, 524 may operate radio access networks (RANs), while the NTOs 526, 528 may operate wireline networks, e.g., the NTOs 526, 528 may be internet service providers (ISPs). The controller 570 may negotiate resource pricing with the NTOs 522-528 on part of users/subscribers. Alternatively, the pricing may fluctuate based on actual or estimated resource availability, e.g., price increases as resources become more scarce. Resource pricing may be calculated/estimated using, inter alia, current loading parameters of the network 500 in accordance with a cost-function. The cost function can be developed by analyzing historical network information to identify correlations between resource availability (e.g., spectral efficiency, etc.) and loading in the network 500. Techniques for developing cost functions and resource cost databases are described in U.S. patent application Ser. No. 14/107,946 entitled "Service Provisioning Using Abstracted Network Resource Requirements" filed on Dec. 16, 2013 and U.S. patent application Ser. No. 14/106,531 entitled "Methods and Systems for Admission Control and Resource Availability Prediction Considering User Equipment (UE) Mobility" filed on Dec. 13, 2013, both of which are incorporated herein by reference as if reproduced in their entireties.

[0035] Aspects of this disclosure provide methods for predicting network costs during path selection and/or user admission. FIG. 6 illustrates an embodiment method 600 for estimating resource pricing during admission/path-selection, as might be performed by a controller. As shown, the method 600 begins with step 610, where the controller receives a service request requesting resources for transporting a service flow to a destination. Thereafter, the method 600 proceeds to step 620, where the controller identifies candidate paths for carrying the service flow to the destination. Subsequently, the method 600 proceeds to step 630, where the controller obtains loading parameters for links in the candidate paths. Next, the method 600 proceeds to step 640, where the controller estimates or determines resource pricing for transporting the service flow over each of the candidate paths based on the loading parameters. In some embodiments, price components may be estimated for each link in the path, and then summed to determine the resource pricing for the path. In some embodiments, the controller may estimate the pricing based on loading information provided by the NTOs. In other embodiments, the controller may determine the pricing by negotiating with the NTOs. Thereafter, the method 600 proceeds to step 650, where the controller assigns the candidate path having the lowest price to transport the service flow. In some embodiments, the service request may be rejected altogether if all prices exceed a threshold.

[0036] All else being equal, links having higher load variations typically exhibit a higher probability of outage than links having lower load variations. When there are alternative links/paths to be chosen, the cost of adding a user

should be increased when the mean load is higher to discourage users/service providers from using highly loaded paths during load balancing. During path selection and/or admission control, the overall cost of multiple paths is searched and the best path is selected. The cost of each path may be a function of the cost of each link in that path. An embodiment represents the cost of a link such that overall cost is additive, but still allows the best path to be chosen from the viewpoint of the network.

[0037] An embodiment provides a method to represent the impact of load variation on service outage when admitting or routing a flow through a path consisting of multiple links. When a data flow is to be added to a link, if the load variation is high, the probability of outage increases. Thus, if the flow is added considering only the increase of mean load, the chance of outage would increase and an embodiment provides a method to represent this outage.

[0038] An embodiment provides a method for a central entity to perform load balancing across a network. The cost of adding a service flow is modeled as a function of current load and load variation using a convex, increasing function, the parameters of which can be changed based on the load variation and other operator needs such as competitiveness or to draw a higher income. When there are alternative links/paths to be chosen, the cost of adding a user is increased when the mean load is higher to discourage users/service providers from using highly loaded paths to do load balancing.

[0039] Embodiment cost functions encourage complete shut-off of a radio node if the user flows can be handled along different paths. An embodiment supports on-demand cost estimation as a function of demand/availability to enable pricing to be adjusted dynamically. Embodiments also may be used for user controlled path selection based on cost, as a central entity to perform admission and routing, load balancing and optimization of the network, and as a network congestion solution to avoid network congestion if demand-based charging is imposed.

[0040] Outage of a link can be modeled as a function of load parameters, e.g., load, loading variation. In some embodiments, loading variation is computed using short-term statistics. In another embodiment, loading variation is computed using long-term statistics. The load itself can be a function of the number of resources used or the power each of these resources used. It also can be a function of the characteristics of traffic flows that go through that link, for example, a total utility. The load also varies with the channel conditions of different traffic flows. In the simplest example, the load could be modeled as a mean and a standard variation, thereby allowing the probability of outage to be modeled as a function of loading on the link.

[0041] The probability of success of the path can be computed by multiplying the individual probabilities of success values for each link, which can be achieved using a logarithm or logarithmic function. In this manner, the probability of success can be used by a central entity for various provisioning activities, such as load balancing across a network, determining whether a certain node can be switched off, supporting on-demand pricing for users, to determine the impact a given service will have on network performance, etc.

[0042] An embodiment method represents the current loading of a link using a cost function that provides a higher system cost for accommodating a flow at higher loading

compared to lower loading taking load variation into account to reduce outage. An embodiment method allows the network operator to increase or decrease charging dynamically (e.g., to address competitive situations), which can be used by the customers to select networks and paths. An embodiment enables a method for different service providers to use the system independently while indirectly balancing the load and minimizing link outage, and allows for automated congestion control if the cost based charging or admission is implemented. An embodiment method allows an operator to charge customers taking resource cost into account.

[0043] Network controllers may maintain a database to keep load-cost dependencies. In cost based control schemes, the cost of each link is evaluated by mapping an expected resource usage to a cost, which can be a factor of many other parameters and may change dynamically.

[0044] FIG. 7 illustrates a cost based model overview. The B and C cost based schemes use a cost function to change the resource usage to network cost. FIG. 8 illustrates a cost function usage overview. Objectives of a cost function may include reducing network congestion call blocking in a manner that permits distributed provisioning (e.g., admission control and path selection). This can be achieved by conservatively admitting or making routing changes when the load of a link changes, e.g., the cost of a link increases with the load. The cost function may also reflect that higher demand increases resource costs on a link. When the variation of load in a link is large (traffic variability) and/or the variation of link capacity (e.g. for a wireless link, the link capacity changes with time), there is a higher probability of outage. At that time the admission control or path selection can be done more conservatively and the cost function can reflect that. In addition, the network operator may be able to dynamically change the cost function parameters for competitive purposes and as per the dynamic per user based requirements. The cost function can also take into account the energy cost of maintaining a link. For example, when a node or transmission link is completely shut down, there is energy saving and a cost proportional to energy usage can be included. The cost function can also be used for dynamic cost based charging. In this case, the cost function could reflect the actual cost of a link to the user so that the user can make a decision to use or not to use a given link depending on the user's needs.

[0045] If loads are added directly (or a linear function) to determine the total cost of the path, the cost of the almost loaded link is compensated by the lightly loaded link. An embodiment function is used to map all the link loads and load variations to a single cost value. The qualities of the function include (1) when load of one link increases, the total load should increase, (2) when load of even one link exceeds the capacity the total cost should exceed the cost threshold, and (3) the increase in cost should be higher at a higher load than a smaller load (if the variance remains the same). The $\text{Cost_per_link} = f(\text{load, variation})$. The Cost per link = Load (or a linear function). The Cost of path = sum (Cost per link). The cost of the almost loaded link is compensated by the lightly loaded link. Therefore, the function can be modified by a cost function such that the cost increases with the likelihood of outage in the link. Then the cost of using the same amount of resources should be higher at higher loads than the smaller loads.

[0046] FIG. 9 illustrates a load variation plot. If it is assumed that a current load variation behaves as a normal distribution, with a mean L_i and standard deviation σ_i , then the probability of outage can be expressed as follows: $\text{Poi} = 0.5 - 0.5 * \text{erf}(1 - L_i) / (\sigma_i * \sqrt{2})$. Cost can be taken as inversely proportional to the probability of success, and let $\text{cost} = \log(1/P_s)$. This cost function is shown in FIG. 10, wherein the curves progress downward from $\sigma = 0.9$ as the uppermost curve, down to $\sigma = 0.1$ as the lowermost curve. FIG. 11 illustrates an embodiment cost function, and shows how it can be used for admission control and routing. The following parameters were used for the cost function: function $\text{cost} = f_cost_func(\text{link, load})$; $\text{power_on_cost} = 100$; $\text{max_load} = 1$; $k = 400$; % value operator uses to compete with other operators in the area; $n = 5$; if $\text{load} \leq 1$, then $\text{cost} = 500 * \text{load}^n$, end; if $\text{load} > 1$, then $\text{cost} = \text{cost} + 1500$, end; if $\text{load} > 0$, then $\text{cost} = \text{cost} + \text{power_on_cost}$, end; end.

[0047] In this scheme, n is changed according to the variation of the load. Operator's prize increase or demand increase can be modeled by a simple proportional parameter. This may also be changed only in a high loading area. FIG. 12 illustrates different cost based submissions.

[0048] FIG. 13 illustrates a block diagram of an embodiment of a network device 1300, which may be equivalent to one or more devices (e.g., controllers, etc.) discussed above. The network device 1300 may include a processor 1304, a memory 1306, a cellular interface 1310, a supplemental interface 1312, and a backhaul interface 1314, which may (or may not) be arranged as shown in FIG. 13. The processor 1304 may be any component capable of performing computations and/or other processing related tasks, and the memory 1306 may be any component capable of storing programming and/or instructions for the processor 1304. The cellular interface 1310 may be any component or collection of components that allows the network device 1300 to communicate using a cellular signal, and may be used to receive and/or transmit information over a cellular connection of a cellular network. The supplemental interface 1312 may be any component or collection of components that allows the network device 1300 to communicate data or control information via a supplemental protocol. For instance, the supplemental interface 1312 may be a non-cellular wireless interface for communicating in accordance with a Wireless-Fidelity (Wi-Fi) or Bluetooth protocol. Alternatively, the supplemental interface 1312 may be a wireline interface. The backhaul interface 1314 may be optionally included in the network device 1300, and may comprise any component or collection of components that allows the network device 1300 to communicate with another device via a backhaul network.

[0049] FIG. 14 is a block diagram of a processing system that may be used for implementing the devices and methods disclosed herein. Specific devices may utilize all of the components shown, or only a subset of the components, and levels of integration may vary from device to device. Furthermore, a device may contain multiple instances of a component, such as multiple processing units, processors, memories, transmitters, receivers, etc. The processing system may comprise a processing unit equipped with one or more input/output devices, such as a speaker, microphone, mouse, touchscreen, keypad, keyboard, printer, display, and the like. The processing unit may include a central processing unit (CPU), memory, a mass storage device, a video adapter, and an I/O interface connected to a bus.

[0050] The bus may be one or more of any type of several bus architectures including a memory bus or memory controller, a peripheral bus, video bus, or the like. The CPU may comprise any type of electronic data processor. The memory may comprise any type of system memory such as static random access memory (SRAM), dynamic random access memory (DRAM), synchronous DRAM (SDRAM), read-only memory (ROM), a combination thereof, or the like. In an embodiment, the memory may include ROM for use at boot-up, and DRAM for program and data storage for use while executing programs.

[0051] The mass storage device may comprise any type of storage device configured to store data, programs, and other information and to make the data, programs, and other information accessible via the bus. The mass storage device may comprise, for example, one or more of a solid state drive, hard disk drive, a magnetic disk drive, an optical disk drive, or the like.

[0052] The video adapter and the I/O interface provide interfaces to couple external input and output devices to the processing unit. As illustrated, examples of input and output devices include the display coupled to the video adapter and the mouse/keyboard/printer coupled to the I/O interface. Other devices may be coupled to the processing unit, and additional or fewer interface cards may be utilized. For example, a serial interface such as Universal Serial Bus (USB) (not shown) may be used to provide an interface for a printer.

[0053] The processing unit also includes one or more network interfaces, which may comprise wired links, such as an Ethernet cable or the like, and/or wireless links to access nodes or different networks. The network interface allows the processing unit to communicate with remote units via the networks. For example, the network interface may provide wireless communication via one or more transmitters/transmit antennas and one or more receivers/receive antennas. In an embodiment, the processing unit is coupled to a local-area network or a wide-area network for data processing and communications with remote devices, such as other processing units, the Internet, remote storage facilities, or the like.

[0054] While this invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications and combinations of the illustrative embodiments, as well as other embodiments of the invention, will be apparent to persons skilled in the art upon reference to the description. It is therefore intended that the appended claims encompass any such modifications or embodiments.

1. A method for resource provisioning, the method comprising:

identifying, by a controller, one or more candidate paths for carrying a service flow to a destination, the one or more candidate paths including at least a first path extending from a core network to a user equipment (UE) accessing a radio access network (RAN), the first path including a first set of links connected in series, wherein at least one link in the first set of links is a wireless link in the RAN;

obtaining load parameters associated with one or more links that either are included in the first set of links or interfere with one or more with wireless links in the first set of links;

determining a network cost for carrying the service flow over the first path in accordance with the load param-

eters and at least one cost function associated at least one link in with the first set of links; and

prompting establishment of the first path when the network cost for carrying the service flow over the first path is less than a threshold.

2. The method of claim 1, wherein the network cost includes the output of the cost function, and wherein the load parameters are input parameters to the cost function.

3. The method of claim 2, wherein the location of the UE is an input parameter to the cost function.

4. The method of claim 2, wherein an energy cost for maintaining a wireless link in the first path is an input parameter to the cost function.

5. The method of claim 2, wherein a charge for transporting the service flow over one or more links in the first path is an input parameter to the at least one cost function.

6. The method of claim 2, wherein the cost function depends a link type associated with one or more links in the first path.

7. The method of claim 1, wherein the at least one cost function and the load parameters are received by the controller from one or more network devices in the RAN.

8. The method of claim 1, wherein the load parameters include a load or a load variation on at least one link that either is in the first set of links or interferes with one or more wireless links in the first set of links.

9. The method of claim 1, wherein the network cost includes an indirect cost component that corresponds to an increase in interference on at least one neighboring wireless link that would result from communicating the service flow over one or more wireless links in the first path.

10. The method of claim 9, wherein the indirect cost component is determined based on one or more current load parameters associated with a wireless link, the one or more current load parameters including at least one of a current load and a current load variation on the wireless link during a current time period.

11. The method of claim 10, wherein the indirect cost component is determined based on both the current load and the current load variation on the wireless link.

12. The method of claim 9, wherein the indirect cost component is determined based on a relationship between one or more historical load parameters and one or more associated costs on a wireless link, the one or more historical load parameters including at least one of a historical load and a historical load variation on the wireless link during a previous time period.

13. The method of claim 12, wherein the indirect cost component is determined based on both the historical load and the historical load variation on the wireless link.

14. The method of claim 9, wherein the indirect cost component is determined based on a location of a wireless receiver on the wireless link.

15. The method of claim 1, wherein the network cost includes a direct cost component that corresponds to an increase in interference on one or more wireless links in the first path from neighboring wireless links that would result from communicating the service flow over the first path.

16. The method of claim 15, wherein the direct cost component is determined based on one or more current load parameters associated with a wireless link, the one or more current load parameters including at least one of a current load and a current load variation on the wireless link during a current time period.

17. The method of claim 16, wherein the direct cost component is determined based on both the current load and the current load variation on the wireless link.

18. The method of claim 16, wherein the direct cost component is determined based on a relationship between one or more historical load parameters and one or more associated costs on a wireless link, the one or more historical load parameters including at least one of a historical load and a historical load variation on the wireless link during a previous time period.

19. The method of claim 18, wherein the direct cost component is determined based on both the historical load and the historical load variation on the wireless link.

20. The method of claim 19, wherein the historical load and the historical load variation are received by the controller from one or more network devices in the RAN.

21. The method of claim 15, wherein the direct cost component is determined based on a location of the UE.

22. The method of claim 1, wherein the network cost corresponds to a total network cost for transporting the service flow over the first path, the total network cost including a network cost component for transporting the service flow over a first portion of the first path extending through the RAN and a second cost component for transporting the service flow over a second portion of the first path extending through the core network.

23. The method of claim 1, wherein the first set of links includes a first radio interface between a first transmitter and the destination.

24. The method of claim 23, wherein the network cost includes a cost component corresponding to interference experienced on a second radio interface as a result of communicating the service flow over the first radio interface.

25. The method of claim 1, wherein determining the network cost comprises:

computing the network cost in accordance with the following formula: $cost = \sum_{i=1}^n C(L_i, \sigma_i, L_{i1}, \sigma_{i1}, L_{i2}, \sigma_{i2}, \dots, L_{im}, \sigma_{im})$, where $C(L_i, \sigma_i)$ is the cost function for the first path, L_i is a loading parameter on an i^{th} link in the first set of links, σ_i is the load variation on an i^{th} link in the first set of links, L_{ij} is a loading parameter of the j^{th} neighbor of the i^{th} link, σ_{ij} is the load variation of the j^{th} neighbor of the i^{th} link, n is the number of links in the first set of links, m is the number of neighbors for the i^{th} link.

26. The method of claim 25, wherein L_i is a mean load on the i^{th} link in the first set of links, and L_{ij} is a mean load on the j^{th} neighbor of the i^{th} link.

27. A controller comprising:
a processor; and

a computer readable storage medium storing programming for execution by the processor, the programming including instructions to:

identify one or more candidate paths for carrying a service flow to a destination, the one or more candidate paths including at least a first path extending from a core network to a user equipment (UE) accessing a radio access network (RAN), the first path including a first set of links connected in series, wherein at least one link in the first set of links is a wireless link in the RAN; obtain load parameters associated with one or more links that either are included in the first set of links or interfere with one or more with wireless links in the first set of links;

determine a network cost for carrying the service flow over the first path in accordance with the load parameters and at least one cost function associated at least one link in with the first set of links; and prompt establishment of the first path when the network cost for carrying the service flow over the first path is less than a threshold.

28. A computer program product comprising a non-transitory computer readable storage medium storing programming, the programming including instructions to:

identify one or more candidate paths for carrying a service flow to a destination, the one or more candidate paths including at least a first path extending from a core network to a user equipment (UE) accessing a radio access network (RAN), the first path including a first set of links connected in series, wherein at least one link in the first set of links is a wireless link in the RAN; obtain load parameters associated with one or more links that either are included in the first set of links or interfere with one or more with wireless links in the first set of links;

determine a network cost for carrying the service flow over the first path in accordance with the load parameters and at least one cost function associated at least one link in with the first set of links; and

prompt establishment of the first path when the network cost for carrying the service flow over the first path is less than a threshold.

* * * * *