Performance of networks containing both MaxNet and SumNet links

Lachlan L. H. Andrew and Bartek P. Wydrowski *

Abstract

Both MaxNet and SumNet are distributed congestion control architectures suitable for the Internet. MaxNet has recently been shown to have better fairness and scaling properties, but the majority of existing Internet links use the SumNet paradigm. If MaxNet links are to be deployed, they will need to be compatible with the existing SumNet infrastructure. This paper investigates the fairness and utilisation of networks consisting of mixtures of MaxNet and SumNet links, in different proportions.

Keywords: Network Flow Control, Congestion Control, Stability, Scalability.

1 Introduction

The problem of network flow control is to control source rates so that link capacities are utilised. For Internet-like networks, where links and sources can only have local information, the challenge is to control the source rates in a fully distributed manner. There are many different possible source rate allocations that fulfill the requirement of utilising the capacity. It has been shown [1] that the source rate allocation achieved by TCP Reno maximises an arbitrary utility function. Most other flow control schemes also maximise (different) utility functions [1]. In [2], we showed that it is possible to achieve a different, fairer rate allocation by altering the way the network signals congestion information.

Models of Internet-like networks control the source rate by a scalar feedback congestion signal. This signal is generated by aggregating the congestion prices of links on the end-to-end connection path of the source. For networks which achieve maximum utility, including TCP Reno networks, the signal is aggregated by summing all of the link prices on the path. We refer to these networks as SumNets. MaxNet was introduced in [2]. Its aggregation function is *Max*, and only the maximum link price along the connection path controls the source rate. MaxNet results in a Max-Min fair rate allocation for sources with homogeneous demand functions [2]. With appropriate parameters, it is stable for arbitrary capacities, delays and routing [3].

Since the difference between MaxNet and SumNet is simply the way the routers calculate the aggregate congestion signal, it is straightforward to build networks consisting of both MaxNet and SumNet links. The price signal forwarded over a MaxNet link is the maximum of the price of the link and the price signal in the incoming packet, while the price signal forwarded over a SumNet link is the sum of the two. All links out of a given router would typically be of the same type (MaxNet or SumNet), but that need not be the case.

*The authors are with the ARC Special Research Centre for Ultra-Broadband Information Networks. Department of Electrical and Electronic Engineering, The University of Melbourne, Australia. This work has been supported by the Australian Research Council.
{l.andrew,b.wydrowski}@ee.mu.oz.au

A major drawback of these hybrid networks is that the fairness may be seriously compromised by the heterogeneity of the network. This paper investigates both the fairness and the stability of such hybrid networks.

In Section 2 we start by briefly describing the MaxNet architecture presented in [2], and Section 3 recalls the steady state rate allocation properties of both MaxNet and SumNet. Section 4 investigates the degree to which the heterogeneity of the links causes unequal rates to be allocated to different connections, while Section 5 considers the impact on the total utility achieved by the steady state rate allocations. Stability issues are briefly discussed in Section 6.

2 MaxNet Architecture

In a MaxNet network, the congestion signal, q_i, communicated to source i is the maximum of all link prices on the end-to-end transmission path, as illustrated in Figure 1. Let p_l be the price at link l and L_i be the set of links connection i uses, then

$$q_i = \max\{p_l : l \in L_i\}.$$

To achieve this, the packet format must include sufficient bits to communicate the complete congestion price. Each link replaces the congestion price in the packet with its own congestion price if its own price is larger than the one in the packet. (This was shown in [4] to be more efficient than the single-bit signalling typically used in SumNets.) The congestion signal is relayed back to the source by the destination host in acknowledgement packets.

The behaviour of connection i is governed by an explicit demand function, $D_i(\cdot)$, such that its transmit rate is

$$x_i = D_i(q_i) \tag{1}$$

for a congestion signal q_i. The link Active Queue Management (AQM) algorithm is the by the well studied integrator process [5]:

$$p_l(t+1) = p_l(t) + (y_l(t) - c_l)\varphi_l \tag{2}$$

where $y_l(t) = \sum_{i:l\in L_i} x_i(t)$ is the aggregate arrival rate for link l at time t, φ_l is the control gain and c_l is the target capacity of link l which is related to its physical capacity C_l by the target utilisation $0 < \mu_l < 1$ such that $c_l = \mu_l C_l$.

3 Rate allocation of homogeneous networks

In steady state, when $p_l(t+1) = p_l(t)$, the rate allocation to connection i, x_i will depend on the magnitude of its demand function relative to those of other connections sharing bottleneck links of connection i.

Using MaxNet, the only bottleneck link which influences the rate is l, the link of connection i with the highest price. If T_l is the set of connections traversing link l and link l has the highest price on path i, then the rate allocation to i is

$$x_i = C_l \frac{D_i(q_i)}{\sum_{k\in T_l} D_k(q_k)} \tag{3}$$

MaxNet can achieve a Max-Min fair rate allocation. A vector of rates, x, is defined as Max-Min fair if, for every feasible rate vector r with $r_i > x_i$ for some connection i, there exists a connection k such that $x_k \leq x_i$ and $x_k > r_k$. Put simply, a rate vector is Max-Min if it is feasible and no flow can be increased while maintaining feasibility without decreasing a smaller or equal flow.

The following proposition is proved in [2].

Proposition 1 *Let $D(\cdot)$ be positive, continuous and decreasing. The rate allocation for a MaxNet network of homogeneous connections with $x_i = D(q_i)$ is Max-Min fair.*

Using SumNet, all bottleneck links traversed by connection i affect the steady state rate. The steady state rate allocate of SumNet is the allocation which maximises the sum over all sources, i, of the utility function, $U_i(x) = \int_0^x D_i(x)$.

4 Fairness of hybrid networks

The degree of unfairness of a rate allocation can be measured as its deviation from the max-min fair allocation. However, fairness *per se* is not the objective; rather it is to ensure that connections get at least their fair share. This motivates the following definition.

Definition 1 *The unfairness index of a rate allocation, x, relative to the max-min fair allocation, y, is*

$$u(x) = \sum_{i:x_i<y_i} \frac{y_i - x_i}{y_i}. \tag{4}$$

In essence, this measures the proportion by which connections fall short of achieving their fair share.

In order to investigate the impact of heterogeneous links on the fairness of the equilibrium rate allocation, networks with various arrangements of SumNet and MaxNet links were simulated. In order to isolate the impact of the flow control from that of the network topology, a simple ring topology was used. Even from this very specific topology, several general conclusions can be drawn.

4.1 Simulation setup

Two regular topologies were studied, with several different traffic patterns and with randomly allocated MaxNet and SumNet links. For each topology, the average over the traffic patterns was recorded.

The first topology studied was a ring with equal capacity on each of its L links. This is a common topology for metropolitan area networks, and represents networks with very sparse interconnections. Both unidirectional and bidirectional rings were simulated. To investigate densely connected networks, a 3×3 mesh-torus with unidirectional links was also simulated.

In both topologies, random combinations of source and destination nodes were used. A route through a ring is uniquely determined by its source and its destination, and whether it is bidirectional or unidirectional. For the mesh-torus, routes were shortest paths going in the "horizontal" direction first and then in the "vertical" direction. Let the number of connections be S. Each connection had the same demand function, $D(p) = x_0 \exp(-kp)$, where x_0 was 1.5 times the link capacity, and the constant $k = 0.005$ controls the sensitivity of the transmission rate to the price.

Figure 3 plots the unfairness index, averaged over all traffic patterns, against the number of SumNet links, s, for both bidirectional and unidirectional rings. Figure 4 plots the maximum unfairness index taken over the same traffic patterns.

Notice that a three node bidirectional ring is always perfectly fair, since all routes are a single hop. Bidirectional rings with two connections are also always perfectly fair, as will be explained below.

4.2 Maximum unfairness

For networks with unidirectional links, the worst case fairness of hybrid MaxNet/SumNet networks is generally worse than pure SumNet. This phenomenon is observable in networks in which two routes use the same pair of links in the same direction but in the reverse order. This is illustrated in Figure 2, with links 12 occurs before link 45 in one path, but after it in the other. If one of these links is costed

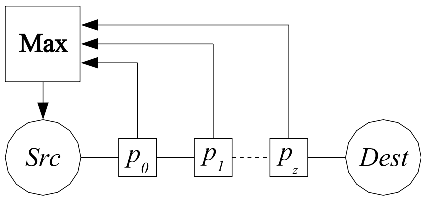

Figure 1: MaxNet logical feedback loop.

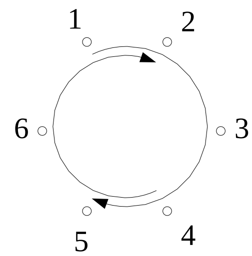

Figure 2: Routes overlapping on two routes in different orders

by a SumNet link and the other is by a MaxNet link, then the symmetry between the two paths which exists in a pure SumNet network is broken in a heterogeneous network, causing unfair rate allocation. If the two connections shown in Figure 2 are the only ones in the network, then both receive equal bandwidth under both MaxNet and SumNet. However, if link 12 is a SumNet link while link 45 is a MaxNet link, then the price signal conveyed by packets destined for node 2 is the sum of the two link prices, while that of packets destined for node 5 is only the cost of one link. The former cost is twice the latter, causing the transmission rates to differ.

This arrangement can only occur for rings with at least four nodes. It only requires one MaxNet link and one SumNet link in the network, and so occurs in heterogeneous unidirectional rings (of more than three nodes), irrespective of the proportion of MaxNet and SumNet links. However, the proportion of MaxNet and SumNet links influences the proportion of configurations which will be of this form, with the maximum proportion being for equal numbers of MaxNet and SumNet links.

Fortunately, it is impossible to have two shortest paths in bidirectional networks of any topology which both traverse the two links, but in the reverse order. This can be stated formally as follows.

Theorem 1 *Let $G = (V, E)$ be a graph with vertices V and undirected edges $E \subset V^2$. Weight the edges by the function $w : E \mapsto (0, \infty)$, with $w(u, v) = w(v, u)$ for all $(u, v) \in E$. Let R_1 and R_2 be two shortest paths in G. That is, R_i is the sequence $(v_{1,i}, v_{2,i}, \ldots, v_{n(i),i})$ where $(v_{j,i}, v_{j+1,i}) \in E$ and there is no other sequence $(u_{1,i} = v_{1,i}, u_{2,i}, \ldots, u_{m(i),i} = v_{n(i),i})$ with $\sum_{j=1}^{m(i)-1} w(u_{j,i}, u_{j+1,i}) < \sum_{j=1}^{n(i)-1} w(v_{j,i}, v_{j+1,i})$. With these definitions, if edge $(v_{k,1}, v_{k+1,1})$ on R_1 is also on R_2 as $(v_{l,2}, v_{l+1,2})$ and edge $(v_{k',1}, v_{k'+1,1})$ on R_1 is also on R_2 as $(v_{l',2}, v_{l'+1,2})$, with $k' > k$, then $l' > l$.*

Proof: The proof uses the fact that a sub-path of a shortest path is also a shortest path. Assume, with a view to obtaining a contradiction, that $l' \leq l$. If $l' = l$, then $v_{k,1} = v_{k',1}$, whence R_1 contains a loop and is not a shortest path. It remains to show that $l' < l$ also yields a contradiction. Let $P(u, v)$ be the length of the shortest path between u and v. Let $A = v_{k,1}$, $B = v_{k+1,1}$, $C = v_{k',1}$ and $D = v_{k'+1,1}$, so that $R_1 = \ldots AB \ldots CD \ldots$ and $R_2 = \ldots CD \ldots AB \ldots$. The cost of the sub-path of R_1 which lies between (and includes) the shared links is then

$$\begin{aligned} P(A, D) &= P(A, B) + P(B, C) + P(C, D) \\ &= P(A, B) + P(C, B) + P(C, D) \\ &> P(C, B). \end{aligned}$$

Considering path R_2 similarly yields $P(C, B) > P(A, D)$, which is the required contradiction. □

4.3 Impact of proportions

Although the maximum level of unfairness does not depend greatly on the ratio of MaxNet to SumNet links in the network, the mean level does. For bidirectional rings, the degree of unfairness is directly proportional to the proportion of SumNet links in the network in all cases investigated. That implies that the benefit gained by deploying MaxNet is directly proportional to the scale of the deployment.

This is particularly important from a practical point of view. If MaxNet only provided a benefit when it constituted a large proportion of the network, then it would be unattractive to large infrastructure providers with a large installed base of SumNet equipment, and also to small providers who have to rely on the rest of the Internet for most of their connections. Since partial deployment provides a partial benefit, both of these groups stand to benefit from upgrading to MaxNet equipment.

Unidirectional rings with five connections also show unfairness approximately proportional to the number of SumNet links, although the total level of unfairness is an order of magnitude higher than for bidirectional rings. However, unidirectional rings carrying only three concurrent connections display unfairness which is concave in the number of SumNet links, especially for larger rings. This indicates that there is little benefit from deploying a small number of MaxNet links in such networks, and that widespread deployment is necessary to obtain significant improvement.

4.4 Scaling

The maximum unfairness increases as the number of connections increases. There are two reasons for this. Firstly, the definition of unfairness includes a sum over the connections. Although most connections will contribute zero to the sum, the number of non-zero contributions also increases with the number of connections.

A second and more subtle effect is the increase of the link prices. Because the demand function is exponential, the ratio of transmission rates depends on the absolute difference between the aggregate prices seen by the connections, rather than on their ratio. However, as the number of connections increases, the absolute value of the link prices also increases. Thus, the difference in total price of a connection traversing two SumNet links and one traversing only one such link will increase.

The second effect is particularly important because it causes the unfairness experienced by an individual user to increase, whereas the first effect is simply a reflection of the fact that there are more users in the system.

5 Utility

MaxNet trades increased fairness for decreased total utility. This section investigates how far short heterogeneous networks fall from the maximum utility, which is achieved by homogeneous SumNets. The measure of wasted utility is the waste index.

Definition 2 *The waste index of a rate allocation, x, relative to the maximum utility allocation, y, is*

$$w(x) = \frac{\sum_i U(y_i) - U(x_i)}{\sum_i U(y_i)}, \tag{5}$$

where $U(x) = \int_0^x D(x)$ is the utility function common to all connections, and the sums are over all connections.

Figures 5 and 6 show the mean and maximum values of waste index of uni- and bi-directional rings, averaged over many traffic patterns.

The maximum waste of unidirectional rings is greater for heterogeneous networks than for either type of homogeneous networks, as was the case for the maximum unfairness. However, this effect decreases as the number of connections increases.

6 Stability of hybrid networks

It is known [3] that MaxNet is locally stable provided that, for each connection, i, the slope of the demand function around the equilibrium value, x_i^*, satisfies $D'(x_i^*) < x_i^*/\tau_i$, i, where τ_i is the round trip time for connection i. Similarly, it is known [5] that SumNet is locally stable provided $D(x_i^*) < x_i^*/M_i\tau_i$, where M_i is the number of bottleneck links on route i. The proofs rely on specific properties of the feedback in each case: for MaxNet it relies on the fact that there is no feedback from any link other than the controlling bottleneck, while for SumNet it relies on the symmetry of a connection receiving feedback from *all* links whose price is influenced by that connection's rate. Heterogeneous networks have neither of those properties, and there is currently no proof that even the stricter condition $D(x_i^*) < x_i^*/M_i\tau_i$ is sufficient for stability.

In particular, for the network in Figure 2, the total delay for the feedback of rate information may be related to the sum of the round trip times for the two connections, rather than either individual round trip time. This is the subject of ongoing investigation.

7 Conclusion

The average behaviour of networks consisting of heterogeneous mixtures of MaxNet and SumNet links is generally between that of homogeneous MaxNets and SumNets. However, in networks with asymmetric links, it is possible for the steady state bandwidth allocation of the heterogeneous network to be significantly less fair than that of either type of homogeneous network. The most common cause, if not the only cause, is routes which overlap in two different parts of the network, and traverse those parts in the same direction, but traverse them in different orders. This cannot occur in bidirectional networks containing only symmetric links.

References

[1] S. H. Low, "A duality model of TCP and queue management algorithms," in *Proc. ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, (Monterey, CA), 2002.

[2] B. Wydrowski and M. Zukerman, "MaxNet: A congestion control architecture for maxmin fairness," *IEEE Commun. Lett.*, vol. 6, pp. 512–514, Nov. 2002.

[3] B. Wydrowski, L. L. H. Andrew, and M. Zukerman, "MaxNet: A congestion control architecture for scalable networks," *IEEE Commun. Lett.*, Accepted for publication.

[4] B. Wydrowski, *Techniques in Internet congestion control*. PhD thesis, University of Melbourne, 2003.

[5] F. Paganini, J. C. Doyle, and S. H. Low, "Scalable laws for stable network congestion control," in *Proc. IEEE Conf. Decision Contr. (CDC)*, (Orlando, FL), 2001, pp. 185–190.

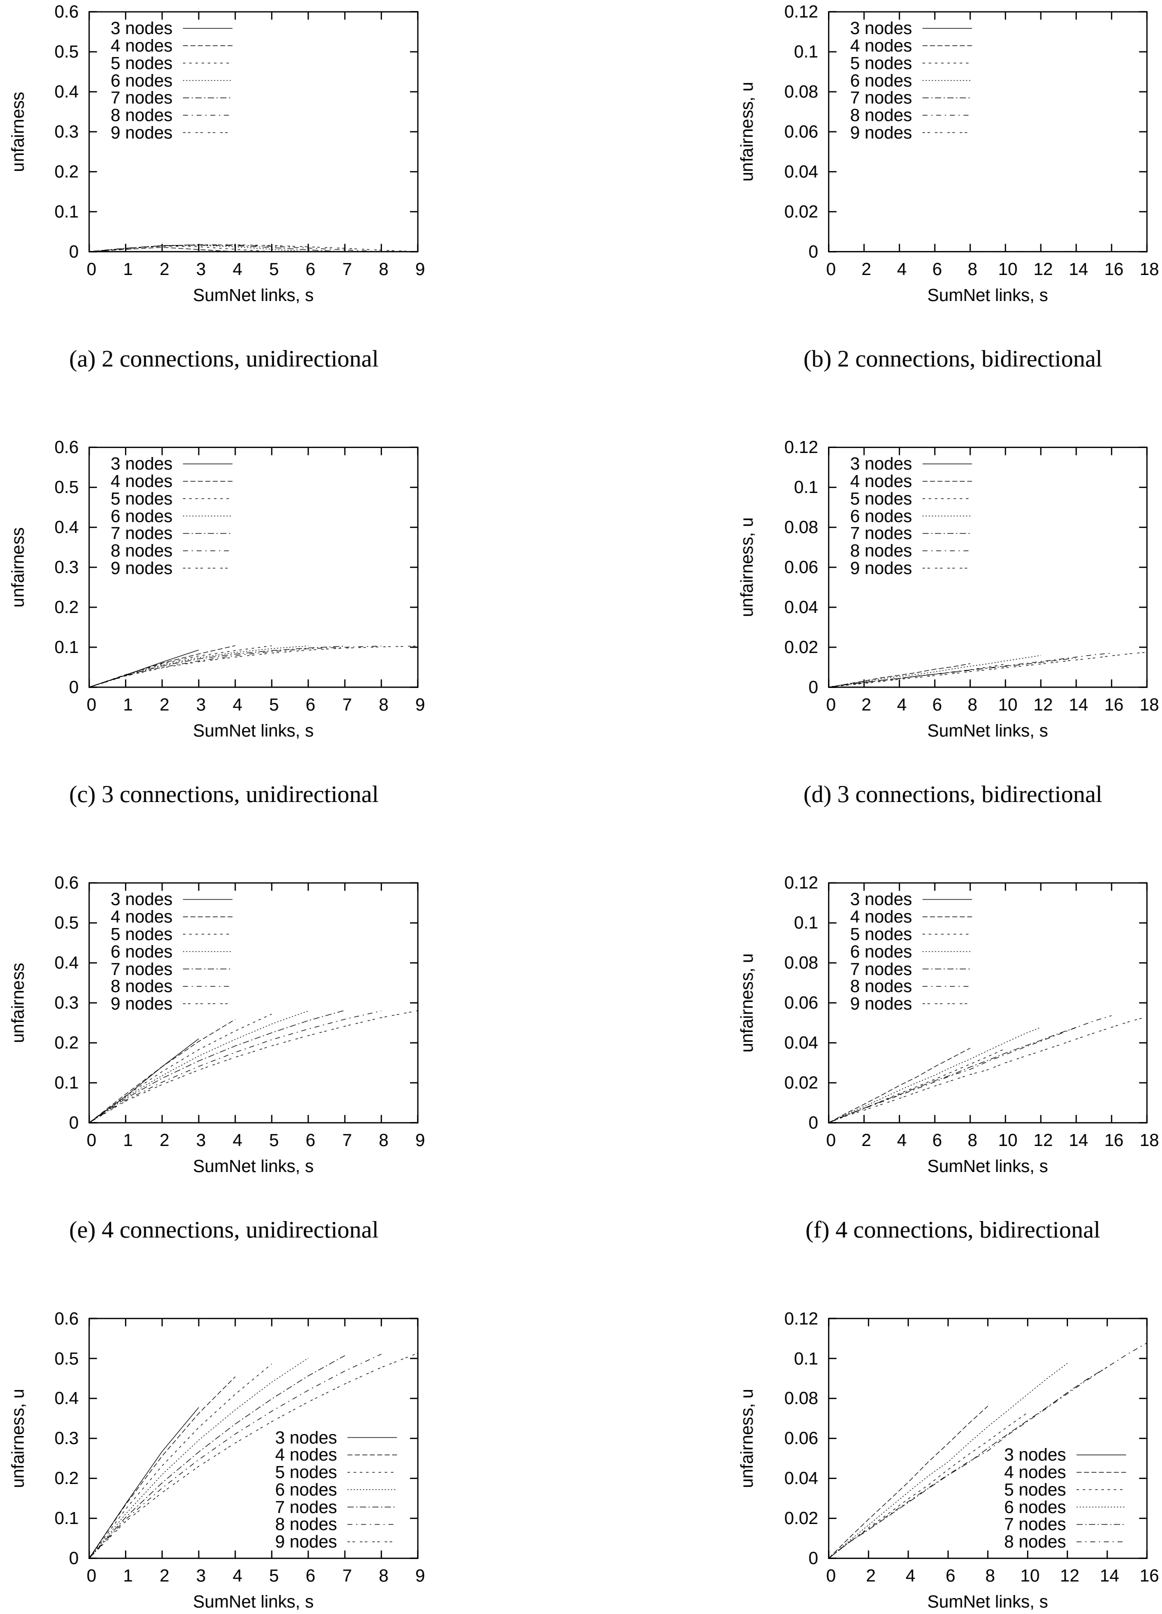

(a) 2 connections, unidirectional
(b) 2 connections, bidirectional
(c) 3 connections, unidirectional
(d) 3 connections, bidirectional
(e) 4 connections, unidirectional
(f) 4 connections, bidirectional
(g) 5 connections, unidirectional
(h) 5 connections, bidirectional

Figure 3: Mean unfairness index for ring networks.

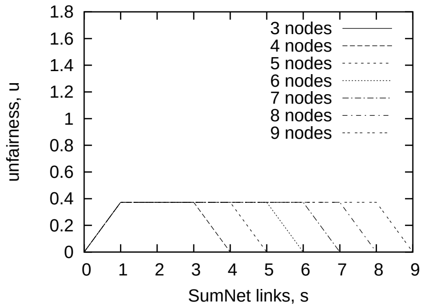

(a) 2 connections, unidirectional

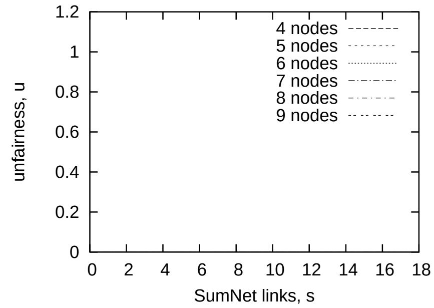

(b) 2 connections, bidirectional

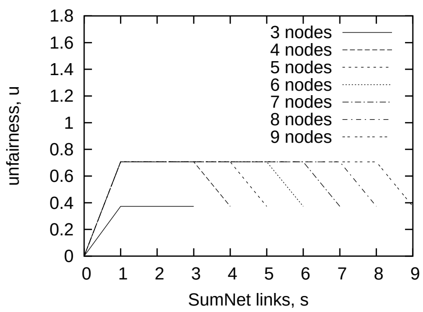

(c) 3 connections, unidirectional

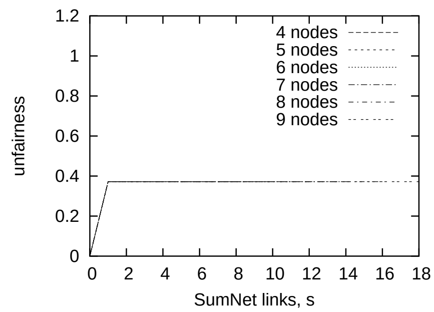

(d) 3 connections, bidirectional

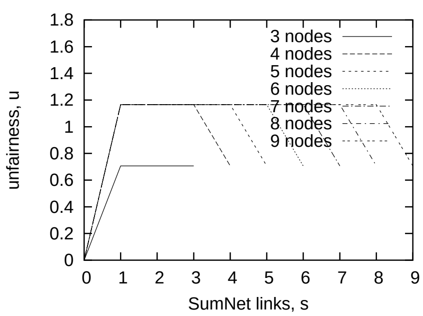

(e) 4 connections, unidirectional

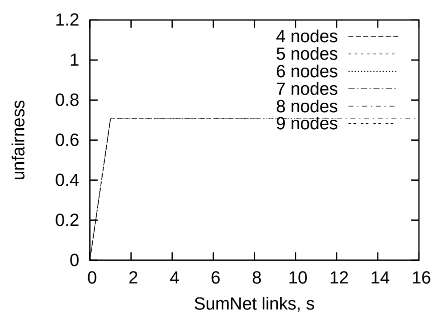

(f) 4 connections, bidirectional

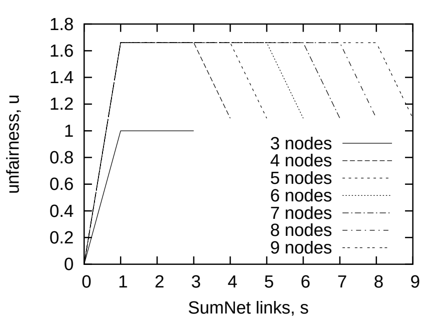

(g) 5 connections, unidirectional

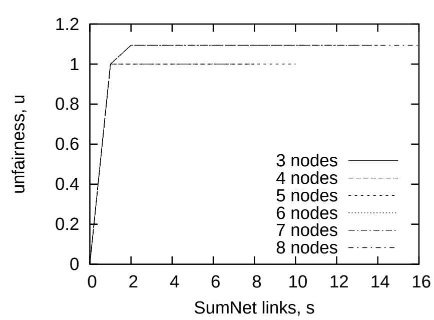

(h) 5 connections, bidirectional

Figure 4: Maximum unfairness index for ring networks.

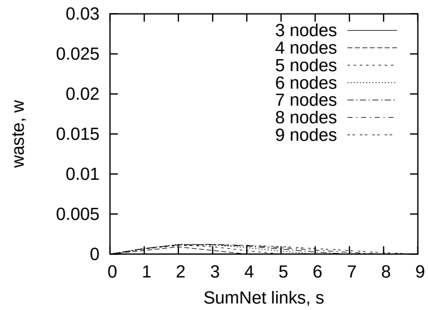

(a) 2 connections, unidirectional

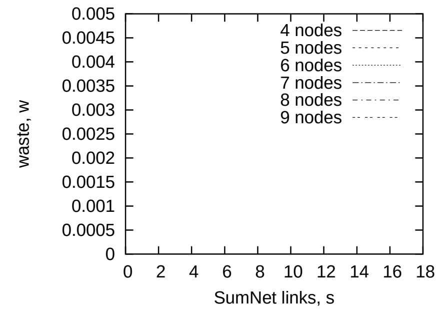

(b) 2 connections, bidirectional

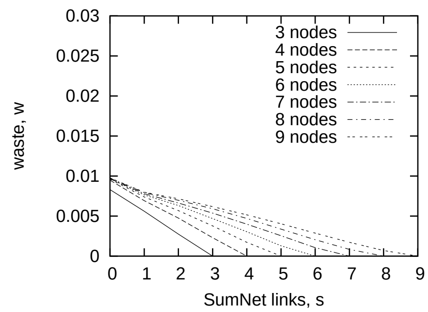

(c) 3 connections, unidirectional

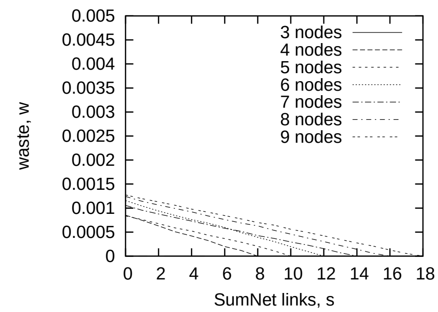

(d) 3 connections, bidirectional

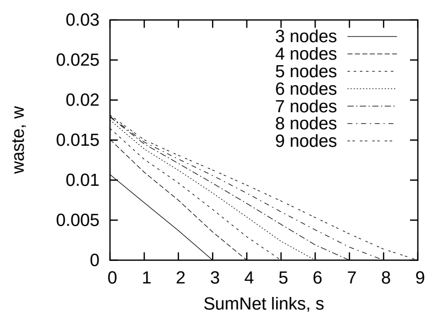

(e) 4 connections, unidirectional

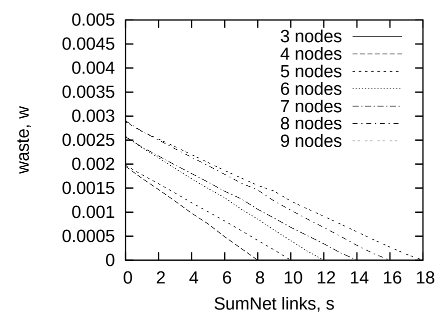

(f) 4 connections, bidirectional

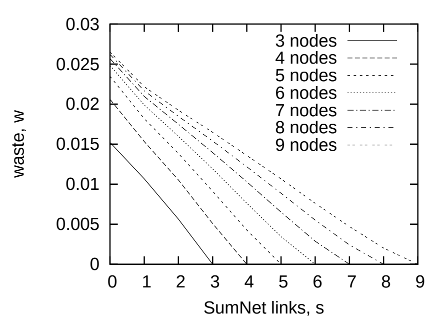

(g) 5 connections, unidirectional

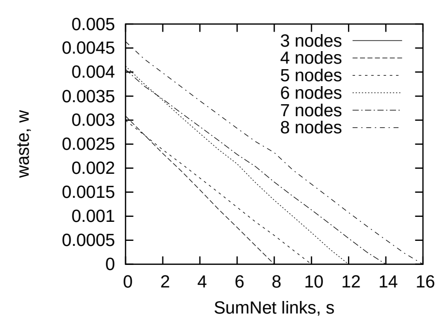

(h) 5 connections, bidirectional

Figure 5: Mean waste index for ring networks.

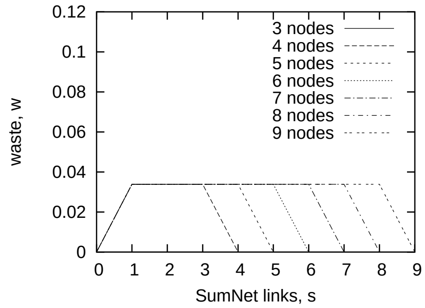

(a) 2 connections, unidirectional

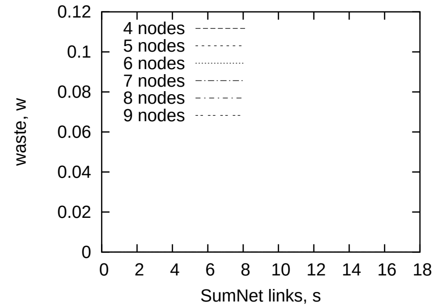

(b) 2 connections, bidirectional

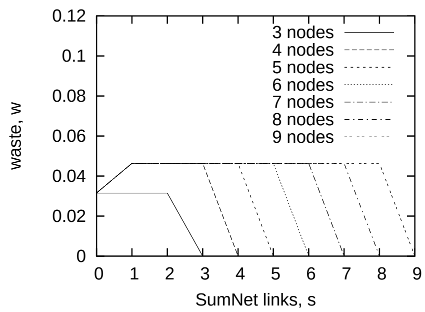

(c) 3 connections, unidirectional

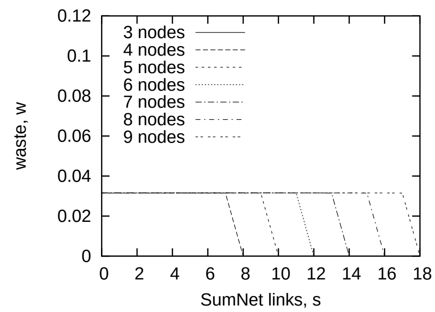

(d) 3 connections, bidirectional

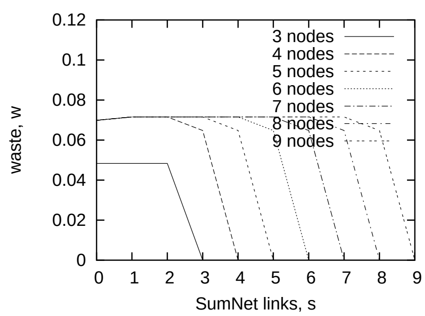

(e) 4 connections, unidirectional

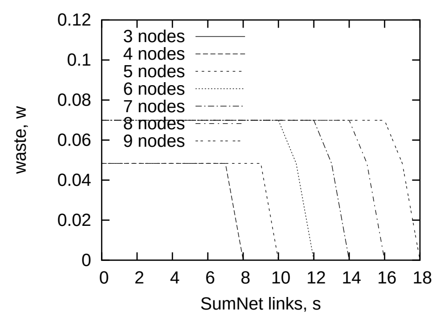

(f) 4 connections, bidirectional

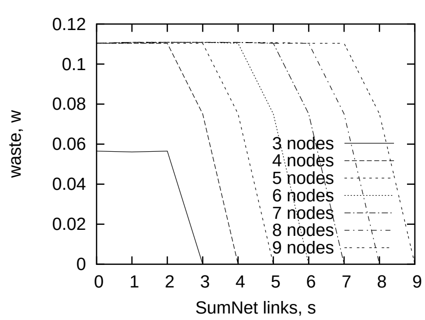

(g) 5 connections, unidirectional

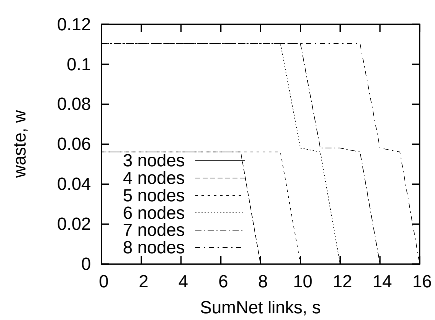

(h) 5 connections, bidirectional

Figure 6: Maximum waste index for ring networks.