Transport Layer Fairness Revisited

Hakim Adhari, Erwin P. Rathgeb
University of Duisburg-Essen
Institute for Experimental Mathematics
Ellernstraße 29, 45326 Essen, Germany
{hakim.adhari, rathgeb}@uni-due.de

Amanpreet Singh, Andreas Könsgen, Carmelita Goerg
University of Bremen
Department of Communication Networks
Otto-Hahn-Allee 1, 28359 Bremen, Germany
{aps, ajk, cg}@comnets.uni-bremen.de

***Abstract*—Fairness amongst the competing flows at the transport layer has always been an important topic, however, the current definition based on the TCP-compatible view is not always suitable. With the increasing deployment of multipath transport protocols such as Multipath TCP (MPTCP) and the Concurrent Multipath Transfer extension of SCTP (CMT-SCTP), the term "fair" can have various interpretations. In this paper, inconsistencies are avoided by classifying fairness definitions according to the resource – bottleneck or network – and the participants – subflow, flow, tariff, etc. that share the resource. With example network scenarios the current (TCP-compatible) fairness view from both the single and multipath perspective is presented and their shortcomings discussed. Alternative definitions are introduced and their benefits are illustrated based on a theoretical analysis. The realization aspects of the discussed fairness definitions are also presented. The evaluations of available coupled congestion control variants for multipath transport are shown to highlight the proximity of the simulated results to the theoretical target values. Due to the complexity of the realization of network-based approaches, bottleneck is chosen as the preferred resource. Tariff is a promising participant as it couples applications and incorporates economic entities for fair resource sharing at the transport layer.** [1]

Keywords: Fairness, Congestion Control, TCP, Multipath

I. Introduction

In the current Internet, it is considered as fair not to push away TCP flows. Therefore, the common definition of fairness in the Internet is called TCP friendliness. TCP-friendly flows are also called TCP compatible and defined by RFC 2309: *"A TCP-compatible flow is responsive to congestion notification, and in steady state it uses no more bandwidth than a conformant TCP flow running under comparable conditions"* [1], [2]. Protocols commonly meet this requirement by using some form of Additive Increase/Multiplicative Decrease (AIMD) congestion window management, or by computing a transmission rate based on equations derived from an AIMD model.

This kind of flow rate fairness centered on a flow, has been criticized for not being based on any respected definitions of fairness from philosophy or social sciences [3]. In fact while handling flows on the transport layer, there is a big architectural vacuum on upper layers. In addition to it, the topic became more confusing with the standardization of the multipath protocols. The standardization community decided to remain by the notion of TCP-friendly flows even if a flow as known from a singlepath environment is different from the new kinds of flows used by the multipath protocols. This leads to design penalization of multipath flows to the advantage of singlepath flows [4] [5].

In this work, fairness is defined in general as the way how available *Resources* are distributed between different *Participants*. The main goal here is not to define a new metric such as Max-Min Fairness [6], Proportional Fairness [7] or Weighted Proportional Fairness [8] but to provide an overview of the notion about participants and resources considered in the past and to propose new ways that still tackle fairness at the transport layer but also considers the higher layer aspects, the end user as well as the network. Thereby this work improves the overall fairness and at the same time clears confusions.

This paper is structured as follows: in Section II the important terms which are essential for the fairness discussion are defined in an abstracted form. The definition of the different fairness mechanisms follows in Section III and various aspects about the mechanisms are discussed in Section IV. After that, the practicability of the different views is considered. In a first step, the realization of the different fairness views is discussed (Section V-A) and in a second step, the existing coupled congestion control variants that realize different fairness aspects are evaluated in Section V-B. Finally, Section VI concludes this paper with a summary.

II. Terminology

In a first step, in order to avoid misunderstandings and confusions, a formal abstraction is given for the fundamental terminology needed for defining and discussing the different views on network resources and participants.

Definition II.1. 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.

L defines a finite set of unique locators. A locator $l \in L$ could e.g. be a Network Interface Card (NIC) or an IP address. A node could include one or multiple locators. The connectivity among the locators is described by the connectivity set $C \subseteq L \times L$. N_{sr} and N_{ds} are the set of source and destination nodes, respectively. n_{sr} and n_{ds} are elements of N_{sr} and N_{ds}, respectively.

Definition II.2. Flow:
A flow f between two nodes n_{sr} and n_{ds} is composed of all

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

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 connectionless and using a path set P .

Definition II.3. Subflow:
A subflow s denotes the subset of PDUs belonging to f and using a specific path $p \in P$. The bandwidth allocated to subflow k is denoted as b_k which is the bandwidth ρ_{p_k} of the path p_k.

Definition II.4. Bottleneck:
Let (λ_i, λ_j) be a link in Γ with the bandwidth $\rho(\lambda_i, \lambda_j)$. The set of subflows crossing this link builds the subflow set $S\langle i, j\rangle$. (λ_i, λ_j) is considered as a bottleneck if

$$\sum_{\forall k \in S\langle i,j\rangle} b_k = \rho(\lambda_i, \lambda_j).$$

There are two kinds of bottlenecks:

- *A simple bottleneck* describes the case where only one subflow is crossing the link ($|S\langle i, j\rangle| = 1$).
- *A shared bottleneck* describes the case where the link is crossed by multiple subflows ($|S\langle i, j\rangle| > 1$).

Definition II.5. Tariff:
A tariff is defined by the cost for using the service of an access network. Depending on the tariff plan this cost can be based on the data volume transferred, the time for which connections were open or the maximum bandwidth.

III. Fairness definitions for Multipath Congestion Control

The Internet of today is dominated by singlepath TCP flows [9] where every flow should have a TCP-compatible behavior [10] i. e., a bottleneck link capacity should be equally shared by the competing flows. Therefore, the resource which should be divided is the *bottleneck* capacity and the entities or participants demanding a fair allocation are *singlepath flows*. This fairness method can be referred as *bottleneck flow fairness*.

With the emergence of multipath protocols, the simple TCP-compatible notion of a flow has become unclear i. e., whether the flow now means a multipath flow or a subflow which in multipath terminology is comparable to a singlepath flow. With the new terminology, a singlepath flow has a single subflow i. e., the ratio of flow to subflows is 1 : 1. But a multipath flow can have several (n) subflows and hence the ratio of a flow to its subflows is $1 : n$. If the TCP congestion control e. g., NewReno is used for multipath transport then every subflow will behave as a singlepath TCP connection. In this case, the realized fairness is not called *bottleneck flow fairness* but *bottleneck subflow fairness* i. e., it considers a *subflow* as a participant and as earlier a *bottleneck* as the resource that should be apportioned.

Thus, the confusion in mapping the singlepath flow to a multipath flow or subflow breaks the existing TCP-compatible fairness definition into two possible cases. At this stage, the standardization community (IETF) felt that the multipath extensions with fairness applied on subflow level may unfairly influence the singlepath dominated network [11], [12]. Therefore, the multipath flow should be TCP-compatible i. e., *bottleneck flow fairness* is the desired goal.

For a multipath flow to be *bottleneck flow fair*, it requires that all its subflows on the respective bottleneck link do not get a combined share bigger than that of a singlepath flow. This led to the idea of coupling the congestion control (CC) of the subflows that share a common bottleneck. In this regard, the Dynamic Window Coupling (DWC) approach and its variant (DWC-DD) have been proposed in [13] and [**?**]. This method couples congestion window growth of subflows that are identified to be sharing a common bottleneck. The two variants differ in the method of grouping subflows together. Bottleneck detection approaches are still not fully reliable and hence they are an important topic in the research community [14].

In order to deal with this issue, an alternative based on the idea of *Resource Pooling (RP)* was introduced in [15]. It is based on the approach of making network resources behave like a single pooled resource. With this idea the two sets of goals: *functional goals* and *compatibility goals* specified in [11] for multipath TCP are extended with an additional goal to *balance congestion* [12], [16]. Therefore, RP intends, by design, to shift traffic from more to less congested paths in order to release resources on the congested paths and hence increase the performance of the whole network. Thus, RP introduced a new perspective on fairness at the resource level, denoted as *network flow fairness*, where the participant is still a *flow* but the resource to be divided is the considered *network* instead of the *bottleneck*. The RP principle has been adopted by several proposed solutions such as *Linked Increases Algorithm (LIA)* [16], *Opportunistic Linked Increases Algorithm (OLIA)* [17] or *Adapted OLIA* [**?**] for MPTCP [18] and *Resource Pooling Multipath version 2 (RP-MPv2)* for CMT-SCTP [19].

A. Issues related with the current fairness views

The impact of the use of the flow as a participant and network as a resource can be explained based on the basic scenarios shown in Fig. 1.

For the topology depicted in Fig. 1(a), a bottleneck subflow/flow fair solution will share the link capacity ρ_β equally between the subflows sf2 and sf3 whereas the network flow fair solution will share the whole network capacity $\rho_\alpha + \rho_\beta$ equally (if $\rho_\alpha \leq \rho_\beta$) between the flows f1 and f2 and thus benefiting the single path flow f1 at the expense of multipath flow f2 even though the multipath flow makes use of more resources.

Fig. 1(b) depicts another multipath scenario where the subflows sf1 and sf2 of flow f1 share a common bottleneck with a subflow of another flow. In this case, the bottleneck subflow fairness will share the bottleneck link capacity ρ_α between the three subflows equally i. e., flow f1 will get double the capacity w. r. t. flow f2. A flow based fairness for this scenario will give the both flows an equal share.

In the scenario shown in Fig. 1(c), two source nodes S1 and S2 are sending traffic to the destination nodes D1 and D2. Two independent flows f1 and f2 are scheduled from S1 and only one flow f3 is scheduled from S2. In this case, the bottleneck capacity ρ_α is apportioned equally through all the existing flows f1, f2 and f3 and this is obviously considered as fair. Now if the scenario depicted in Fig. 1(b) is re-considered then, in both cases (Topology 1(c) and 1(b)) S1 is scheduling two subflows, where in the first case the two subflows belong to different flows (f1 and f2) and in the second case both subflows belong to the same flow f1. Thus, a flow-based fairness approach which is also TCP-friendly will give different allocations for the considered two cases. This reflects a conflict between the fairness established on the transport layer and the behavior of applications on higher layers i. e. if an application initiates multiple flows then it can get a larger share whereas an application that initiates a multipath flow is not entitled for the larger share.

A subflow-based fairness approach will remove all the discrepancies highlighted so far in this section i. e. preference to single path flows at the expense of multipath flows. But the question remains which approach is the most fair. Fairness metrics such as Max-Min Fairness [6] cannot be used here as the fundamental definition of a participant need to be clarified and fairness metrics are both dependent on the participant/resource combination as well as the network topology. Considering the scenario shown in Fig. 1(d), neither sharing the bandwidth ρ_α equally between the three flows nor between the 4 subflows seems to be fair to all the participants. At this point a simple question from the design point of view needs to be addressed, should the participant view be made broader than a flow and involve the application and utilized resources as well?

B. Proposed alternative fairness definitions

In order to deal with dubious fairness definition issues due to the traditional participant flow, new participants are defined in this work. In addition, methods to define the set of subflows to be coupled together are proposed. Table I shows an overview of all the fairness views considered in this work.

For multipath flows it is highly likely that not all the corresponding subflows have the same end locator. A new proposed approach is to share the resource amongst the different locators fairly. Thus, for *network locator fairness* all the (sub)flows initiated from the same locator are coupled together and for *bottleneck locator fairness* all (sub)flows that share the same bottleneck and locator are coupled together even if they belong to different flows. For the topologies depicted in Fig. 1(b) and 1(d), the bottleneck link between the routers R3 and R4 is shared between 3 locators. Therefore in Fig. 1(b), the share of flow f1 is $2 \times \frac{\rho_\alpha}{3}$. In Fig. 1(d), the first locator is shared by two subflows (belonging to a singlepath flow and a multipath flow) and therefore these two subflows share their corresponding locator share ($\frac{\rho_\alpha}{3}$). Basing the resource allocation on the locator instead of flow as the participant rewards the extra effort made in using more than one path (locator) for the flow as well as restricts any misuse made on upper layers by applications that start multiple flows simultaneously through the same locator.

However, even if a locator based fairness view is a huge improvement, the one-to-one liaison between the locator and the resource allocation reflects only a minimized view on how networks nowadays are handled. The reality at every Internet Service Provider (ISP) is that users are supposed to get what they are paying for and that at every ISP, different tariffs are offered based on the different characteristics of the connection. In fact, inside the network of one ISP, a provider has the task to make the users accountable and has the obligation to deliver specific characteristics promised by the tariff. In this work, a view that could easily be adopted by the ISPs in order to involve the economical entities into their calculation is proposed. Here the ISP is free to define its utility function and to associate weights to its different pricing models and based on it to its users. Each tariff, whether inside the scope of bottleneck or the ISP, is then assigned a capacity share proportional to the weighting factor. Thus, for *network tariff fairness* all the (sub)flows initiated from the same tariff and being part of the same network are coupled together while for *bottleneck tariff fairness* all (sub)flows that share the same bottleneck and tariff are coupled together. In general, care must be taken that only those (sub)flows which share a common resource and participant are coupled together and not all of them.

Outside the ISP network, the complete network could be seen as a hierarchical graph where the current ISP network Γ_{child} is a child of a parent network Γ_{parent} and where inside Γ_{parent} also a weighting factor is associated to Γ_{child}. Based

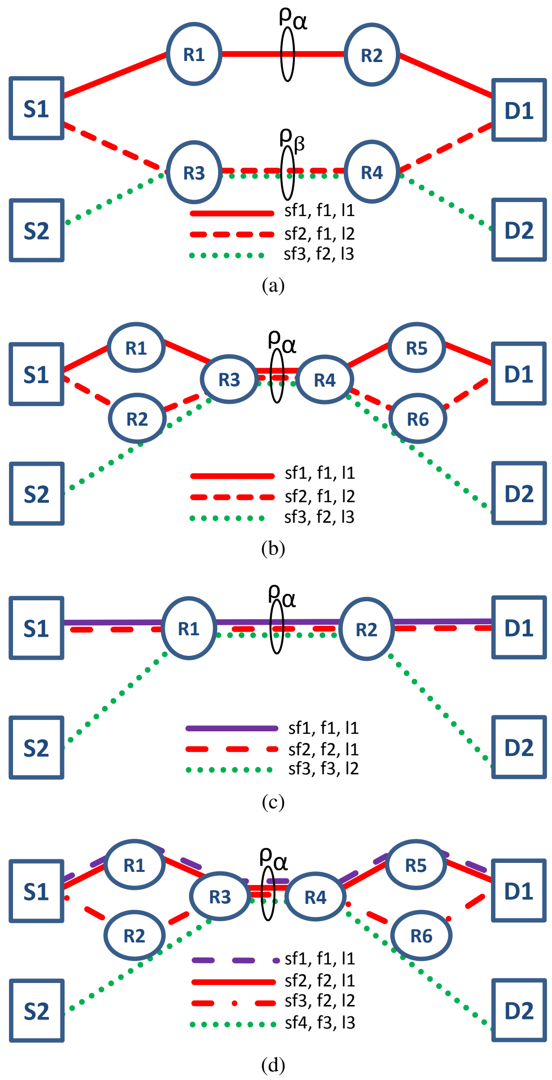

Figure 1. Scenarios to highlight the shortcomings of traditional fairness definitions

Table I. OVERVIEW OF FAIRNESS VIEWS

Participant/Resource	Network (N)	Bottleneck (B)
Subflow (sf)	Network Subflow Fair (NS)	Bottleneck Subflow Fair (BS)
Flow (f)	Network Flow Fair (NF)	Bottleneck Flow Fair (BF)
Locator (l)	Network Locator Fair (NL)	Bottleneck Locator Fair (BL)
Tariff (t)	Network Tariff Fair (NT)	Bottleneck Tariff Fair (BT)

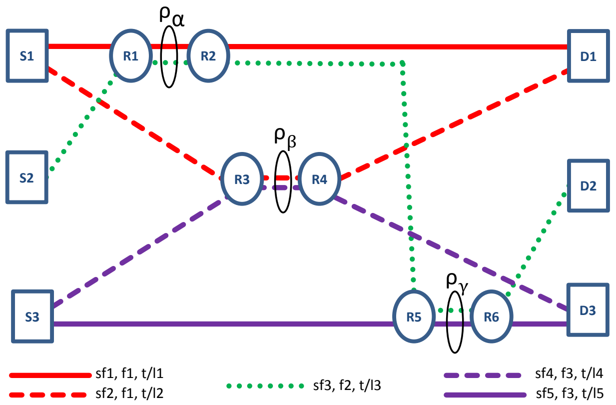

Figure 2. Consecutive bottlenecks scenario (3 flows, 5 tariffs/locators and 5 subflows)

Table II. CAPACITY SHARE IN MBIT/S FOR THE SCENARIO IN FIG. 2 ρ_α=ρ_β=ρ_γ=12 MBIT/S

sf#	BS/BF/BL/BT/NS/NL/NT	NF
1	ρ_α-min(ρ_α,ρ_γ)/2 = 6	3
2	ρ_β/2 = 6	6
3	min(ρ_α,ρ_γ)/2 = 6	9
4	ρ_β/2 = 6	6
5	ρ_γ-min(ρ_α,ρ_γ)/2 = 6	3
c^{fc}_{max}	30	27

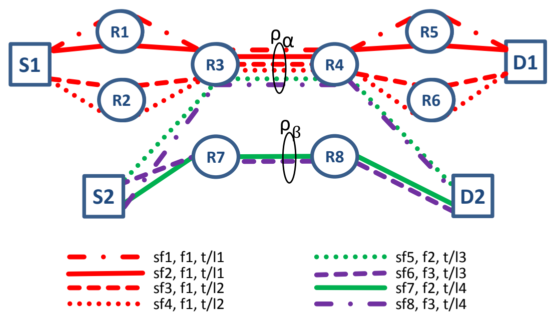

Figure 3. Parallel bottlenecks scenario with several participants (3 flows, 4 tariffs/locators and 8 subflows)

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_{parent} 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}.

IV. DISCUSSION OF THE DIFFERENT FAIRNESS VIEWS

In this section, two important points, related with the different fairness views introduced in Section III, are further discussed w. r. t. the scope of identified participants and resources.

A. Quantifying the network as a resource

In this paragraph, the case is considered where the network is the resource. The network capacity c in general is defined as the pooling capacity of a network section which basically results from the capacities of the bottleneck links in the network. For example, in the scenario shown in Fig. 1(a), in the case that no limitations are set on the access links, $c = \rho_\alpha + \rho_\beta$. In general, this is in most of the case considered as an obvious case in the literature such as in [15].

However, determining the network capacity is not always that obvious. As an example, the topology shown in Fig. 2 is considered where $\rho_\alpha = \rho_\beta = \rho_\gamma = 12$ Mbit/s. Let $c_{\max}$ be the fairness independent network capacity which is the maximum throughput that can be reached between the source and destination node sets $N_{sr} = \{S1, S2, S3\}$ and $N_{ds} = \{D1, D2, D3\}$, respectively, regardless if any fairness criterion is considered or how much capacity is assigned to a subflow. Obviously, $c_{max} = 36$ Mbit/s can be reached in the case subflow sf3 is completely excluded ($b_3 = 0$ Mbit/s).

c^{fc}_{max} is defined as the fairness criterion (fc) dependent network capacity which is the maximum throughput that can be reached between N_{sr} and N_{ds} based on the fairness criterion fc.

Table II gives the capacity (Mbit/s) assignments in case of different fairness methods for the scenario depicted in Fig. 2. The numerical results are given assuming that each of the bottleneck links in the figure has a speed of 12 Mbit/s. A multipath transport enabled participant distributes its total load across multiple subflows over different paths. A bottleneck link on the path may be shared by subflows which may belong to the same or different participants. Hence there are various mutual dependencies between capacity assignments for the different participants and subflows.

If a network fairness view is considered, a dilemma is faced. In order to allocate the resources, the network capacity needs to be calculated. On the other hand, it must be known how much resources are allocated to each subflow in order to be able to determine the network capacity. One possible approach to solve this dependency is forming a linear set of equations if an equal share is possible. Otherwise linear programming methods that allow inequality may be used. The network flow fair allocation leads in this case to an equal share for every flow $A = 9$ Mbit/s and the $c^{\text{network-flow}}_{\max} = 27$ Mbit/s.

Thus, it can be seen from Table II that the resource pooling idea of sharing the network as the resource can lead to different overall throughputs depending on the participant and hence is not always the most efficient solution.

B. Hierarchical participants

Since the entity that is actually composed of data is the subflow, the calculation of the subflow share due to different

Table III. BOTTLENECK CAPACITY SHARE IN MBIT/S FOR SCENARIO IN FIG. 3
SF# SHARE = (CAPACITY/NUMBER OF PARTICIPANTS)/ NUMBER OF SUBFLOWS PER PARTICIPANT

Sf#	BS	BF	BL/BT
1	ρ_α/6	(ρ_α/3)/4	(ρ_α/4)/2
2	ρ_α/6	(ρ_α/3)/4	(ρ_α/4)/2
3	ρ_α/6	(ρ_α/3)/4	(ρ_α/4)/2
4	ρ_α/6	(ρ_α/3)/4	(ρ_α/4)/2
5	ρ_α/6	(ρ_α/3)/1	(ρ_α/4)/1
6	ρ_β/2	(ρ_β/2)/1	(ρ_β/2)/1
7	ρ_β/2	(ρ_β/2)/1	(ρ_β/2)/1
8	ρ_α/6	(ρ_α/3)/1	(ρ_α/4)/1

Table IV. NETWORK CAPACITY SHARE IN MBIT/S FOR SCENARIO IN FIG. 3,
NF: $2\rho_\alpha \leq \rho_\beta$, NT/NU: $\rho_\alpha \leq \rho_\beta$

Sf#	NS	NF	NL/NT
1	ρ_α/6	(ρ_α/1)/4	(ρ_α/2)/2
2	ρ_α/6	(ρ_α/1)/4	(ρ_α/2)/2
3	ρ_α/6	(ρ_α/1)/4	(ρ_α/2)/2
4	ρ_α/6	(ρ_α/1)/4	(ρ_α/2)/2
5	ρ_α/6	0	0
6	ρ_β/2	(ρ_β/2)/1	(ρ_β/2)/1
7	ρ_β/2	(ρ_β/2)/1	(ρ_β/2)/1
8	ρ_α/6	0	0

fairness methods in a hierarchical way is the focus in this subsection. Fig. 3 shows a scenario which is suitable to highlight the difference between various fairness methods with subflows, flows and locators/tariffs as the participants and bottlenecks or network as the resource. There are two source/destination node pairs S1/D1 and S2/D2, respectively. S1 is connected to two points of attachment R1 and R2 with different tariff plans. S1 runs one flow, consisting of four subflows where two are routed via R1 and the other two via R2. S2 has two flows with two subflows each, one subflow of each flow is connected to R3, the other one to R7 which are access points with different tariff models. Hence bottleneck link R3–R4 with capacity ρ_α is shared by six subflows, bottleneck link R7–R8 with capacity ρ_β is shared by two subflows. Due to the topology of the network, each fairness method yields different participant shares as shown in Tables III to V. In contrast to the scenario in Fig. 2, it is possible in this special case to give the capacity assignments for the different participants as closed expressions which give a good overview of how the available capacity is shared between different participants.

Table III gives the results for bottleneck fairness, corresponding to the fairness cases shown in the right column of Table I. In case of bottleneck subflow fairness, the capacity ρ_α and ρ_β are equally shared by the number of subflows sharing the respective link. In case of the other fairness views,

Table V. NETWORK CAPACITY SHARE IN MBIT/S FOR SCENARIO IN FIG. 3,
NF: $2\rho_\alpha > \rho_\beta$, NT/NU: $\rho_\alpha > \rho_\beta$

Sf#	NS	NF	NL/NT
1	ρ_α/6	(($\rho_\alpha + \rho_\beta$)/3)/4	(($\rho_\alpha + \rho_\beta$)/4)/2
2	ρ_α/6	(($\rho_\alpha + \rho_\beta$)/3)/4	(($\rho_\alpha + \rho_\beta$)/4)/2
3	ρ_α/6	(($\rho_\alpha + \rho_\beta$)/3)/4	(($\rho_\alpha + \rho_\beta$)/4)/2
4	ρ_α/6	(($\rho_\alpha + \rho_\beta$)/3)/4	(($\rho_\alpha + \rho_\beta$)/4)/2
5	ρ_α/6	((2$\rho_\alpha - \rho_\beta$)/6)/1	(($\rho_\alpha - \rho_\beta$)/4)/1
6	ρ_β/2	(ρ_β/2)/1	(ρ_β/2)/1
7	ρ_β/2	(ρ_β/2)/1	(ρ_β/2)/1
8	ρ_α/6	((2$\rho_\alpha - \rho_\beta$)/6)/1	(($\rho_\alpha - \rho_\beta$)/4)/1

the capacity is shared on two levels – the higher level is the participant which is the goal of the respective fairness scheme, i.e. flow, tariff, and the lower level is the number of subflows belonging both to the same participant and sharing the same bottleneck. For example, in case of BF, three flows share the bottleneck ρ_α, and therefore the flow share of the link R3–R4 is ρ_α/3. Since subflows 1 to 4 belong to f1, their share of the total capacity is (ρ_α/3)/4.

For network fairness, again the above-mentioned participants have to be considered, which results in the cases given in the left column of Table I. Further, there are subcases e.g., for NT $\rho_\alpha \geq \rho_\beta$ or $\rho_\alpha \leq \rho_\beta$, shown in Tables IV and V, respectively. For NS, the results are the same as for BS – for subflows it is irrelevant whether they share a bottleneck or a network. The total network capacity ρ_α+ρ_β is on the higher level equally shared by the number of participants – 3 flows for NF and 4 tariffs for NT. On the lower level, the upper participant-level share is then equally divided by the number of subflows both belonging to the same parent participant as well as sharing the same bottleneck, e.g. in case of subflows 1 to 4, all four subflows belong to the same flow, two of them each belong to the same tariff and all four subflows share the same bottleneck. This two-level capacity sharing yields the double denominators which can be seen in columns NF and NT. Subflows 6 and 7 share the lower bottleneck, the denominators are determined analogously to subflows 1 to 4. Subflows 5 and 8 are a special case which is determined as follows. Each of the participants in the network gets an equal share of the network capacity ρ_α+ρ_β as described earlier. Only flow 2, tariffs 3 resp. 4 uses the bottleneck link R7–R8 with capacity ρ_β. If these subflows cannot provide the capacity which the participants should get according to the fair network share, the remaining parts of the share are transported through subflows 5 and 8 via link R3–R4 with capacity ρ_α.

V. THE DIFFERENT FAIRNESS VIEWS IN PRACTICE

Aspects that are important for realizing the discussed fairness mechanisms are briefly mentioned in this section. The congestion control variants proposed in the literature for the multipath transport are also evaluated with the help of simulations.

A. *Realization of the fairness views*

One of the important points related with this discussion is the level where the fairness is established. In fact, which instance is responsible for establishing fairness is a decisive criterion which sets limits for the amount of share that can be obtained. In order to make this clear, the notion of the resource and participant is considered separately.

Using the *subflow* as a participant turns out to be an easy task and can obviously be realized on the source nodes with minimum of efforts. Considering *flow* as a participant is more complicated. Since a flow, in case of using multipath could include more than one subflow, mechanisms have to be involved in order to be able to couple subflows belonging to a single flow. In this case, this could be done by resource pooling based CCs.

The *locator* view of the participant is based on coupling not only subflows belonging to one flow but subflows that may belong to different flows initiated by different applications. A way to implement such a view of the participant is to use a solution based on sharing information between different flows such as the congestion manager [20] or TCB interdependence [21]. The congestion manager for example maintains congestion parameters (bandwidth, round-trip times, etc.) and makes it possible for applications to learn about network characteristics, to pass information to the congestion manager, to share congestion information with each other and to schedule data transmissions [20].

However, the biggest leap is moving to the *tariff* participant view. While all other notions of participants could be realized at the source nodes, this might not be the case with the *tariff* view. The tariff based fairness solution is hard to realize without an active participation of the infrastructure. For establishing this fairness view, adaptations are needed at least at the access routers connecting the end hosts in order to mark the different subflows and enable other routers to identify the right weighting factor for every subflow. Obviously a management entity [22] having full knowledge of all the active flows belonging to a tariff would involve a major simplification of the deployment of such a fairness view.

Implementing fairness from the resource view is also complicated. Reliable bottleneck detection mechanisms are needed to identify shared bottlenecks in order to ascertain the respective bottleneck and network capacity. In addition, it is important to identify subflows that share a bottleneck or network to be coupled together. Existing solutions at the source nodes are able to approach the real values however in multiple cases, this turns out to be difficult. Here also information from the network operator or a central management entity might be needed to achieve fairness in all cases.

B. Evaluation of the current CC methods

In this section the behavior of the different existing coupled congestion (CC) control mechanisms based on bottleneck/network as the resource and subflow/flow as the participant are analyzed by means of simulations to highlight the gap between the simulated and the theoretical values obtained in Section IV. For this analysis, the QualNet-4.0 simulator has been used with the MPTCP implementation [**?**] [**?**].

In the first part of the evaluation, we consider the scenario shown in Fig. 3 ($\rho_\alpha = \rho_\beta = 10\,\text{Mbit/s}$), with focus on which level the capacity assignment is performed. Fig. 4 to 6 show how well the different CCs are able to fulfill the theoretical fairness goals which were introduced in Section III.

The bottleneck subflow fairness properties of the different CC mechanisms are investigated with Fig. 4, which depicts the capacity assignments to the individual subflows. All CC mechanisms assign equal share to the subflows which belong to the same participant (i.e., flow) and share the same bottleneck. The equal share between subflows on the same bottleneck but belonging to different participants achieved by the schemes optimized for network fairness, i.e. (O)LIA is just by chance and not by design as is the case for NewReno.

In Fig. 5, the capacities of all subflows belonging to the same flow and sharing the same bottleneck have been summed up, so the figure shows the bottleneck flow fairness properties. All CC schemes use the link R7–R8 with capacity ρ_β which is

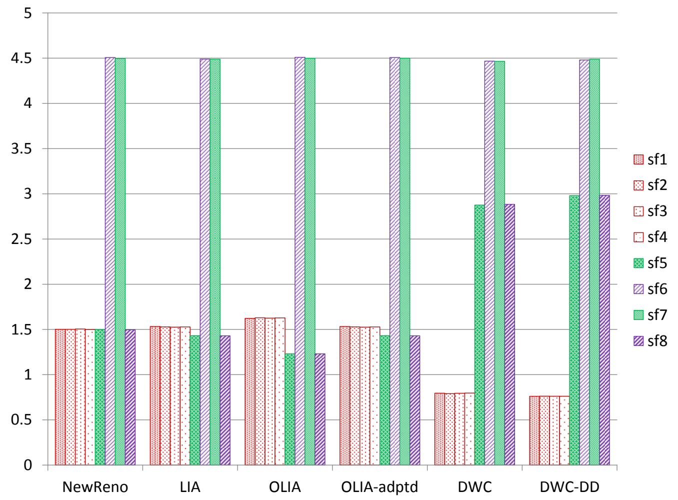

Figure 4. Capacity assignment in Mbit/s of the subflows for the different CCs applied on scenario in Fig. 3. ρ_α=ρ_β=10 Mbit/s.

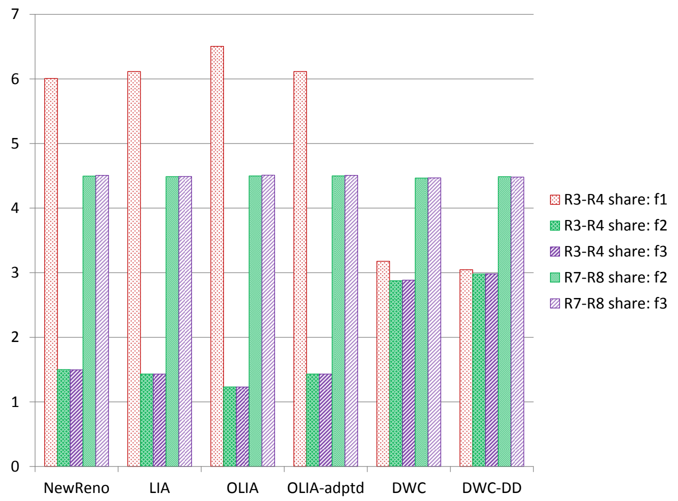

Figure 5. Capacity assignment in Mbit/s of the flows at the bottleneck links for the different CCs applied on scenario in Fig. 3. ρ_α=ρ_β=10 Mbit/s.

shared by flows f2 and f3 to the maximum possible extent, in order to remove load from link R3–R4 with capacity ρ_α. The bottleneck-oriented schemes perform well w. r. t. the bottleneck flow fairness with DWC-DD being the better variant.

The total capacity assignment for the entire flows is shown in Fig. 6 which gives an idea about the network flow fairness. It can be seen that the network fairness oriented approaches LIA, OLIA and adapted OLIA assign approximately the same capacity to all flows with a small negative bias towards flows that share distinct bottlenecks. The equal share of flows obtained by NewReno is just coincidental.

In the next part of the evaluation, the scenario shown in Fig. 2 is considered. The focus of this experiment is to find out to which extent the RP principle goals can be reached by the corresponding CCs. In fact RP introduced the feature of *balance congestion* [15] and motivated with it the shift of the resource view from the bottleneck to the network. The *balance*

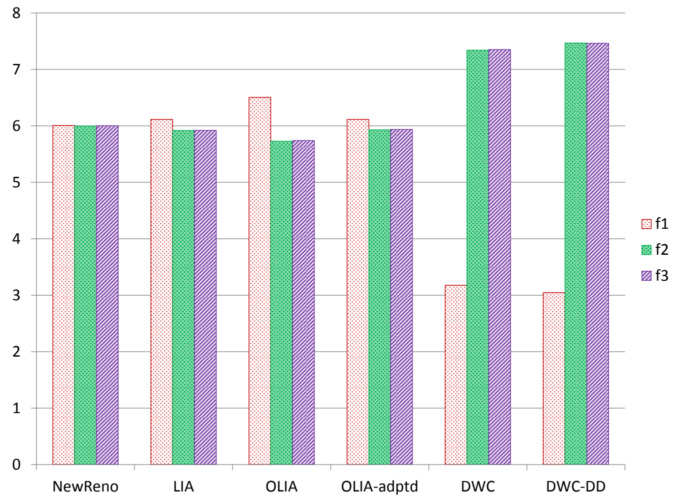

Figure 6. Capacity assignment in Mbit/s of the flows within the network for different CCs applied on scenario in Fig. 3. ρ_α=ρ_β=10 Mbit/s.

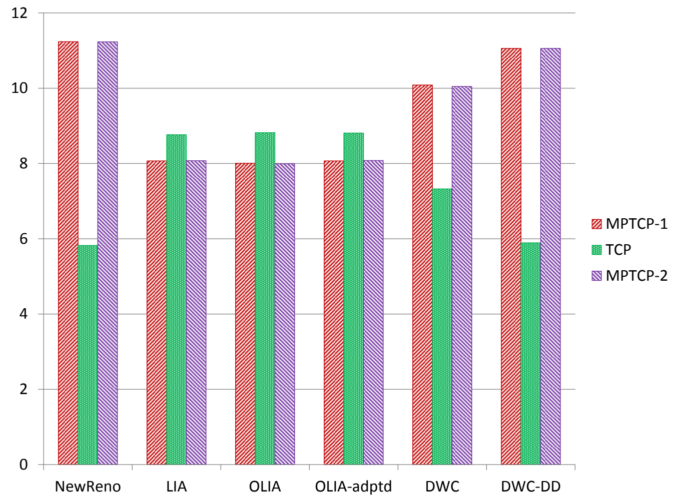

Figure 8. Capacity assignment in Mbit/s of the flows for the different CCs applied on scenario in Fig. 2. ρ_α=ρ_β=ρ_γ=12 Mbit/s.

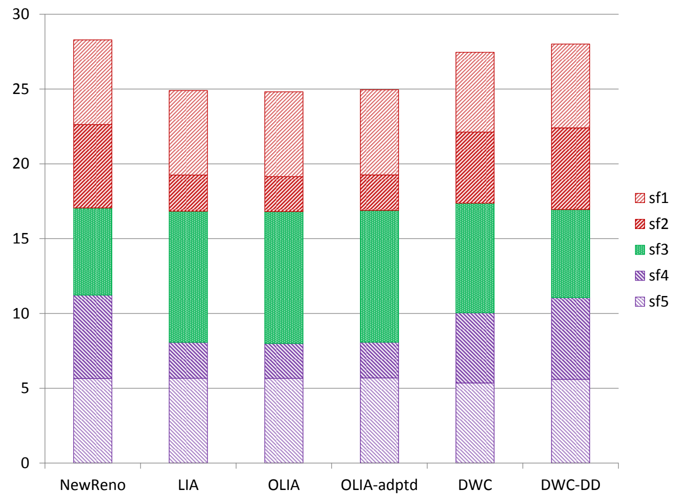

Figure 7. Network capacity in Mbit/s with different CCs applied on scenario in Fig. 2. ρ_α=ρ_β=ρ_γ=12 Mbit/s.

congestion goal was thought to lead to an equal share of the network resources and to improve the network utilization. This definitely holds for multiple cases, especially for the typical multipath scenarios presented in [15]. However, the results depicted in Table II for the scenario described in Fig. 2 show that the use of resource pooling can also lead to an under-utilization of the network compared to the TCP-compatible fairness criteria bottleneck flow fair or other fairness methods. This result is confirmed by the simulation results depicted in Fig. 7. For each CC used, the total sum of the throughput values reached by all subflows in the network is displayed. The results confirm the statement that depending on the CC chosen, different levels of network efficiency can be reached. The theoretical values are not reached but there is a tendency that the uncoupled CC as well as the DWC are able to reach a higher network capacity than the RP CCs LIA and OLIA in this scenario. It is observed that the network flow fair solution aims to give all the flows an equal share of the network resource but flow f2 is part of two bottlenecks and hence reduces the overall network throughput. These results do not mean that RP always leads to lower network capacity but should be considered as a counterexample for the claim of RP to always reach a better network efficiency [23]. The maximum network throughput can only be obtained, if the flow f2 gets no share and this will be against the *do no harm* goal of the CC mechanisms. The bottleneck flow fair solution on the contrary can satisfy both the goals defined in [11] i.e., *improve throughput* and *do no harm*.

Fig. 8 shows the simulation results of the different CC mechanisms in terms of flow share. It can be seen that the RP-based mechanisms give advantage to the singlepath flow at the expense of the multipath subflows. Due to the coupling of the multipath subflow's congestion window, their respective increase is less aggressive in comparison to the competing singlepath flow over a bottleneck. In this case, the bottleneck subflow fair solution given by NewReno gives the ideal solution in terms of both network throughput and fairness amongst different flows. DWC-DD gives the best performance from the group of CCs that provide bottleneck flow fairness as depicted in Fig. 7 and 8.

In a summary, how well the current CCs are able to handle a multitude of hierarchical participants is depicted in this section. This is remarkable w. r. t. to the small knowledge they have about the internals of the network, i. e. unlike the theoretical capacity assignments based on a perfect network knowledge which were demonstrated in Section IV-A, they are based on rough estimations. However, it can be said that existing CC schemes for multipath transport provide, even in the similar paths case (i.e. paths having similar QoS characteristics) deficits realizing the current fairness views. Especially the RP based alternatives penalizes multipath flows to the benefit of singlepath flows. In addition to this, it has been shown that the use of RP CCs not always implies a higher network capacity.

VI. Conclusions

In this paper, aspects of fairness in legacy and multipath transport were discussed. The existing fairness views have been categorized according to different resources and participants

Based on a theoretical analysis and on simulations, multiple issues have been identified with the current choices of participant and resource. Therefore, to bridge the gap between the fairness aspects on the transport layer and further higher layers, new participants in the form of a locator and tariff have been introduced. It has been shown that the locator/tariff view of the participant is able to deal with the typical flow based issues such as the disadvantage faced by multipath flows and possible misuse started from layers over the transport layer. In addition, the tariff as a participant is seen as the means for the new fairness definitions to involve the economical entities which reflects a philosophical and social background.

On the other side, with the help of simple example scenarios, the difficulties related with choosing the network as a resource have been illustrated. With the newly introduced participants that extend over multiple flows, sharing of the network resource becomes even further complicated. Considering Bottleneck as the resource localizes the problem and hence simplifies the allocation process but determining the bottlenecks in an operational network has its own challenges. An active participation of the network would aid in determining the bottlenecks as well coupling the appropriate subflows with their respective tariff weights.

Therefore, *Bottleneck tariff fairness* is proposed as the most suitable fairness approach for all kind of flows, single or multipath flows initiated by one or different applications. This fairness approach is beneficial to both the end-user as well as the network provider.

References

[1] 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.

[2] M. Welzl, *Network Congestion Control: Managing Internet Traffic*. John Wiley & Sons, 2005.

[3] B. Briscoe, "Flow Rate Fairness: Dismantling a Religion," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 37, 2007.

[4] 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.

[5] 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.

[6] E. L. Hahne, "Round-Robin Scheduling for Max-Min Fairness in Data Networks," *IEEE Journal on Selected Areas in Communications*, vol. 9, 1991.

[7] R. Mazumdar, L. G. Mason, and C. Douligeris, "Fairness in Network Optimal Flow Control: Optimality of Product Forms," *IEEE Transactions on Communications*, vol. 39, 1991.

[8] F. Kelly, "Charging and Rate Control for Elastic Traffic," *European Transactions on Telecommunications*, vol. 8, 1997.

[9] J. B. Postel, "Transmission Control Protocol," IETF, Standards Track RFC 793, 1981.

[10] S. Floyd, "Congestion Control Principles," IETF, RFC 2914, 2000.

[11] A. Ford, C. Raiciu, M. Handley, S. Barré, and J. R. Iyengar, "Architectural Guidelines for Multipath TCP Development," IETF, Informational RFC 6182, 2011.

[12] C. Raiciu, M. Handley, and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols," IETF, RFC 6356, 2011.

[13] S. Hassayoun, J. Iyengar, and D. Ros, "Dynamic Window Coupling for multipath congestion control," in *19th IEEE International Conference on Network Protocols (ICNP)*, 2011.

[14] D. A. Hayes, S. Ferlin, and M. Welzl, "Practical Passive Shared Bottleneck Detection using Shape Summary Statistics," in *Proceedings of the 39th IEEE Conference on Local Computer Networks (LCN)*, 2014.

[15] D. Wischik, M. Handley, and M. B. Braun, "The Resource Pooling Principle," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 38, no. 5, 2008.

[16] C. Raiciu, D. Wischik, and M. Handley, "Practical Congestion Control for Multipath Transport Protocols," University College London, Tech. Rep., 2009.

[17] R. Khalili, N. Gast, M. Popović, and J.-Y. L. Boudec, "MPTCP is not Pareto-Optimal: Performance Issues and a Possible Solution," *IEEE/ACM Transactions on Networking*, vol. 21, no. 5, 2013.

[18] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses," IETF, RFC 6824, 2013.

[19] 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.

[20] H. Balakrishnan and S. Seshan, "The Congestion Manager," IETF, RFC 3124, 2001.

[21] J. Touch, "TCP Control Block Interdependence," IETF, RFC 2140, 1997.

[22] M. Becke, T. Dreibholz, H. Adhari, and E. P. Rathgeb, "A Future Internet Architecture supporting Multipath Communication Networks," in *Proceedings of the 13th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2012.

[23] D. Wischik, C. Raiciu, A. Greenhalgh, and M. Handley, "Design, Implementation and Evaluation of Congestion Control for Multipath TCP," in *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI'11)*, 2011.