Two-rate three-color marker (trTCM)

- The trTCM meters an IP packet stream and marks its packets based on two rates
  - Peak Information Rate (PIR) and
  - Committed Information Rate (CIR)
- and their associated burst sizes to be either
  - Green
  - Yellow
  - Red
- A packet is marked red if it exceeds the PIR
- Otherwise it is marked either yellow or green depending on whether it exceeds or doesn't exceed the CIR

Two-rate three-color marker (trTCM)

- The trTCM is configured by setting its mode (color-blind or color-aware) and by assigning values to four traffic parameters:
  - a Peak Information Rate (PIR) and its associated Peak Burst Size (PBS)
  - a Committed Information Rate (CIR) and its associated Committed Burst Size (CBS)
- The PIR and CIR are measured in bytes of IP packets per second, i.e., it includes the IP header, but not link specific headers
- The PIR must be equal to or greater than the CIR
- The PBS and the CBS and are measured in bytes and both of them must be configured to be greater than 0
- It is recommended that they be configured to be equal to or greater than the size of the largest possible IP packet in the stream

Two-rate three-color marker (trTCM)

- The behavior of the Meter is specified in terms of its mode and two token buckets, P and C, with rates PIR and CIR, respectively
- The maximum size of the token bucket P is PBS and the maximum size of the token bucket C is CBS
- The token buckets P and C are initially (at time 0) full, i.e., the token count Tp(0) = PBS and the token count Tc(0) = CBS
- Thereafter, the token count Tp is incremented by one PIR times per second up to PBS
- The token count Tc is incremented by one CIR times per second up to CBS

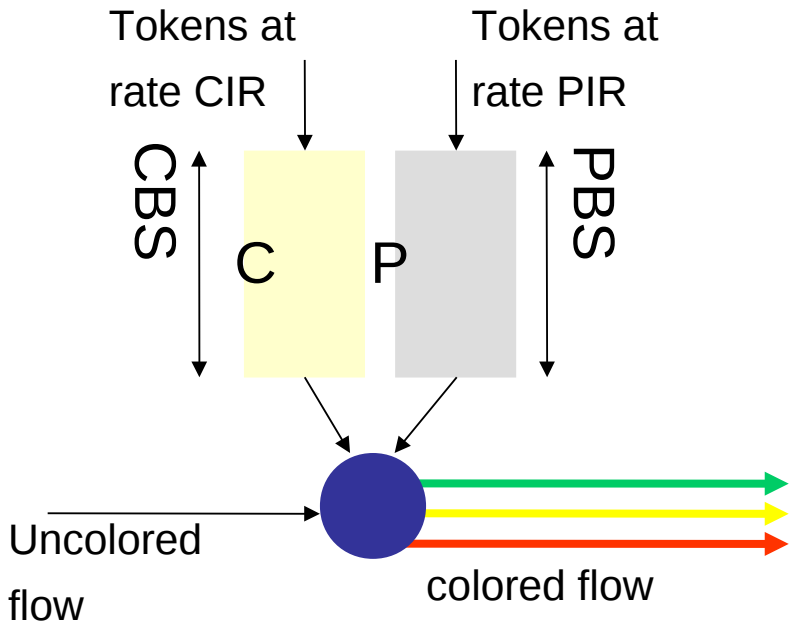

trTCM operating in the Color-blind mode

Two-rate three-color marker (trTCM)

- When a packet of size B bytes arrives at time t, the following happens if the trTCM is configured to operate in the Color-Blind mode:
  - If Tp(t)-B < 0, the packet is red, else
  - if Tc(t)-B < 0, the packet is yellow and Tp is decremented by B, else
  - the packet is green and both Tp and Tc are decremented by B

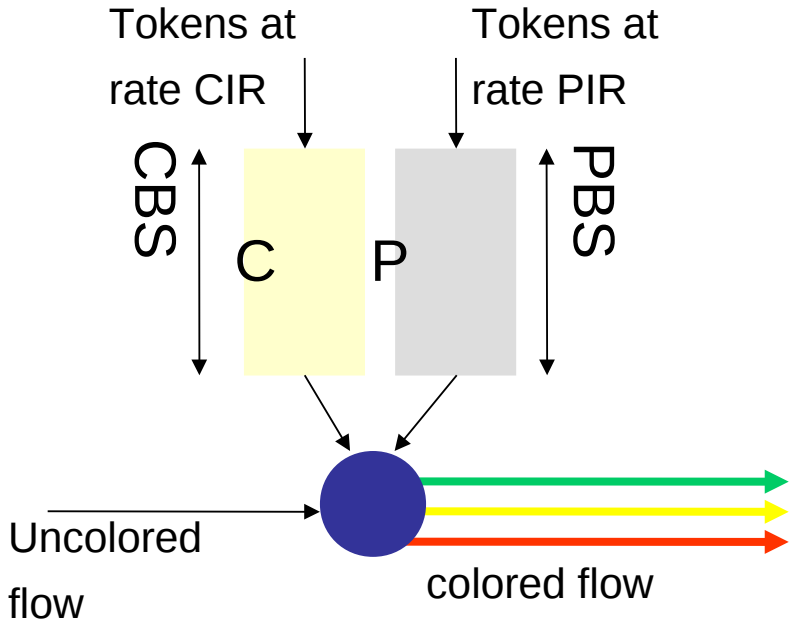

trTCM operating in the Color-blind mode

Two-rate three-color marker (trTCM)

- When a packet of size B bytes arrives at time t, the following happens if the trTCM is configured to operate in the Color-Aware mode:
  - If the packet has been precolored as red or if Tp(t)-B < 0, the packet is red, else
  - if the packet has been precolored as yellow or if Tc(t)-B < 0, the packet is yellow and Tp is decremented by B, else
  - the packet is green and both Tp and Tc are decremented by B

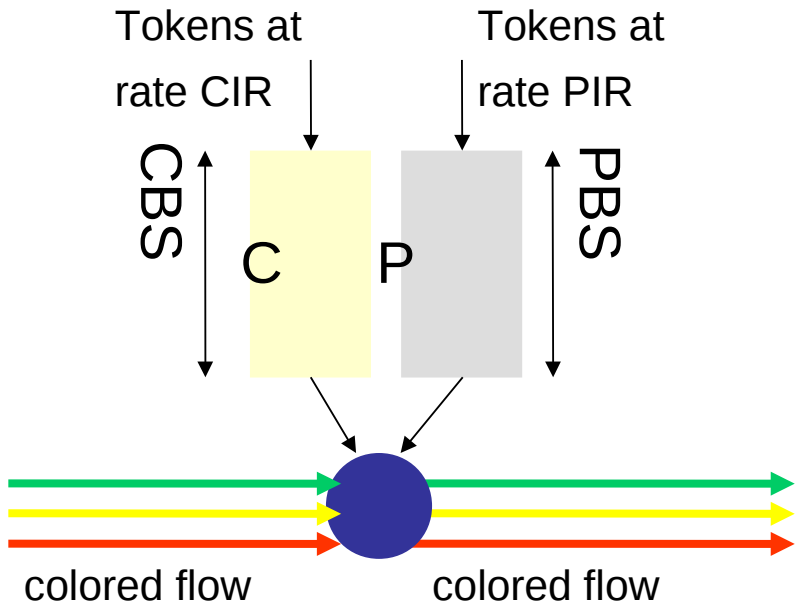

trTCM operating in the Color-aware mode

Two-rate three-color marker (trTCM)

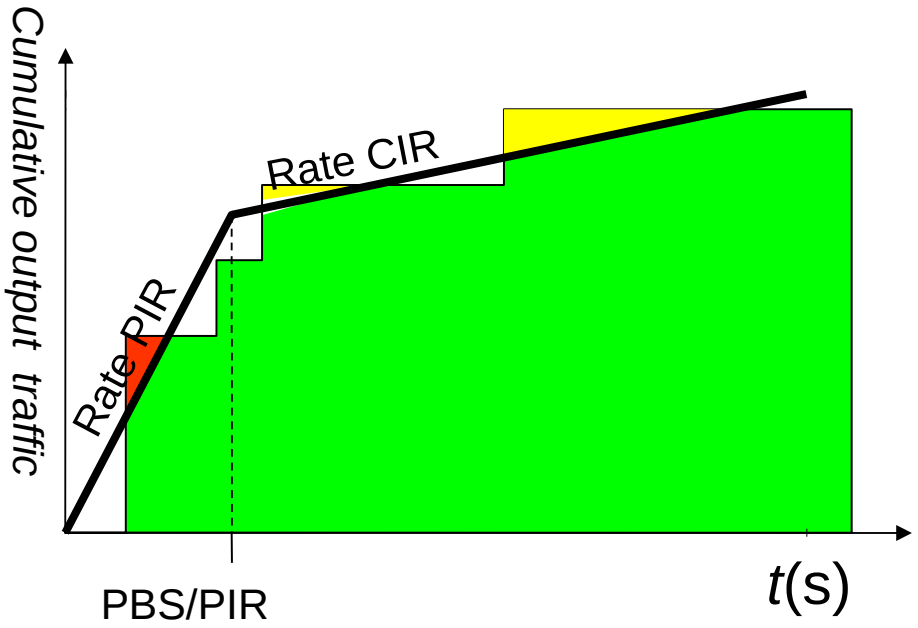

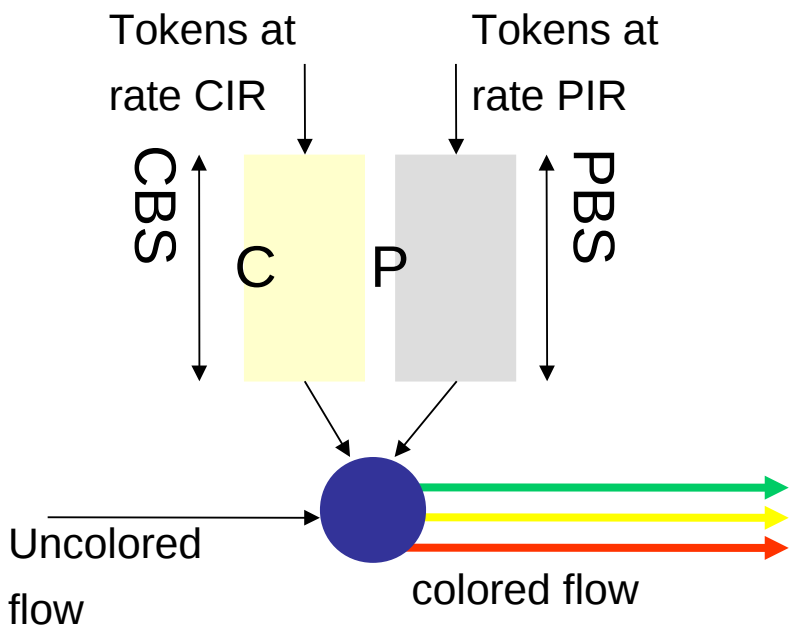

trTCM operating in the Color-blind mode

Traffic contract

A traffic contract between a provider and a customer includer a TCA and a SLA

The TCA specifies the traffic profile

The SLA specifies the QoS requirements

a delay threshold

the maximum fraction of packets that may exceed the delay threshold

- If the fraction of packets is zero, the service is GUARANTEED while if the fraction of packets is greater than zero, the service is STATISTICAL

a maximum fraction of packets that can be lost

Guaranteeing QoS

When a request is issued by the customer, a traffic contract (TCA + SLA) is established

Admission control checks if this traffic contract can be fulfilled, without affecting the SLA of already established contracts

If this is possible, resources are allocated for this contract

The connection becomes active and the real-time management of traffic starts
Scheduling at each network node
Regulation at the network ingress
Active queue management in nodes

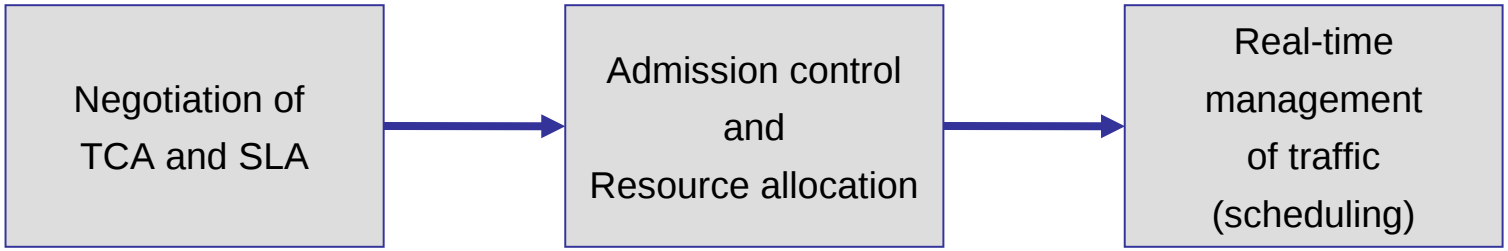

Admission control

The admission control function accepts or denies user requests

A request is accepted only if it

is possible to do it fulfilling the requested and

Accepting the new request does not degrade the SLA of already accepted requests

For example, admitting a VoIP flow can be done only if delay can be granted

After all, it is in the nature of the telephone service to block requests if resources are not available

Admission control is strictly related to the real-time management of traffic(scheduling)

In fact, it is possible to determine if the SLA can be met only if the scheduling policy implemented by nodes is known, as well as the amount of available resources in each node

Scheduling

In a network node, each output link is managed by a scheduler

The scheduler organizes the transmission of packets, by choosing on a per-packet base which service calss is to be served

In particular, as soon as the transmission of a packet is over, the scheduler selects, according to a suitable algorithm, which service class is to be served next

In some architectures, scheduling is based on service classes, i.e., traffic flows are aggregated into service categories and, in the scheduler, all packets belonging to the same service category are stored in the same buffer