

Algorithms for Theoretical Investigation of Fairness in Multipath Transport

Amanpreet Singh¹, Andreas Könsgen¹, Hakim Adhari², Carmelita Görg¹,
and Erwin P. Rathgeb²

¹ University of Bremen
{aps, ajk, cg}@comnets.uni-bremen.de

² University of Duisburg-Essen
{hakim.adhari, rathgeb}@uni-due.de

Summary. With the onset of multipath transport protocols such as MPTCP and multihomed mobile devices, fairness considerations which have been widely analyzed for legacy TCP need to be re-investigated. A practical realization of fairness amongst different participants is known to be difficult but even the theoretical calculation of the resource capacity and its allocation is not a trivial task. Therefore in this work, resource allocation algorithms are presented to thoroughly evaluate the impact of the fairness definitions. For a rigorous analysis, existing fairness definitions are identified according to the resources (bottleneck or network) and the competing participants (flow, tariff or user). Tariff as the participant, provides a realistic option to comply with the service level agreement between the operator and the user where as flow as the participant leads to TCP-compatible allocation. From the obtained results, it can be seen that if fairness is applied at the bottleneck then it is absolutely fair to the individual participants w.r.t. the bottleneck. On the other hand, fairness mechanisms considering the whole network as a single resource exploit the freedom of resource allocation (due to multipath flows) to achieve an overall similar allocation for the different participants (irrespective if the participant is composed of singlepath or multipath flows) but are still restricted by the topological constraints and might even result in a lower overall network throughput. ³

1.1 Introduction

The Internet of today is dominated by singlepath TCP flows [1] and it is considered as fair not to push away TCP flows. Therefore, the common definition of fairness in the Internet is called TCP friendliness [2] i. e. if n TCP sessions share the same bottleneck link, each should get $1/n$ of the bottleneck link capacity. Thus, a transport layer protocol is fair if it displaces no more TCP traffic than a TCP stream itself would displace. TCP-friendly flows are

³ This work has been funded by the German Research Foundation (Deutsche Forschungsgemeinschaft – DFG).

also called TCP-compatible and defined by RFC 2309 [3]. Mechanisms that protocols commonly use to meet the TCP-friendly requirement use some form of *additive increase multiplicative decrease* (AIMD) congestion window management or compute a transmission rate based on equations derived from an AIMD model [4]. Therefore, the TCP-friendly or -compatible view basically controls the rates of the flows in such a way that during congestion the bottleneck link capacity (resource) is shared equally by the competing flows (participants). This fairness view can also be called *flow rate fairness* [5] or *bottleneck flow fair (BFF)* [6]. RFC 6077 [7] outlines fairness issues as part of the open research issues in the Internet congestion control without favoring any particular fairness definition.

With the emergence of multipath transport protocols like CMT-SCTP [8] and MPTCP [9], a flow can have several (k) subflows (which are comparable to singlepath flows) to increase its overall throughput by utilizing the multi-homing capability of the endpoints and at the same time make the network more efficient. With this new terminology, a singlepath flow can be seen as a flow composed of a single subflow. If any of the standard TCP congestion control methods (e.g., NewReno) is used for multipath transport then every subflow will behave as an individual TCP connection. Hence the realized fairness mechanism is called *bottleneck subflow fair (BSfF)* i.e., it considers *subflows* as the participant and *bottlenecks* as the resource that should be shared.

The multipath transport standardization group at the IETF felt that since the current Internet is based on the principle of “do no harm” to existing singlepath flows, any multipath transport protocol design should satisfy this goal [10]. Bottleneck subflow fairness is deemed too aggressive on a bottleneck link, if two or more subflows of a flow share the same bottleneck link, therefore bottleneck flow fairness is the desired mechanism. In order to have bottleneck flow fairness, all subflows of a multipath flow sharing the same bottleneck should not get a combined share more than that of a singlepath flow [11].

Since a flow can have several subflows using different paths, the overall throughput is not limited by a single bottleneck link anymore. In addition, different subflows will observe different degrees of congestion on their respective paths. The idea of *Resource Pooling (RP)* that makes network resources behave like a single pooled resource, exploits this feature of the multipath transport to *balance congestion* [11, 12, 13] in the network. RP aims at shifting traffic from more to less congested paths and thereby decreasing the overall congestion in the network and increasing the performance. Thus, RP brings in a new fairness perspective by considering the whole network as a resource while the participants are still the flows. This new fairness view is denoted as *network flow fair (NFF)*.

The flow rate fairness approach means that it is considered fair that two flows from the same application or endpoint can get double share of the bottleneck link capacity at the expense of another flow from another endpoint (Fig. 1(a)), but it is deemed unfair if a multipath flow with two subflows

gets a double share (Fig. 1(b)). In previous work [6, 14, 15], issues related with the current fairness methods proposed for multipath transport have been highlighted. In order to deal with this limited scope of (*bottleneck/network*) *flow fairness*, alternative methods w. r. t. the additional participants that still tackle fairness at the transport layer but also consider the higher layer aspects, the end user as well as the network are discussed in section 1.3.

Independent of the way where/how fairness methods can be deployed, a comprehensive analysis is required to determine their impact. In addition it has been shown in [6] that even for simple topologies, neither the theoretical allocation of the resources nor the calculation of the resource capacity is an obvious task. This work presents resource assignment algorithms for the discussed fairness methods as a first step towards such an analysis.




Fig. 1.1. Issue with fairness definition for multipath

This paper is structured as follows: in Section 1.2 the important terms which are essential for the fairness discussion are defined in an abstracted form. The definition of the alternative fairness mechanisms follows in Section 1.3. In Section 1.4 resource allocation algorithms that determine the theoretical share of the participants are presented and the corresponding impact analysis of the fairness definitions as well as the validation of the algorithms is discussed in Section 1.5 with the help of an example scenario. Finally, Section 1.6 concludes this paper with a summary and a short outlook.

1.2 Terminology

For a better overview and systematic analysis, the different fairness methods discussed in this work are categorized based on the choice of the resource and participant. Therefore, a formal abstraction of the different networking resources and participants is given in this section.

Network: A *network* $\Gamma := (L, N, C)$ can be abstracted as:

- L – a finite locator set,
- $N \subseteq \mathfrak{P}(L)$ – a node set, $\mathfrak{P}(L)$ is the powerset of L ,
- $C \subseteq L \times L$ – a connectivity set.

Locator: A locator $l \in L$ is a network interface where L defines a finite set of unique locators. The connectivity among the locators is described by the connectivity set C which also describes the individual link capacity ρ .

User: A user u_q is defined as an entity which can simultaneously use more than one terminal node for its applications/connections.

Source: A source $n_{sr} \in N_{sr}$ is the source node in a communication e.g., it represents a computer or mobile device. N_{sr} and N_{ds} are the set of source and destination nodes, respectively used to define the demand matrix $D := N_{sr} \times N_{ds}$.

Tariff: A tariff t is defined by the access network with which the user has a contract. There can also be a shared-tariff model where the subscription is shared amongst several locators by a user e.g., a mobile operator also provides a WLAN hotspot or partner cards under a common tariff.

Flow: A flow f_x between two nodes $n_{sr} \in N_{sr}$ and $n_{ds} \in N_{ds}$ is composed of all Protocol Data Units (PDU) belonging to the same communication (e.g. an individual file transfer between n_{sr} and n_{ds}) irrespective of whether it is connection-oriented or connection-less and using a path set P_x . The bandwidth allocated to a flow $f_{q,x}$ of user u_q is denoted as $b_{q,x}$.

Subflow: A subflow $s_{q,x}^a$ denotes the subset of PDUs belonging to $f_{q,x}$ initiated by user u_q and using a specific path $p_{q,x}^a \in P$. The bandwidth allocated to this subflow is denoted by $b_{q,x}^a$.

Bottleneck: Let $\langle i, j \rangle$ be a link in Γ with the bandwidth $\rho\langle i, j \rangle$. The set of subflows crossing this link builds the subflow set $S\langle i, j \rangle$. A link is considered to be bottlenecked if no spare capacity is left after allocation of the subflow capacities that share this link i.e.,

$$\sum_{s_{q,x}^a \in S\langle i, j \rangle} b_{q,x}^a = \rho\langle i, j \rangle$$

1.3 Alternative Fairness Definitions

To design a fair multipath transport protocol, the standardization community has decided to remain with the notion of TCP-compatible flows even if a flow as known from a singlepath environment is different from the new kinds of flows used by the multipath protocols [10]. With the use of RP, the decision was taken to couple all subflows belonging to a single multipath flow. However, in multiple situations, multipath flows appear to be penalized to the advantage of singlepath flows [14, 16, 17]. In order to deal with these issues, in this section alternative ways to define the set of subflows that should be coupled together are presented.

For multipath subflows it is highly likely that not all the subflows have the same end locator. A simple case is to share the resource amongst the different locators fairly. Thus, for *network locator fairness (NLF)* all the (sub)flows initiated from the same locator are coupled together and for *bottleneck locator fairness (BLF)* all (sub)flows that share the same bottleneck and locator are coupled together.

A user may use multiple tariffs, e.g. one for LTE and one for WLAN. In this case, it is fair if the source node/user gets a share of the bottle-

neck/network capacity w. r. t. its tariff. Thus, for *network tariff fairness (NTF)* all the (sub)flows initiated from the same tariff and being part of the same network are coupled together while for *bottleneck tariff fairness (BTF)* all (sub)flows that share the same bottleneck and tariff are coupled together.

There are various tariff models which grant different guarantees e. g. different capacities in the access link. In order to reflect this property of tariffs in fairness methods, weighting factors could be used. Each tariff, whether inside the scope of bottleneck or network fairness, is then assigned a capacity share proportional to the weighting factor. An important aspect of weighted fairness is the mapping between the tariff plan and the weighting factor. In addition, a weighting factor could be adopted by other participants as well e. g. by flows to prioritize different types of applications or subflows to prioritize an interface.

Defining the source node as the participant leads to *network source fairness (NSrF)* or *bottleneck source fairness (BSrF)* where for the former, all subflows and flows initiated from the same source node are coupled together and for the latter, subflows and flows sharing the same bottleneck as well as the same source node are coupled together. Since a user may use multiple devices simultaneously the fairness can be taken to even a higher participant level wherein all the (sub)flows initiated by a user should be coupled depending on which resource they are sharing – a bottleneck (*bottleneck user fairness (BUF)*) or the network (*network user fairness (NUF)*).

1.4 Algorithms for the considered fairness views

Different fairness types can be realized practically by applying various methods of coupled congestion control. But these realizations have imperfect knowledge of the network conditions and are limited by their respective protocol design [16] [17]. Therefore to determine the theoretical optimum share of the participants, resource allocation algorithms are presented in this section. In addition the results from these algorithms can be used to thoroughly evaluate the impact of the different fairness methods introduced in sections 1.1 and 1.3 on both the network and the user. Thus these algorithms enable a systematic comparison of the various fairness mechanisms but do not have any influence on the practical realization in form of coupled congestion control variants.

1.4.1 Bottleneck scope

This section describes the means to achieve the theoretical allocation of subflow capacities within the network at a given time with respect to the different fairness definitions that share the bottleneck as a resource. The subflow capacities are obtained in an iterative process to consider the bottlenecks and spare capacity within the network. Due to elastic traffic, subflow capacity is bounded by the minimum share with respect to the links on its path i. e. , other non bottlenecked links of its path will have spare capacity left. Therefore *fair plus*

spare share allocation considers the spare capacity made available by the participants that are bottlenecked by other links on their respective paths. Hence in the algorithm, multiple iterations are needed to identify all bottlenecked flows. In each iteration, bottlenecked subflows are identified by searching for links that are fully utilized by their respective fair share constraint. In each iteration a new set of links and corresponding subflows becomes bottlenecked. Once all the subflows are bottlenecked, the final allocation of subflow capacities is obtained.

A user may initiate multiple flows which in turn might consist of several subflows. Thus, a bottleneck link may be shared by multiple flows or subflows from a single user. By definition of the *Bottleneck user fair*, all flows or subflows belonging to a single user should get a combined share equal to that of other sharing users. Since the subflow is the smallest entity, its capacity share is calculated dependent on the fairness policy with multiple iterations based on which subflows get bottlenecked first.

The link that will bottleneck an unbottlenecked subflow has to be identified which not only depends on the user share of the link but also the number of subflows that share the bottleneck link. To obtain the unbottlenecked subflow capacity, first a fair share of the respective user on each link of the subflow's path is calculated. Therefore, the remaining bandwidth of link $\langle i, j \rangle$ is determined by considering the link-specific bottlenecked users $\mathcal{Y}\langle i, j \rangle_b$ (i.e., all the flows of the user on the link are bottlenecked) and then shared amongst non-bottlenecked user's $\mathcal{Y}\langle i, j \rangle_{ub}$ to provide the fair+spare user share of the link. The fair share of user v_q on link $\langle i, j \rangle$ is denoted by $ub_q\langle i, j \rangle$ and can be further apportioned equally amongst its unbottlenecked subflows and single-path flows or first at the level of flows ($b_{q,x}\langle i, j \rangle$) and then subflows ($b_{q,x}^a\langle i, j \rangle$). In this way, the subflow capacity share over each link of its path is obtained. Finally, the subflow capacity which is the minimum capacity over all links that constitutes its path is determined.

Once the capacity of all the subflows is obtained, the bottleneck link(s) can be identified i.e, a link which does not have any spare bandwidth is bottlenecked and all the subflows that share this link have reached their maximum capacity. The bottlenecked subflows are now made part of the bottlenecked subflows set over all users S_b as well as the bottleneck subflow set $S_{q,x}\langle i, j \rangle_b$ for a particular user and flow on link $\langle i, j \rangle$. If all the subflows of a flow $f_{q,x}$ on link $\langle i, j \rangle$ are bottlenecked then the flow becomes part of the user-specific bottlenecked flow set $F_q\langle i, j \rangle_b$ over link $\langle i, j \rangle$. Similarly, the link-specific bottlenecked user set $\mathcal{Y}\langle i, j \rangle_b$ is also updated. The capacity of the remaining unbottlenecked subflows needs to be determined again with the extended set of bottlenecked subflows, flows and users until all the subflows are bottlenecked.

For multipath flows, different subflows might use different tariffs and therefore there is an inter-dependency between the different tariffs that might even go across multiple source nodes/users depending on the shared tariffs. For tariff as the participant, only the network that is affected by the tariffs is considered i.e., a subset $\Gamma'_{child} = (L', N', C')$ of the whole network $\Gamma = (L, N, C)$.

Algorithm 1 Bottleneck weighted user fair allocation algorithm

• Input: network $\Gamma := (L, N, C)$, demand matrix $D := N_{sr} \times N_{ds}$, user set \mathcal{U} , flow set F , user-specific flow sets F_q , non-bottlenecked subflow set S_{ub} , bottlenecked subflow set S_b , user and flow specific subflow sets $S_{q,x}$ and path set P .

• Output: user (weighted) share ub_q , flow (weighted) share $b_{q,x}$ and subflow (weighted) share $b_{q,x}^a$.

• Initialization: $ub_q := 0 \forall v_q \in \mathcal{U}; b_{q,x} := 0 \forall f_{q,x} \in F; b_{q,x}^a := 0 \forall s_{q,x}^a \in S; S_{ub} := S; S_b := \emptyset; \mathcal{Y}\langle i, j \rangle := \emptyset \forall \langle i, j \rangle \in C; F_q\langle i, j \rangle := \emptyset \forall v_q \in \mathcal{U}, \langle i, j \rangle \in C; S_{q,x}\langle i, j \rangle := \emptyset \forall f_{q,x} \in F_q, \langle i, j \rangle \in C;$

foreach $\langle i, j \rangle \in C$ **do**

foreach v_q in \mathcal{U} **do**

foreach $f_{q,x}$ in F_q **do**

foreach $s_{q,x}^a$ in $S_{q,x}$ **do**

$S_{q,x}\langle i, j \rangle := S_{q,x}\langle i, j \rangle \cup \{s_{q,x}^a \mid \langle i, j \rangle \in p_{q,x}^a, p_{q,x}^a \in P\}$

$F_q\langle i, j \rangle := F_q\langle i, j \rangle \cup \{f_{q,x} \mid S_{q,x}\langle i, j \rangle \neq \emptyset\}$

$\mathcal{Y}\langle i, j \rangle := \mathcal{Y}\langle i, j \rangle \cup \{v_q \mid F_q\langle i, j \rangle \neq \emptyset\}$

$S_{q,x}\langle i, j \rangle_b := \emptyset; S_{q,x}\langle i, j \rangle_{ub} := S_{q,x}\langle i, j \rangle; F_q\langle i, j \rangle_b := \emptyset, F_q\langle i, j \rangle_{ub} := F_q\langle i, j \rangle$

$\mathcal{Y}\langle i, j \rangle_b := \emptyset; \mathcal{Y}\langle i, j \rangle_{ub} := \mathcal{Y}\langle i, j \rangle.$

• Computation: **while** $S_{ub} \neq \emptyset$ **do**

foreach $s_{q,x}^a$ in S_{ub} // for each non-bottlenecked subflow //

do

foreach $\langle i, j \rangle \in p_{q,x}^a$ // for each link on the subflow's path //

do

$$ub_q\langle i, j \rangle := \left[\frac{\left(\rho\langle i, j \rangle - \sum_{v_r \in [\mathcal{Y}\langle i, j \rangle]_b} \left(\sum_{f_{r,y} \in F_r\langle i, j \rangle} \left(\sum_{s_{r,y}^d \in S_{r,y}\langle i, j \rangle} b_{r,y}^d \right) \right) \right) \cdot \Xi_q}{\sum_{\mathcal{Y}\langle i, j \rangle_{ub}} \Xi_r} \right]$$

$$b_{q,x}\langle i, j \rangle := \left[\frac{\left(ub_q\langle i, j \rangle - \sum_{f_{q,y} \in F_q\langle i, j \rangle_b} \left(\sum_{s_{q,y}^d \in S_{q,y}\langle i, j \rangle} b_{q,y}^d \right) \right) \cdot \Psi_{q,x}}{\sum_{F_q\langle i, j \rangle_{ub}} \Psi_{q,x}} \right]$$

$$b_{q,x}^a\langle i, j \rangle := \left[\frac{\left(b_{q,x}\langle i, j \rangle - \sum_{s_{q,x}^d \in S_{q,x}\langle i, j \rangle_b} b_{q,x}^d \right) \cdot \psi_{q,x}^a}{\sum_{S_{q,x}\langle i, j \rangle_{ub}} \psi_{q,x}^d} \right]$$

 where,

Ξ_q is the weight of user \mathbf{U}_q , $\Psi_{q,x}$ is the weight of flow $f_{q,x}$ that belongs to user q and $\psi_{q,x}^a$ is the weight of subflow $s_{q,x}^a$ that belongs to flow $f_{q,x}$.

$b_{q,x}^a = \min_{\langle i, j \rangle \in p_{q,x}^a} (b_{q,x}^a\langle i, j \rangle)$ // intermediate subflow capacity //

 // update the respective non-/bottlenecked sets //

$S_b := \left\{ s_{q,x}^a \in S \mid \exists \langle i, j \rangle \in p_{q,x}^a, \sum_{s_{q,x}^a \in S\langle i, j \rangle} (b_{q,x}^a) = \rho\langle i, j \rangle \right\}, S_{ub} := S \setminus S_b$

foreach $\langle i, j \rangle \in C$ **do**

foreach v_q in \mathcal{U} **do**

foreach $f_{q,x}$ in F_q **do**

$S_{q,x}\langle i, j \rangle_b := S_b \cap S_{q,x}\langle i, j \rangle$

$S_{q,x}\langle i, j \rangle_{ub} := S_{q,x}\langle i, j \rangle \setminus S_{q,x}\langle i, j \rangle_b$

$F_q\langle i, j \rangle_b := F_q\langle i, j \rangle_b \cup \{f_{q,x} \mid S_{q,x}\langle i, j \rangle_b = S_{q,x}\langle i, j \rangle\}$

$F_q\langle i, j \rangle_{ub} := F_q\langle i, j \rangle \setminus F_q\langle i, j \rangle_b$

$\mathcal{Y}\langle i, j \rangle_b := \mathcal{Y}\langle i, j \rangle_b \cup \{v_q \mid F_q\langle i, j \rangle_b = F_q\langle i, j \rangle\}$

$\mathcal{Y}\langle i, j \rangle_{ub} := \mathcal{Y}\langle i, j \rangle \setminus \mathcal{Y}\langle i, j \rangle_b$

Outside the tariff-specific network, the complete network could be seen as a hierarchical graph where inside Γ also a weighting factor is associated to Γ_{child} . Based on this weighting factor, the resources associated to the sum of all subflows going out of Γ_{child} and crossing for example a bottleneck in Γ is determined. Amongst the subflows belonging to Γ_{child} , the Γ_{child} weighting factors are still valid even if the bottleneck is outside Γ_{child} .

A user might be interested to use its different interfaces/flows in a particular way and therefore might have different weights attached to it. Similarly a user might have a common tariff over multiple interfaces and hence also have weights depending on the tariff cost. A general algorithm with weights that allows a user to give weights to specific users, flows or interfaces (subflows) is given in Algorithm 1. *Bottleneck source fair*, *bottleneck flow fair*, *bottleneck locator fair* and *bottleneck subflow fair* mechanisms share the bottleneck link capacity fairly amongst the participants: source node, flow, locator and subflow, respectively. Thus, this is either a subset of the algorithm presented in Algorithm 1 or can be derived from it.

1.4.2 Network scope

In this section the network is seen as a single resource which can be shared by different participants. The dependency between the calculation of the network capacity ρ^n and subflow capacity $b_{u,f}^{sf}$ can be solved by forming a linear set of equations. The linear equation system can be classified into two parts:

- Bottleneck bound: the subflows that are bounded by the bottleneck link capacity,
- Fairness bound: the participants (i.e. flows/users) that should get a fair share.

If the number of equations fits the number of variables, there is exactly one solution. If the number of equations is less (the system is underdetermined), there are multiple solutions, i.e. in the network, distributing capacities to different subflows is arbitrary within a specific range dependent on the network topology and bottleneck link capacities.

In some cases a perfectly fair solution may not be possible due to the unavailability of sufficient capacity corresponding to a flow that cannot get a fair share. An iterative method is then applied which systematically performs allocations to flows that are restricted due to the limited bottleneck capacity and identifies its dependencies on the other flows. Any allocation update of a flow share may introduce new dependencies. Once all the dependencies are taken care of then a final allocation of subflow capacities is obtained. For complex and large networks such an iterative method might become too difficult to handle and therefore, an alternative method of mixed integer (non-)linear programming may be used to identify the network fair allocation of flow capacity. In this method, the fairness requirement is represented in the form

of constraints while the optimum solution is obtained by the objective function. Here multiple objectives can be defined such as maximizing the network throughput while at the same time minimizing the difference between the share of the participants.

Due to the fact that subflows may share different bottlenecks, it is highly unlikely that all the subflows can get an equal share. Since a flow is defined as a set of subflows, a *network flow fair* solution is feasible wherein the allocation of the different subflows corresponding to a flow can be tuned. To obtain a *network locator fair*, *network tariff fair*, *network source fair* or *network user fair* solution, the *fairness bound* equations have to be formulated with respect to locator, tariff, source node or user, respectively. The bottleneck bound equations are always applied at the subflow level.

1.5 Validation of the Algorithms and Analysis of Fairness Mechanisms

To validate the algorithms implemented in MATLAB as well as highlight the advantage of considering tariff as the participant, an example scenario shown in Fig. 1.2 is used. For better overview, the notation used in this scenario differs from the one used in the algorithm previously described e. g., subflow $sf_k, 1 \leq k \leq 7$. The scenario consists of two users where user u_1 has two source nodes sr_1 and sr_2 while user u_2 has a single source node sr_3 . Source node sr_1 initiates two multipath flows, each composed of two subflows initiated from two different locators (l_1 and l_2) under different tariffs (t_1 and t_2). The different tariffs are shown with different colors while subflows of the same flow have the same line style. Source node sr_2 is also under tariff t_2 and initiates a singlepath flow. Source node sr_3 initiates a multipath flow also composed of two subflows via two different locators (l_4 and l_5) having independent tariffs (t_3 and t_4). In the scenario's topology, there are three bottleneck links where ρ_1, ρ_2 and ρ_3 are defined as the capacities of these bottleneck links.

Table 1.1 depicts the results obtained from the bottleneck and network resource allocation algorithms (introduced in Section 1.4) for the scenario shown in Fig. 1.2 with $\rho_1 = \rho_2 = \rho_3 = 100$. In this scenario, flows 1, 2 and 4 are multipath flows, each composed of two subflows while flow 3 is the only single path flow and part of two bottleneck links ρ_1 and ρ_3 . As a metric to compare the different fairness mechanisms, the overall throughput and Jain's fairness index (FI) [18] is used. In order to highlight that fairness can be perceived differently with respect to the different participants, Jain's fairness index is shown for the following participants - subflow (SfFI), flow (FFI) and tariff (TFI).

The *bottleneck subflow fair* allocation algorithm identifies the link between routers r_1 and r_2 to be bottlenecked first as this link has four subflows sharing its capacity $\rho_1 = 100$. Thus, the four subflows sf_1, sf_2, sf_3 and sf_5 get bottlenecked with a share of 25 each. Subflows sf_4 and sf_6 get bottlenecked next on

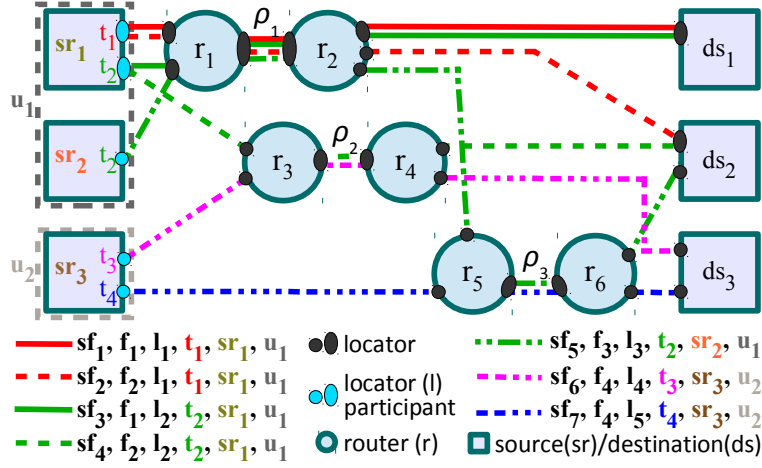


Fig. 1.2. Example scenario for fairness methods analysis (participants - sf: subflow, f: flow, l: locator, t: tariff, sr: source, u: user)

Table 1.1. Capacity share in Mbit/s for scenario in Fig. 1.2, (N: Network, B: Bottleneck, Sf: Subflow, F: Flow, L: Locator, T: Tariff, Sr: Source and U: User)

sf#	share $b_{u,f}^{sf}$	$\rho_1 = \rho_2 = \rho_3 = 100$										
		BSfF	BFF	BLF	BTF	BSrF	BUF	NFF	NLF	NTF	NSrF	NUF
sf ₁	$b_{1,1}^1$	25.0	16.67	16.67	25.0	16.67	16.67	25.0	50.0	33.33	12.5	21.5
sf ₂	$b_{1,2}^1$	25.0	33.33	16.67	25.0	16.67	33.33	00.0	00.0	33.33	00.0	14.0
sf ₃	$b_{1,1}^2$	25.0	16.67	33.33	25.0	16.67	16.67	25.0	00.0	0.00	12.5	21.5
sf ₄	$b_{1,2}^2$	50.0	50.00	50.00	50.0	50.00	50.00	75.0	50.0	33.33	50.0	29.0
sf ₅	$b_{1,3}^1$	25.0	33.33	33.33	25.0	50.00	33.33	50.0	50.0	33.33	75.0	43.0
sf ₆	$b_{2,1}^1$	50.0	50.00	50.00	50.0	50.00	50.00	25.0	50.0	66.67	50.0	71.0
sf ₇	$b_{2,1}^2$	75.0	66.67	66.67	75.0	50.00	66.67	50.0	50.0	66.67	25.0	58.0
sum		275.0	266.67	266.67	275.0	250.00	266.67	250.0	250.0	266.67	225.0	258.0
Jain's SfFI		0.823	0.8312	0.8312	0.823	0.8242	0.8312	0.741	0.714	0.7619	0.625	0.778
Jain's FFI		0.776	0.7805	0.8205	0.776	0.8654	0.7805	0.961	0.893	0.7273	0.880	0.750
Jain's TFI		0.917	0.9143	0.8205	0.917	0.7922	0.9143	0.595	0.893	1.0000	0.570	0.850

link $\langle r_3, r_4 \rangle$ with limited capacity $\rho_2 = 100$ with a share of 50 each. The third bottleneck link $\langle r_5, r_6 \rangle$ with capacity $\rho_3 = 100$ is shared between subflows sf_5 and sf_7 but since sf_5 is already bottlenecked with a share of $b_{1,3}^1 = 25$, the remaining capacity of the link is allocated to subflow sf_7 i.e., $b_{2,1}^2 = 75$.

The *bottleneck flow fair* allocation shares the available bottleneck link capacity between flows and subflows. Therefore, $\rho_1 = 100$ capacity is shared between 3 flows (not 4 subflows) giving a share of 33.33 to each of the three flows (f_1, f_2 and f_3) initiated from the same user u_1 . Since both the subflows of flow f_1 are part of this bottleneck, the flow allocation is shared equally between the two subflows sf_1 and sf_3 . The higher-level participants such as locator or tariff share their respective allocations equally amongst their flows and then

the flow share is shared equally between its respective subflows (hierarchical equal share policy). For example, the *bottleneck tariff fair* allocation first shares the bottleneck link capacity $\rho_1 = 100$ equally between the two tariffs (t_1 and t_2). Both the tariffs t_1 and t_2 have two subflows of different flows over this bottleneck link and therefore they further share the tariff allocations equally. Even though flow f_1 has two subflows over this bottleneck, each subflow gets a share equal to the share of the single subflows belonging to other flows by the virtue of being part of two different tariffs.

The hierarchical equal share policy used by the bottleneck resource allocation algorithm is not trivial for the network resource because all the participants and their inter-dependency have to be considered. For example, a flow can allocate its share equally between its two subflows only if they are part of the same bottleneck and participant. In the investigated scenario (Fig. 1.2) both subflows of flow f_1 share the same bottleneck link and for participants such as flow (Eq. nf5), source (Eq. nsr5) and user (Eq. nu3) the two subflows can have equal allocation. If a subflow is part of more than one bottleneck then the network capacity ρ^n is not just the sum of the bottleneck link capacities. Therefore, the network capacity is defined to be the sum of all subflow capacities in the network (Eq. n1) where $b_{q,x}^a$ is the capacity assigned to the subflow $s_{q,x}^a$ of flow $f_{q,x}$ that is initiated by user u_q . The linear equation system representation of the scenario considered in Fig. 1.2 implemented in MATLAB is as follows.

- Network capacity equation:

$$\rho^n = b_{1,1}^1 + b_{1,1}^2 + b_{1,2}^1 + b_{1,2}^2 + b_{1,3}^1 + b_{2,1}^1 + b_{2,1}^2 \quad (\text{n1}).$$

- Bottleneck bounded set of equations:

$$b_{1,1}^1 + b_{1,1}^2 + b_{1,2}^1 + b_{1,3}^1 = \rho_1 \quad (\text{b1}); \quad b_{1,2}^2 + b_{2,1}^1 = \rho_2 \quad (\text{b2}); \quad b_{1,3}^1 + b_{2,1}^2 = \rho_3 \quad (\text{b3}).$$

Depending on the participant i. e., flow (nf1-nf5), tariff (nt1-nt4), etc. different set of fairness equations are used to obtain the respective allocation of subflow capacities.

- Fairness bounded set of equations for network flow fair (NFF):

$$b_{1,1}^1 + b_{1,1}^2 = \rho^n/4 \quad (\text{nf1}); \quad b_{1,2}^1 + b_{1,2}^2 = \rho^n/4 \quad (\text{nf2}); \quad b_{1,3}^1 = \rho^n/4 \quad (\text{nf3}); \\ b_{2,1}^1 + b_{2,1}^2 = \rho^n/4 \quad (\text{nf4}); \quad b_{1,1}^1 - b_{1,1}^2 = 0 \quad (\text{nf5}).$$

- Fairness bounded set of equations for network locator fair (NLF):

$$b_{1,1}^1 + b_{1,2}^1 = \rho^n/5 \quad (\text{nl1}); \quad b_{1,1}^2 + b_{1,2}^2 = \rho^n/5 \quad (\text{nl2}); \quad b_{1,3}^1 = \rho^n/5 \quad (\text{nl3}); \\ b_{2,1}^1 = \rho^n/5 \quad (\text{nl4}); \quad b_{2,1}^2 = \rho^n/5 \quad (\text{nl5}).$$

- Fairness bounded set of equations for network tariff fair (NTF):

$$b_{1,1}^1 + b_{1,2}^1 = \rho^n/4 \quad (\text{nt1}); \quad b_{1,1}^2 + b_{1,2}^2 + b_{1,3}^1 = \rho^n/4 \quad (\text{nt2}); \quad b_{2,1}^1 = \rho^n/4 \quad (\text{nt3}); \\ b_{2,1}^2 = \rho^n/4 \quad (\text{nt4}).$$

- Fairness bounded set of equations for network source fair (NSrF):

$$b_{1,1}^1 + b_{1,1}^2 + b_{1,2}^1 + b_{1,2}^2 = \rho^n/3 \text{ (nsr1)}; b_{1,3}^1 = \rho^n/3 \text{ (nsr2)};$$

$$b_{2,1}^1 + b_{2,1}^2 = \rho^n/3 \text{ (nsr3)}; b_{1,1}^1 - b_{1,1}^2 = 0 \text{ (nsr4)}; b_{1,1}^1 - b_{1,1}^2 = 0 \text{ (nsr5)}.$$

- Fairness bounded set of equations for network user fair (NUF):

$$b_{1,1}^1 + b_{1,1}^2 + b_{1,2}^1 + b_{1,2}^2 + b_{1,3}^1 = \rho^n/2 \text{ (nu1)}; b_{2,1}^1 + b_{2,1}^2 = \rho^n/2 \text{ (nu2)};$$

$$b_{1,1}^1 - b_{1,1}^2 = 0 \text{ (nu3)}.$$

The *network flow fair* allocation based on the linear set of equations n1, b1-b3 and nf1-nf5, as desired, results in an equal allocation of 60 to the competing four flows within the network. But in order to do so, it gave an allocation of -20 and 80 to subflows $s_{1,2}^1$ and $s_{1,2}^2$ respectively of flow $f_{1,2}$. A negative allocation implies that the competing flows on the bottleneck link have got an overall larger share than the bottleneck link capacity. Thus, with the help of a simple corrective algorithm the additional share is reduced equally from the competing flows $f_{1,1}$ and $f_{1,3}$. Thus a maximum possible allocation of 50 for the two flows due to topological restriction is obtained. A reduced set of linear equations related to the remaining network capacity and unallocated participants is solved to obtain an equal share of 75 between flows $f_{1,2}$ and $f_{2,1}$. The *network tariff fair* allocation results in each tariff getting an equal share of 66.66. Due to topological constraints i.e., tariff t_1 only shares the bottleneck link with capacity $\rho_1 = 100$, the tariff t_2 can only get a share of 33.33 on this link. Thus, the tariff t_2 cannot be shared equally amongst its three flows i.e., an equal share of 22.22. In addition, priority is given for a further fair share amongst flows leading to an allocation of 33.33 to flow $f_{1,3}$ and nothing to subflow $s_{1,1}^2$ as the other subflow $s_{1,1}^1$ of flow $f_{1,1}$ gets a share of 33.33 from tariff t_1 . A similar allocation discrepancy due to the topological limitation is observed for the *network source fair* allocation where all the three source nodes get an equal share of 75 each, but due to flow $f_{1,3}$ share of 75 from source sr_2 , flow $f_{1,1}$ can be allocated only 25 (limited by ρ_1) from the share of source sr_1 . The *network user fair* allocation can achieve an equal share of the user u_1 allocation amongst its 3 flows, each getting a share of 43.

From the obtained results it can be seen that if fairness is applied at the bottleneck then it is absolutely fair to the individual participants w.r.t. the bottleneck. Fairness mechanisms considering the whole network as a single resource is restricted by the topological constraints and might even result in a lower overall network throughput when a flow is part of multiple bottleneck links (Fig. 1.2). Flows can be chosen as participants if TCP compatible fairness is preferred while tariffs combine the individual flows and couple their share together w.r.t. economical aspects thereby allowing for a better implementation of service level agreements between the user and the service provider. The network flow fair solution for the investigated scenario leads to a highly unfair distribution of network capacity between the four different tariffs, which highlights the issue in selecting the flow as a participant for resource allocation.

1.6 Conclusion and Outlook

In this work, the investigated fairness mechanisms are not limited to universally accepted TCP-friendly notions but defined w. r. t. available resources (e. g. bottleneck and network) and competing participants (e. g. subflow, flow, locator, tariff, source and user). The discussed alternative fairness mechanisms extend the scope of fairness beyond a transport flow to include the higher layer aspects, the end user as well as the network. Furthermore, theoretical resource (bottleneck/network) allocation algorithms with regard to different fairness goals are presented as a means for a comprehensive analysis of the various fairness definitions. With the help of a carefully chosen scenario, different aspects of fairness in multipath transport are highlighted with the help of results obtained from the introduced theoretical resource assignment algorithms. In addition, the investigated scenario is kept small enough to validate the operation of the proposed algorithms.

To achieve a fair end-to-end solution without the aid of the network, congestion windows of subflows belonging to the same participant can be coupled. Internal weighting of the subflows corresponding to the same participant may be achieved at the end host but for weights to work at the participant level, network elements need to collaborate. Exchange of signalling information between the end host and the network can further enhance the performance of the congestion control mechanisms in achieving the desired goals specified in [10]. Thus, these theoretical allocation algorithms can be used not only to classify but also validate the performance of the multi/single-path congestion control algorithms for complex scenarios.

The theoretical algorithm for bottleneck fair allocation is presented as a flexible method which assumes that due to the elastic traffic, every subflow will be bottlenecked by the network and not restricted by the application. The algorithm can also be extended to include the application limited cases (real-time traffic). The algorithm allocates the fair share to every subflow by utilizing the subflow path mapping information. The functionality of choosing a limited set of subflows out of all the available subflows can also be added, if needed, but will introduce large computational complexity due to the number of possible combinations.

With a linear equation system, a fair share solution might not be obtained if the bottleneck link does not allow a participant to get its fair share. In this case, some of the subflows competing for the bottleneck link capacity will be assigned a negative allocation. In this work, an iterative method is used to overcome this shortcoming by systematically updating the capacity allocations. Also care is taken that no flow gets starved if the participant is a higher level entity such as tariff. In future work, a comprehensive solution which is capable of both fair resource allocation as well as optimum routing will be worked out based on mixed-integer (non-)linear programming.

References

1. J. B. Postel, "Transmission Control Protocol," IETF, Standards Track RFC 793, 1981.
2. M. Welzl, *Network Congestion Control: Managing Internet Traffic*. John Wiley & Sons, 2005.
3. R. Braden, D. D. Clark, J. Crowcroft, B. Davie, S. E. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K. K. Ramakrishnan, S. Shenker, J. Wroclawski, and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet," IETF, Informational RFC 2309, 1998.
4. R. H. Tse, J. H. Jannotti, and R. H. Tse, "TCP fairness in multipath transport protocols," Department of Computer Science, BROWN UNIVERSITY, Bachelor's thesis, 2006.
5. B. Briscoe, "Flow Rate Fairness: Dismantling a Religion," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 37, 2007.
6. H. Adhari, E. P. Rathgeb, A. Singh, A. Könsgen and C. Goerg, "Transport Layer Fairness Revisited," Submitted, 2015.
7. D. Papadimitriou, M. Welzl, M. Scharf, and B. Briscoe, "Open Research Issues in Internet Congestion Control," IRTF, RFC 6077, 2011.
8. T. Dreibholz, M. Becke, H. Adhari, and E. P. Rathgeb, "On the Impact of Congestion Control for Concurrent Multipath Transfer on the Transport Layer," in *Proceedings of the 11th IEEE International Conference on Telecommunications (ConTEL)*, 2011.
9. A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses," IETF, RFC 6824, 2013.
10. A. Ford, C. Raiciu, M. Handley, S. Barr, and J. R. Iyengar, "Architectural Guidelines for Multipath TCP Development," IETF, Informational RFC 6182, 2011.
11. C. Raiciu, M. Handley, and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols," IETF, RFC 6356, 2011.
12. D. Wischik, M. Handley, and M. B. Braun, "The Resource Pooling Principle," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 38, no. 5, 2008.
13. C. Raiciu, D. Wischik, and M. Handley, "Practical Congestion Control for Multipath Transport Protocols," University College London, Tech. Rep., 2009.
14. M. Becke, T. Dreibholz, H. Adhari, and E. P. Rathgeb, "On the Fairness of Transport Protocols in a Multi-Path Environment," in *Proceedings of the IEEE International Conference on Communications (ICC)*, 2012.
15. H. Adhari, M. Becke, and T. Dreibholz, "On the Fairness of Transport Protocols in a Multi-Path Environment," in *Proceedings of the 83rd IETF Meeting*, 2012.
16. A. Singh, M. Xiang, A. Könsgen, and C. Goerg, "Performance and Fairness Comparison of Extensions to Dynamic Window Coupling for Multipath TCP," in *9th International Wireless Communications and Mobile Computing Conference (IWCMC)*, 2013.
17. A. Singh, M. Xiang, A. Könsgen, C. Goerg, and Y. Zaki, "Enhancing Fairness and Congestion Control in Multipath TCP," in *6th Joint IFIP Wireless and Mobile Networking Conference (WMNC)*, 2013.
18. R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," in *CoRR*, vol. cs.NI/9809099, 1998. [Online]. Available: <http://arxiv.org/abs/cs/9809099>.