Utility Max-Min Fair Flow Control for Multipath Communication Networks

Jiong Jin, Wei-Hua Wang and Marimuthu Palaniswami
Department of Electrical and Electronic Engineering
The University of Melbourne
Victoria 3010, Australia
E-mail: {j.jin, weihw, swami}@ee.unimelb.edu.au

***Abstract*— This paper considers flow control and resource allocation problem as applied to multipath communication networks. We propose a novel distributed algorithm, show and prove that among all the sources with generic increasing and bounded utilities (no need to be concave) in steady state, the utility max-min fairness is achieved, which is essential in providing application QoS (Quality of Service) guarantee. In addition, by combining a first order Lagrangian method and filtering mechanism, the resulted approach eliminates typical oscillating behavior for the multipath network and possesses a rapid convergence property.**

I. INTRODUCTION

Current communication network, like the prevailing Internet, has made a great success in providing efficient data transmission services, e.g., web browsing and electronic mail, but it is not sufficient to support the increasing demand on real-time services, such as audio, video and multimedia delivery through the network. These real-time applications usually have stringent Quality of Service (QoS) requirements, and are sensitive with allocated bandwidth, time delay and packet loss ratio, which are generally not easy to be guaranteed in the TCP based Internet services nowadays. Therefore, future communication networks are expected to support applications with various QoS requirements.

To provide a better traffic management in computer networks than the traditional TCP does, an extensive study has been carried out in the literature. Among them, the most successful result in the area of network congestion control and resource allocation is the "Optimal Flow Control"(OFC) approach proposed by Kelly [1]. This pioneer work was further advanced by the researches in single path networks [2]–[8], multipath networks [9], [10] and multirate multicast networks [11], [12].

The main idea of OFC is essentially the same to formulate flow control as an optimization problem and then maximize the total utilities under the network bandwidth constraint. The utility function of the bandwidth associated with each application mathematically models its QoS performance. Following that, OFC algorithm is derived by solving the optimization problem distributively. It consists of a link algorithm to measure the congestion (link price) in the network and a source algorithm to adapt the transmission rate according to the feedback congestion signals. This optimization approach not only leads to a social utility maximization at the convergence, what is more important, also the resulting bandwidth allocation in equilibrium is in a fair manner.

With a popularity to select utility as logarithmic function, Kelly [1] shows that the OFC approach achieves a *proportional fairness* of bandwidth allocation. Using the OFC strategy, another important fairness criterion called *max-min fair* allocation [13] (which emphasizes an equal sharing compared with *proportional fairness*) is also studied by Mo and Walrand [4] and La and Anatharam [6]. In their work, the authors use a family of utility functions to approximate arbitrarily close to a max-min fair allocation. But the selected utility function becomes *ill conditioned* when the max-min fairness is reached, and the related link prices at congested links either turn to 0 or diverge to ∞. Thus their max-min fair flow control algorithms are impractical for engineering purpose. Meanwhile, in order to deal with different users with different QoS requirements, Cao and Zegura [14] define a new criterion named *utility max-min fairness* and propose an allocation algorithm. In their approach, the links require the information of utility functions from all the traversed sources, which makes network implementation difficult.

Even though the optimal flow control approach has made a great success in dealing with congestion control and resource allocation, it also exposes serious limitations as pointed out in our paper [15].

- At current stage, OFC approach is only suitable for *elastic* traffic, where each application attains a strictly increasing and concave utility function to ensure the feasible optimal solution and convergence of utility maximization process. It can not deal with congestion control and resource allocation for communication networks where real-time applications are engaged.
- In the utility maximization approach, if each user selects different utility function based on their real performance requirement, then OFC approach usually leads to a totally unfair resource allocation for practical use, in particular, an application with low demand is usually allocated with a high bandwidth.

On the other hand, multipath communication networks attract significant attention recently due to scalability and robustness. With the help of MPLS technology, even the most common IP networks, which more or less require single path routing previously, enable the traffic to split across several paths. For these reasons, in this paper, we propose a novel distributed flow control algorithm for multipath communication networks to achieve utility max-min fair resource allocation. By taking into account of different QoS requirements, the new flow control algorithm is friendly with both elastic traffic and real-time applications.

The rest of the paper is organized as follows. In Section II, we describe and formulate the problem. Section III proposes the utility max-min fair flow control algorithm. After that, we present the simulation results to illustrate performance of the algorithm in Section IV. Finally, the conclusions are drawn in Section V.

II. PROBLEM FORMULATION

For a practical network application, people may concern about the bandwidth allocation, but a more important and direct factor that the application really cares about is the QoS performance or the utility it achieves in the network. The utility function of an application is a measurement of its QoS performance based on provided network services such as bandwidth, transmission delay and loss ratio. In this paper, we deal with the utility as a function of the allocated bandwidth only, which is a common assumption in most optimal flow control literatures.

As pointed out in the paper [16], the traditional data applications such as file transfer, electronic mail, and web browsing are rather tolerant of throughput and time-delays. This class of applications are called *elastic* applications, and their utility functions can be described as a strictly concave function as shown in Fig. 1(a). The utility (performance) increases as the increasing of bandwidth, but the marginal improvement is decreased. This class of applications has been well studied in OFC literatures.

Nowadays due to the development of multimedia technologies, real-time applications, such as audio and video delivery, become ubiquitous. These applications are generally delay sensitive and have a strict Quality of Services (QoS) requirement. Unlike the *elastic* traffic, they have an intrinsic bandwidth requirement because the data generation rate is independent of the network congestion. Thus the degradation in bandwidth may result in serious packet drops and severe degradation of the performance. A reasonable description of the utility of this class applications is close to a single step function as shown in Fig. 1(b) (solid line), which is convex but not concave at the lower bandwidths. For some hard real-time applications, they may require an exact step utility function as in Fig. 1(b) (dash line).

In this context, we consider a network that consists of a set $\mathcal{L} = \{1, 2, \dots, L\}$ of links of capacity c_l, $l \in \mathcal{L}$. The network is shared by a set $\mathcal{S} = \{1, 2, \dots, S\}$ of sources. Each source s attains a non-negative QoS utility $U_s(x_s)$ when it transmits at a rate $x_s \in [m_s, M_s]$ where m_s and M_s are the minimum and maximum transmission rates required by

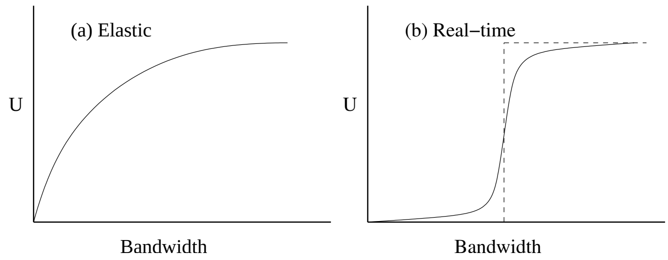

Fig. 1. Utility functions for different classes of applications.

source s respectively. The utility function $U_s(x_s)$ is assumed to be continuous, strictly increasing and bounded (no need to be concave) in the interval $[m_s, M_s]$. Without the loss of generality, it can be assumed that $U_s(x_s) = 0$ when $x_s < m_s$ and $U_s(x_s) = U_s(M_s)$ when $x_s > M_S$.[1]

Different from single path network, now each source s has n_s available paths or routes from the source to the destination. Denote the $L \times 1$ vector $R_{s,i}$ the set of links used by source $s \in S$ for its path $i \in \{1, 2, \dots, n_s\}$, whose $l\mathrm{th}$ element is equal to 1 if and only if the path passes through link l, and 0 otherwise. Then the set of all the available paths of user s is defined by

$$R_s = [R_{s,1}, R_{s,2}, \dots, R_{s,n_s}]$$

and the total paths in the network are defined by a $L \times N$ routing matrix R,

$$R = [R_1, R_2, \dots, R_S]$$

where $N = n_1 + n_2 + \dots + n_S$ is the total number of the paths.

For each source s, define $x_{s,i}$ be the rate of source s on path $R_{s,i}$, and naturally the total source rate $x_s = \sum_{i=1}^{n_s} x_{s,i}$. Let

$$x = [x_{1,1}, \dots, x_{1,n_1}, x_{2,1}, \dots, x_{2,n2}, \dots, x_{n,1}, \dots, x_{n,n_S}]^{\mathrm{T}} \in \mathcal{R}_+^N$$

be the vector of all path rates of all sources. In order to formulate the flow control problem, we first define the notion of feasible (or attainable) path rate allocation.

Definition 1: A path rate allocation x for all available paths is *feasible* or *attainable* if and only if the corresponding total source rate x_s for each source s is within the range $[m_s, M_s]$, and no links in the network are congested, i.e.:

$$m_s \le x_s \le M_s, \qquad x_s = \sum_{i=1}^{n_s} x_{s,i}, \quad s \in \mathcal{S} \tag{1}$$

$$Rx \le c, \qquad x \ge 0 \tag{2}$$

where $c = [c_1, c_2, \dots, c_L]^{\mathrm{T}}$ is the vector of link capacities.

Since we are considering the different QoS requirements among network users, it may not be appreciative for the network to simply share the bandwidth equally as max-min fairness does. Instead, the network should allocate the bandwidth to the competing users according to their different

[1]For the scalability, it can be further assumed that $0 \le U_s(x_s) \le 1$ and $U_s(M_s) = 1$.

QoS utilities. This motivates the proposal of the criterion of *utility max-min fairness* [14].

Definition 2: A source rate allocation is utility max-min fair, if it is feasible and for each user s, its utility $U_s(x_s)$ cannot be increased while maintaining feasibility, without decreasing the utility $U_{s'}(x_{s'})$ for some user s' with a lower utility $U_{s'}(x_{s'}) \leq U_s(x_s)$.

It is even more complicated in the environment of multipath networks, where the source rate is made up of constituted path rates. Our objective is to guide traffic to a feasible path rate allocation, in such a way that the summing source rate is utility max-min fair. In other words, each source is treated in a fair manner and guaranteed high utility performance. In the following section, we will develop a new flow control algorithm to achieve utility max-min fairness within a given multipath network and study its properties in detail.

III. UTILITY MAX-MIN FAIR FLOW CONTROL ALGORITHM

Consider the flow control problem formulated in Section II. Now, we propose a distributed algorithm that achieves utility max-min fairness for multipath communication networks.

A. A Distributed Utility Based Flow Control Algorithm

The utility max-min fair flow control algorithm uses the similar flow control structure as the optimal flow control approach [3] does, with the help of pricing scheme. There are three price vectors $\alpha \in \mathcal{R}_+^S$, $\beta \in \mathcal{R}_+^S$ and $p \in \mathcal{R}_+^L$ associated with constraint (1) and (2) (Regard constraint (1) $m_s \leq x_s \leq M_s$ as two separated constraints $x_s \geq m_s$ and $x_s \leq M_s$) respectively. A link algorithm is deployed at each link to update the link price depending on the severity of link congestion, and a source algorithm is implemented at each source edge to adapt the transmission rate based on these three prices.

Both link algorithm and source algorithm are iterative. At time $t+1$, each link l updates its link price p_l according to:

$$p_l(t+1) = [p_l(t) + \gamma(x^l(t) - c_l)]^+ \tag{3}$$

where $\gamma > 0$ is a small step size, and $x^l(t) = R_l.x$ is the aggregate path rate at link l. Equation (3) implies that if the aggregate path rate at link l exceeds the link capacity c_l, the link price will be increased; otherwise it will be decreased. The projection $[z]^+ = \max\{0, z\}$ ensures that the link price is always non-negative.

For each source s, we use the following first-order Lagrangian algorithm to update its ith path rate:

$$x_{s,i}(t+1) = [x_{s,i}(t) + \gamma(\frac{1}{U_s(x_s(t))} + \alpha_s(t) - \beta_s(t) - p_{s,i}^r(t))]^+ \tag{4}$$

and then calculate the source rate:

$$x_s(t+1) = \sum_{i=1}^{n_s} x_{s,i}(t+1) \tag{5}$$

where

$$\alpha_s(t+1) = [\alpha_s(t) + \gamma(m_s - x_s(t))]^+ \tag{6}$$

$$\beta_s(t+1) = [\beta_s(t) - \gamma(M_s - x_s(t))]^+ \tag{7}$$

are the lower bound and upper bound price to restrict the source rate constraint $m_s \leq x_s \leq M_s$, and

$$p_{s,i}^r(t) = \max_{l \in R_{s,i}} p_l(t) \tag{8}$$

is the path price, which is the maximum of the link prices along the particular route. Combining them all together, the utility max-min fair flow control algorithm for multiple paths is summarized as follows:

Algorithm

- **Link l's algorithm:**
  At time $t = 1, 2, \ldots$, link l:
  1) Aggregates flow rates $x_{s,i}(t)$ for all paths $R_{s,i}$ that contain link l.
  2) Computes a new link price
  $$p_l(t+1) = [p_l(t) + \gamma(x^l(t) - c_l)]^+.$$
  3) Communicates new price $p_l(t+1)$ to all the sources whose path $R_{s,i}$ contains link l.
- **Source s's algorithm:**
  At time $t = 1, 2, \ldots$, source s:
  1) Receives from the network the path prices
  $$p_{s,i}^r(t) = \max_{l \in R_{s,i}} p_l(t)$$
  for all its paths $R_{s,i}$, $i = 1, 2, \ldots, n_s$.
  2) Updates the path rate $x_{s,i}(t+1)$ and the source rate $x_s(t+1)$ using Equation (4) and (5).
  3) Computes the new lower bound and upper bound price $\alpha(t+1)$ and $\beta(t+1)$ for the next step according to Equation (6) and (7).
  4) Communicates the new flow rate $x_{s,i}(t+1)$ to all the links which contained in paths $R_{s,i}$.

As we know, for multipath networks, the set of feasible path rates $x_{s,i}$ may not be unique, such that the first-order Lagrangian algorithm usually oscillates. In order to eliminate this undesirable effect and further improve the convergence speed, we introduce another augmented variable $\overline{x}_{s,i}$, called the optimal estimation of path rate $x_{s,i}$. Borrowing the concept of low-pass filtering, we slightly modify Equation (4) as

$$\begin{aligned} x_{s,i}(t+1) &= [(1-\gamma)x_{s,i}(t) + \gamma\overline{x}_{s,i}(t) \\ &\quad + \gamma(\frac{1}{U_s(x_s(t))} + \alpha_s(t) - \beta_s(t) - p_{s,i}^r(t))]^+ \\ \overline{x}_{s,i}(t+1) &= (1-\gamma)\overline{x}_{s,i}(t) + \gamma x_{s,i}(t). \end{aligned} \tag{9}$$

By applying the filtering theory, at optimality, $x_{s,i} = \overline{x}_{s,i}(t+1)$, so that the augmented variable is only used to remove the oscillation without changing the optimal solution of $x_{s,i}$. It is clearly figured out by the simulation in Section IV.

B. Utility Max-Min Fairness

Recalling *Definition 1*, the interested region of the source rate is $[m_s, M_s]$. The associated utility for the source rate outside this region is scaled to 0 or 1. In this scenario, the lower bound and upper bound price (α and β) are both equal to 0. The path rate algorithm of Equation (4) simplifies to

$$x_{s,i}(t+1) = [x_{s,i}(t) + \gamma(\frac{1}{U_s(x_s(t))} - p^r_{s,i}(t))]^+. \quad (10)$$

From Equation (10), it is observed that either $\frac{1}{U_s(x_s(t))} = p^r_{s,i}(t)$ or $x_{s,i}(t) = 0$ at convergence. If we define $p_s^{r^*} = \frac{1}{U_s(x_s(t))}$ for every source s, the latter case can be interpreted in another way, i.e., when the path price $p^r_{s,i}(t)$ is greater than $p_s^{r^*}$, this particular path is too "expensive" to carry any flow ($x_{s,i}(t) = 0$). The above fact establishes *Theorem 1*.

Theorem 1: For multipath communication networks, in steady state, the prices on paths $R_{s,i}$ that carry positive flows $x_{s,i} > 0$ must be minimum, and hence equal, among all the paths R_s of source s. Moreover, the optimal source rates are given by

$$x_s^* = \sum_{R^*_{s,i} \in R^*_s} x^*_{s,i} = U_s^{-1}\left(\left[\frac{1}{p_s^{r^*}}\right]^{U_s(M_s)}_{U_s(m_s)}\right)$$
$$\text{and} \quad x_{s,i} = 0 \quad \text{if} \;\; p^r_{s,i} > p_s^{r^*}$$

where $[z]^b_a = \max(a, \min(b, z))$, path $R^*_{s,i}$ has the minimum path price $p^{r^*}_{s,i} = p_s^{r^*}$, and R^*_s defines the set of all minimum price paths $R^*_{s,i}$ of source s.

From this theorem, it is clear that in steady state, the associated utility U_s of source s is equal to $\frac{1}{p_s^{r^*}}$ when $p_s^{r^*} \in [\frac{1}{U_s(M_s)}, \frac{1}{U_s(m_s)}]$, otherwise, it attains a utility $U_s(m_s)$ which is greater than $\frac{1}{p_s^{r^*}}$ for minimum rate requirement (which cannot be decreased anymore due to QoS requirement), or a utility $U_s(M_s)$ which is less than $\frac{1}{p_s^{r^*}}$ for maximum rate requirement (which needs not to be increased any further).

For this reason, here we only consider the resource allocation among the sources who attain a normal utility $U_s^* = \frac{1}{p_s^{r^*}}$.

Let $\mathcal{S}_l$ be the set of sources which have at least one path traversing link l. We first select the link with the highest link price in the network. Suppose it is link l_1 and its link price is p_{l_1}, then all the sources $s \in \mathcal{S}_{l_1}$ attain the same utility $U_s = 1/p_{l_1}$, which are the smallest allocated utilities compared with other sources. If there is a source $s \in \mathcal{S}_{l_1}$ that increases its utility U_s by increasing the transmission rate x_s, then there must be another source $s' \in \mathcal{S}_{l_1}$ to decrease its rate $x_{s'}$ and further decrease its utility $U_{s'}$ which is equal to U_s. This violates the law of utility max-min fairness.

Next, we select the link l_2 with the second highest link price p_{l_2}. Then all the sources $s \in \mathcal{S}_{l_2} \setminus \mathcal{S}_{l_1}$ have the same utility $U_s = 1/p_{l_2}$. If there is a source $s \in \mathcal{S}_{l_2} \setminus \mathcal{S}_{l_1}$ that increases its rate and utility, then there must be another source $s' \in \mathcal{S}_{l_2}$ to decrease its rate which already has a lower utility $U_{s'} \leq U_s$. This again violates the utility max-min fairness.

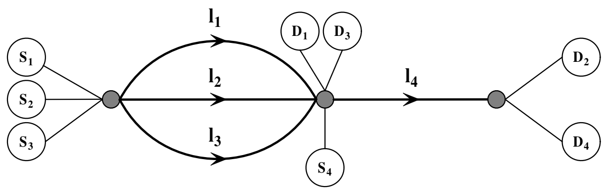

Fig. 2. The network topology

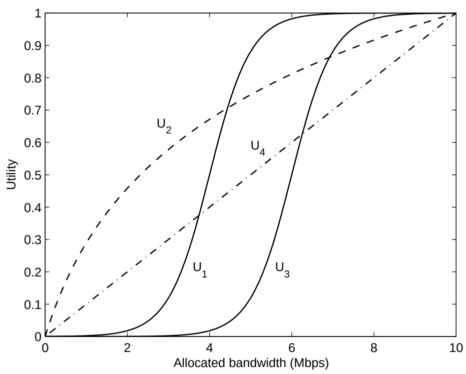

Fig. 3. Source utility functions

Continuing in this way, selecting all the links with positive link price, it is concluded that all the source allocation rates are utility max-min fair and global fairness is achieved.

IV. SIMULATION RESULTS

In this section, we evaluate through simulations the performance of our proposed utility max-min fair flow control algorithm for multipath communication networks. Fig. 2 depicts the topology of the network. It consists of four unidirectional links labelled l_1, l_2, l_3 and l_4 with capacities $c = (4, 6, 8, 10)$ in Mbps and shared by 4 sources S_1, S_2, S_3 and S_4. S_1 with a total rate of x_1 uses the paths: l_1 with rate $x_{1,1}$ and l_2 with rate $x_{1,2}$. S_2 with a total rate of x_2 uses the paths: $l_2 \to l_4$ with rate $x_{2,1}$ and $l_3 \to l_4$ with rate $x_{2,2}$. S_3 with a total rate of x_3 uses the paths: l_3 with rate $x_{3,1}$ and l_1 with rate $x_{3,2}$. S_4 with a rate of x_4 uses a single path: l_4 with rate $x_{4,1}$ i.e. $x_4 = x_{4,1}$.

The utility function of each source is given as: $U_1(x_1) = 1/(1 + e^{-2(x_1-4)})$, $U_2(x_2) = \log(x_2 + 1)/\log 11$, $U_3(x_3) = 1/(1 + e^{-2(x_3-6)})$ and $U_4(x_4) = 0.1x_4$. All the sources have their maximum rate requirement at 10 Mbps. Fig. 3 illustrates these utility functions. The logarithmic utility function represents an elastic data flow application such as FTP whereas the sigmoidal function approximates the real-time application. The linear utility function corresponds to the application whose satisfaction increases linearly.

In the simulation, we run the original algorithm with $\gamma = 0.2$. The simulation results are given in Fig. 4. As expected, the oscillation is observed, which motives the modification replacing Equation (4) with Equation (9) in the algorithm.

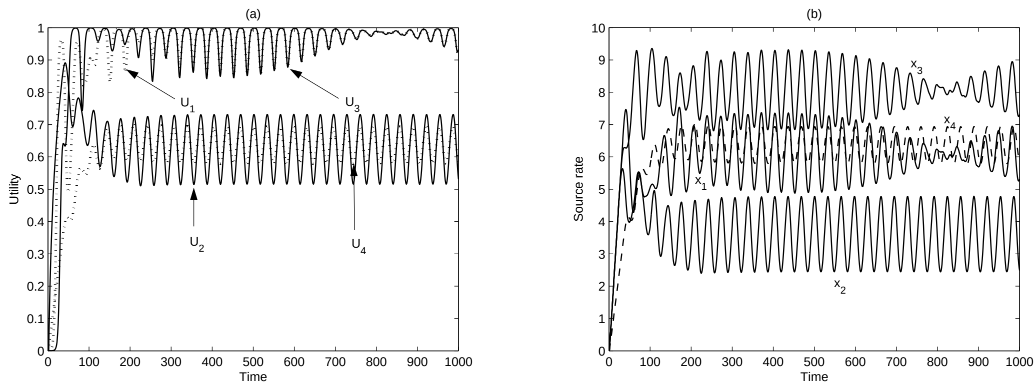

Fig. 4. Simulation results of original flow control algorithm (a) Source utilities (b) Source rates

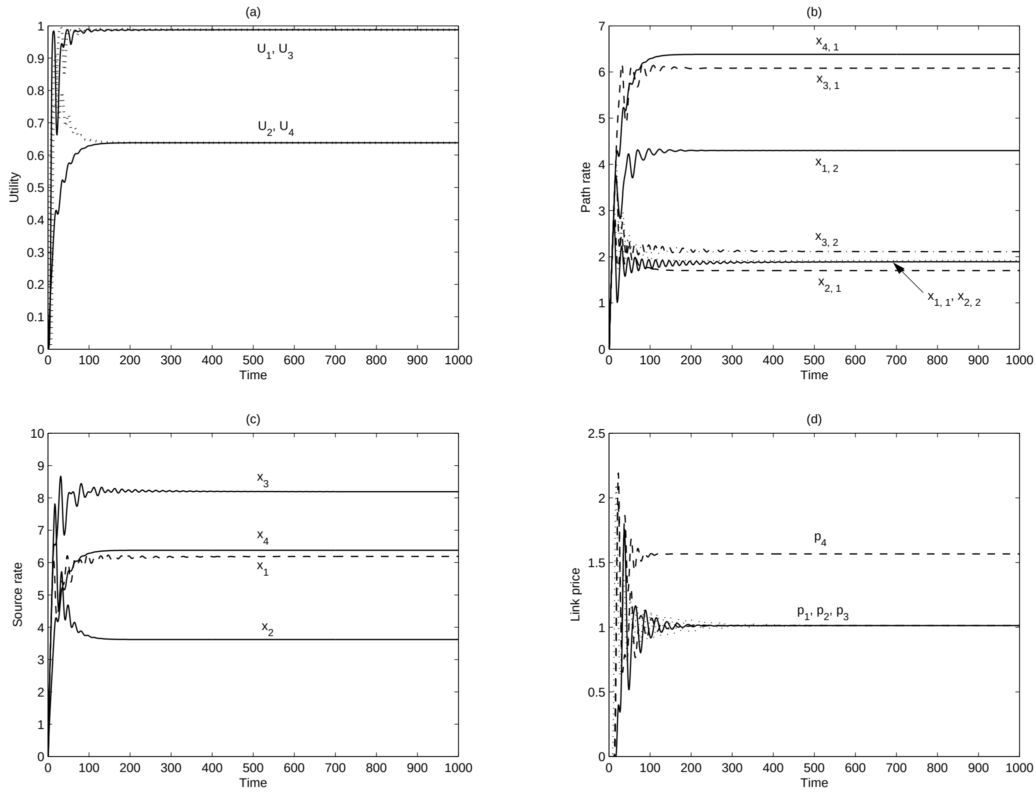

Fig. 5. Simulation results of modified flow control algorithm (a) Source utilities (b) Path rates (c) Source rates (d) Link prices

Fig. 5 shows the behavior of the modified algorithm. S_2 and S_4 share the bottleneck link l_4 ($p_4 = 1.5671$) with source rate $(3.6188, 6.3812)$. Both achieve a utility $U = U_2 = U_4 = 1/p_4 = 0.6381$. S_1 and S_3 then equally share the remaining "cheaper" network resource ($p_1 = p_2 = p_3 = 1.0125$) with a utility of 0.9877.

This confirms that the flow control algorithm given in this paper can provide an efficient utility max-min fair resource allocation for multipath communication networks among different applications, moreover, their utility functions (i.e., $U_1(x_1)$ and $U_3(x_3)$) may not need to satisfy the critical strictly concave condition which is required by optimal flow control approach.

V. CONCLUSIONS

In this paper, we have developed a distributive flow control algorithm for networks with multiple paths between source-destination pairs, and the objective is to achieve the utility max-min fair resource allocation among completing users. We have shown that in steady state, the algorithm does meet the goal for all choices of utility functions. It leads to a very desirable result. The utility max-min fair flow control algorithm presented in this paper only requires that each source utility function is positive, strictly increasing and bounded over the bandwidth. It is more suitable for practical networks where the QoS utility functions of real-time applications usually do not satisfy the strict concavity condition that is strongly desired by the standard optimal flow control approach. Furthermore, the simulation reveals that the means we have taken to speed up the convergence and remove the oscillation effect in multipath networks is effective. For our future work, the dynamic behavior such as stability will be studied and analyzed.

ACKNOWLEDGMENT

This work was supported by the Australian Research Council under Grant DP0559131 and ARC Research Networks on Intelligent Sensors, Sensor Networks and Information Processing.

REFERENCES

[1] F. P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, no. 1, pp. 33–37, January 1997.

[2] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *Journal of Operations Research Society*, vol. 49, no. 3, pp. 237–252, March 1998.

[3] S. H. Low and D. E. Lapsley, "Optimal flow control, I: basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, December 1999.

[4] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, October 2000.

[5] S. H. Low, F. Paganini, and J. C. Doyle, "Internet congestion control," *IEEE Control Systems Magazine*, vol. 22, no. 1, pp. 28–43, February 2002.

[6] R. La and V. Anantharam, "Utility based rate control in the Internet for elastic traffic," *IEEE/ACM Transactions on Networking*, vol. 10, no. 2, pp. 272–286, April 2002.

[7] T. Alpcan and T. Basar, "A utility-based congestion control scheme for Internet-style networks with delay," in *Proceedings of IEEE INFOCOM 2003*, March-April 2003.

[8] F. P. Kelly, "Fairness and stability of end-to-end congestion control," *European Journal of Control*, vol. 9, pp. 149–165, 2003.

[9] K. Kar, S. Sarkar, and L. Tassiulas, "Optimization based rate control for multipath sessions," Institute for Systems Research, University of Maryland, Tech. Rep. 2001-1, 2001.

[10] W. H. Wang, M. Palaniswami, and S. H. Low, "Optimal flow control and routing in multiple paths networks," *Performance Evaluation*, vol. 52, no. 2–3, pp. 119–132, 2003.

[11] K. Kar, S. Sarkar, and L. Tassiulas, "A scalable lower-overhead rate control algorithm for multirate multicast sessions," *IEEE Journal of Selected Areas in Communications*, vol. 20, no. 8, pp. 1541–1557, August 2002.

[12] W. H. Wang, M. Palaniswami, and S. H. Low, "Necessary and sufficient conditions of optimal flow control in multirate multicast networks," *IEE Proceedings — Communications*, vol. 150, no. 5, pp. 385–390, October 2003.

[13] D. Bertsekas and R. Gallager, *Data Networks*, 2nd ed. Prentice-Hall Inc., 1992.

[14] Z. Cao and E. W. Zegura, "Utility max-min: an application oriented bandwidth allocation scheme," in *Proceedings of IEEE INFOCOM 1999*, March 1999, pp. 793–801.

[15] W. H. Wang, M. Palaniswami, and S. H. Low, "Application-oriented flow control: fundamentals, algorithms and fairness," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1282–1291, December 2006.

[16] S. Shenker, "Fundamental design issues for the future Internet," *IEEE Journal of Selected Areas in Communications*, vol. 13, no. 7, pp. 1176–1188, July 1995.