25

A Fairness Index for Communication Networks

Wing Shing Wong and Fengjun Li*

Department of Information Engineering, Chinese University of Hong Kong
{wswong;fjli2}@ie.cuhk.edu.hk

Summary. This paper introduces a new fairness index in open architecture networks. The concept can be used to balance the load of self-authority servers and keep them operating in a fair manner. Properties, such as existence and uniqueness, of this index are investigated for some typical network structures. By connecting to von Neumann's equilibrium concept, the proposed fairness concept can be related via a pricing duality to an equilibrium index, which uniquely exists in general. We also investigate the problem of how distributed users can achieve a given set of target indices. A distributed, low data rate control algorithm is introduced and its convergence property is discussed.

25.1 Introduction

The altruistic spirit of routing sharing is a key factor that contributes to the rapid growth of Internet. Nodes in an open architecture network such as the Internet are assumed to participate in the routing of third party traffic whenever the demand is within their resource capacity limit. However, the hierarchical architecture of Internet dictates that it cannot be as flat and altruistic as one may imagine. Only nodes of similar caliber can peer with each other, sometimes through Bi-Lateral Peering Agreements (BLPA). Smaller nodes have to aggregate their traffic via backbone nodes. This kind of arrangement can be interpreted as a scheme to ensure fairness in routing contribution. However, multiple servers could participate in a fair peering without satisfying pair-wise symmetric conditions. The fairness issue is even more critical for wireless ad-hoc networks [1]. Unlike hierarchical networks, these networks are self-organizing and totally distributed in administration. Failure to address fairness issue could lead to congestions that may prove to be impossible to recover from. Moreover, unlike servers in a wired network, battery power is a critical resource for mobile servers. Routing for third party traffic demands

*The work reported in this paper was partially supported by a Direct Grant of the Chinese University of Hong Kong.

W.S. Wong and F. Li: *A Fairness Index for Communication Networks*, LNCIS **321**, 389–400 (2005)
www.springerlink.com

not only communication channel resources but also consumes battery power. In [2]fairness issues in such a setting are investigated. In this paper, we report some of the key results in [2].

We present here a new fairness concept from [2]. An early version of this concept was first proposed in [3]. We begin by defining concepts of a *fairness index* and a *perfectly fair* solution in Sect. 25.2. The fairness index of a node is an indicator of its contribution to network routing. The ideal case where all the nodes have identical fairness indices is an important baseline reference point, and the corresponding network is called perfectly fair. One can connect the perfectly fair index with the von Neumann equilibrium by introducing a duality concept of pricing. The von Neumann equilibrium always exists for a network.

For implementation considerations, one needs to consider the problem how nodes in the network can achieve a perfectly fair state. Since each node can have only a local view of the network status, it is important that the desired state should be achieved through distributed algorithms. The problem is further complicated by the fact that in such a network, system parameters are basically unknown. In Sect. 25.3, a simple distributed controller is proposed to achieve a given set of fairness targets. A main result is to establish the convergence of such a class of distributed controllers. The algorithm can be extended to provide a heuristic solution to achieve a perfectly fair state without knowing the value of the perfectly fair index beforehand.

25.2 Fairness Index and Neumann Equilibrium

Nodes in a network tend to send as much traffic as possible into networks in order to achieve their maximum throughput unless being regulated. This behavior could lead to severe network congestion and unfairness in common resource usage. Many investigations use pricing and game theory to find efficient or fair operating states. By realizing the max-min fairness defined in [4] or the proportional fairness defined in [5] and [6], network link resources can be used efficiently. On the other hand, network traffic routing consumes resources at nodes as well as at links. For nodes peering in a network, it makes sense to demand that they are contributing to the routing function in a fair manner. We propose to use a *fairness index* for each node in a peer-to-peer network as a measure of whether routing load is shared fairly.

We conceptualize a communication network as a graph (V, E), with the nodes representing peering servers and an edge connects two nodes whenever there is a direct, duplex data link between two servers. Each node generates and consumes traffic data; it also routes traffic on behalf of other nodes.

Label the nodes in (V, E) from 1 to K. For each j, represent by r_j the rate of total traffic generated from node j to the network. Denote the vector of originating traffic rates, $(r_1, ..., r_K)^T$ by $\mathbf{r}$. A node, j, controls the network by means of r_j.

For any $i \neq j$, let $M_{ij}r_j$ represent the rate of the traffic generated from node j that is ultimately destined for node i. For all i, define

$$M_{ii} = 0. \tag{25.1}$$

Denote by $\mathbf{M}$ the K-by-K traffic distribution matrix with non-negative entries and column sums equal to 1,

$$\begin{pmatrix} 0 & M_{12} & \cdots & M_{1K} \\ M_{21} & 0 & \cdots & M_{2K} \\ \ldots & \ldots & \ldots & \ldots \\ M_{K1} & M_{K2} & \cdots & 0 \end{pmatrix} \tag{25.2}$$

The traffic from any source-destination pair can be routed over a variety of paths. We assume that the distributions of traffic into these alternative paths are arbitrary but known a priori and remain unchanged in a sufficiently long enough period for the consideration of this problem. Given a traffic distribution matrix and a set of routing schemes, the distribution of transitory traffic among intermediate nodes in the networks is fixed accordingly. Let $L_{ij}r_j$ denote the traffic rate of transitory traffic passing through node i that originates from node j, and $\mathbf{L}$=(L_{ij}) denote the corresponding K-by-K matrix. $\mathbf{L}$ is also a non-negative matrix with all entries bounded by 1. Moreover, $\mathbf{Lr}$ is the column vector representing the total transitory data traffic rates passing through each node.

A *fairness index* for each node is, roughly speaking, the ratio of traffic directly attributed to the node as a source or destination to the amount of total traffic it handles. Depending on the economical model one adopts to account for the utility or the revenue from the traffic, three classes of fairness indices are defined below, all taking values between 0 and 1.

Definition 1: For a network, (V, E), suppose $r_i+\sum_{j=1}^{K}(M_{ij}+L_{ij})r_j > 0$, then the *source-weighted fairness index* for node i is defined to be the ratio

$$l_i = \frac{r_i}{r_i + \sum_{j=1}^{K}(M_{ij} + L_{ij})r_j}. \tag{25.3}$$

Otherwise, the index is defined to be 0.

Suppose $r_i+\sum_{j=1}^{K}(M_{ij}+L_{ij})r_j > 0$, then the *destination-weighted fairness index* for node i is defined to be the ratio

$$o_i = \frac{\sum_{j=1}^{K} M_{ij}r_j}{r_i + \sum_{j=1}^{K}(M_{ij} + L_{ij})r_j}. \tag{25.4}$$

Otherwise, the index is defined to be 0.

Suppose $r_i + \sum_{j=1}^{K}(M_{ij} + L_{ij})r_j > 0$, then the *combined fairness index* or simply the *fairness index* for node i is defined to be the ratio

$$\chi_i = \frac{r_i + \sum_{j=1}^{K} M_{ij}r_j}{r_i + \sum_{j=1}^{K}(M_{ij} + L_{ij})r_j}. \tag{25.5}$$

Otherwise, the index is defined to be 0.

The denominator of these indices accounts for the total traffic handled by node i. For the source-weighted fairness index, the numerator accounts for the total data rate generated by node i to the network; traffic received by a node as final destination is assumed to have no economical benefit to it. One can interpret the meaning of the numerator term for the other indices accordingly.

Definition 2: A *perfectly source-fair* solution exists if there is a non-negative, non-zero rate vector $\mathbf{r}=(r_1, ..., r_K)^T$ so that the source-weighted fairness indices for all nodes are equal.

The concepts of a *perfectly destination-fair* and *perfectly fair* solution are defined similarly.

A perfectly fair solution exists if and only if there is a positive γ_C and a non-negative, non-zero rate vector so that

$$\gamma_C(\mathbf{I}+\mathbf{L}+\mathbf{M})\mathbf{r} = (\mathbf{I}+\mathbf{M})\mathbf{r}, \tag{25.6}$$

where $\mathbf{I}$ stands for the K-by-K identity matrix. One can establish similar equations for *perfectly source-fair* or *perfectly destination-fair* solutions:

$$\gamma_S(\mathbf{I}+\mathbf{L}+\mathbf{M})\mathbf{r} = \mathbf{r}, \tag{25.7}$$

$$\gamma_D(\mathbf{I}+\mathbf{L}+\mathbf{M})\mathbf{r} = \mathbf{M}\mathbf{r}. \tag{25.8}$$

Notice that perfectly fair solutions can be scaled uniformly without affecting their fairness properties. Hence with a proper scalar, a given set of link and node capacity constraints can always be satisfied.

The existence and uniqueness property of a perfectly fair solution is a natural question for investigation. For source-fair solution, the question can be settled by using the Perron-Frobenius Theorem on irreducible non-negative matrices [7]. The case for the other two types of indices is much more complicated. Mathematically, the issue hinges on finding generalized non-negative eigenvectors for a pair of nonnegative matrices. However, very little results have been reported in the literature on this subject. In [2], some existence and uniqueness properties of these solutions are reported. Furthermore, specific characterizations are found for networks with special topology.

Examples show that the existence of a perfectly fair solution depends on system parameters. The non-existence of a perfectly fair solution can be interpreted as an indication that nodes in the network are not suitable candidates for peering agreement. An alternative is to bring in a concept of pricing to compensate for the lack of perfect fairness. This latter approach makes contact with an equilibrium concept proposed by von Neumann [8].

A non-negative, non-zero price vector, $\mathbf{p}=(p_1, ..., p_K)$, is introduced for a network as a dual to the rate vector. The interpretation of the price vector is that a node can charge according to the traffic it handles, including traffic it generates, it receives as destination, and traffic it routes for other nodes.

Denote (**I**+**L**+**M**) by **B** and let **A** represent **I**, **M**, or **I**+**M**, depending on what type of fairness index is being considered. Introduce $1/\beta$ as the *minimum fairness ratio* so that for each node i,

$$\beta \sum_{j=1}^{K} a_{ij} r_j \geq \sum_{j=1}^{K} b_{ij} r_j. \tag{25.9}$$

Similarly, for each node, j, define a *maximum payment return ratio*, $1/\alpha$, so that for each node j,

$$\alpha \sum_{i=1}^{K} a_{ij} p_i \leq \sum_{i=1}^{K} b_{ij} p_i. \tag{25.10}$$

For a more concrete discussion, assume that both source and destination traffic have economic value, so that perfectly fair index is considered and **B** represents **I**+**M**. The discussion can be extended to the other two types of indices as well.

From a fairness consideration, each node, i, sends and receives data traffic of interest to it, amounting to $\sum_{j=1}^{K} a_{ij} r_j = r_i + \sum_{j=1}^{K} M_{ij} r_j$, while it handles the total amount $\sum_{j=1}^{K} b_{ij} r_j = r_i + \sum_{j=1}^{K} (M_{ij} + L_{ij}) r_j$ for the network. Normalized by the amount of useful traffic it receives, the contribution to a network by a node is maximal if equality holds in (25.9) for that node. A strict inequality implies the node enjoys more benefit from the network then nodes with equality.

For traffic originating from node j, (with rate r_j), a total charge of $\sum_{i=1}^{K} b_{ij} p_i = p_j + \sum_{i=1}^{K} (M_{ij} + L_{ij}) p_i$ per unit rate is incurred by the network for handling it. On the other hand $\sum_{i=1}^{K} a_{ij} p_i = p_j + \sum_{i=1}^{K} M_{ij} p_i$ can be viewed as the economic benefit generated by the traffic flow to the network. (One can assume the originating and destination nodes can charge end-users for traffic handling at a rate that is proportional to the shadow prices.) Hence, if equality holds in equation (25.10) for a node, traffic from that node can be viewed as enjoying the most efficient payment return ratio in the network; otherwise the payment return ratio is not optimal.

It is natural to define a rule requiring that nodes enjoying more routing benefit from the network than others should not require charges for traffic handling. Similarly, traffic rate from nodes that do not enjoy the optimal payment return ratio should be set to zero. Hence, one can adopt von Neumann's concept of an equilibrium solution here. That is an equilibrium solution to (25.9) and (25.10) is a pair of non-negative, non-zero vectors, **r** and **p**, satisfying both set of equations for some positive constants, α and β, with the additional property that if an inequality holds for index i in (25.9) then $p_i = 0$ and if an inequality holds for index j in (25.10) then $r_j = 0$. That is:

$$\begin{cases} \alpha \sum_{i=1}^{K} a_{ij} p_i \leq \sum_{i=1}^{K} b_{ij} p_i, & \text{and} \quad r_j = 0 \quad \text{if } < \text{ applies;} \\ \beta \sum_{j=1}^{K} a_{ij} r_j \geq \sum_{j=1}^{K} b_{ij} r_j, & \text{and} \quad p_i = 0 \quad \text{if } > \text{ applies;} \end{cases} \tag{25.11}$$

From Neumann's result (see [8]), one can claim that there always exists a unique equilibrium value γ for the model discussed here so that $\alpha = \beta = \gamma$. If a perfectly fair index also exists for the system, then it is equal to or less than $1/\gamma$.

25.3 Distributed Controller and Its Convergence

Based on the approach first proposed in [9] and further extended in [10], we present here a tri-state distributed control algorithm that can achieve a pre-defined fairness index target. This algorithm assumes that each node only has a local view of the network: it knows the traffic rate it generates, the total rate of traffic ultimately destined to it, and the total rate of transitory traffic passing through it. Note that none of the nodes is assumed to know any global system-wide parameters. Each node computes its fairness index individually, and updates its originating traffic rate according to the following rule:

$$r_i(n+1) = \begin{cases} r_i(n)\delta & \text{if} \quad \gamma_i(n) > \varepsilon\lambda_i, \\ r_i(n)\delta^{-1} & \text{if} \quad \gamma_i(n) < \varepsilon^{-1}\lambda_i, \\ r_i(n) & \text{else.} \end{cases} \tag{25.12}$$

In the above algorithm, δ and ε are scalar parameters which control the speed of convergence. They should satisfy the conditions $\delta > 1$ and $\delta^2 \leq \varepsilon$. Moreover, $\{\lambda_1, ...\lambda_K\}$ denotes a set of feasible performance targets and $\gamma_i(n)$ represents the fairness index for node i at iteration n. For example, if one considers the destination-weighted fairness index, then $\gamma_i(n)$ is defined as:

$$\gamma_i(n) = \frac{\sum_{j=1}^{K} M_{ij} r_j(n)}{r_i(n) + \sum_{j=1}^{K} (M_{ij} + L_{ij}) r_j(n)}. \tag{25.13}$$

The algorithm can start from any positive initial state $(r_1(0), ..., r_K(0))$. The traffic rates, $(r_1(n), ..., r_K(n))$, are adjusted according to 25.12 at each iteration. Therefore, for any server i the control levels are of the form $r_i(0)\delta^k$ for some integer k. It is shown in [2] that under suitable technical conditions, the simple algorithm defined by 25.12 always converges to $(r_1^*, ...r_K^*)$ and $\{\gamma_i^*\}$ such that

$$\gamma_i^* = \frac{\sum_{j=1}^{K} M_{ij} r_j^*}{r_i^* + \sum_{j=1}^{K} (M_{ij} + L_{ij}) r_j^*}, \quad \text{and} \quad \varepsilon^{-1}\lambda_i \leq \gamma_i^* \leq \varepsilon\lambda_i. \tag{25.14}$$

Numerical computations were carried out to examine the convergence property of the tri-state algorithm under different network topologies, system parameters, initial values, scalar parameters, etc. Some typical examples are shown below.

First consider a 6-node network with topology shown in Figure 25.1.

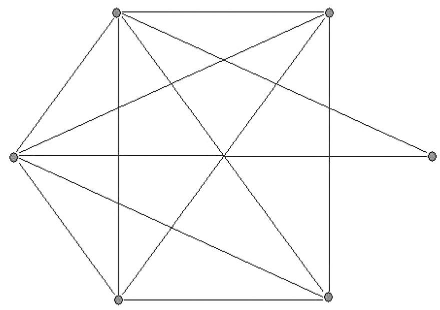

Fig. 25.1. Topology of a 6-node network.

Assume the traffic distribution, $\mathbf{M}$, and the transitory matrix, $\mathbf{L}$, are set as follows:

$$\mathbf{M} = \begin{bmatrix} 0 & 0.3433 & 0.1622 & 0.2240 & 0.2740 & 0.0724 \\ 0.1991 & 0 & 0.1615 & 0.2006 & 0.2915 & 0.1893 \\ 0.1772 & 0.2063 & 0 & 0.0626 & 0.2072 & 0.2564 \\ 0.1276 & 0.3604 & 0.4192 & 0 & 0.1923 & 0.1229 \\ 0.1489 & 0.0605 & 0.1479 & 0.1900 & 0 & 0.3589 \\ 0.3473 & 0.0295 & 0.1093 & 0.3229 & 0.0350 & 0 \end{bmatrix} \tag{25.15}$$

$$\mathbf{L} = \begin{bmatrix} 0 & 0.4332 & 0.5532 & 0.5270 & 0.4876 & 0.6196 \\ 0.6417 & 0 & 0.6585 & 0.6401 & 0.5477 & 0.7089 \\ 0.6584 & 0.6100 & 0 & 0.7452 & 0.6119 & 0.6458 \\ 0.5912 & 0.4218 & 0.3820 & 0 & 0.5433 & 0.5849 \\ 0.5766 & 0.6217 & 0.5628 & 0.5502 & 0 & 0.4227 \\ 0.1650 & 0.1941 & 0.1831 & 0.1756 & 0.2459 & 0 \end{bmatrix} \tag{25.16}$$

For these parameters, there exists a perfect destination-weighted index, λ, with value 0.2159. Set δ equal to 1.001, $\varepsilon = \delta^2$ and set the initial traffic rate to be $[1, 1, 1, 1, 1, 1]^T$. Under the tri-state distributed control algorithm, the fairness indices of the nodes converge to γ^* after 928 iterations, where $\gamma^* = [0.2162, 0.2163, 0.2155, 0.2155, 0.2162, 0.2155]^T$. The trajectories of fairness indices of the six nodes are shown in Figure 25.2.

As stated previously, the tri-state algorithm can also be used to achieve feasible performance targets other than perfectly fair solutions. As an example, consider again the previous network. Let a set of performance targets, $\{\lambda_i, 1 \leq i \leq K\}$, be specified by the vector,

$$\lambda = [0.2462, 0.1958, 0.2186, 0.2896, 0.2054, 0.1290]^T. \tag{25.17}$$

This set of targets can be attained by setting the traffic rates to $\mathbf{r}_\lambda = [0.2, 0.5, 0.3, 0.1, 0.6, 0.8]^T$. Starting from an arbitrarily chosen initial state, the fairness index trajectories of all the nodes converge to this feasible target value under the tri-state distributed algorithm as shown in Figure 25.3.

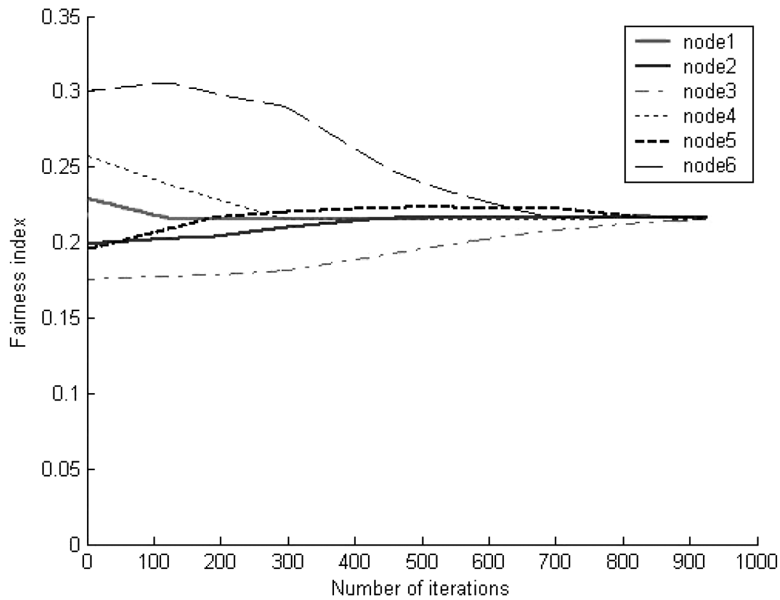

Fig. 25.2. Convergence to a perfectly destination-fair solution for a 6-node network.

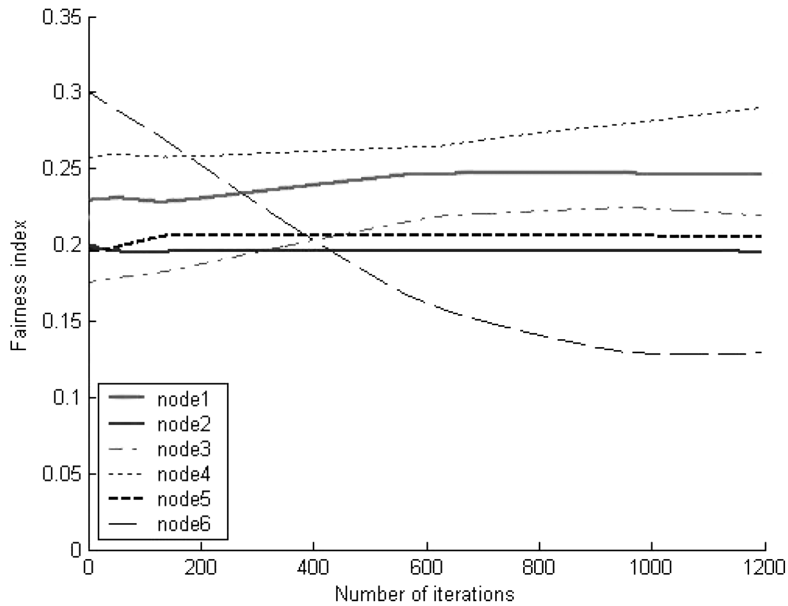

Fig. 25.3. Feasible performance target convergence for a 6-node network.

Consider now the 7-node network shown in Figure 25.4. By choosing the matrices **M** and **L** properly, one can find cases where the fairness index exists. Simulation results show that for a feasible set of performance targets, the tri-state algorithm always converges to the targeted performance band. Some simulation results for this network are listed in Table 25.1. (In these simulations, $\delta = 1.001$ and $\varepsilon = \delta^2$).

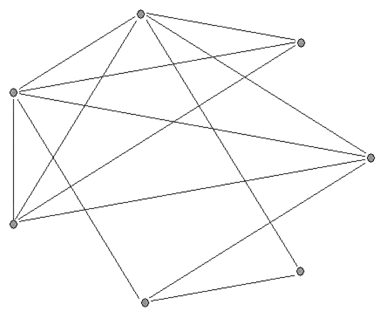

Fig. 25.4. Topology of a 7-node network.

Table 25.1. Convergence results of a 7-node network.

Target	0.1912	0.1986	0.1953	0.1979	0.1921
	0.1909	0.1983	0.1949	0.1982	0.1919
	0.1913	0.1984	0.1949	0.1979	0.1925
Converged	0.1916	0.1990	0.1956	0.1982	0.1925
	0.1908	0.1990	0.1956	0.1975	0.1919
Indexes	0.1909	0.1983	0.1949	0.1982	0.1917
	0.1916	0.1983	0.1949	0.1975	0.1922
	0.1916	0.1990	0.1956	0.1975	0.1917
No. of Iterations	1348	1305	2953	615	1223

The convergence process is affected by parameters such as the initial traffics rate and the scalars, δ and ε. Choosing these values correctly can speed up the convergence. Simulation runs were conducted for the previous 7-node network example. For one set of runs, the initial traffic rates were chosen independently from an identical distribution. For a second set of runs, the initial rates were chosen to be identically for all the nodes. Simulation results show that the algorithm converges faster for runs starting with identical initial rates, if other conditions are set identically. Effects of the adjusting the parameters, δ and ε, were also studied. In general, the lager these parameters are, the faster the convergence rate.

Table 25.2. Convergence rate under different conditions.

Target	δ	ε	Initial state chosen	No. of Iterations
0.1921	1.001	1.002001	Identically	1223
			Independently	1954
0.1937	1.001	1.002001	Identically	4949
			Independently	5602
0.1921	1.01	1.03	Identically	89
			Independently	243
0.1937	1.01	1.03	Identically	389
			Independently	536

From the table, the tri-state control algorithm appears to converge reasonably fast. Nevertheless, a weakness for the algorithm is that it requires the feasible performance target be known beforehand and any invalid target value will result in a process that never converges. It is difficult to determine a priori whether a set of performance targets is feasible or not in a large-scale network. However, a heuristic algorithm exists for reaching a perfectly fair solution, even if the index value is not known a priori. The basic idea is to apply the previous distributed algorithm using an estimated index value and updates the target value if the algorithm does not converge. The target update procedure implies that the algorithm is only partially distributed, since a central server is needed to exchange and updated target value to all the nodes. The heuristic algorithm is defined below:

1. Set proper initial values: Set δ and ε, $1 < \delta^2 \leq \varepsilon$, with relatively large values; set initial target value, γ_0, and initial rate vector, $\mathbf{r}_0$.
2. Perform the tri-state algorithm until it converges or remains unchanged at a value outside the target band after a maximal iteration number, κ.
3. If the maximum iteration number is reached, the nodes send their current index value to the central server. The mean of the current index values is set as the target value for the next turn.
4. If the tri-state algorithm stops due to convergence to currently set target zone, the error control parameters are decreased. (One approach is to set the new value to the square root of the current value.)
5. Re-do the target tracking process using the tri-state algorithm until δ and ε are decreased to predefined acceptable value δ_0 and ε_0 respectively.

Numerical computations were carried out to examine the convergence behavior of the heuristic algorithm. A 5-node network with a fairness index value of 0.2462 was studied. Key convergence stages of the heuristic algorithm are

Table 25.3. Convergence rate of the heuristic algorithm.

Target	δ	ε	Lower bound	Upper bound	Status	No. of Iterations
0.3	1.7320	3.0000	0.1	0.9	Converged	1
0.3	1.3160	1.7320	0.17321	0.5196	Converged	1
0.3	1.1471	1.3160	0.2280	0.3948	Converged	6
0.3	1.0710	1.1471	0.2615	0.3441	Not converged	281
0.2460	1.0710	1.1471	0.2145	0.2822	Converged	4
0.2460	1.0348	1.0710	0.2297	0.2635	Converged	11
..	...	...	...	...	...	...
0.2460	1.0005	1.0011	0.2457	0.2463	Not converged	761
0.2462	1.0002	1.0005	0.2459	0.2461	Converged	1499

shown in Table 25.3. The target value was initially set to be 0.3. It took 11 outer loops for algorithm to converge to the perfectly fair index value.

Conclusion

A new fairness index for communication networks is introduced in this paper. This index can be used as a reference point to measure the contribution of independent, self-authority servers to the networks. Some initial properties of this index are presented here. However, much remains to be done. Issues such as how the index can be applied to practical network control and traffic pricing are of great interest.

Acknowledgment

The authors would like to thank Professor Dah Ming Chiu for helpful discussions on this work. In particular, the discussions helped us reformulate the definition of fairness index from the form stated in [3] to the form stated here.

References

1. Andy Oram. *Peer-to-Peer: Harnessing the Power of Disruptive Technologies.* O'Reilly & Associates, 2001.
2. Wing Shing Wong and Fengjun Li. Fairness index and distributed control in communication networks. Preprint, 2004.
3. Wing Shing Wong. Fairness issues in peer-to-peer networks. In *Proceedings of the IEEE 2002 Conference on Decision and Control*, Las Vegas, Nevada, Dec 2002.
4. Jean-Yves Le Boudec. Rate adaptation, congestion control and fairness: a tutorial. Avaliable: http://ica1www.epfl.ch/PS_files/LEB3132.pdf.

5. Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, Jan 1997.
6. Richard J. Gibbens and Frank Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35:1969–1985, 1999.
7. E. Seneta. *Non-negative Matrices and Markov Chain*. Springer-Verlag, New York, second edition, 1981.
8. John von Neumann. A model of general economic equilibrium. *The Review of Economic Studies*, 13(1):1–9, 1945-1946.
9. Chi Wan Sung and Wing Shing Wong. A distributed fixed-step power control algorithm with quantization and active link quality protection. *IEEE Transactions on Vehicular Technology*, 48(2):553–562, 1999.
10. Wing Shing Wong and Chi Wan Sung. Robust convergence of low data rate distributed controllers. *IEEE Transactions on Automatic Control*, 49(1):82–87, 2004.