

## 1 **9. MAC fairness**

### 2 **9.1 Overview**

#### 3 **9.1.1 Scope**

4 This clause defines the fairness algorithm for RPR MACs. The MAC uses the algorithm  
5 to enforce fairness among stations on the ring. The RPR fairness algorithm RPR-FA  
6 handles the fairness eligible traffic (Class B and Class C traffic).

#### 7 **9.1.2 Goals and objectives**

8 The fairness protocol has the following objectives:

- 9 a) Source-based weighted fairness—on any given segment on the ringlet, the  
10 available bandwidth is allocated to each station in proportion to its relative  
11 weight. For example, if every station has an equal weight, then the available  
12 bandwidth on the segment should be shared equally by all stations. On the other  
13 hand, if one station has a higher weight, the bandwidth allocated to that station  
14 should be in proportion to the station's weight divided by the sum of the weights  
15 of all the active stations.
- 16 b) Fast response time—In order to ensure maximum ring bandwidth utilization and  
17 to ensure that the protocol is responsive to instantaneous changes in traffic load, it  
18 must have a fast response time.
- 19 c) High bandwidth utilization on the ring—the protocol should be able to achieve  
20 very high levels of bandwidth utilization even under heavy load approaching  
21 100% of the ring capacity.
- 22 d) Scalability—the protocol should be scalable and should be able to function  
23 predictably for all ringlet speeds and ring diameters allowed by this standard.

#### 24 **9.1.3 Relationship to other clauses**

25 The RPR-FA is implemented within a control entity called the Fairness Control Unit  
26 (FCU) located in the MAC Control Sublayer, as described in Clause 5.

### 27 **9.2 Acronyms**

28 This clause contains the following acronyms:

29		
30	FA	Fairness Algorithm
31	FCU	Fairness Control Unit
32	MTU	Maximum Transmission Unit
33	RTT	Round Trip Time
34	FCM	Fairness Control Message

### 35 **9.3 Variables and terminology used**

36 This clause contains the following definitions and variables in alphabetical order:

1 **9.3.1 calculation round**

2 The information collected in the information round is used by all stations in the  
3 calculation round. This round is where the fair rates are calculated by the fairness  
4 algorithm. Each station starts its new cycle after it has done its fair rate calculation in this  
5 round.

6 **9.3.2 cycle**

7 The intervals over which the fairness algorithm schedules and allocates the fair rates.

8 **9.3.3 designated station**

9 The station that initiates the information rounds, and therefore also the new cycles.

10 **9.3.4 destinationTraffic**

11 This array holds information about the local traffic load to each destination. It is a local  
12 array, i.e., each station holds one for each ringlet. The station fills this array during the  
13 information round, it is used during the calculation round.

14 **9.3.5 estimator**

15 Array holding the current estimated traffic to each destination

16 **9.3.6 greedy mode**

17 Source-destination flows not flowing over one or more bottleneck links are in greedy  
18 mode, i.e., free access in the current cycle. See also “reservation mode”.

19 **9.3.7 information round**

20 A round trip of the fairness control message just before the start of a new cycle, with the  
21 goal to collect information about the traffic demand on each link. This round is started by  
22 a designated station.

23 **9.3.8 linkTraffic**

24 This array holds information about the local traffic load on each link. It is a local array,  
25 i.e., each station holds one for each ringlet. The station fills this array during the  
26 information round, it is used during the calculation round.

27 **9.3.9 numLinks**

28 The number of links on the ringlet (equals numStation).

29 **9.3.10 numStations**

30 The number of stations on the ringlet (equals numLinks).

31 **9.3.11 reservation mode**

32 Flows not in greedy mode are in reservation mode, i.e., they flow over one or more  
33 bottleneck links.

### 1 **9.3.12 table**

2 Table is an array in the FCM that holds the following information for each link:

- 3 • demand: Total demand on the link
- 4 • remainingCap: Capacity available on the link

### 5 **9.3.13 useSourceFairness**

6 The fairness algorithm uses source fairness when this variable is set to true. Flow based  
7 fairness is used when this variable is set to false.

## 8 **9.4 MAC fairness operation**

9 The fairness algorithm implemented within the FCU consists of the following functions:

- 10 a) sourcing and consuming fairness messages
- 11 b) calculation of fair rates for each source destination pair
- 12 c) determining the state for each flow: reserved or free access

13  
14 Each station is assigned a weight, which allows the user to allocate more ring bandwidth  
15 to certain stations as compared with other stations.

16  
17 The FA uses a proactive method that assigns fair rates to each source-destination flow on  
18 a ringlet. It can be used for source-fairness as well as source-destination (or flow)  
19 fairness.

20 The algorithm uses three rounds of the FCM for each cycle. The first round is used for  
21 collecting the traffic demand for each source destination pair, and in the second round  
22 each station computes its own fair rate based on the information in the fairness message.  
23 Instead of using all flow information on the ringlet, the algorithm uses aggregate flows,  
24 keeping the fairness message small and the algorithm scalable. The third round is used to  
25 inform all stations about the available bandwidth that still can be used for greedy traffic.

26  
27 The control information needed to accomplish this, flows in the same direction (i.e.  
28 ringlet) as the data flow, which simplifies the protocol in a single ring topology and any  
29 configuration of multiple rings.

### 31 **9.4.1 FCM Processing**

#### 32 **9.4.1.1 Fairness Control Message**

33 The designated station creates fairness control messages. This message travels three  
34 rounds, thereby visiting each station on the ringlet three times. Apart from the normal  
35 packet header fields, the FCM contains a round counter (2 bits) and two arrays, each of  
36 size numLinks (the number of links on the ringlet), one for the total traffic demand on all  
37 links for fairness eligible traffic, and one array for the remaining capacity (total capacity  
38 minus provisioned traffic) on all links (Figure 1).

39




Figure 1. Table structure contained in the FCM

1  
2  
3

### 9.4.1.2 Overview




Figure 2. The fairness control message is input to the fairness algorithm

5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22

The designated station holds a timer for each ringlet that fires every cycle interval. At this event the designated station creates a fairness control message, initializes it and starts the first round of this message. In this “information gathering” round each station writes the amount of bytes that are available for its outgoing link, in the remaining capacity field. Additionally it also adds its own flows (information comes again from the waiting traffic demand) to the sum of all flows on all links. Once the control message arrives back at the designated station, the control message starts its second round where each station performs the fairness algorithm and immediately can start sending its fair share. In this second round the stations can modify the contents of the FCM. The third round is used to notify all stations about the amount of remaining capacity that can be used for free-access traffic. The control message is taken from the ringlet at the time it returns for the third time at the designated station.

Since the control message is relative small in size, it produces a small overhead even for short calculation intervals.




Figure 3. Cycle timeout for the designated station




Figure 4. Flow diagram for non-designated stations

Each station starts its new cycle directly after calculation of the fair rates, i.e., stations can start transmission of the packets according to the just calculated fair rates.

The flow diagram for designated stations (Figure 5) is similar to the diagram for non-designated stations. The difference is that in the designated station the transition from one round to the other takes place. Additionally, the FCM will be deleted after the remaining capacity round.




Figure 5. Flow diagram for the designated station

1  
2  
3  
4  
5  
6




Figure 6. Time diagram for the designated station

7  
8  
9  
10

## 1 **9.4.2 Information collection round**

2 The fairness algorithm is a proactive mechanism that uses the traffic demand to schedule  
3 the traffic in the next cycle. Ideally, the information about the waiting traffic demand  
4 comes from the MAC client. The rate will be estimated (Section 9.4.2.2), for clients  
5 unable to tell their traffic demand.

### 6 **9.4.2.1 Getting the queued traffic demand**

7 Upon request, the MAC-client notifies the MAC about the amount of bytes currently  
8 waiting in the destination queues. If queues are nearly or completely full, the MAC uses  
9 rate estimation, the topic of the next section.

### 10 **9.4.2.2 Rate estimation**

11 Two arrays are used to estimate the traffic flows: Estimated[] holds the current estimation  
12 to all destinations, and Measured[] holds the total amount of fairness eligible traffic  
13 sourced to each destination in the current cycle.

14 At the end of each cycle, all nodes perform the following update of the Estimated array:  
15 If there are no packets to transmit for a destination  $d$ , then Estimated[ $d$ ] is set to  
16 Measured[ $d$ ]. Else, we take the maximum of three values: the previous estimation, the  
17 measured value and a constant value B. This maximum multiplied by a constant A gives  
18 us the new estimation: Estimated[ $d$ ] = A \* max( Estimated[ $d$ ], Measured[ $d$ ], B).

19 The constants A and B control a trade-off between throughput and response time during  
20 transitions.

### 21 **9.4.2.3 Weights**

22 As long as there are no bottlenecks, weights are not used. When there are bottlenecks  
23 however, the estimated or real traffic demand is multiplied by the weight of each station.  
24 This value will be written in the FCM by each station in the information round.

### 25 **9.4.2.4 Source and flow fairness**

26 The fairness algorithm is a flow or source-destination fairness algorithm. Source fairness  
27 can be easily achieved with the same algorithm by limiting the sum of all flows leaving  
28 each station to the link capacity. These adjusted values are then written in the FCM.  
29

## 30 **9.4.3 Code**

31 This section describes two functions: cycleTimeOut and handleFCM. CycleTimeOut is  
32 executed at each cycle timeout at the designated station. It generates a fairness control  
33 message and passes the message to the handleFCM function. This function is also  
34 executed at each station when a fairness message arrives.

```
35  
36 void cycleTimeOut() {  
37     FairnessControlMessage *fcm = new FairnessControlMessage;  
38     fcm->round = -1;  
39     handleFCM(fcm);  
40 }
```

```
41  
42 void handleFCM(FairnessControlMessage *fcm)
```

```
1  {
2  double *allowForDest;
3  int i;
4
5  if (designatedStation()){
6      fcm->round++;
7      if (fcm->round == 3){
8          delete fcm;
9          return;
10     }
11 }
12
13 switch (fcm->round){
14 case 0: // information round
15     infoRound(fcm);
16     break;
17 case 1: // calculation round
18     allowForDest= new double[getSize()];
19     makeFair(fcm,allowForDest);
20     for (i=0;i<getSize();i++){
21         setBitsAllow(i,allowForDest[i]);
22     }
23     delete allowForDest;
24     break;
25 case 2: // update remaining capacity round
26     for (i=0;i<getSize();i++){
27         remainingCap[i] = fcm->table.available[i];
28     }
29     break;
30 }
31
32 // schedule the forwarding
33 forwardFairnessControlMessage (fcm);
34 }
35
36
37 bool designatedStation(){
38     // returns true if this station is the designated station
39     // on the current ringlet, false otherwise;
40 }
41
42
43 int getSize(){
44     // returns the number of stations on the ringlet;
45 }
46
47
48 void infoRound(...){
49     // see Section 9.4.3.1
50 }
51
52
53 void makeFair(...){
54     // see section 9.4.3.2
55 }
56
57
58 void setBitsAllow(int dst,double value){
59     // Sets the amount of bytes that can be transmitted by this station
60     // to destination dst.
61 }
62
```

```
1
2 void forwardFairnessControlMessage(FairnessControlMessage *fcm) {
3     // function that forwards the fcm to the next downstream neighbor as
4     // quick as possible
5 }
6
7
```

### 8 9.4.3.1 Information Round

9 Each station in the information round executes the following function. The only argument  
10 is the fairness control message.

```
11
12 void infoRound(FairnessControlMessage *fcm)
13 {
14     double sourceFairFactor      = 1.0;
15     int i,dest,link;
16
17     // set the available bandwidth to 100%
18     fcm->table.available[atStationID()]=getCycleMaxLoad();
19
20
21     // make a copy of the current fill sizes
22     for (i=0;i<getSize();i++) {
23         destinationTraffic[i] = min(getLoadForDestination(i), getCycleMaxLoad());
24     }
25
26     if (useSourceFairness)
27         sourceFairFactor = calculateSourceFairFactor();
28
29     // now loop through all links and add the amount of bytes
30     link = downStreamLinkId();
31     for (i=0;i<getSize();i++){
32
33         // now find all destinations over "link"
34         dest = stationIdAtEndOfLink(link);
35         while (dest!=atStationID()){
36             double Iwant = min(getCycleMaxLoad(),getLoadForDestination(dest));
37             fcm->table.demand[link] += Iwant * sourceFairFactor;
38             linkTraffic[link]      = fcm->table.demand[link];
39             dest = next(dest);
40         }
41         link = next(link);
42     }
43 }
44
```

---

```
45
46 double getCycleMaxLoad(){
47     // return the maximum number of bits that can be transmitted on
48     // the outgoing link in one cycle:
49     // e.g. return cycleInterval * getLinkSpeed();
50 }
51
```

---

```
52
53 long getLoadForDestination(int dest){
54     // see section 9.4.2.1
55     double MAXpossible = getCycleMaxLoad();
56
57     if (isQueueEmpty(dest))
58         estimator[dest] = getBytesSend(dest);
59     else
60         estimator[dest] = A*max( estimator[dest], max(getBytesSend(dest), B));
```

```
1
2     if (estimator[dest]>MAXpossible/2) estimator[dest]=MAXpossible/2;
3
4     return estimator[dest];
5 }
6
7 bool isQueueEmpty(int dest){
8     // returns true iff the MAC client has no fairness eligible traffic
9 }
10
11 int getBytesSend(int dest){
12     // returns the amount of bytes sources to destination dest in the
13     // current cycle
14 }
15
16
17 int stationIdAtEndOfLink(int lnk){
18     // function that returns the id of the station at the end of link "lnk"
19     // on the current ringlet
20 }
21
22
23 double calculateSourceFairFactor(){
24     // this function computes a factor that is used by the fairness algorithm
25     // to limit the flows of a single station. Limiting all flows from
26     // one station to the link capacity, results in source fairness (in the
27     // used fairness algorithm)
28
29     double MAXpossible = getCycleMaxLoad();
30     double ret      = 1.0;
31     double total    = 0;
32     int i;
33     for (i=0;i<getSize();i++)
34         total += min(MAXpossible,getLoadForDestination(i));
35
36     if (total>MAXpossible)
37         ret = MAXpossible / total;
38
39     return ret;
40 }
41
```

### 42 9.4.3.2 Calculation Round

```
43
44 void makeFair(FairnessControlMessage *fcm, double *allowForDest){
45     // This is the function that computes the fair rates for the station
46     // where this function is being called. The FCM is input to this function,
47     // allowForDest is the resulting array with the fair amount of bytes
48     // for each destination.
49     // Note that the contents of the fcm can be modified by this function.
50
51     int strongestBottleNeckLink,i;
52     bool *bottleNeckDone = new bool[getSize()]; // array indicating whether or not
53                                                // a bottleneck link is processed
54
55     // initiliaze:
56     for (i=0; i<getSize(); i++) {
57         bottleNeckDone[i] = false;
58         allowForDest[i]   = 0;
59     }
60
```

```
1 do {
2     // look for the strongest bottleneck where this station is involved...
3     strongestBottleNeckLink=getStrongestValidBottleneck(bottleNeckDone, fcm);
4
5     if (strongestBottleNeckLink!=-1){
6
7         // yes we are involved in a bottleneck
8         bottleNeckDone[strongestBottleNeckLink] = true;
9
10        // ok, reduce all flows over this bottleneck
11        // we start at the destination just over the bottleneck
12        // and loop through all destinations from there on
13        int toDest = stationIdAtEndOfLink (strongestBottleNeckLink);
14        while (toDest != atStationID()){
15            // do I have something for this destination?
16            if (destinationTraffic[toDest]>0){
17                double ratio      = fcm->table.ratio(strongestBottleNeckLink);
18                double oldValue   = destinationTraffic[toDest];
19                double newValue    = oldValue/ratio;
20                destinationTraffic[toDest] = 0;
21                allowForDest[toDest]      = newValue;
22
23                // now update the table in the control message
24                // and our local linkTraffic table accordingly.
25                // all links between this station and toDest need
26                // to be updated.
27                double diff = oldValue-newValue;
28                int link = downstreamLinkId();
29                while (link != downstreamLinkIdAtStation(toDest)){
30                    fcm->table.available[link] -= newValue;
31                    fcm->table.demand[link] -= oldValue;
32                    linkTraffic[link] -= diff;
33                    link = next(link);
34                }
35            }
36            toDest = next(toDest);
37        }
38    }
39    } while (strongestBottleNeckLink!=-1); // as long as there are bottlenecks
40
41    // copy remeainig traffic load since this
42    // traffic is not involved in any bottleneck
43    for (i=0; i<getSize(); i++) allowForDest[i] += destinationTraffic[i];
44
45    delete bottleNeckDone;
46 }
47
48
49
50 int getStrongestValidBottleneck(bool *done, FairnessControlMessage *fcm){
51     // Returns the id of the strongest bottleneck, that is not yet
52     // processed (done).
53     // Returns -1 if no such bottlenecks exists.
54
55     int i,ret =-1;
56     for (i=0;i<getSize();i++)
57         if (!done[i] && fcm->table.isBottleNeck(i) &&
58             ((ret==-1) || (fcm->table.ratio(ret)<fcm->table.ratio(i))))
59             ret = i;
60     return ret;
61 }
62
63
```


```
1
2 int downstreamLinkId(){
3     // Returns the downstream link id at the current station on the
4     // current ringlet
5 }
6
7 -----
8 int downstreamLinkIdAtStation(int s){
9     // Returns the downstream link id at station "s" on the current ringlet
10 }
11
12 -----
13 int atStationID(){
14     // Returns the station ID of the station
15 }
16
17 -----
18 bool table::isBottleNeck(int lnk){
19     // We have a bottleneck if the demand is larger than what is available,
20     // and there is a positive non-null demand:
21     return (demand[lnk]>0) && (demand[lnk]>available[lnk]);
22 }
23
24 -----
25 double table::ratio(int lnk){
26     // Returns the ration demand/available for the specified link
27     // To avoid division be zero, a very large constant "BIG" is
28     // returned if available equals zero.
29
30     if (available[lnk]==0) return BIG;
31     else return demand[lnk]/available[i];
32 }
33
```

### 34 9.5 Example

35 This section gives an example of the operation of the algorithm. For simplicity a single  
36 ringlet is used with one priority class and only 4 stations. Furthermore, the following  
37 assumptions are made:

- 38 • 100 units (e.g. bytes) can be transmitted on each link, in one cycle
- 39 • Station 0 is the designated station
- 40 • Weights are all 1
- 41 • The traffic demand from and to each station is shown in Figure 7.

42  
43



44  
45

Figure 7. Traffic demand

1 **9.5.1 Information round**

2 The information round starts at the designated station 0 after receiving the cycle timeout.  
3 It creates a new FCM and adds all of its flows to the summed traffic fields. The  
4 remaining capacity is set to full capacity (100).  
5

Link	Summed Traffic	Remaining Capacity
0	20+80=100	100
1	80	100
2	80	100
3	-	100

6 Station 0 forwards this FCM to station 1, which also adds it own flows to the summed  
7 traffic fields.  
8

Link	Summed Traffic	Remaining Capacity
0	100	100
1	80+30 =110	100
2	80+30 =110	100
3	-	100

9 FCM leaving station 1  
10  
11

12 Similar for station 2 and station 3:

Link	Summed Traffic	Remaining Capacity
0	100	100
1	110	100
2	110+60 = 170	100
3	-	100

13 FCM leaving station 2  
14

Link	Summed Traffic	Remaining Capacity
0	100+90 = 190	100
1	110	100
2	170	100
3	90	100

15 FCM leaving station 3  
16  
17

18 **9.5.2 Calculation round**

19 A new round is started whenever the designated station receives the FCM from its  
20 upstream neighbor, in the example the new round will be the calculation round. This is  
21 the round where each station calculates its own fair rates and immediately can start to  
22 transmit these fair rates.  
23

1 The highest bottleneck where station 0 is involved is at link 0. All flows leaving station 0  
2 over this bottleneck will be reduced by a factor  $100/190$ . For flow 0->1 with a demand of  
3 20, the assigned value will become  $20 * 100/190 = 10.5$ .

4 The FCM has to be updated to reflect this: The summed traffic fields involved should be  
5 decreased by 20, which is in this case only the field for link 0. Since 10.5 is assigned to  
6 this flow, the remaining capacity fields should be decreased by 10.5.

7 The new FCM:

8

Link	Summed Traffic	Remaining Capacity
0	$190-20=170$	$100-10.5=89.5$
1	110	100
2	170	100
3	90	100

9 Station 0 still has a flow over a bottleneck, which is still link 0. Flow 0->3 with a demand  
10 of 80 will get a value of  $80 * 89.5 / 170 = 42.1$ .

11

12 The new FCM:

Link	Summed Traffic	Remaining Capacity
0	$170-80 = 90$	$89.5-42.1=47.4$
1	$110-80 = 30$	$100-42.1=57.9$
2	$170-80=90$	$100-42.1=57.9$
3	90	100

13

14 Equals:

15

Link	Summed Traffic	Remaining Capacity
0	90	47.4
1	30	57.9
2	90	57.9
3	90	100

16 FCM leaving station 0

17

18

19 Station 1 follows the same procedure: highest bottleneck is at link 2, the assigned value  
20 will be  $30 * 57.9 / 90 = 19.3$ .

21

Link	Summed Traffic	Remaining Capacity
0	90	47.4
1	$30-30=0$	$57.9-19.3=38.6$
2	$90-30=60$	$57.9-19.3=38.6$
3	90	100

22 FCM leaving station 1

23

24

25 Station 2: flow gets  $60 * 38.6 / 60 = 38.6$

Link	Summed Traffic	Remaining Capacity
0	90	47.4
1	0	38.6
2	60	38.6
3	90	100

0	90	47.4
1	0	38.6
2	60-60=0	38.6-38.6=0
3	90	100

1 FCM leaving station 2

2

3

4

5 Station 3: flow gets  $90 * 47.4/90 = 47.4$

6


Link	Summed Traffic	Remaining Capacity
0	90-90=0	47.4-47.4=0
1	0	38.6
2	0	0
3	90-90=0	100-47.4=52.6

7 FCM leaving station 3

8

9

10 All flows are now assigned, as can be seen in the following figure:



11

12

Figure 8. Result after the calculation round

13

### 14 9.5.3 Remaining Capacity Round

15 The purpose of this round is to inform all station on the ringlet about the amount of  
16 capacity available on all links. This is used for greedy traffic.

17

Link	Summed Traffic	Remaining Capacity
0	0	0
1	0	38.6
2	0	0
3	0	52.6

18

19 When the FCM returns back to the designated station, the FCM will be deleted. The timer  
20 in the designated station will trigger the start of the next cycle.