Network Provisioning as a Game Against Nature

A Multicommodity Network Flow Model under Uncertain Demands

V. Marbukh
National Institute of Standards and Technology
100 Bureau Drive, Stop 8920
Gaithersburg, MD 20899-8920
E-mail: marbukh@nist.gov

Abstract—**Traditional approaches to network provisioning assume availability of the reliable estimates for the expected demands. This assumption, however, oversimplifies many practical situations when some incomplete information on the expected demands is available, and proper utilization of this information may improve the network performance. In a case of traffic engineering the uncertainty in the expected demands may be a result of sudden changes in the demand pattern, when significant statistical uncertainty in determining the varying demand pattern and possible undesirable transient effects make continuous adjustment of the routing algorithm to varying demands difficult. A long-term network provisioning, e.g., capacity planning, is a subject to uncertainties in the overall economic conditions. Despite the network may be capable of controlling demands through pricing, the overall economic conditions affect the price-demand curve. As the recent sharp downturn in the demand for communication bandwidth demonstrated, making long-term network planning decisions without assessing reliability of the underlying assumptions on the expected demands may lead to disastrous results. Assuming that the expected demand is an unknown mixture of some known scenarios, i.e., demand matrices, this paper proposes a framework for robust network provisioning by guarding against the worst case scenario with respect to the future demands. This framework can be interpreted as a game between the network, e.g., service provider, and nature. The service provider makes the network provisioning decisions in an attempt to minimize losses due to the uncertain future demands, while the nature selects a feasible demand matrix. Solution to this game balances risks of over and under provisioning of the network.**

Keywords-robust network provisioning; uncertain demands; game against nature.

I. INTRODUCTION

Traditional approaches to short term network provisioning, e.g., flow assignments, as well as long term network provisioning, e.g., capacity and topology planning, assume availability of the reliable estimates for the expected demands [1]-[4]. In economic models the network, i.e., service provider, can influence the demand by controlling the price of the service [5]-[7]. The purpose of the network provisioning is to optimize certain performance criterion, typically to maximize the revenue generated by the network. Network provisioning decisions may be very sensitive to the estimates of the future demands in models with fixed prices, or, future price-demand curves in models with varying prices. However, in many practical situations reliability of the corresponding point estimates may be insufficient for planning purposes.

Initial results produced by the recent surge in activities related to measurements of the Internet traffic suggest high degree of volatility and uncertainty in the traffic pattern [8]-[9]. This conclusion raises question of reliability of the comparatively short-term network provisioning, e.g., traffic engineering, decisions based on the corresponding point estimates, since significant statistical uncertainty in determining the traffic pattern and possible undesirable transient effects make continuous adjustment of the routing algorithm difficult. A long-term network provisioning, e.g., capacity and topology planning, is a subject to uncertainties in the overall economic conditions. Despite the network may be capable of controlling demands through pricing, the overall economic conditions affect the price-demand curve. As the recent sharp downturn in the demand for communication bandwidth demonstrated, making long-term network planning decisions based on the point estimates for the long-term demand forecast without assessing reliability of these estimates may lead to disastrous results [10].

Several attempts have been made to evaluate and incorporate the concept of robustness into the network provisioning process. A measure of network lifetime intended to quantify the growth and shifts in the load (traffic demand perturbations) that a network can sustain has been introduced in [11]. A Bayesian framework, based on the assumption that uncertainties can be described by some fixed probability distribution, has been applied in [12] to *QoS* routing in presence of unstable routes, and in [13] to traffic engineering under varying by "business hour" external demands. Robust network management schemes [14]-[16], attempting to minimize the worst-case scenario loss in performance due to various uncertainties, have been motivated by decision theory [17] and complexity theory of algorithms [18]. Possible applications of this approach to *QoS* routing and traffic engineering under assumption that the available to the protocol network state information can be more reliably quantified in terms of the "confidence" intervals rather than point estimates have been discussed in [19]-[21].

This paper attempts to cast the problem of robust network provisioning as a game of the network, e.g., service provider, against the nature. The service provider makes the network provisioning decisions in an attempt to minimize losses due to

This work was supported in part by DARPA NMS Program

the uncertain demands, while the nature selects a feasible scenario with respect to the future demands. In a case of traffic engineering over or under provisioning causes mismatch between the loads and capacities in the different parts of the network. In a case of long-term network planning over provisioning results in low or negative return on capital investment, while under provisioning results in inability to meet demand and potential lost in the market share. Once the relevant potential losses are quantified (not a simple task!), solution to this game balances risks of over and under provisioning of the network.

The paper is organized as follows. Sections II and III describe some possible multicommodity network flow optimization model under fixed and, respectively, priced demands, assuming complete knowledge of the relevant information on the future demands. Section IV describes models for uncertain demands and casts network provisioning as a game against nature. Finally, conclusion summarizes results and outlines directions for future research.

II. Network Provisioning under Fixed Demands

A. Multicommodity Network Flow Model

Consider a network with set of nodes $n \in N$ and set of directed links $l \in L$ with capacities c_l. The vector of traffic flows carried on all possible routes r in the network $x = (x_r)$ satisfy the following capacity constraints:

$$\lambda_l \overset{def}{=} \sum_{r:l\in r} x_r \le c_l, \ \forall l, \tag{1}$$

$$y_{ij} \overset{def}{=} \sum_{r\in R_{ij}} x_r \le \mu_{ij}, \ \forall (i,j), \tag{2}$$

$$x_r \ge 0, \ \forall r \in R_{ij}, \ \forall (i,j), \tag{3}$$

where the total flow carried on a link l is λ_l, the set of feasible routes with origin-destination (i,j) is R_{ij}, the matrix of accepted demands for all origin-destination pairs (i,j) is $y = (y_{ij})$, and the matrix of external demands for all origin-destination pairs (i,j) is $\mu = (\mu_{ij})$. Constrains (1) state that the total flow carried on each link l cannot exceed the capacity of this link, and constraints (2) state that the accepted load for each origin-destination (i,j) cannot exceed the corresponding offered load. Note that formulation (1)-(3) allows for rejection of some portion of offered traffic.

Sometimes it is convenient to separate admission and routing strategies. The admission strategy is characterized by the fraction of the offered load with origin-destination (i,j) to be rejected:

$$L_{ij} = 1 - y_{ij} / \mu_{ij} \tag{4}$$

Routing strategy is characterized by the vector $\xi = (\xi_r)$ where the fraction of the accepted load y_{ij} with origin-destination (i,j) carried on a feasible route $r \in R_{ij}$ is

$$\xi_r = x_r / \mu_{ij} \tag{5}$$

Capacity constraints (1)-(3) can be rewritten in terms of the losses (4) and fractions (5) as follows:

$$\lambda_l \overset{def}{=} \sum_{(i,j)} (1 - L_{ij}) \mu_{ij} \sum_{r:l\in r\subseteq R_{ij}} \xi_r \le c_l, \ \forall l, \tag{6}$$

$$0 \le L_{ij} \le 1, \ \forall (i,j), \tag{7}$$

$$\sum_{r\in R_{ij}} \xi_r = 1, \ \xi_r \ge 0, \ \forall r \in R_{ij}, \ \forall (i,j) \tag{8}$$

The purpose of the network provisioning is to optimize certain criterion, typically to maximize the net revenue generated by the network subject to the capacity and Quality of Service (*QoS*) constraints, given demands and prices. In this paper, due to limited space, we consider the following very simple model for the *QoS* constraint:

$$T \le T^* \tag{9}$$

where

$$T = \sum_l T_l \tag{10}$$

and the penalty associated with deterioration of the quality of service for traffic of rate λ_l carried on link l of capacity c_l is characterized by some increasing and convex in $\lambda_l \in [0, c_l)$, function $T_l(\lambda_l, c_l)$, $\forall l$. We also assume that $\lim T_l = \infty$ as $\lambda_l \to c_l - 0$. Under these assumptions *QoS* constraints (9) imply capacity constraints (1), (6). The following penalty function, inspired by the average delay in a $M/M/1$ queueing system, is often used for a packet network: $T_l = \lambda_l/(c_l - \lambda_l) + h_l\lambda_l$, where the processing and propagation delay associated with link l is h_l [22]-[23]. For this specific penalty function the total penalty (10) after normalization approximates the overall average delay in the network under hypothesis of independence [22]. Note that much more sophisticated models for *QoS* requirements, based on the utility functions of the sources are possible [5]-[7].

Given price p_{ij} charged by the network for a unit of accepted demand with origin-destination (i,j), the total revenue generated by the network is $P = \sum_{(i,j)} (1 - L_{ij}) \mu_{ij} p_{ij}$. Assuming that the price of the link l capacity c_l is $q_l = q_l(c_l)$, the total price of the bandwidth is $Q = \sum_l c_l q_l(c_l)$, and the net revenue generated by the network is $W = P - Q$.

B. Network Provisioning

If the matrix of expected demands $\mu = (\mu_{ij})$ is fixed and known, the optimal admission, routing and capacity provisioning policies can be determined by solution of the following optimization problems.

Flow Assignment (FA) problem. Given link capacities $c = (c_l)$, find optimal flow assignment $\xi = \xi^*$ by maximizing revenue P subject to constraints (6)-(10).

Capacity Assignment (CA) problem. Given route assignments ξ, find optimal link capacities $c = c^*$ and rejection probabilities $L = L^*$ by maximizing net revenue W subject to constraints (6)-(10).

Capacity and Flow Assignment (CFA) problem. Given network topology, find optimal link capacities $c = c^*$, rejection probabilities $L = L^*$ and route assignments $\xi = \xi^*$ by maximizing net revenue W subject to constraints (6)-(10).

These optimization problems are reformulations of the well-known optimization problems for network provisioning [22]-[23]. The purpose of this reformulation is to accommodate demand pricing and uncertainty issues into the network provisioning process. These models can be generalized to cover cases of multiple *QoS* classes, Virtual Private Network (*VPN*) provisioning, stochastic nature of the offered load, etc. The optimization problem for topological design can be also reformulated to accommodate demand pricing and uncertainty issues. Due to limited space we do not consider more sophisticated optimization models integrating pricing with user utilities. In the rest of this subsection we briefly discuss approaches to solving stated optimization problems.

Solving *FA* and *CA* problems is comparatively easy [6]-[7]. *CFA* problem can be solved iteratively as follows. First, given link capacities c, find optimal route assignments ξ by solving the *FA* problem. Second, given these fractions ξ, find new link capacities c and losses L by solving the *CA* problem. Third, repeat the process. Note, however, that since *CFA* problem typically has multiple local optima [22]-[23], the space of possible initial conditions for this iterative process should be explored.

III. NETWORK PROVISIONING UNDER PRICED DEMANDS

A. Priced Demands

In economic models [4]-[6] the network is capable of controlling demands μ_{ij} by varying prices p_{ij}:

$$\mu_{ij} = \mu_{ij}(p_{ij}) \tag{11}$$

Sometimes, price-demand curve is obtained by maximizing user utility function as follows. If the user utility of transmitting from node i to node j at rate μ is $u_{ij}(\mu)$, it is natural to assume that the user maximized its total utility:

$$\mu_{ij}(p_{ij}) = \arg\max_{\mu \geq 0}[u_{ij}(\mu) - p_{ij}\mu] \tag{12}$$

In a case of monotonously increasing, concave utility function $u_{ij}(\mu)$, rate (12) is the unique solution to the following equation: $du_{ij}/d\mu = p_{ij}$. The price elasticity is defined as follows: $\varepsilon_{ij} = -(p_{ij}/\mu_{ij})(d\mu_{ij}/dp_{ij})$. In a case of constant elasticity $\varepsilon = \varepsilon_{ij}$, demand (11) takes form $\mu_{ij} = A_{ij}/p_{ij}^{\varepsilon_{ij}}$, where constant A_{ij} may be interpreted as a demand potential for origin-destination pair (i, j). Since $d(\mu_{ij}p_{ij})/dp_{ij} = (\varepsilon - 1)\mu_{ij}$, the effect of increase in price p_{nk} on the revenue is

$$\frac{\partial P}{\partial p_{nk}} = (1 - L_{nk})(1 - \varepsilon_{nk})\mu_{nk} - \sum_{(i,j)}\left(p_{ij}\mu_{ij}\frac{\partial L_{ij}}{\partial p_{nk}}\right)$$

In a case of lightly loaded network $L_{ij} << 1$, $\partial L_{ij}/\partial p_{nk} << 1$, $\forall (i,j),(n,k)$, and thus

$$\frac{\partial P}{\partial p_{nk}} \approx \frac{d(\mu_{nk}p_{nk})}{dp_{nk}} = (1 - \varepsilon_{nk})\mu_{nk} \tag{13}$$

It follows from (13) that lightly loaded network has incentive to increase utilization by reducing price p_{nk} if $\varepsilon_{nk} < 1$, and decrease utilization by increasing price p_{nk} if $\varepsilon_{nk} > 1$ [6].

B. Network Provisioning

If the price-demand relation (11) is known, the network may attempt to price accepted demands in order to maximize the revenue generated by the network. In this subsection we formulate some typical optimization problems.

Price Assignment (PA) problem. Given link capacities $c = (c_l)$ and fractions $\xi = (\xi_r)$, find prices $p = (p_{ij})$ and admission policy $L = (L_{ij})$ that maximize revenue P subject to constraints (6)-(10).

Price and Flow Assignment (PFA) problem. Given link capacities $c = (c_l)$, find prices $p = (p_{ij})$, admission policy $L = (L_{ij})$, and fractions $\xi = (\xi_r)$ that maximize revenue P subject to constraints (6)-(10).

Price and Capacity Assignment (PCA) problem. Given fractions $\xi = (\xi_r)$, find prices $p = (p_{ij})$, capacities $c = (c_l)$ and admission policy $L = (L_{ij})$ that maximize the net revenue W subject to constraints (6)-(10).

Price, Capacity and Flow Assignment (PCFA) problem. Find prices $p = (p_{ij})$, capacities $c = (c_l)$, admission policy $L = (L_{ij})$, and fractions $\xi = (\xi_r)$ that maximize the net revenue W subject to constraints (6)-(10).

Solution to the optimization problems *PA* and *PFA* can be formulated in terms of the link costs [6]. However, solving *PCA* and *PCFA* are to a large degree open problems.

IV. Network Provisioning under Uncertainty

A. Models of Uncertainty

All, so far formulated in this paper optimization problems for network provisioning assumed that the performance criterion Φ was a known function of the set of design parameters, i.e., controlled variables $z \in Z$: $\Phi = \Phi(z)$. The optimal provisioning decision $z = z^*$ was determined by solution to the optimization problem $\Phi^* = \max \Phi(z)$ over $z \in Z$. Long-term network provisioning decisions, i.e., long-term price contracts, capital investment into capacity expansion, etc., are more prone to uncertainties than short-term network provisioning decisions, i.e., short-term price contracts, admission and routing strategies, etc. To account for this difference we separate the set of design variables z into variables $\omega \in \Omega$ describing long-term network provisioning decisions and variables $\theta \in \Theta$ describing short-term network provisioning decisions: $z = (\omega, \theta) \in Z = \Omega \otimes \Theta$. Of course, in absence of uncertainty, the optimization problem can be simply rewritten in terms of the long and short term design variables as follows: $\Phi^* = \max \Phi(\omega, \theta)$ over $(\omega, \theta) \in \Omega \otimes \Theta$. Note that separation of the design variables into long and short-term design variables depends on the particular uncertainties to be incorporated into the model. It is also possible to have design variables describing intermediate-term network provisioning decisions.

We propose to model uncertainty by replacing criterion $\Phi(\omega, \theta)$ with the following mixture

$$\mathrm{M}_\gamma(\omega,\theta) = \sum_s \gamma_s \Phi_s(\omega,\theta) \qquad (14)$$

of known functions $\Phi_s(\omega, \theta)$ with unknown weights

$$\gamma \in \Gamma = \left\{ \gamma_s \middle| \sum \gamma_s = 1, \gamma_s \geq 0, s = 1,..,S \right\} \qquad (15)$$

We will refer to criteria $\Phi_s(\omega, \theta)$ as scenarios and interpret criterion (14) as a mixture of these scenarios with weights $\gamma = (\gamma_s)$. If the scenario s is known, the optimal network provisioning decisions (ω^s, θ^s) are determined by solution to the following optimization problem $\Phi_s^* = \max \Phi_s(\omega, \theta)$ over $(\omega, \theta) \in \Omega \otimes \Theta$. However, this solution cannot be directly implemented since scenario s is not known at the moment of making the network provisioning decision.

B. Examples of Uncertainty

Any projection for the future demands is a subject to statistical uncertainty. If the underlying statistical model does not change, the statistical uncertainty can be reduced by increase in the amount of historical data used for the projections. However, inevitable variations in the external conditions affecting the demand pattern may severely reduce the amount of historical data to be reliably used for the projection of the future demands. In this situation the point estimates for the expected demands $\mu_{ij} \approx \tilde{\mu}_{ij}$ may be meaningless. Instead, the confidence intervals $\mu_{ij} \in [\breve{\mu}_{ij}, \hat{\mu}_{ij}]$ should be used. In this situation one may interpret the end points of the confidence intervals as possible scenarios. The set of possible scenarios is determined by the correlations between expected demands with different origin-destinations. This model of uncertainty can be also used in a case of variable by "business hour" demands, when the pattern of variability is known, in an attempt to identify a single traffic engineering scheme suitable for all feasible demand matrices. In this case different scenarios correspond to demand matrices at different business hours, and thus the number of scenarios S is comparatively small.

If the network is capable of controlling demands μ_{ij} by varying prices p_{ij}, uncertainty in the overall economic conditions affect the price-demand curve (11). One can model this uncertainty by assuming the price-demand curve is a known function of the origin-destination (i, j), service price p_{ij} and scenario $s = 1,..,S$: $\mu_{ij} = \mu_{ij}^s(p_{ij})$. Uncertainty is modeled by assuming that the scenario $s = 1,..,S$ is unknown at the moment of making the network provisioning decision. If a scenario $s = 1,..,S$ is formulated in terms of the user utility $u_{ij}^s(\mu)$ of transmitting from node i to node j at rate μ, one may assume that

$$\mu_{ij}^s(p_{ij}) = \arg\max_{\mu \geq 0}[u_{ij}^s(\mu) - p_{ij}\mu] \qquad (16)$$

C. Network Provisioning

Broadly speaking, there are two possible frameworks for decision making under uncertainty. The *Bayesian* framework is concerned with the "average" performance by assuming that uncertain parameters follow some probability distribution [17]. This framework has been applied to *QoS* routing in [12] and to *OSPF* link weight optimization in [13]. Note that [13] implicitly assumed that all scenarios are equally likely: $\gamma_s = 1/S$, $\forall s = 1,..,S$. The Bayesian framework, however, may not be adequate if one is concerned with the worst rather than "average" case scenario performance. Such robustness concerns can be addresses within the *minimax* or game theoretic framework [17]-[18] by identifying the network provisioning decisions that minimize the worst-case scenario losses in performance resulted from the uncertainty. However, this approach is typically beyond reach due to computational complexity even for moderate size networks. The following framework offers some relief to this burden.

Consider a zero sum game G of two players (s) and (t). Player (s) represents nature and controls actual scenario

$s=1,..,S$. Player (t) represents the network, i.e., provider, and attempts to guess scenario $s=1,..,S$ selected by the nature. To obtain the payoff function for this game Δ_{st} consider loss in performance incurred by the network if player (t) selects strategy $t=1,..,S$, and player (s) selects strategy $s=1,..,S$. Obviously $\Delta_{st}=0$ if $t=s$. In an unfortunate for the network case $t \neq s$,

$$\Delta_{st} = \max_{(\omega,\theta)\in\Omega\otimes\Theta} \Phi_s(\omega^s,\theta) - \max_{\theta\in\Theta}\Phi_s(\omega^t,\theta) \geq 0 \qquad (17)$$

where ω^t is determined by solution to $\max \Phi_t(\omega,\theta)$ over $(\omega,\theta)\in\Omega\otimes\Theta$.

Obviously $\max_{s=1,..,S}\min_{t=1,..,S}\Delta_{st}=0$. If $\min_{t=1,..,S}\max_{s=1,..,S}\Delta_{st}=0$ this game has solution in pure strategies, the value of the game is zero, and the uncertainty does not result in any loss in the performance [24]. If $\min_{t=1,..,S}\max_{s=1,..,S}\Delta_{st}>0$ this game has solution in mixed strategies, the value of the game is positive, and the uncertainty results in the performance loss quantified by the game value [24]. Let $\zeta=(\zeta_1,..,\zeta_S)$ and $\eta=(\eta_1,..,\eta_S)$ be the optimal mixed strategies for the nature and network respectively, i.e., the optimal strategy for nature is to select scenario t with probability ζ_t, and the optimal strategy for the network is to bet on scenario s with probability η_s. Then, the optimal value of the long-term design parameters $\omega=\omega^{opt}$ is $\omega^{opt}=\sum_s \eta_s\omega^s$, assuming that $\omega^{opt}\in\Omega$. Since at the moment of selecting the short-term design parameters θ scenario s is already known, the network can select $\theta=\theta^{opt}$ by maximizing $\sum_s \zeta_s\Phi_s(\omega^{opt},\theta)$.

V. CONCLUSION

This paper has proposes framing network provisioning problem as a game against nature, where the network makes the network provisioning decisions and nature selects a feasible scenario with respect to future demands. Various extensions of the described network provisioning models, e.g., by incorporating quality of service, protection and restoration, etc., are possible. Despite this paper assumed a multicommodity network flow model, the same approach can be applied to a flow level network model with random flow arrivals. A promising extension of the proposed game theoretic framework appears to be framing the network provisioning process as a sequential (multi-stage) game [25]. This formulation may provide a natural framework for cognitive process of network provisioning. However, the most immediate problem is developing computationally effective methods for solving the corresponding optimization problems. Solutions for some particular cases and relation to regularization of ill-conditioned problems have been discussed in [19]-[21].

REFERENCES

[1] D.O. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus, "Requirements for traffic engineering over MPLS," Network Working Group, RFC 2702, http://search.ietf.org/rfc/rfc2702.txt, September 1999.

[2] K.W. Ross, Multirate Loss Models for Broadband Telecommunications, Springer, New York, 1995.

[3] F.P. Kelly, "Routing in circuit-switched networks: optimization, shadow prices and decentralization," *Adv. Appl. Prob*., 20, 1988, pp. 112-144.

[4] D. Mitra and K.G. Ramakrishnan, "Techniques for traffic engineering of multiservice, multipriority networks," *Bell Labs Tech. J.*, Vol. 6, No. 1, 2001, pp. 139-151.

[5] S. Shenker, "Fundamental design issues for the future Internet," *IEEE JSAC*, 13 (1995) 1176-1188.

[6] D. Mitra, K.G. Ramakrishnan and Q. Wang, "Combined economic modeling and traffic engineering: joint optimization of pricing and routing in multi-service networks," *Proc. 17th ITC*, Brazil, 2001.

[7] S.G. Lanning, D. Mitra, Q. Wang and M.H. Wright, "Optimal planning for optical transport networks," *Phil. Trans. Royal Soc*., London, Vol. 358, No. 1773, 2000, pp. 2183-2196.

[8] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: methodology and experience," *IEEE/ACM Transactions on Networking*, June 2001, pp. 265-279.

[9] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, "Traffic matrix estimation: existing techniques and new directions," *ACM SIGCOMM*, Pittsburgh, PA, 2002, available at http://www.acm.org/sigcomm/sigcomm2002/

[10] A. Odlyzko, "Internet growth: myth and reality, use and abuse," available at http://www.dtc.unm.edu/~odlyzko/doc.

[11] N. Maxemchuk, I. Ouveysi and M. Zukerman, "A quantitative measure for telecommunications networks topology design," available at http://comet.ctr.columbia.edu/~nick/.

[12] R.A. Guerin and A. Orda, "QoS Routing in Networks with Inaccurate Information: Theory and Algorithms," *IEEE/ACM Trans. on Networking*, 7 , 1999, pp. 350-364.

[13] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world," *IEEE JSAC*, Vol. 20, No. 4, May 2002, pp. 519-528.

[14] S. Plotkin, "Competitive routing of virtual circuits in ATM networks," *IEEE JSAC*, 13, 1995, pp. 1128-1136.

[15] V. Marbukh, "Network Management under Incomplete Information on the Operational Environment," *International Symposium on Information Theory and Its Applications (ISITA2000)*, Hawaii, pp. 637-640.

[16] V. Marbukh, "Minimum Regret Approach to Network Management under Uncertainty with Applications to Connection Admission Control and Routing," *Int. Conf. on Networking.(ICN'01)*, France, 2001.

[17] D. Blackwell and M. Girschick, Theory of Games and Statistical Decisions, Wiley, New York, 1954.

[18] E. Koutsoupias and C.H. Papadimitriou, "Beyond competitive analysis," *Foundations of Computer Science*, 1994. (Updated version available at http:/www.cs.berkeley.edu/christos/)

[19] V. Marbukh, "On Shortest Random Walks under Adversarial Uncertainty", *Fortieth Annual Allerton Conference on Communication, Control, and Computing*, Monticello, Illinois, 2002.

[20] V. Marbukh, "Minimum Cost Routing under Adversarial Uncertainty: Robustness through Randomization," *Int. Symp. on Information Theory (ISIT2002)*, Lausanne, Switzewland, 2002.

[21] V. Marbukh, "QoS Routing under Adversarial Binary Uncertainty: Solution for a Symmetric Case," *Int. Conf. on Communications (ICC2002),* New York, 2002.

[22] Kleinrock, Queueing Systems. Volume II: Computer Applications, John Wiley, 1976.

[23] D. Bertsekas and R. Gallager, Data Networks, Prentice-Hall, New Jersey, 1992.

[24] N.N. Vorob'ev, Game Theory, Springer-Verlag, New York, 1977.

[25] D. Fudenberg and J. Tirole, Game Theory, MIT Press, Cambridge, Massachusetts, 2000.