The Expedited Forwarding PHB

- The Expedited Forwarding (EF) PHB is a basic building block in the Differentiated Services architecture
- EF is intended to provide a building block for low delay, low jitter and low loss services by ensuring that the EF aggregate is served at a certain configured rate
- Or example, circuit emulation services could be transported over the EF PHB

The Expedited Forwarding PHB

- The dominant causes of delay in packet networks are fixed propagation delays (e.g. those arising from speed-of-light delays) on wide area links and queuing delays in switches and routers
- Since propagation delays are a fixed property of the topology, delay and jitter are minimized when queuing delays are minimized
- For our purposes, we define jitter as the variation between maximum and minimum delay
- The intent of the EF PHB is to provide a PHB in which suitably marked packets usually encounter short or empty queues
- Furthermore, if queues remain short relative to the buffer space available, packet loss is also kept to a minimum

The Expedited Forwarding PHB

- To ensure that queues encountered by EF packets are usually short or empty, it is necessary to ensure that the service rate of EF packets on a given output interface exceeds their arrival rate at that interface over long and short time intervals, independent of the load of other (non-EF) traffic
- Thus, EF packets must be guaranteed to receive service above a configured rate
- Moreover, a means to quantify the accuracy with which this service rate is delivered over any time interval must be provided

Traffic and service curves

- In order to perform basic performance evaluation, two basic notions must be introduced:
  - Traffic curve
  - Service curve

Traffic curve

- Consider a specific network interface and assume to be able to count the number of bits $X(t_1, t_2)$ crossing the interface in a given time interval $[t_1, t_2]$
- $X(t_1, t_2)$ is the traffic curve at the reference interface
- If, by convention, $t_1=0$, the traffic curve is indicated as $X(t)$

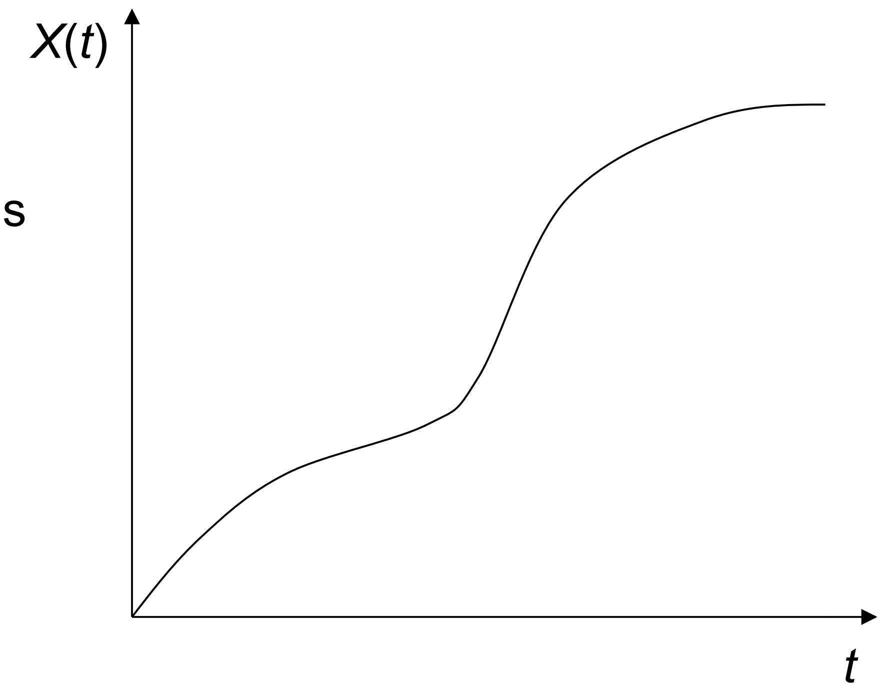

Traffic curve

- Since traffic is packetized, the actual shape of a traffic curve is shown in this graph
- When a packet is crossing the reference interface, the curve increases linearly with slope *C* [bit/s], where x is the link capacity
- When no packet is present, the traffic curve is steady

L_1 [bit]: length of packet 1
L_2 [bit]: length of packet 2
L_3 [bit]: length of packet 3
C [bit/s]: link capacity

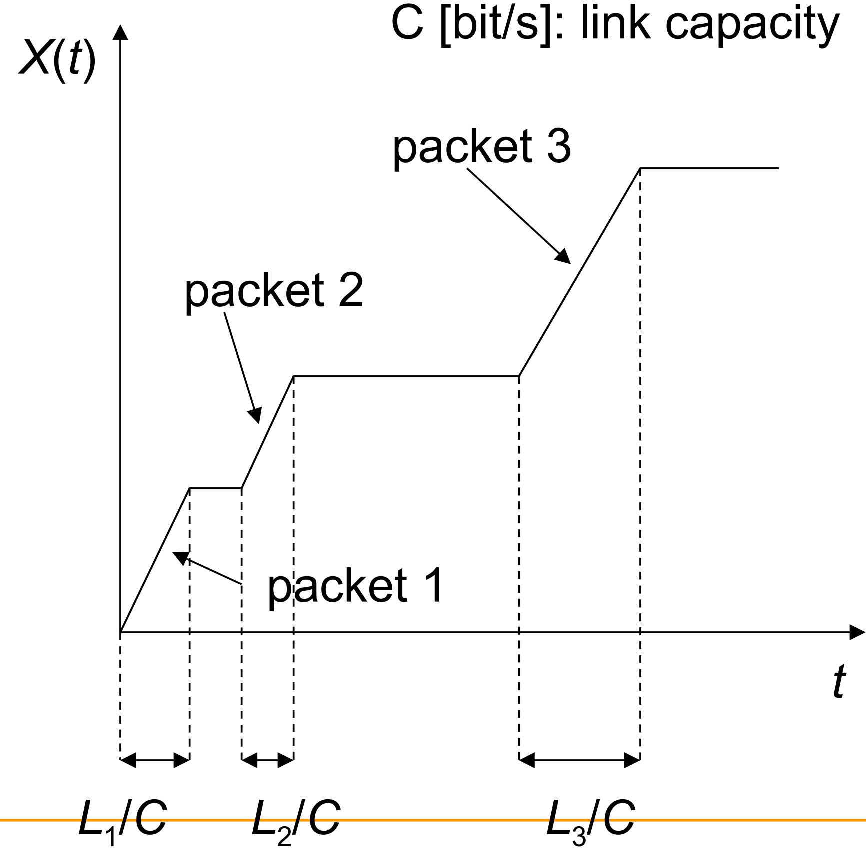

Potential service curve

- Given a scheduler (in the example, a FIFO scheduler), the potential service curve $S(t)$ is the maximum number of bits that the scheduler can serve in a time interval of duration t
- Given a FIFO scheduler provided with an output link with capacity equal to C, the potential service curve is equal to Ct

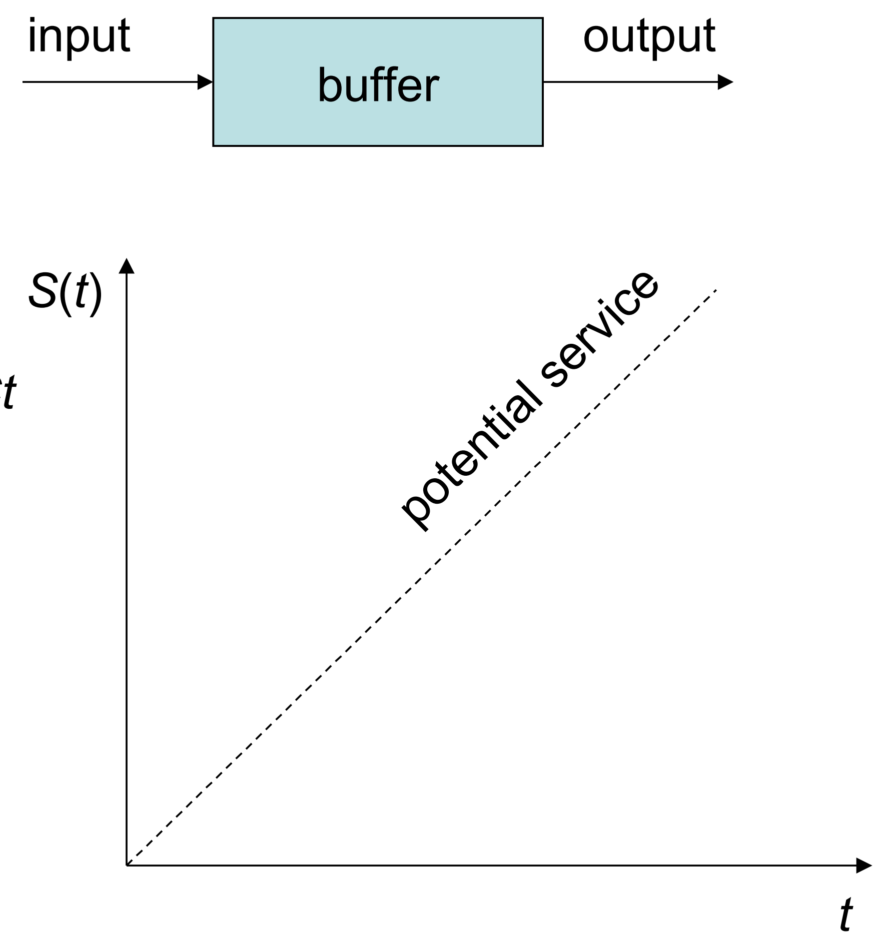

Actual service curve

- Given a scheduler (in the example, a FIFO scheduler), the actual service curve is the traffic curve at the output of the scheduler, given an input traffic with traffic curve $X(t)$

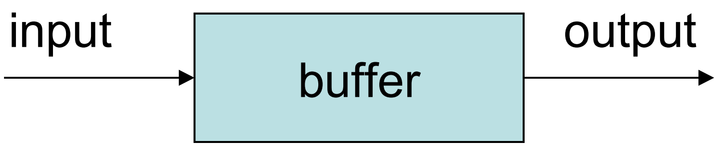

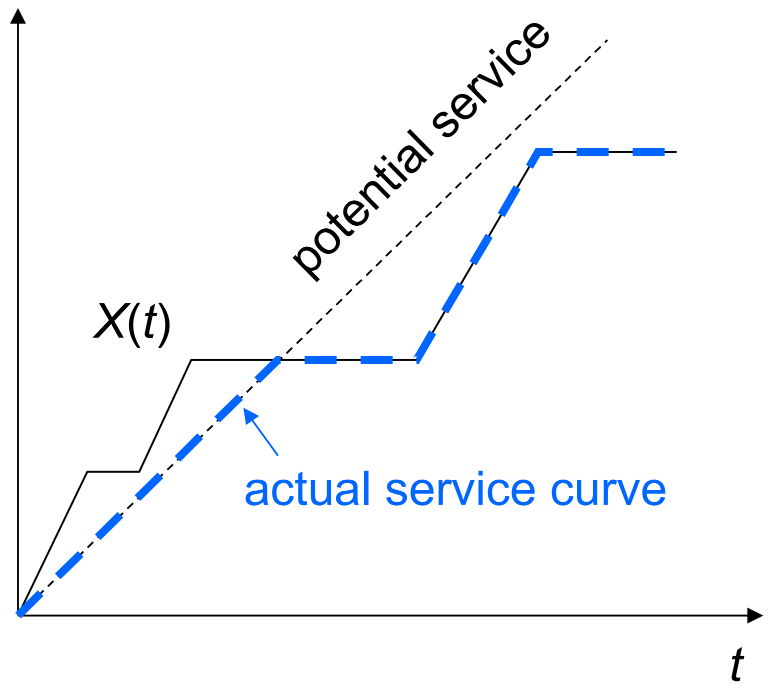

EF PHB

- Usually, a fresh EF traffic flow is characterized with a periodic arrival process of equal-length packets
- All packets have length L (bit) and the arrival rate is 1/T (packet/s)

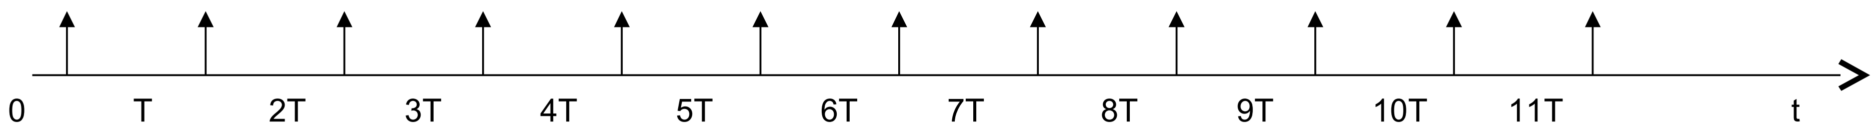

Packet arrival process: $$r(t) = \sum_k L\delta(t - kT)$$

EF PHB

- The traffic curve relative to the previous example is shown in the figure

$$X(t)=\int_0^t r(t)\,dt=\int_0^t \sum_k L\delta(t-kT)\,dt=\sum_k LU(t-kT)$$

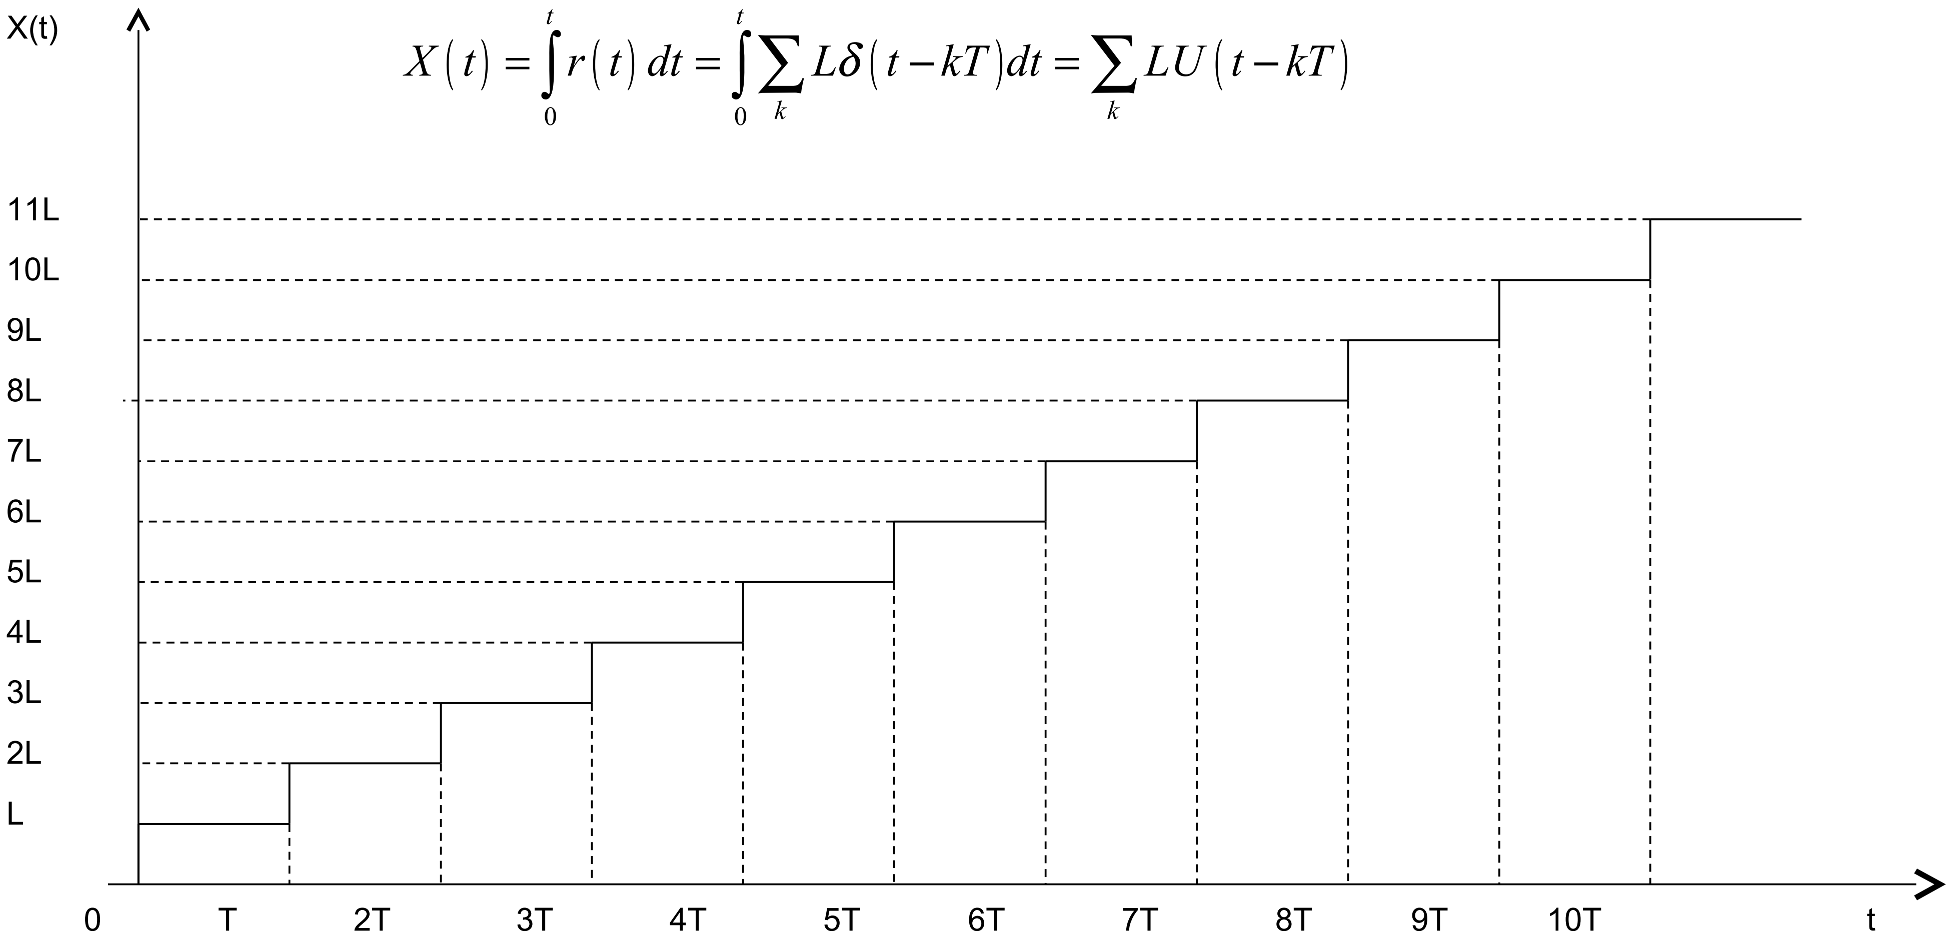

EF PHB

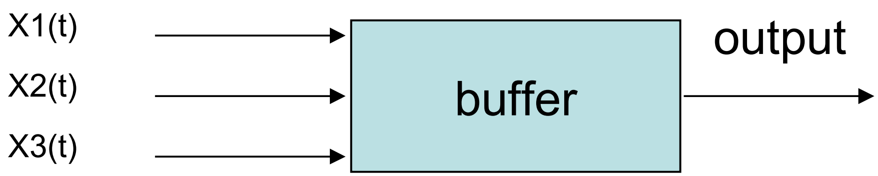

- Let us consider the case of three identical EF flows offered to a FIFO scheduler
- The worst possible case is that in which the phases of the three flows coincide:

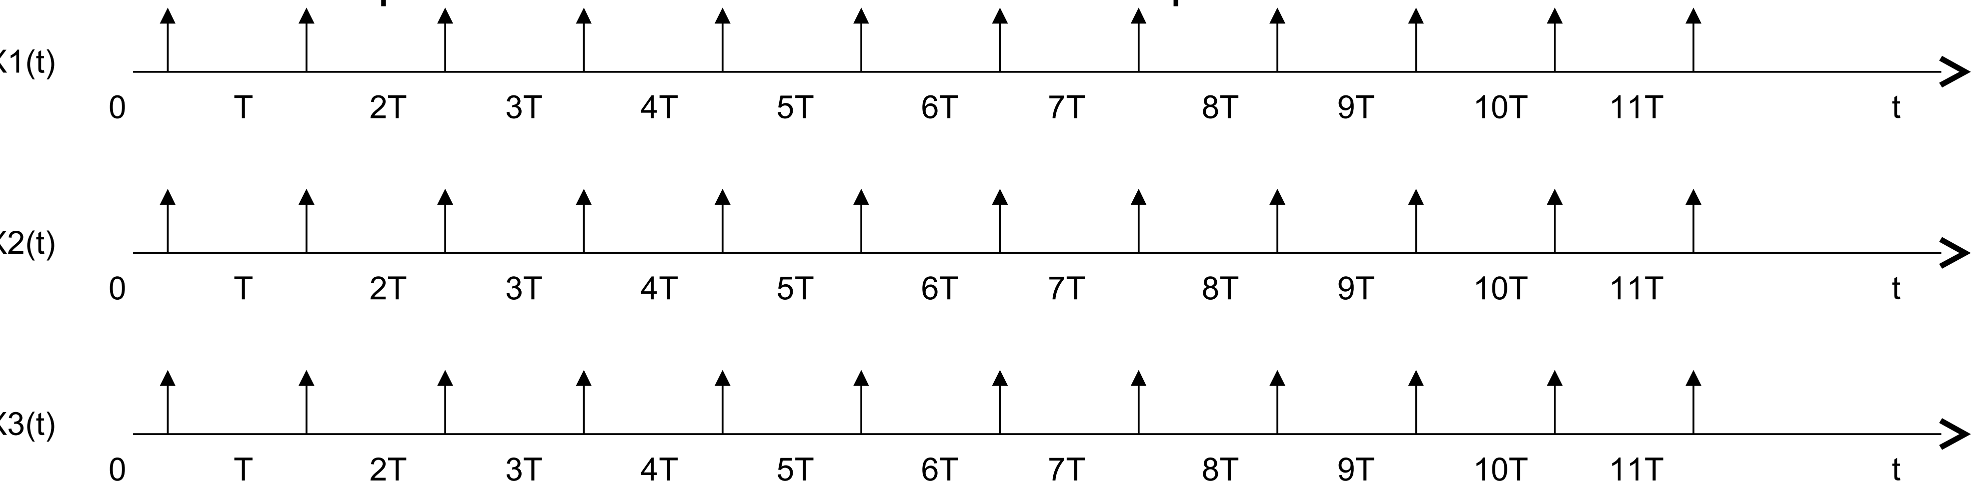

EF PHB

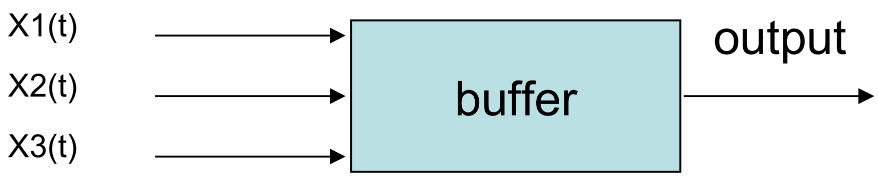

- Let us consider the following example, where the packet length L is 1000 (bit):
  - X1(t): 100 kbit/s; 1 packet every 10 ms
  - X2(t): 200 kbit/s; 1 packet every 5 ms
  - X1(t): 250 kbit/s; 1 packet every 4 ms

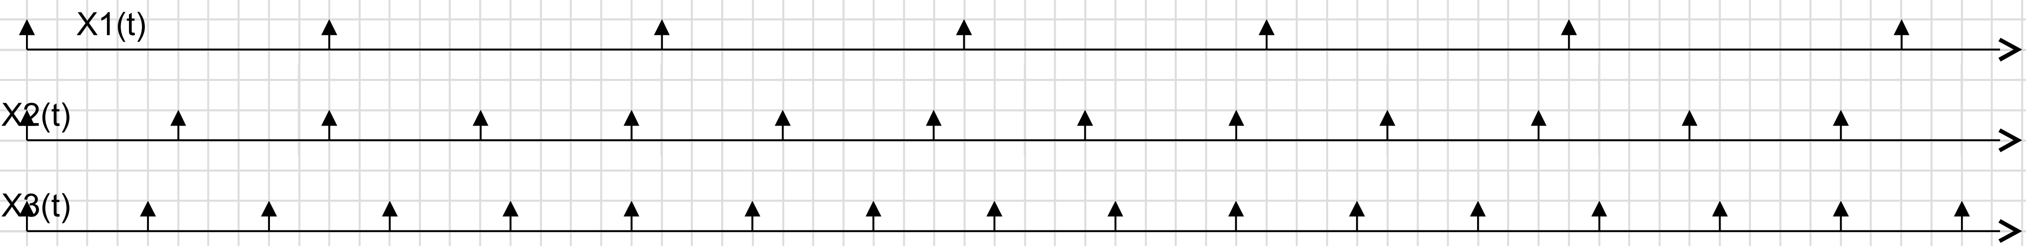

EF PHB

The three traffic curves

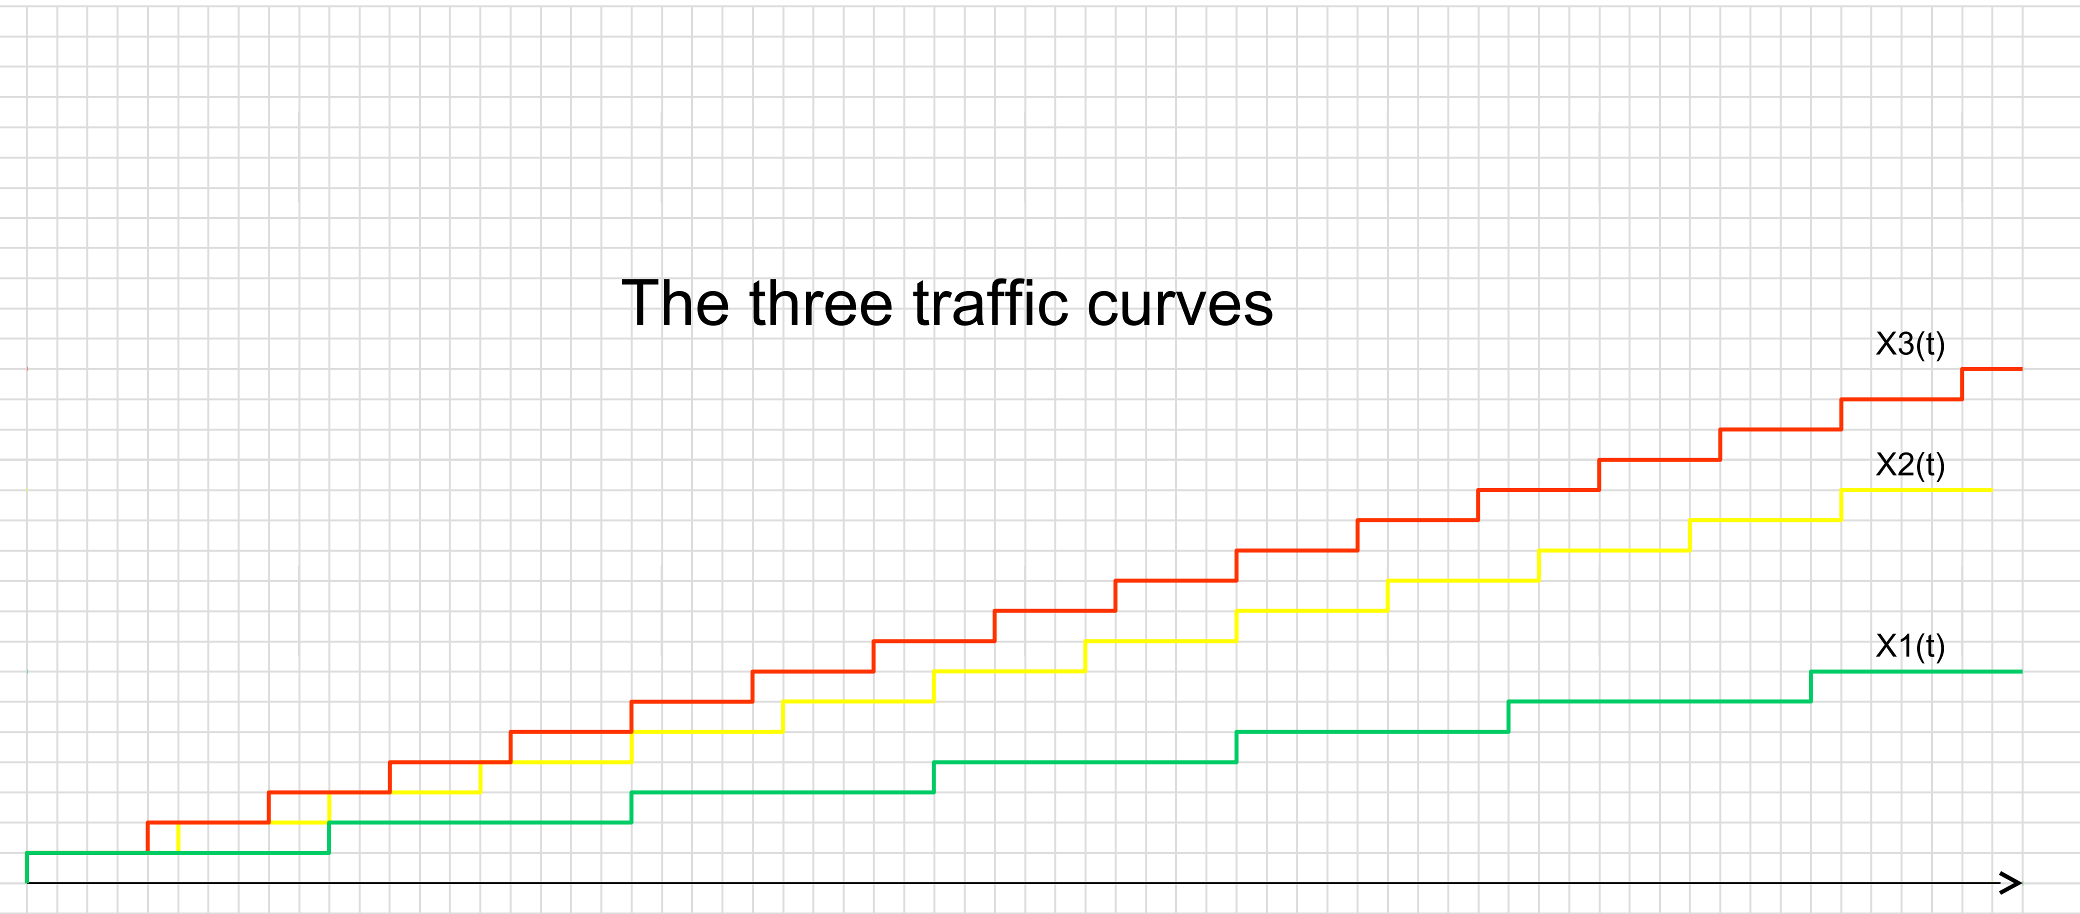

EF PHB

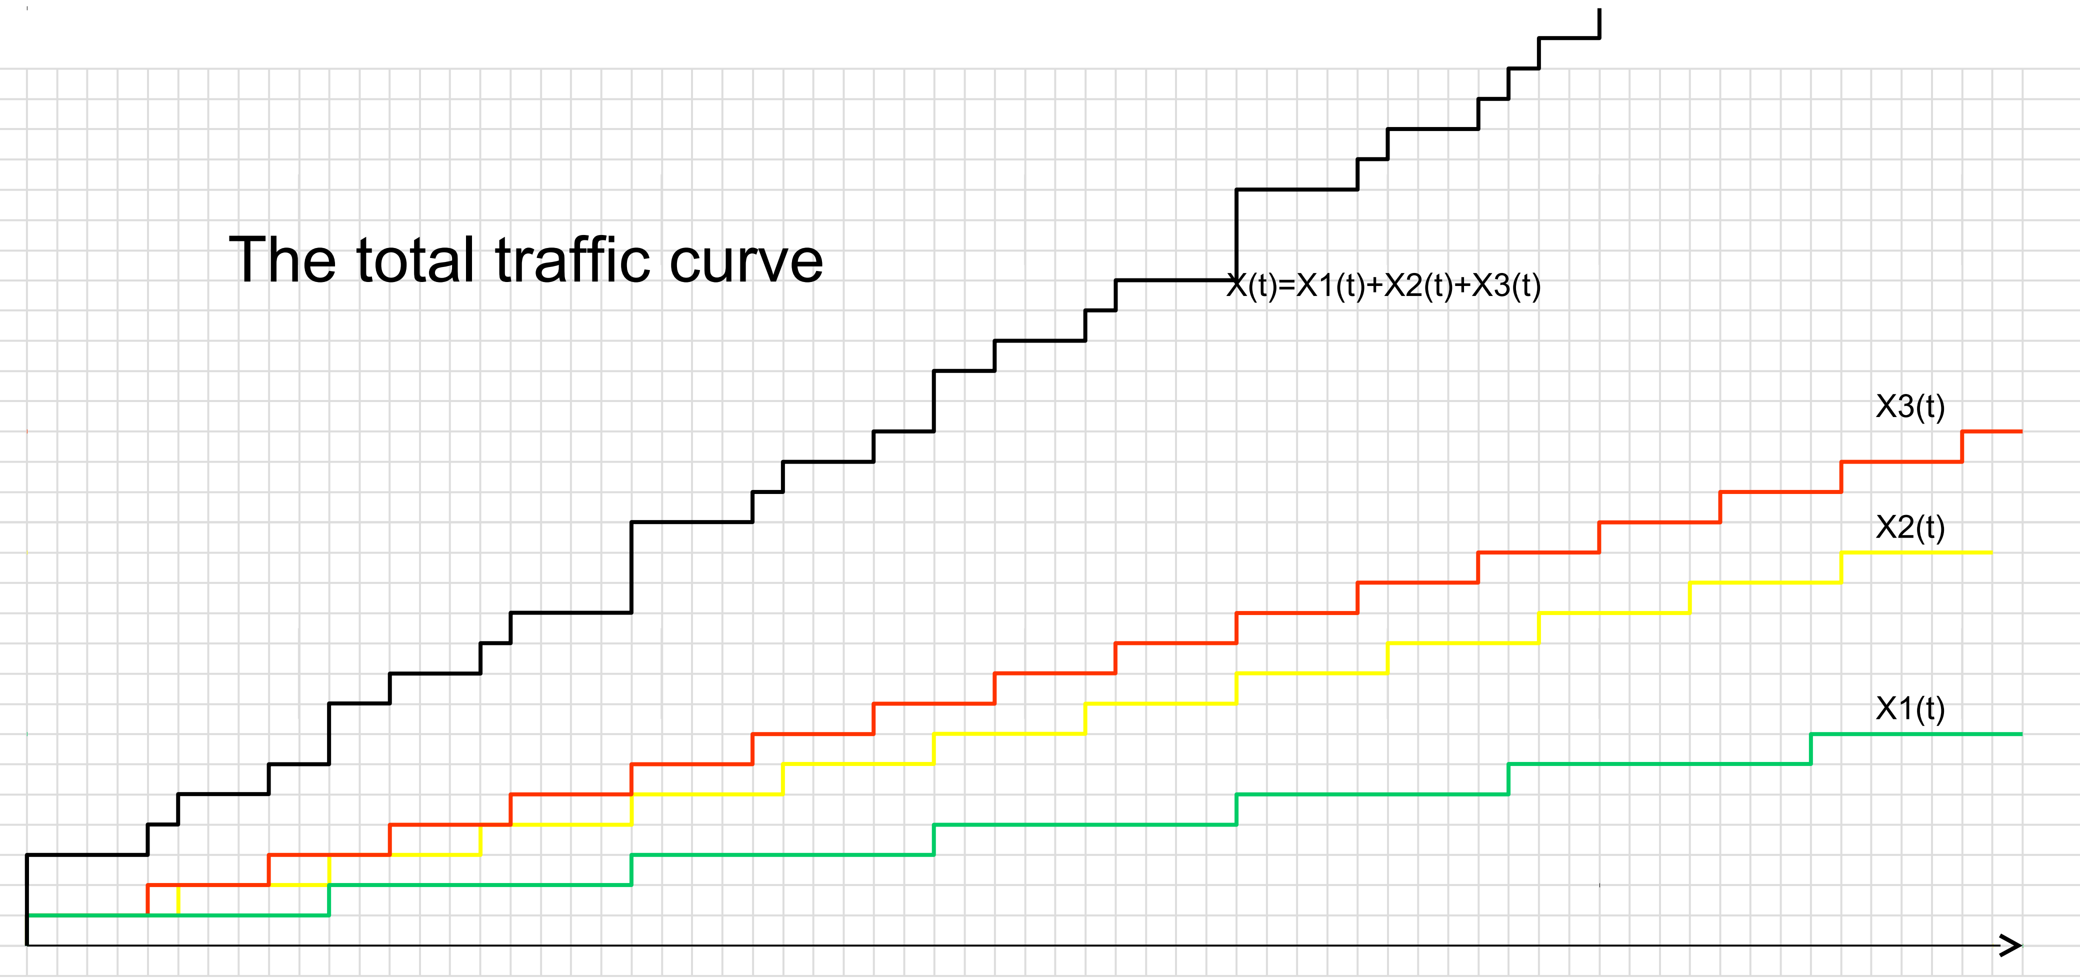

EF PHB

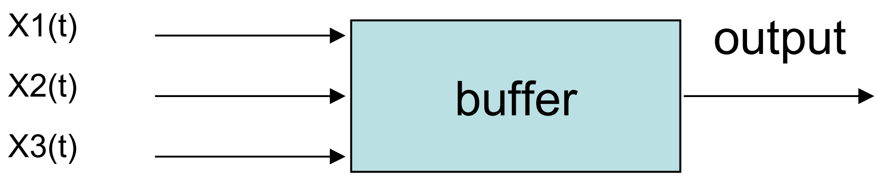

- Now, assumer that the line capacity is equal to 1 Mbit/s
- In this way, the service time of one packet is equal to 1 ms
- The service of one packet is over when the last bit of the packet is transmitted
- Therefore, if the first bit of a packet is transmitted at time t, the last bit of that packet will be transmitted at time t + 1 ms
- This is why the actual service curve grows by one packet as soon as the packet's service is over (i.e., 1 ms after the start of its service)

EF PHB

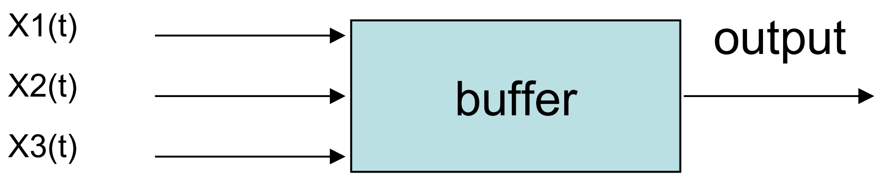

- The maximum delay is calculated as the maximum horizontal difference between the traffic curve and the actual service curve
- The maximum backlog is calculated as the maximum vertical difference between the traffic curve and the actual service curve

EF PHB

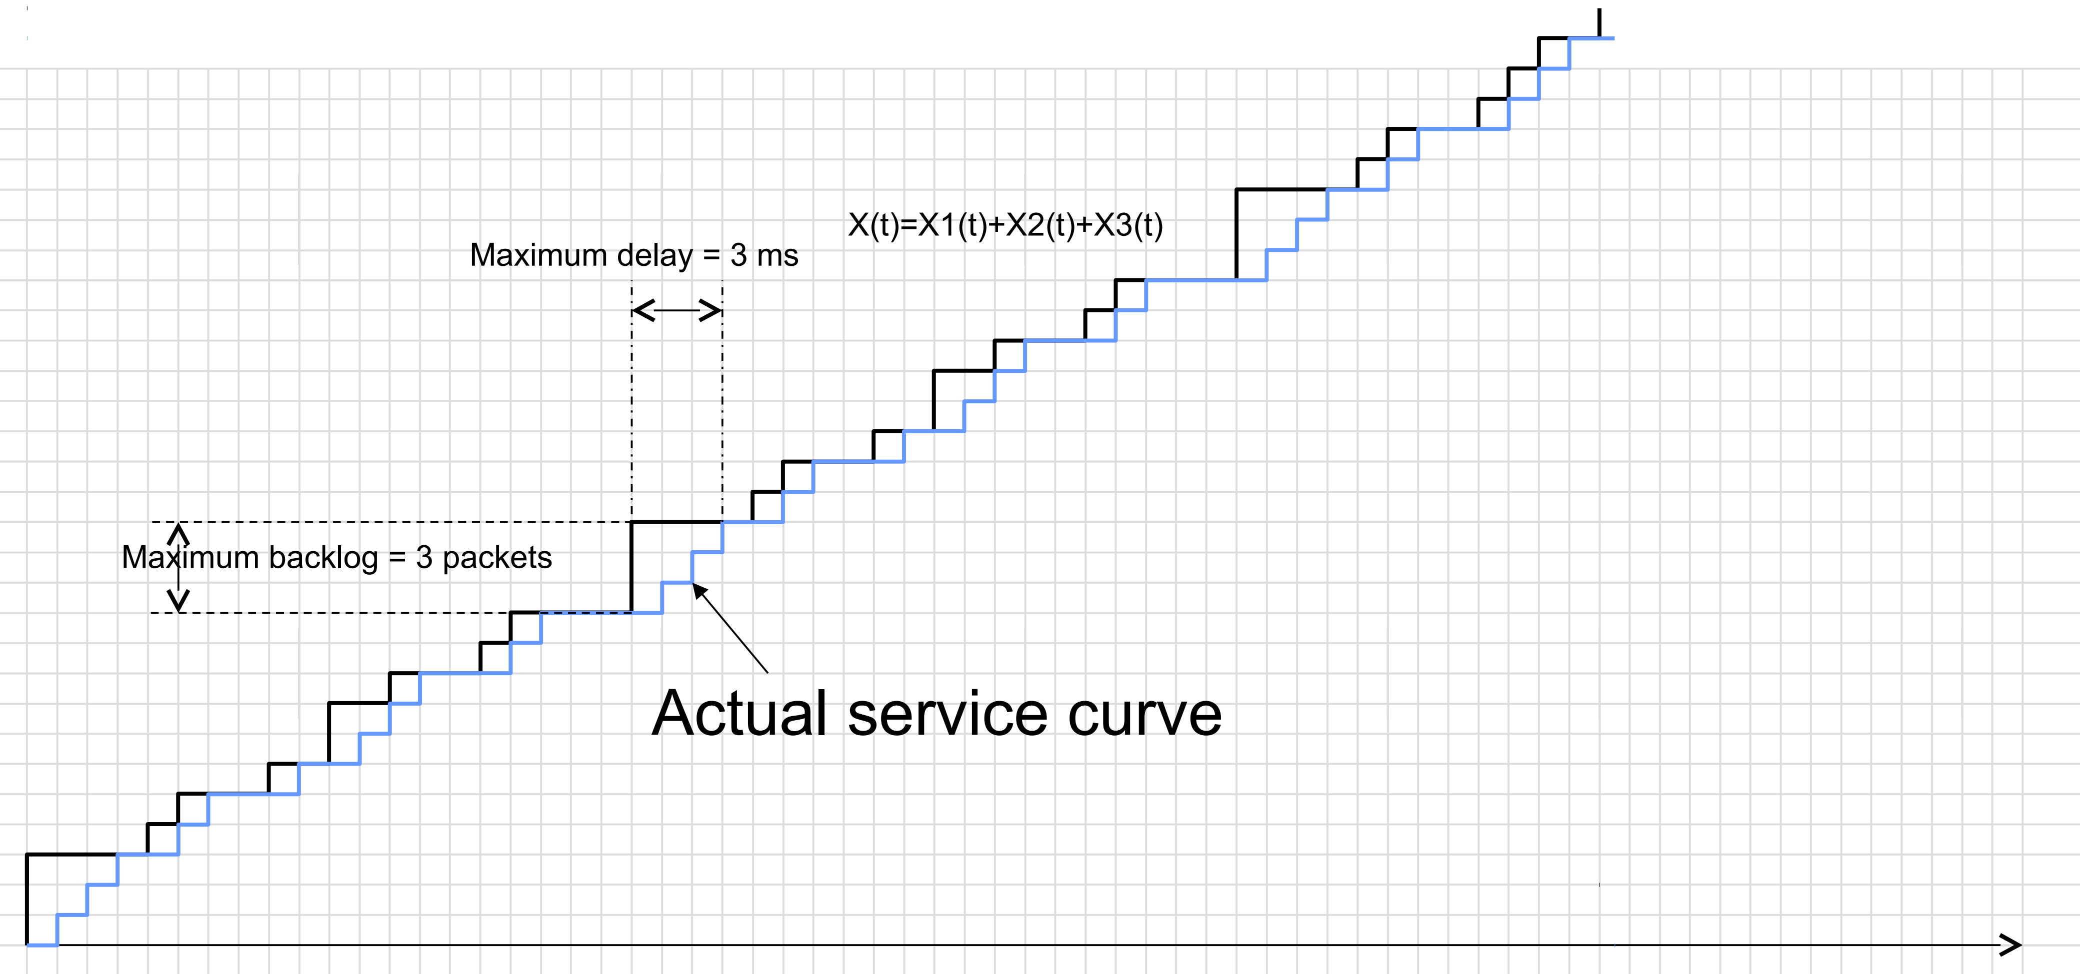
X(t)=X1(t)+X2(t)+X3(t)
Maximum delay = 3 ms
Maximum backlog = 3 packets
Actual service curve

EF PHB

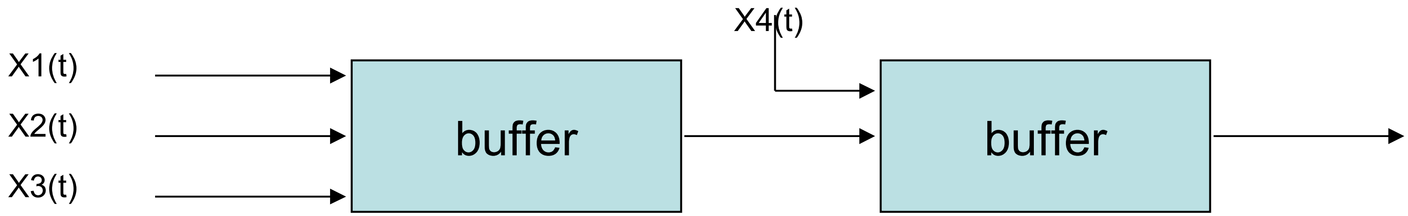

- Now, let us consider a multihop case, in which the output of the first scheduler is fed to a second scheduler, that receives an additional flow, X4(t)
- X4(t) has the same statistical features of the fresh flow X3(t)
- We want to calculate the maximum end-to-end delay

- The curve of the traffic X1(t) + X2(t) + X3(t), fed to the second scheduler, is the actual service curve of the first scheduler, previously calculated

EF PHB

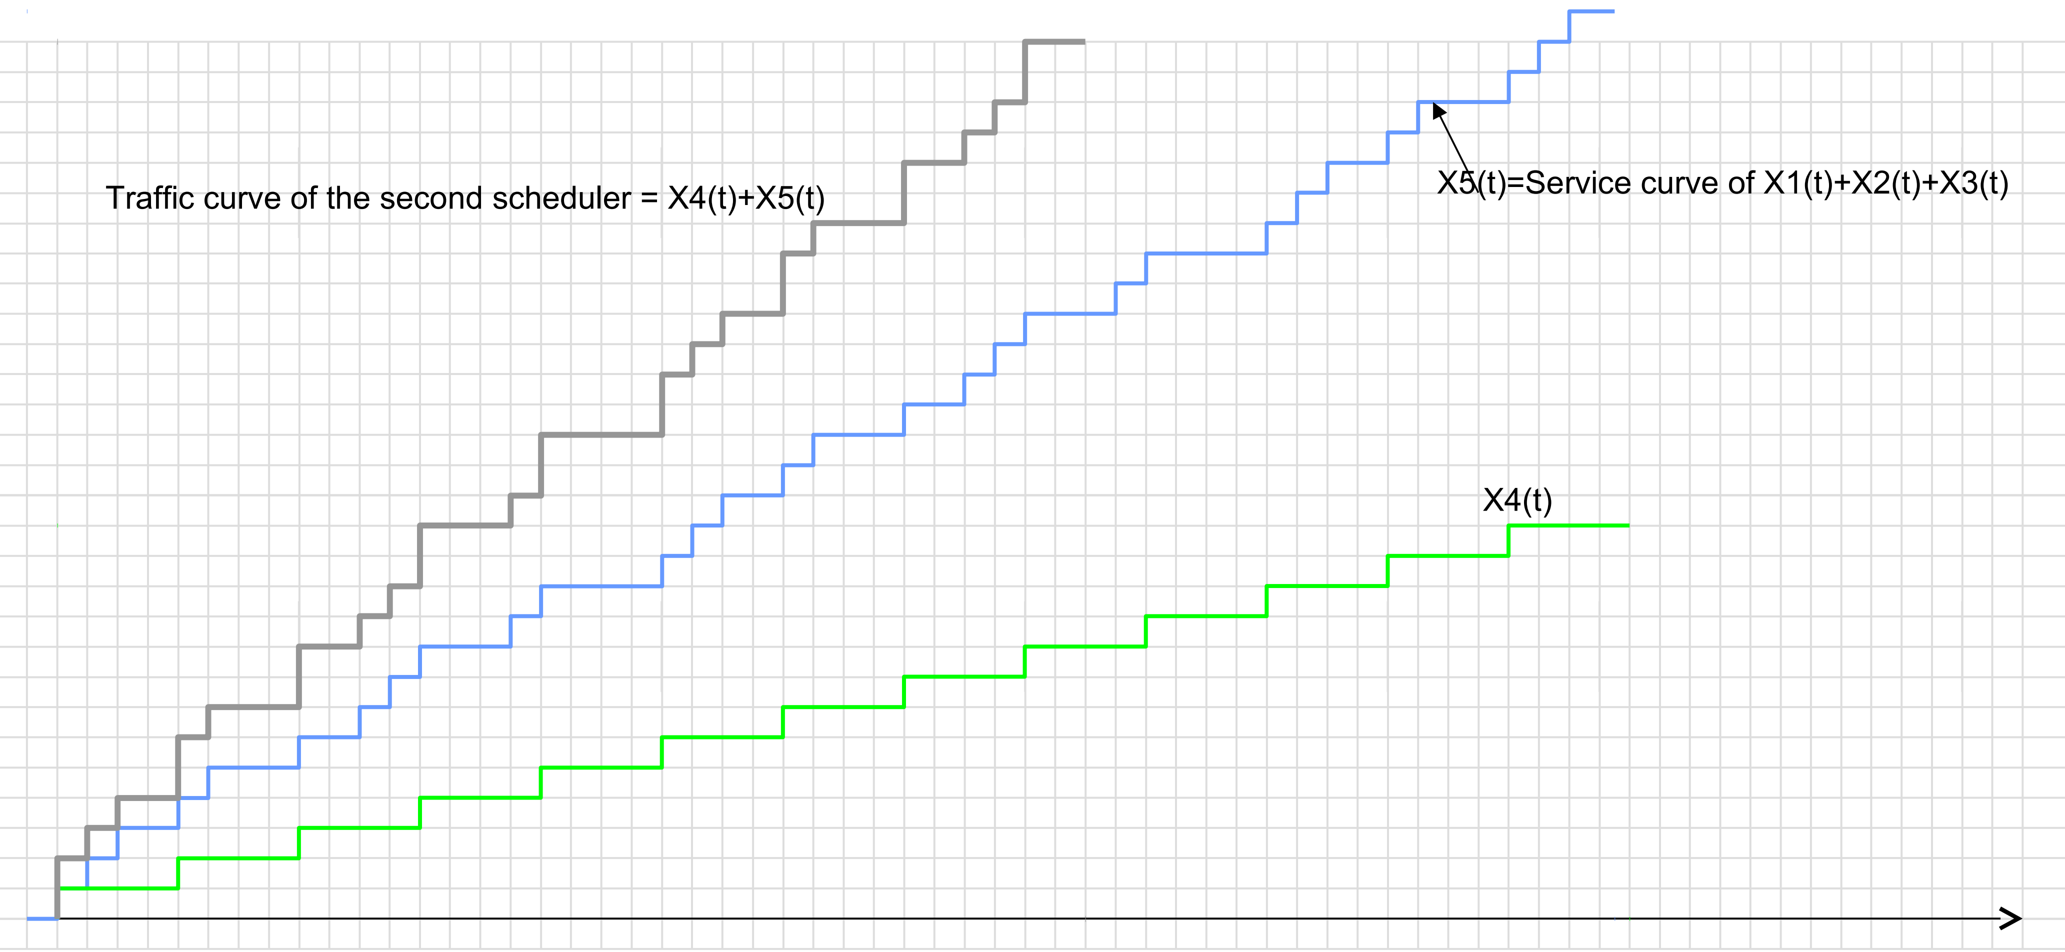
Traffic curve of the second scheduler = X4(t)+X5(t)
X5(t)=Service curve of X1(t)+X2(t)+X3(t)
X4(t)

EF PHB

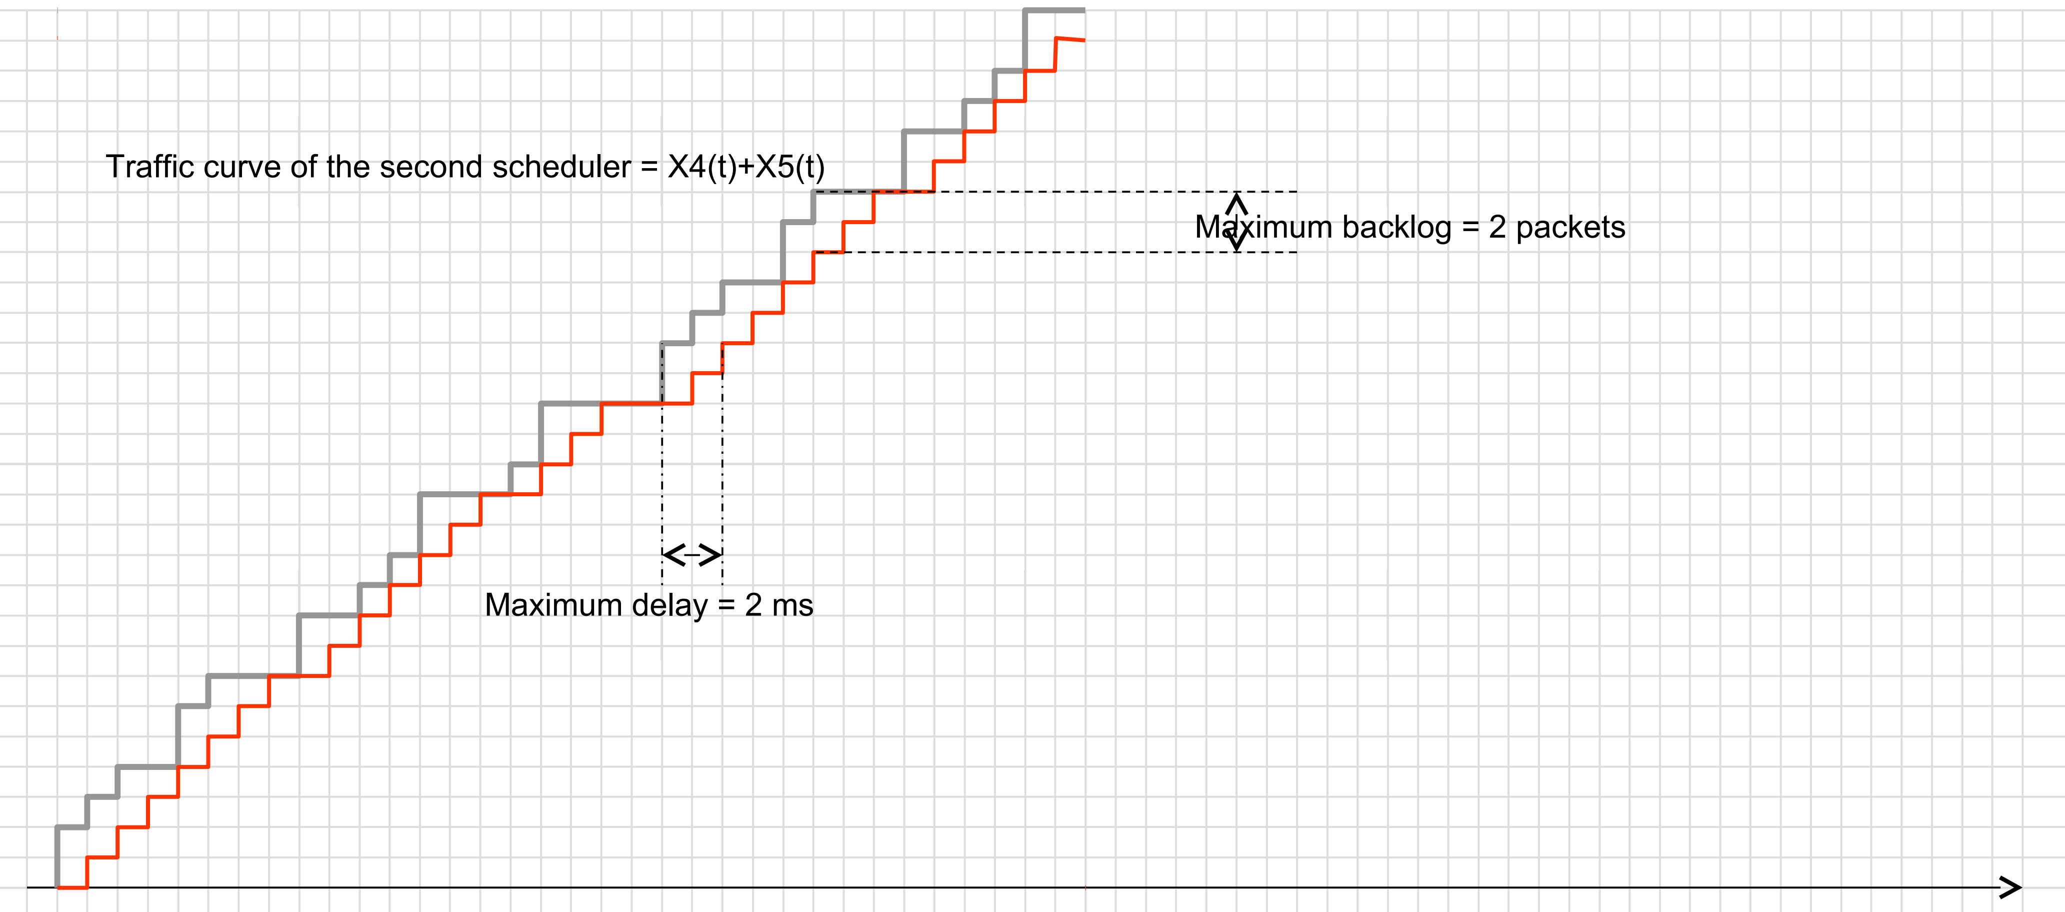
Traffic curve of the second scheduler = X4(t)+X5(t)
Maximum backlog = 2 packets
Maximum delay = 2 ms

EF PHB

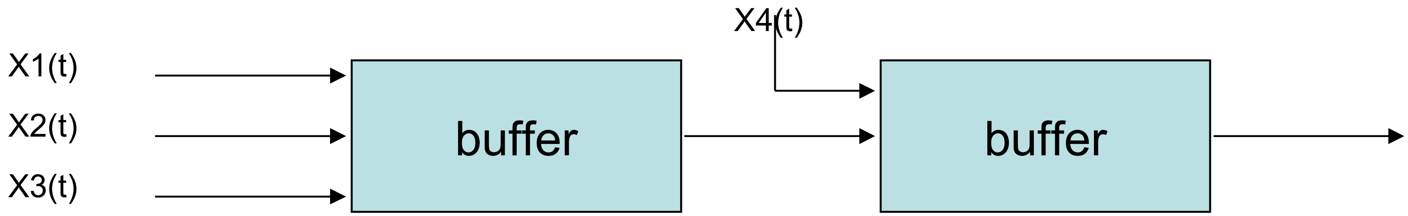

- Thus, the maximum end-to-end delay of flows X1(t), X2(t), X3(t) is equal to 3 ms + 2 ms = 5 ms

EF PHB

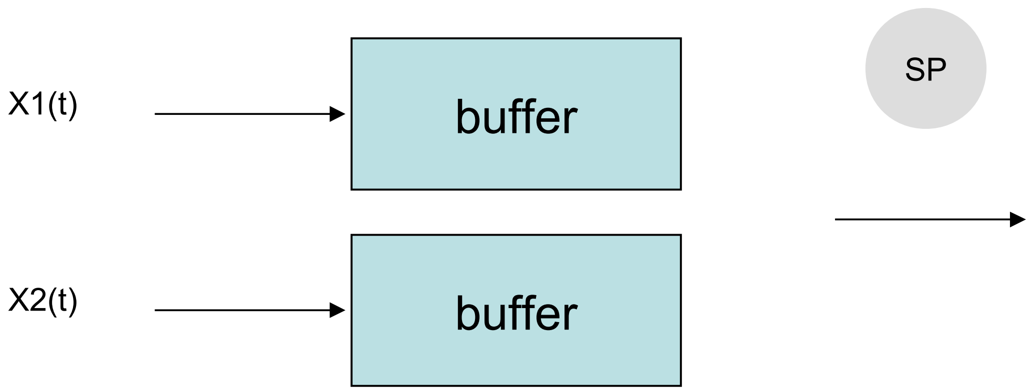

- Now, consider a multi-class scheduler (in particular, a Strict Priority scheduler) where a fresh EF flow (X1(t)) has higher service priority than a not-EF flow, X2(t)
- X2(t) might be a Best Effort flow
- In this case, the worst case delay for EF traffic is calculated by assuming that, when an EF packet is offered to the scheduler, a not-EF packet has just started its service
- The EF packet must wait the end of the not-EF packet's service, before starting transmission
- In the following example, let us assume that X1(t) has the same features of the X3(t) flow of the previous example
- Moreover, we assume that the packet length is constant and identical for the two flows

EF PHB

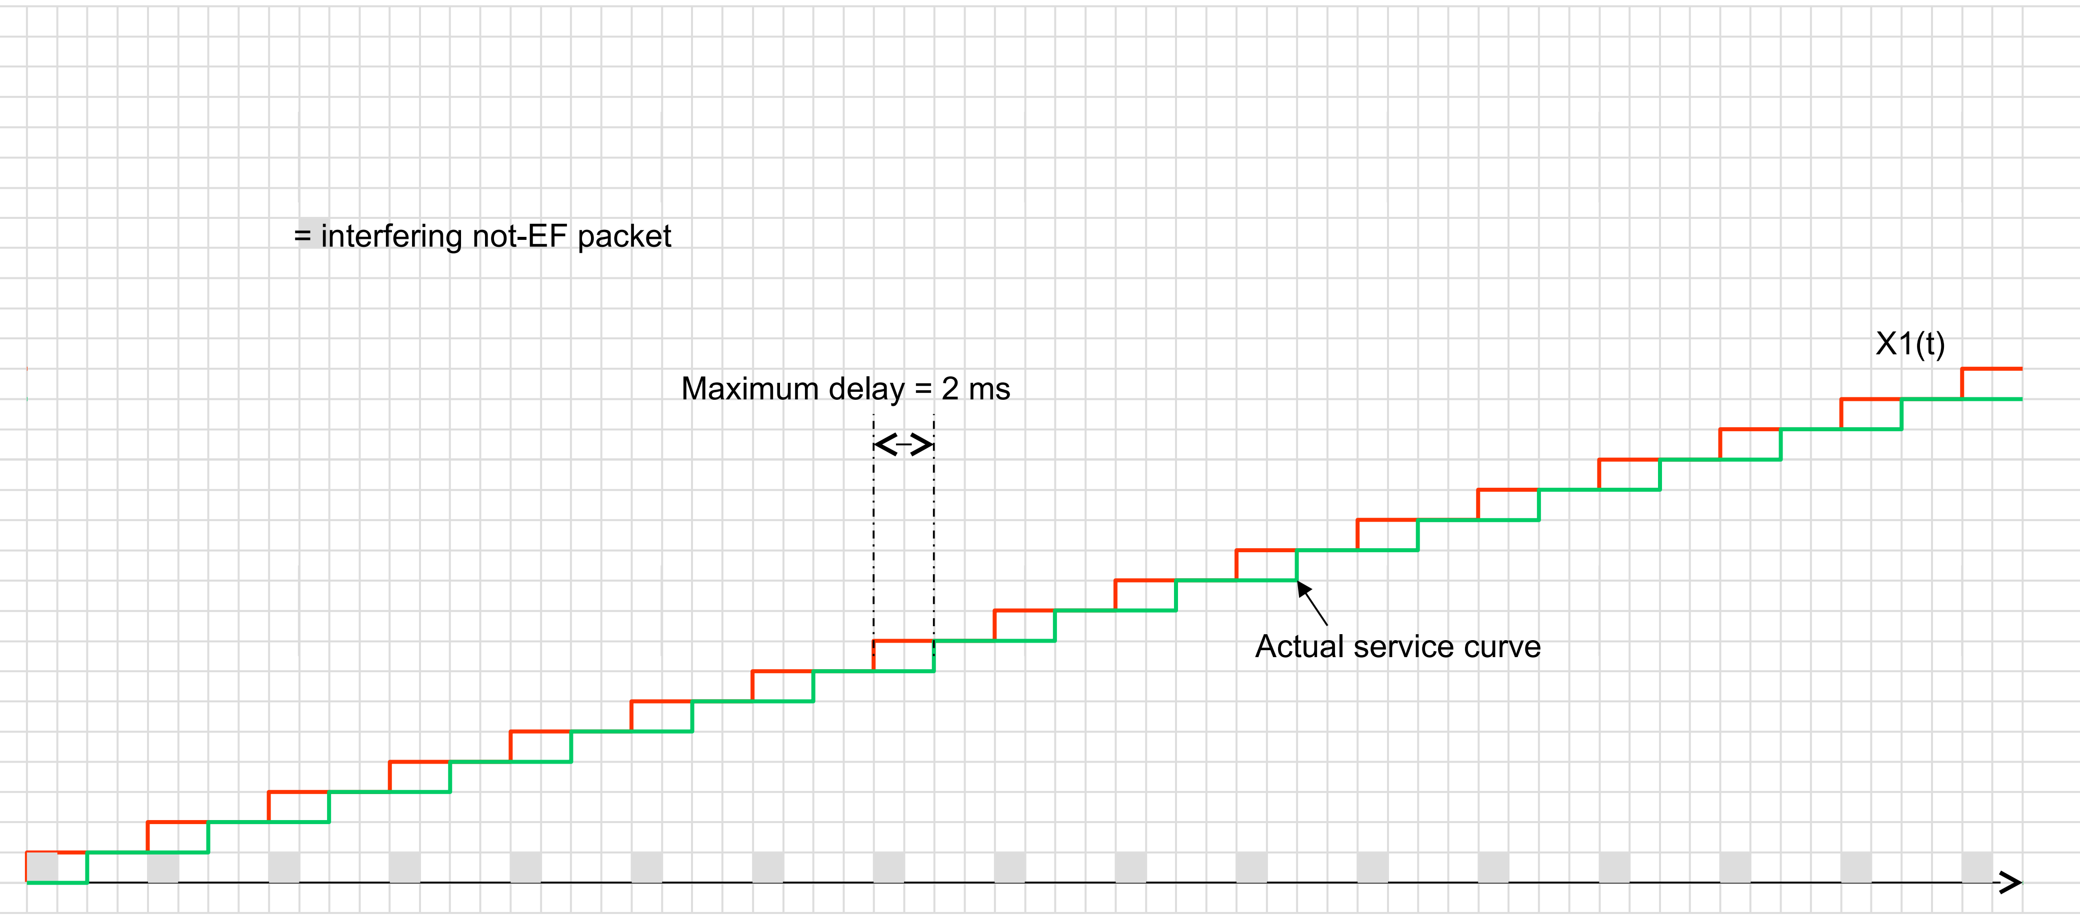
= interfering not-EF packet
Maximum delay = 2 ms
X1(t)
Actual service curve

Maximum configurable rate of EF traffic

- There exists a maximum configurable rate of EF traffic that can be allocated on a node's output line
- This is because all EF packets may find, at their arrival, another non-EF MTU-sized packet, this halves the sustainable EF rate
- In general, given an EF flow with packet length L1 and an interfering not-EF flow with packet length L2, with a line rate equal to C, the maximum configurable rate of EF traffic is obtained by considering the diagram below, where for each EF packet a non-EF interfering packet causes an additional delay

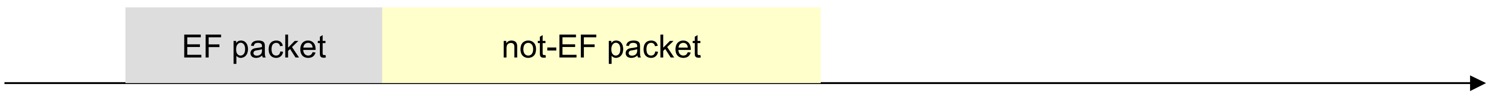

Maximum configurable rate of EF traffic

- The maximum configurable rate of EF traffic is C_{EF}
- The largest the not-EF packet length, the smaller the maximum configurable rate of EF traffic

$$C_{EF} = \frac{L_1}{L_1 + L_2} C$$

Final (tiny) detail: the DS codepoint of the EF PHB

- Codepoint 101110 is recommended in RFC 3246 for the EF PHB