Fairness on the Internet and its Importance in Development Context

Elly Amani Gamukama
Department of Computer and System Sciences
Stockholm University/KTH, Forum 100, Isafjordsgatan 39, SE-16440, Kista, Sweden
Tel: +46-70-7568891, Fax: +46-8- 703 90 25
gamukama@dsv.su.se

Abstract. *The notion of fairness with respect to resource sharing among the competing flows is one of the important considerations in network design. This is true in IP network where the service model is based on best effort and any possible distortion of it may lead to flow starvation and eventually system imbalances. Fairness is one of the major objectives both on a network and transport layer. This is evident in the case of elastic flows where fairness has a major impact on congestion resolution. On a network layer, fairness mechanisms combined with scheduling and queuing policies lead to equitable service, which may also induce higher router utilization and hence better network performance. The paper investigates the current trends in understanding and applying the fairness concept on the Internet. Then it studies and examines the extension of the fairness concept in the context of development and developing regions, where both the traditional lack of infrastructure and costly communication services have also affected the penetration of the Internet and more even distribution of its benefits. The key question is whether or not it is plausible to identify a framework for the evaluation of efficiency-fairness tradeoffs that may provide a sound basis for a model of a more equitable access to the Internet to a diversity of users with different needs and financial possibilities representing mainly developing regions and emerging economies.*

Keywords: Fairness, efficiency, Internet, resource sharing, developing regions

1 Introduction

Although fairness is conceptually associated with allocation of limited resources among competing users, it is an important goal for effective designing, operating and maintaining of a system. It instills a cooperative behavior among participating users that leads to achieving an economically, socially and environmentally sustainable system. In an environment of competitive individualists users, the critical factor of cooperation relies on the underlying notion of fairness as well as incentives for adopting certain behaviors [1].

In respect to the Internet architecture, fairness adoption has far reaching in the design of TCP congestion control mechanism. For example, the implementation of AIMD in TCP was on the basis that multiple flows competing for network resources at the same bottleneck would converge to fair sharing in the steady state[2]. Fairness qualifies a network resource allocation where no user is penalized compared to others that share the same bottleneck. This philosophy has enabled the best-effort service model (that basically relies on the fair sharing of network resources due to the lack of explicit admission control and qualities service assurances) to lead the growth of the Internet from a network of handful researchers to a social institution of substantial import.

The concept of fairness has been further extended into various other mechanisms for fair scheduling and fair queuing management at the network layer [3-7] to ensure that no flow starves another.

The viable economic aspects of fair resource allocation was theoretically proved using utility optimization methods [8, 9]. In [8] the problem of rate allocation is posed as one achieving maximum aggregate utility for users, with assumption that users' resource demands are elastic and their rates are adjustable basing on the network feedback. Because of the decentralization nature of networks, it further proposes the use of pricing as to decompose the overall systems problem into sub-problems for the network and for the individual users. This is seen as an incentive to the users in requesting for the proper service classes for their applications. In this way the systems optimal can be achieved when both the users' choices of allocation rates that they are willing to pay for and the network's allocated rate are in equilibrium.

This paper investigates the current trends in understanding and applying fairness in IP based networks. It strives to identify a framework for evaluation of efficiency-fairness tradeoffs that would be a basis for developing a model for provision of the Internet to a diversity of users with different needs and varying scales of income possibly representing developing regions and emerging economies where poor infrastructures and costly communication services have hindered the community from accessing the Internet benefits.

The rest of the paper is organized as follows. In section II, we present a brief background of the fairness concept and common fairness criteria in use on the present Internet. Section III, presents the current trends of furthering the fairness concept. We discuss its relevance and its importance in context of developing regions and present the envisioned research directions for developing a framework under study in section IV. We sum up our views in section V the conclusion.

2 Fairness Concept and Models in IP Networks

A large variety of fairness notions exists in the literature. The fairness notion spans several disciplines - philosophy, social sciences and science. From a societal perspective, fairness notion is appealing because it: provides an intuitive basis for anayzing distribution issues; exhibits symmetry across individuals; is consistent with on ordinal representation of individual prefrences; is free of interpersonal comparison of utility; requires each indivual to evaluate others' bundle using own preferences. In summary it strives for equality in treatment of all participating agents.

The more grounded theoretical approaches that lay a basis of fairness notion on the Internet stem from the economic theories that uses the concepts of *utility* and *welfare functions* in defining fairness [10-12].

The idea of utility is brought in context as a way of describing consumer preferences. It is a measure of relative satisfaction or desiredness from consumption of goods.

A utility function is formalized as $u(x_1,x_2, ..., x_n)$, where $x_1,x_2, ..., x_n$ are the quantities of each of n bundles. It describes the ordering of individuals' preferences but its outcome value has no absolute meaning. Thus, a bundle x_i is preferred to bundle $x_i^{'}$ if and only if $u_i(x_i)) > u_i(x_i^{'})$. Consequently, under any monotonic transformation of a utility function will represent the same preferences.

2.1 Formalization of utility functions in IP networks

In networks *utility functions* are defined for networking applications as functions that map a service delivered by the network into the performance of the application for that service. This notion was formalized by Shenker in [13]. From this definition, the magnitudes of utility levels have a meaning since they also describe the relative differences in applications performance between distinct levels of services. In [13] Shenker gives qualitative examples of utility functions for different classes of applications (elastic, real-time, rate-adaptive, and step) as function of bandwidth.

2.2 The Common Fairness models in use

The goal of a fairness model is fair sharing of network recourses. A model has to define a formal objective that has to be used as a fairness performance target in a network, i.e. there are two aspects in model's definition – a resource to be shared and a policy. For example the fairness model could consider bandwidth as the shared resource. Then, the formal objective to be used as a fairness performance target in a network would be throughput. Therefore the policy would be that every flow gets an equal share of the bandwidth. In case of a bottleneck on the path, such a policy will penalize flows with high rates while those with low rates are given the greatest possible allocation. The common fairness criteria are discussed below.

2.2.1 Max-Min fairness

This is the classic and most popular fairness criterion in data networks [14]. It roots from the welfare theory with a corresponding welfare function; $W(u_1,u_2, ... ,u_n) = min(u_1,u_2, ... , u_n)$ and individualistic utility functions $u_i(x_i) = x_i \forall i \in \{1,2,...,n\}$. That is, for all users, the utility of the resource bundle allocated to user i is represented by the value of the resource bundle itself. Maximin fairness thus yields a solution $x^{'} = (x_1^{'},....x_n^{'})$ for $\max(\min(x_1,....,x_n))$. The interpretation of this solution is that for all $\forall i, x_i^{'}$ cannot be increased without simultaneously decreasing $x_j^{'}$ for some j with $x_j^{'} \leq x_i^{'}$. Its fair share policy can be summarized as:

- The shared resource is allocated in order of increasing demand
- No user receives a share of the resource large than its demand
- Users with unsatisfied demands receive equal shares of the resource

As a consequence, with this criterion the most poorly treated users are given the greatest possible allocation. The scheme protects them from the bad behaved users and can be implemented with local information only.

2.2.2 Proportional fairness

The proportional fairness criterion formal objective is to maximize the overall utility of rate allocations of the flows in progress [9, 15]. It roots from the welfare theory with a function corresponding to the Nash criterion as; $W(u_1,u_2,.......,u_n) = \sum_i \log(u_i)$ and with individualistic utility functions $u_i(x_i)$.

A proportional fair allocation is a solution to the welfare maximization problem. Taking Kelly's model in [9], an allocation of rate x is proportional fair iff for any other feasible allocation $x^{'}$, satisfies the condition $\sum_{i=1}^{S} \frac{x_i^{'} - x_i}{x_i} \leq 0$

The idea behind proportional fairness is to maximize the overall performance. With proportional fairness, a worse treated entity may see its utility decreased if this allows a large enough increase to an already better treated entity.

3 Current trends in understanding the fairness concept on the Internet

3.1 Causes of the divergence in the fairness trends

The state of concern is that the concept of fairness on the Internet has been used as a criterion to guide the design of traffic controls and its implementation has been successful in the best effort networks. The successes achieved have resulted into a brief that the Internet connectionless model needs not to change as long as there is adequate provisioning of the network resources. That what is needed is a relaxed end system congestion control that can co-exist with the widely used TCP congestion control. All efforts therein invested have resulted into a compromise of TCP-friendly congestion controls. However, the current Internet is driven by a congregate of ideas from the society. Thus, the society's dynamics have led to a congregation of heterogeneous applications on the Internet, of which some of their flows cannot adapt to network congestion in a TCP or TCP-friendly style (gentle adaption that leads to a convergence to point similar to that of TCP flows). Evenif the adaption was possible, for inelastic (multimedia) flows the offered TCP-fair rate might not meet the minimum required quality of service rate. Consequently, this has raised divergent views that contravene the TCP/TCP-friendly congestion control compromise.

Furthermore, other applications use TCP (for example the P2P applications – Gnutella, BitTorrent, etc.) and their use at run-time have distorted the essence of fairness[16, 17]. According to the end-to-end philosophy, TCP fair sharing of network resources is between the flows not the users. With P2P, an application can open a hundred of multiple connections concurrently at their run time that can congest links with less capacity. We demonstrate this perception of TCP distortion at run-time with an example as follows. Assume a network with a bottleneck of capacity C is being used by N users. Each has a one unit flow, then TCP would allocate the available capacity fairly and each flow gets a capacity of C/N. But if one of the users runs an application (e.g. P2P application) that starts N flows in the network, then the total flows in the network will be 2N. TCP will fairly allocate the bottleneck capacity to each flow with C/2N capacity. Therefore this implies that the one user with an application that started N flows is getting a total of N*C/2N = C/2 capacity of the network to her/himself, which is as much as the capacity for the rest of the N-1 users. Therefore from the users' perspective, this way of resource allocation is unfair. Consequently, TCP fairness is distorted[1] at run-time though at protocol design-level it perfectly fulfills the ethical fairness notion.

Consequently these problems have created other research directions of revisiting the fairness concept on the Internet. We discuss them these tracks below:

3.2 Introduction of admission control proponents

The direction here is to let each type of flow have its own type of controls. It argues for the abandoning of the traditional notion of fairness in the designing of end-system flow controls for different applications. It advances the idea of flow co-existence with each implementing its own flow controls. It proposes the introduction of admission controls [13, 18, 19] for inelastic flows and argues for the application of TCP-friendliness at aggregate level for elastic flows instead of the present per-flow TCP-friendly control.

Unlike past works that proposes distributed implementation of admission control based on probing the network with assumption that the network supports Integrated or differentiated services that help in isolating elastic and inelastic flows, this work assumes that inelastic traffic simply performs distributed controls without any network assistance. This approach proposes the use of stochastic process to model the co-existence of the different flows under different controls. It further strives to extend the standard network utility maximization [8] by considering non-concave utility functions to model inelastic traffic as well considering both the elastic and inelastic flows as stochastic processes with finite sizes.

The findings already show that from maximizing the social welfare point of view, it is better to be TCP-friendly at the traffic class aggregate level rather than at the per-flow level. These findings coupled with the user fairness concepts reviewed in section II can lead in establishing new optimal levels at which networks can serve a diversity of users with different scales of income that depicts those from developing regions.

---

[1] Note only that P2P applications distort the state of fairness through misuse of protocols like TCP, but they are also security threats in poor and slow Internet infrastructures. The million multiple concurrent connections can exhaust the routers' resources. In such a state a router can't accept any more connections, thus creating a new problem of DOS form of congestion.

3.3 Shift controls to application run-time

The research trend challenges the concept that fairness can be achieved by controlling flow rates alone. It argues that flow rate fairness was the goal behind the design of fair resource allocation mechanism that is widely deployed in protocols like weighted fair queuing, TCP congestion control, TCP-friendly rate control, etc, but up-todate issues of resource allocation and accountability are still on the list of the Internet architecture requirements. Hence, it proposes that relative flow rates should be the outcome of fairness mechanism but not a mechanism itself [20] Furthermore, from theoretical economics fairness concerns comparing benefits, costs or both. In this context fairness is concerned with the benefit, cost or both for data transfer over the network. Though the benefits of data transfer can be assumed to increase with flow rate, the utility function relating the two is unknown or is unclear and even is subjective and private to individual users. In addition there is a great difference in the magnitude of flow rates (bits/sec) among applications that realizes users' benefit. For example consider a multimedia application with an ordinary web application (SMS, e-mail, web browsing) delivering services to the users. The users' benefits flow rates are in magnitude of differences. Therefore there is no where one can compare the benefits delivered by different applications at equal flow rates.

Considering the cost of data transfer across a network, it is well known that the cost increases with the path capacity exhaustion. If the offered rate is higher than the available capacity then congestion occurs and the cost becomes higher, otherwise the user can send more data at no increase cost, i.e. cost is dependent on the level of congestion not only the flow rate. But the present network layer mechanism does not restrict the cost to the only user or users causing congestion instead the cost is sent to all active users, eventhough through queue management the higher the flow is the higher loss notifications the user/source gets from the router. The underlying argument is that users should be accountable for the cost of the congestion they causes, if flow rate allocation is to be a measure of fairness. However this can be achieved through router policing every user connected to the Internet.

The overall basis of the research trend is that the design ideology of controlling fairness at protocol design-time has been misused and distorted by the applications in use, users and network operators in the present Internet. Consequently resource shares are now primarily determined at run-time. The shift has brought divergent views among users and network operators on the issues of protocol fairness. The use of P2P applications or P2P application implementation is a manifesto of such claim that have distorted the gentile and well crafted TCP fair resource sharing; hence making the flow rate equality insufficient/inappropriate in technical and policy terms. Similarly the use of Deep Packet Inspection (DPI) tools by operators to discriminate among traffic types and throttle their heaviest users' traffic mainly at peak periods overrides the resource sharing decisions of TCP. In addition it contravenes the end-2-end principle of the Internet architecture and it is unethical to throttle the users' traffic without their consent because operators don't know the users priorities among ones traffic flows.

Therefore the research trend argues that the challenge in instituting fair resource sharing is not a design-time issue of the protocols but a control and determination of resource shares at run-time. The remedy to the problem at hand is to create a run-time policy framework within which users and operators can control relative resource shares. A policy that principled and enforceable bring an agreement between users and operators on which fair use policy they want locally to use at run-time.

3.4 Bounded fairness compromise

The bounded fairness compromise ((α-β) – fairness) concept is mainly coined by Zukerman *et al* [21] and references therein. It defines a resource allocation to be ((α-β) – fair if a source is allocated a rate neither less than $0 \leq \alpha \leq 1$ nor more than $\beta \geq 1$ times its fair share. The overall thesis of the study is to investigate to what extent fairness can be compromised by maximizing the efficiency. In this context efficiency refers to a value of a general utility function of flows rate allocation. Such utility functions are interpreted to be profit, revenue, investment costs, utilization, throughput, etc. Unlike earlier studies, the utility function that defines efficiency is unrelated to fairness criterion, i.e. by making variations in α, β (decreasing α and increasing β) values increases the feasible set from which an optimal point is achieved for maximizing the efficiency and an efficiency-fairness function can be obtained to further analysis of the benefit in efficiency as a function of the extent to which fairness is compromised.

Taking this framework under a broad generalization, the efficiency – fairness tradeoff can be used as a basis to determine the types of services to be offered and the corresponding beneficiaries.

4 Efficiency-Fairness framework in development context

4.1 Basis for developing a framework

Having looked the fairness concept and its trends on the Internet we hereby point at possible research directions that would lead in developing a framework appropriate for developing regions. Our basis for the argument in this work stands on the following pillars:

- The present level of the Internet technology is sufficient enough to foster development in developing regions. However the challenge is to guard the perpetuating distortion of the fairness notion on the Internet while maintaining a recommendable level of efficiency in networks. Through this approach we foresee an efficient and fair use of the available limited network resources for connecting to the global Internet. Else no steady growth of the Internet usage which brings real benefits that manifests development in society will be realized in a near future.

- There are social inequities of access between urban and rural Internet users partly due to the current ambiguities in the costing models that do not consider users' level of income and applications in use.

The term "development" used is in context of the border concept of human development as defined by the United Nations Development (UNDP) [22]. In this context, we define the Internet for development as services to an individual being able to access information (using the IP based networks infrastructure) and use it for the betterment of one's life/living [23]. We enumerate the types of the required basic services and evaluate their corresponding levels of the Internet infrastructure requirements (in respect to transport and network layer) that would be appropriate to deliver the services. Then we use the findings as a basis of pointing at possible research directions in formulating a framework that would appropriately evaluate efficiency-fairness tradeoffs that may provide a sound basis for a model of a more equitable access to the Internet services in developing regions.

From the last mile user's perspective, the Internet for development can be summaries (but not limit to) in the following access domains [22, 23]. Access to:

- Market information (news, item/products prices, business partners, market places and their restrictions, etc.)
- Medical information/awareness (type of medical/health services – where/places, who, cost; policies, guidelines, epidemic outbreak warnings, counseling services, etc.)
- Affordable and quick communication services (e.g. voice, e-mail, SMS).
- Human skills upgrade; able to access web-based self tutors/guides both online or offline that would enhance one's skills e.g. in ICTs, agriculture, cookery, health/hygiene, etc.
- Existing national policies and contributing to news ones in formation for promotion of democracy and better governance
- Local and international news and keep embraced with current affairs.

The basic required Internet services for development in developing regions can be well covered under the best-effort services model. Then, the issue is why the Internet penetration and use has been and is still poor to date in developing regions?

We believe that among the many other issues that may contribute to the failure of the Internet penetration in developing regions includes models in use that are likely to be not favorable to both operators and users. To the operators, the Internet (telecommunication) Infrastructure require a heavy investment that puts it at a high risk of delivering services to poor communities whose population of consumption/use is very low. While to the users, the present Internet costing models do not put in considerations issues like the users';

- Level of incomes
- Billing only the active time online (time of data transfer)
- Application type in use.

The costing model commonly used in developing regions agitates for flat rates. Consequently subjecting rural users to pay the same user-costs like those from urban areas even-if their levels of income are not comparable. The ambiguities and the unfairness of these costing models can be ascertained by the cost of access in developing countries vis-à-vis developed countries, whereby the cost of access to the Internet in developing countries is much higher than the cost in developing countries. The point of contention is not how much should a rural Internet user pay for a bit-rate vis-à-vis the urban user, but it is how Internet can penetrate rural areas and users get its benefits. The present flat costing model does not favor the rural user with low income unless subsidies are introduced as a government policy or as an operators' business model for marketing. These costing models do not have such in built options to favor such users neither governments from developing regions have considered the Internet as utility in a human's life to qualify for national subsidies that can offset the investment risks faced by network operators. Consequently we are still faced with issues of social inequities of access to the Internet. This is a challenge still faced in developing regions that calls for social equity of access. Our framework will also look into research directions that would address such issues.

The unfairness caused by cost signals resulting from congested routers to the transport layer has been discussed in section 3.3 above. We also noted that the trend has now intensified by the distortion of fairness notion by applications like P2P. The implication is that costing models have become inefficient and calls for new models that can properly account for the user's bit-rate as well as the congestion the user cause to others users. That is to say one should pay for the congestion that he/she causes to other users. However, we note as discussed in the next section that bounded fairness can be achieved by establishing proper utility functions that delivers users and operators benefits for specified developmental targets. Consequently, one operating within such limits will not only attain his/her optimal throughput but also not cause congestion along the path.

In order to address the central issue of fairness importance on the Internet in developing context, we propose to direct out studies on developing a framework under which efficiency – fairness tradeoff of the network can be analyzed.

4.2 Efficiency versus fairness in developing context

Efficiency and Fairness with respect to resources sharing are main objectives in any system design. Many choices in life are made based on their tradeoffs. In respect to this work we strive to identify a framework to evaluate this concept that may provide a sound basis for a model of providing Internet services to users with a diversity level of income in developing regions. The pertinent questions under consideration in developing the framework are:

- What are the efficient and fair levels needed for an operational network in developing regions that can enable an operator (government owning the public network, a private telecom company, and ISP, etc.) extend the network to rural areas and provide the Internet services to rural users based on the users' level of income without bearing investment risks and or jeopardizing sustainability and upgrade of the network.
- What principled and enforceable policies are required to equip with the user (individual person, institution, company, civil societies proponents, non-governmental organizations, public office, etc.) and operator that can uniquely bring an agreement on fly between them on what fair resource sharing policy they want locally to use at run-time.

4.2.1 Redefining the network efficiency

To achieve the intended framework, there is a need to redefine the network efficiency in respect of the benefits it delivers to the users and operators. Although traditionally network efficiency is strongly related to the bandwidth sharing objectives and measured through overall network goodput, we extend the definition to include benefits accrued to the users and operators. In respect to the operator, network efficiency would be interpreted in view of revenue, profit, (investment returns) that can be generated while to users is how the network leads them in fulfillment of their benefits. Evenif there are some divergences in the definition or extension of the term efficiency in our context, the benefits can be seen as a function of the throughput /goodput of the network. Consequently nay level changes in the throughput/goodput will monotonically change the benefits' level of the user and operator.

4.2.2 Identifying the fairness bounds

The central issue is to identify the fairness bounds under which the set (investment returns in case of the operator and user benefit) targets can be achieved given constraints that uniquely define the area in study (developing regions). In order to overcome the fundamental conflicts between fairness and efficiency [24] (and see references therein that gives elaborative examples on resource allocation in wired networks, wireless networks and economics), we propose to adopt the bounded fairness compromise ((α-β) – fairness) concept similar to that of Zukerman *et al* [21]. As discussed in section 3.4 above, this concept is not biased to any particular fairness criteria and can be applied to any topology of network. In addition it enables one to set minimum and upper allowable bounds that meet specific target benefits.

However, to clearly make sense out of the tradeoff between efficiency and fairness we have to define utility functions that uniquely define or have a direct interpretation of the level of satisfaction attained by the users/operator in the study area. With respect to types of access required in development context as indicated in section 4.1 above, what might be required is to maximize aggregate benefit to the users. It is evident that maximizing aggregate utility subject to linear capacity constraints would lead to the required benefits as longer as the utility functions have been clearly formulated. However to maximize the benefit of the operator, it might require taking a more central approach that puts both considerations of the linear and nonlinear capacity constraints. That is, nonlinear to cater for those users who need more or are beyond the requirements of development context. This requires establishing the degree the network efficiency can be attained by compromising on fairness within given α and β bounds.

Hence this leads in defining an efficiency-fairness function that will qualify the tradeoff.

4.2.3 Control of social dynamics on the Internet in development context

Having noted that the society's dynamics drives the present Internet applications' development and use, as discussed in section 3.1 above, then a challenge is to implement new policies and controls that can effectively protect the network while at the same time fairly meeting the users' dynamics. Integration of such measures in the frameworks would be able to address issues of intra and inter protocol fairness and implementation of applications based costing. The implication is that access to the Internet services not tailored in context of development, would cost a bit higher hence bringing more profit and revenue to the operator and the nation, indirectly introducing subsides to users whose Internet access services are classified under development context. This would lead to Internet penetration to rural areas and enabling users getting the highest level of Internet benefits.

5 Conclusions

This paper has presented the current trends in understanding and applying the fairness concept on the Internet. It also discussed the level to which the fairness concept is being distorted and the consequent effects to Internet costing. In general fairness has been one of the main objectives in the design of resources allocation in IP networks.

Evaluation studies based on efficiency-fairness tradeoffs can lead in development of models that gives more insight in making choices in life. Consequently, we presented a way-forward for identifying a framework for evaluation of efficiency-fairness tradeoffs that may provide a sound basis for a model of a more equitable access to the Internet to a diversity of users with different needs and financial possibilities representing mainly developing regions and emerging economies.

6 Acknowledgments

We wish to thank the Swedish International Development Agency through its department for Research Co-operation (Sida/SAREC) which is supporting the ongoing research study.

References

[1] P. Gevros, J. Crowcroft, P. Kirstein, and S. Bhatti, "Congestion control mechanisms and the best effort service model," *IEEE Network,* vol. 15, pp. 16 - 26, 2001.

[2] D. M. Chiu and R. Jain, "Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks," *Journal of Computer Networks and ISDN,* vol. 17, pp. 1-14, 1989.

[3] J. Nagle, "Congestion control in IP/TCP internetworks," *ACM SIGCOMM Computer Communication Review,* vol. 14, pp. 11 - 17 1984.

[4] M. Shreedhar and G. Varghese, "Efficient fair queueing using deficit round-robin," *IEEE/ACM Transactions on Networking (TON),* vol. 4, pp. 375 - 385 1996.

[5] Y. Zhoua and H. Sethu, "On achieving fairness in the joint allocation of buffer and bandwidth resources: Principles and algorithms," *Computer Networks,* vol. 50, pp. 2239-2254, 2006.

[6] A. Demers, S. Keshavt, and S. Shenker, "Analysis and simulation of a fair queueing algorithm," in *Symposium proceedings on Communications architectures & protocols*, 1989.

[7] S. Floyd and V. Jacobson, "Link-sharing and resource management models for packet," *IEEE/ACM Transactions on Networking (TON),* vol. 3, pp. 365 - 386 1995.

[8] F. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications,* vol. 8, pp. 33-37, 1997.

[9] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," *The Journal of the Operational Research Society,* vol. 49, pp. 237-252, 1998.

[10] E. A. Pazner, "Pitfalls in the theory of fairness," *Journal of Economic Theory,* vol. 14, pp. 458-466, 1977.

[11] H. R. Varian, "Equity. Envy, and Efficiency," *Journal of Economic Theory,* vol. 9, pp. 63-91, 1974.

[12] H. R. Varian, "Distributive Justice, Welfare Economics, and the Theory of Fairness," *Philosophy and public Affairs,* vol. 4, pp. 223-247, 1975.

[13] S. Shenker, "Fundamental Design Issues for the Future Internet," *IEEE Journal on selected area in Communications,* vol. 13, pp. 1176-1188, 1995.

[14] D. Bertsekas and R. Gallager, *Data Networks*, 2nd ed.: Prentice Hall, Upper Saddle River, NJ 07458, 1992.

[15] L. Massoulie and J. Roberts, "Bandwidth sharing: objectives and algorithms," *IEEE/ACM Transactions on networking,* vol. 10, pp. 320 - 328, 2002.

[16] N. Basher, A. Mahanti, A. Mahanti, C. Williamson, and M. Arlitt, "A Comparative Analysis of Web and Peer-to-Peer Traffic " in *17th International World Wide Web Conference (WWW2008)*, Beijing, China,2008.

[17] B. Briscoe, T. Moncaster, and A.-L. Burness, "Problem Statement: We Don't Have To Do Fairness Ourselves, draft-briscoe-tsvwg-relax-fairness-00draft-briscoe-tsvwg-relax-fairness-00," 2007.

[18] D. M. Chiu and A. S.-W. Tam, "Fairness of traffic controls for inelastic flows in the Internet," *Computer Networks,* vol. 51, pp. 2938 - 2957, 2007.

[19] D. M. Chiu and A. S.-W. Tam, "Network Fairness for Heterogeneous Applications " Department of Information Engineering, The Chinese University of Hong Kong, Hong Kong 2004.

[20] B. Briscoe, "Flow rate fairness: Dismantling a religion," *ACM SIGCOMM Computer Communication Review,* vol. 37, pp. 63 - 74, 2007.

[21] M. Zukerman, M. Mammadov, L. Tan, I. Ouveysi, and L. Andrew., "To be fair or efficient or a bit of both " *Computer and Operations Research,* 2007.

[22] UNDP, "United Nations Development Programme on Human Development Report," 2007.

[23] L. Press, "The role of computer networks in development," *Communications of the ACM,* vol. 39 pp. 23 - 30, 1996.

[24] A. Tang, J. Wang, and S. H. Low, "Is fair allocation always inefficient," in *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, 2004, pp. 45-55.