Advanced Computer Networks
Professor Neminath Hubballi
Department of Computer Science Engineering
Indian Institute of Technology, Indore

Lecture 18
Traffic Management - Part 5

(Refer Slide Time: 0:36)

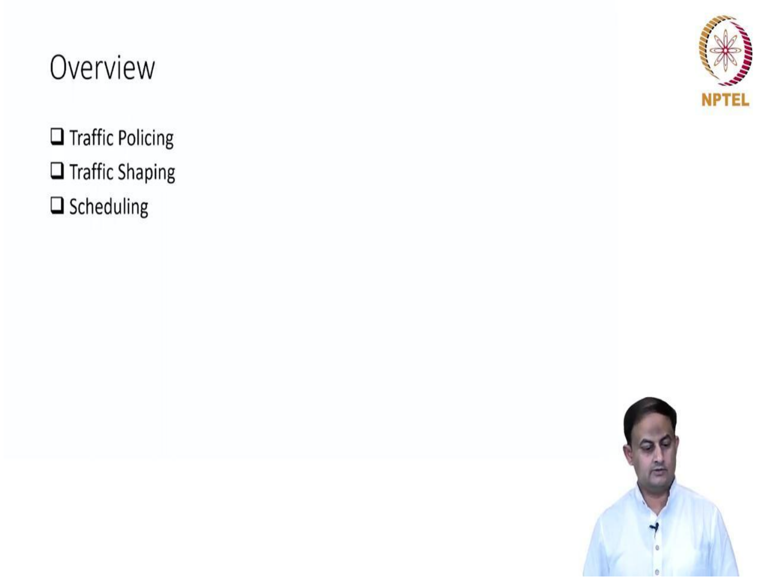

Overview

□ Traffic Policing
□ Traffic Shaping
□ Scheduling

Welcome back. In the last couple of lectures, we were discussing about the traffic management techniques. We will continue that discussion on traffic management today as well. So, I want to discuss in a little more detail these 3 things, we understood that traffic management is primarily done through these three techniques called policing, shaping, and scheduling. So, what I want to cover today is basically how exactly the policing, shaping, and scheduling work is done inside the router and what are the algorithms or the mechanisms that are used for this shaping, policing, and the scheduling.

(Refer Slide Time: 1:04)

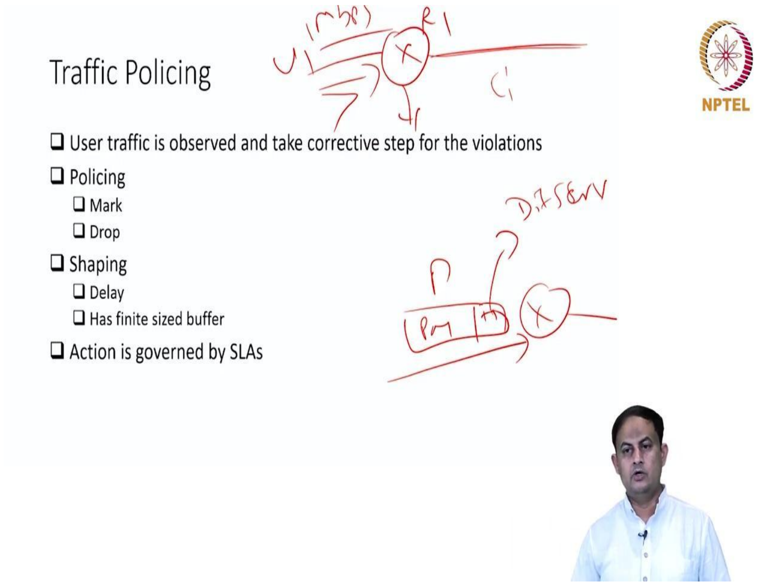

Traffic Policing

- ❑ User traffic is observed and take corrective step for the violations
- ❑ Policing
  - ❑ Mark
  - ❑ Drop
- ❑ Shaping
  - ❑ Delay
  - ❑ Has finite sized buffer
- ❑ Action is governed by SLAs

So, let us begin with the first method, called as traffic policing, policing is required to make the end user traffic rate to some agreed upon traffic rate. So, whenever the transmission rate is violating the condition or the agreed upon transmission rate, then we need to take the policing actions. So, what I mean by that is, if we have a router, let me call this router as R1, and there is a user, let me call that a user U1, he is transmitting the traffic to this router or to any other destination, and that is flowing through this router.

And let us assume that the user has agreed to transmit at a peak rate of 1 Mbps, and if the traffic arrival rate at this router R1 exceeds this, then you need to take the corrective action; that corrective action is called as the policing so you observe the rate of the arrival or any other parameter that you agreed upon and then you take the corrective action for the violations.

So what is the corrective action? So, we understood that there are primarily two things that the policing action can do: one is if the arrival is exceeding the agreed upon rate, then the simple thing that you can do is to drop the packet, up to 1 Mbps you are supposed to transmit, you admit and pass on to the next link. If it exceeds that, then you drop the packet right here. So, that is one thing, and the second thing is if there is a capacity available, remember this is the shared link it has got a capacity, if no other user is presently utilizing that capacity, if the spare capacity is available and although at this router R1, the arrival rate of the user U1 is exceeding the agreed rate then since the spare capacity is available, you can do the transmission.

You do the transmission with the condition that you mark the packet, what I mean by that is that the packet is arriving at this router, let me call that packet as P, and it has got two things one is the header portion, the second one is the payload and the router is here and the packet is arriving at the router and is exiting and you mark this packet, marking by that I mean you set the differentiated service field inside the IP header, it has got a bunch of headers, be it the link layer header, the IP header, and the TCP header and maybe the application layer header itself. So, you pick up the IP header from this packet and set the differentiated service field, and send it. So how does this help? Although the link between R1 and R2 probably is underutilized, you are able to transmit in excess of 1 Mbps, but R2 to some other next hop router R3, the capacity may not be available. So it is an indication for the downstream routers to drop the packet if the capacity is not available.

So if it has to make a choice for dropping a certain number of packet then probably the packets exceeding the 1 Mbps capacity are the candidates for dropping that is the meaning of that. So that is how the policing action is done, primarily to take the corrective action; corrective action here is the drop or mark. When the router R1 is receiving the packet from the user U1. So when one router is sending a packet to the second router, there might be an agreed-upon rate at which the router R1 is supposed to send it to the R2.

Although on an individual router, we may not do it, on an aggregated level for one network sending the traffic to another network, there might be some contract at the peak rate at which the first network is supposed to transmit to the second network then you actually do the shaping. So, in order to impose the constraint that I will not be transmitting more than what is agreed upon, you do the shaping.

So, what are the things that you do inside of the shaping, you basically delay the transmission, and you put the packet inside a queue, which has got a finite buffer and you make it wait, and then you transmit these other two things that happen inside the router. So, these two things, be it the policing or the shaping are governed by something called as the SLAs, which we discussed earlier as well. So, which packets, how much, when to mark, what to mark, when to drop, and when to delete those decisions are governed by these SLAs.

(Refer Slide Time: 6:10)

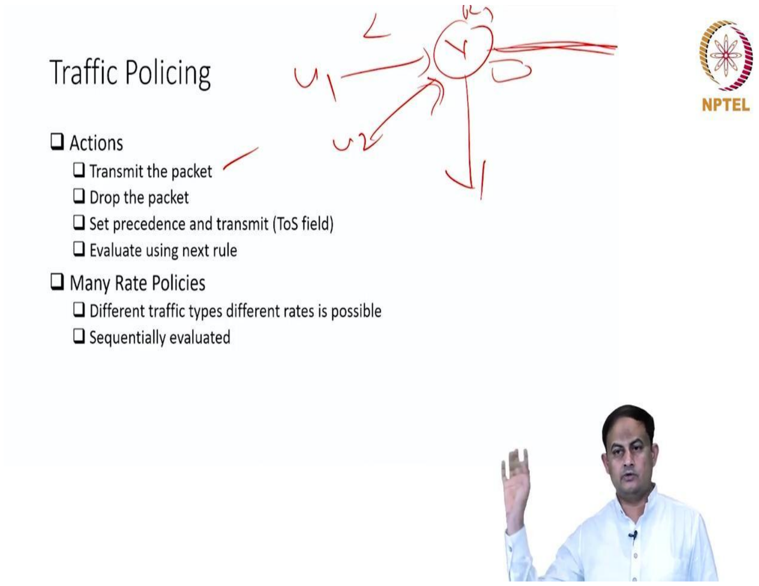

Traffic Policing

- ❑ Actions
  - ❑ Transmit the packet
  - ❑ Drop the packet
  - ❑ Set precedence and transmit (ToS field)
  - ❑ Evaluate using next rule
- ❑ Many Rate Policies
  - ❑ Different traffic types different rates is possible
  - ❑ Sequentially evaluated

So, the actions for the policing is basically you either decide to transmit the packet. So, when you transmit it, when the arrival rate from user 1 is less than the agreed-upon condition or the agreed rate, and you decide to drop the packet when it is in excess and you do not have the capacity to accommodate the excess capacity that time you drop it.

So, if there is a spare capacity available at the outgoing link and in spite of the user transmitting the traffic in excess of the agreed-upon rate, then you mark the packet and send it to the next hop. And these things, whether to drop, mark, and who is transmitting the choice of determining what action to apply for the particular series of the packets is done through the classifier, which we studied earlier, and the classifier has got a collection of the rules, and it might happen that more than one rule needs to be evaluated against a particular packet.

So, you go one rule at a time and then subsequently go and evaluate if any other rule is actually applicable. So, all of this, whether you decide to mark it, transmit it, drop it, or do this you need to go through a series of rules inside your classifier to arrive at a particular choice. And it is possible that many different transmission rates are possible. So even for a particular user itself, you might say that I am going to give you this peak rate if this is the condition, this peak rate if this is the condition, something like that can be defined.

So as long as the SLAs are very clear about what to happen, when it should happen then it is actually quite very clear. So there are a bunch of users transmitting to this router. So they are separated, or there is a different set of rules defined for them, user 1, user 2, and you can apply those rules to the router. So user 1's traffic is going at a certain rate, and user 2 traffic is going at a different rate; that is possible. So this policing action, as I said, although we saw the picture of the user transmitting the traffic to the first hop router, usually the first hop router is the one which actually does this policing action, but the notion of the first hop router can be different.

So one action is where the end user is transmitting to the router, and this router R1 has got an outgoing link, there you are doing the policing action, this is the router where you are doing the policing action. Or it might also happen that the end user is somewhere here, and his traffic is going to the first hop network, and this network has got internally a series of routers, and the traffic passes through these routers; the policing action is happening here. And then, when it is exiting one network and entering into the second network.

(Refer Slide Time: 9:00)

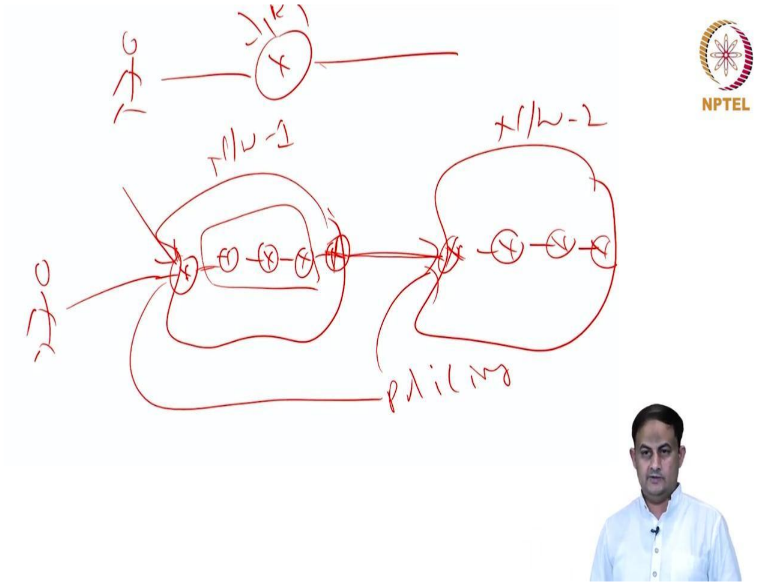

So one action is where the end user is transmitting to the router and this router R1 has got an outgoing link, there you are doing the policing action, this is the router where you are doing the policing action. Or it might also happen that the end user is somewhere here and his traffic is going to the first hop network and this network has got internally a series of router and the traffic passes through these routers, the policing action is happening here. And then when it is exiting one network and entering into the second network.

So from this border router to the border router of the second network, this is the network number 1 and let me call this network as a network number 2, and again, you put you do a marking you do whatever the rules you want to apply right here which is an edge router and we will come to… in between this one this one nothing actually happening these routers which are internal to the this network 1 are basically primarily looking into that marking which is already done and then just transmitting the packet to the next hop router.

So, when the packet arrives at the border router when it is transmitting at the next hop router, this router is receiving the packets from the border router of a previous network there also the policing action can be done. So, here at this point of time and at this location, so, the policing action is happening.

So when I say that incoming traffic is actually policed, so, if the network one violates the transmission rate agreed transmission rate then the network 2 can decide to do the policing action basically it can drop network 1 you were supposed to transmit at so and so rate and you are exceeding that rate I am going to take the corrective action. So, at every… let us say, the network 1 we can think of as an autonomous system network 2 as a different autonomous system and the source and destination are apart are separated by 10 different autonomous system, at 10 different places the policing action is happening during the transmission.

(Refer Slide Time: 11:34)

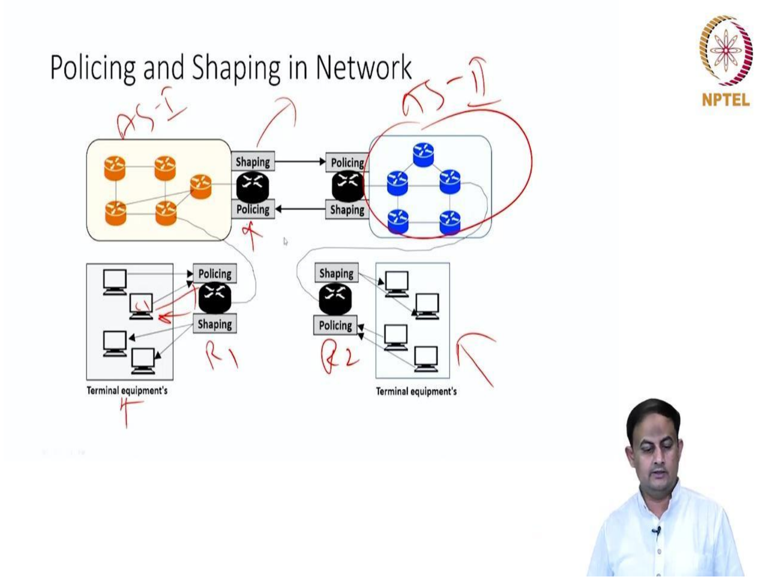

Policing and Shaping in Network

So, here is a picture that is actually showing in the network where exactly the policing and shaping actions are taking place. So, as I said, every time a router receives the packet or receives the traffic from either the end-user or the network there policing action is required. And every time a router is sending traffic to some other network, that time shaping is required. So, what you see on this diagram is a bunch of computers connected to this router, let me call this router as R1, and a bunch of computers connected on the right-hand side to this router, let me call this router as R2.

And any computer on the left-hand side transmitting the packet to some computers on the right inside is going through this router R1 So, the packets of, let me call this computer a C1, C1 traffic is arriving at the router R1. Now R1 is the place where the policing action is happening.

So if C1 is supposed to send at a particular rate, let us say 1 Mbps, any excess traffic in excess of 1 Mbps will be dropped or marked at the R1. And anything that is going to the C1 received from the other end would be shaped at the R1, so traffic exiting the C1 and going to R1 is policed, and traffic exiting the R1 and received at C1 is shaped that is the best way to understand what is shaped and what is policed. And again, you can see that this traffic assuming R1 is part of the same autonomous system as these routers and R2 is part of this autonomous system number 2 as that of these bunch of the routers. So, across these routers, whatever is marked in the yellow color, no policing and shaping is happening. So when the traffic is actually exiting the network, I am going to the second autonomous system shaping is happening, and whatever is received from the autonomous system 2 is actually policed at this router. So that is, between the autonomous systems or the networks you do the policing and shaping operation. So that is the big picture. And policing and shaping are the two things of the traffic management action. So, the third one is scheduling.

(Refer Slide Time: 14:20)

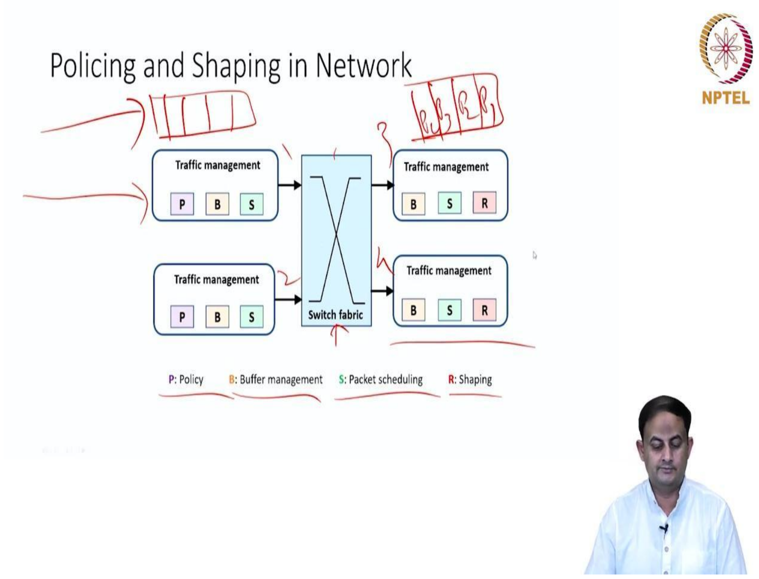

Policing and Shaping in Network

P: Policy B: Buffer management S: Packet scheduling R: Shaping

So, here is another picture. You can recollect our previous discussion previous lecture discussion, what we said is any router in the network has got a line card, and it has got a switch fabric and the switch fabric is here, and the line card or whatever the packets are coming they are arriving at the line card, and they are on a particular port.

Let me call this port number as 1, 2, 3, and 4 and any packet which is arriving at input port. Now the header information of that packet is consulted, and then you look into the FIB information available at the input line card and you decide on which output port number that particular packet should go through, and you request the switch fabric I want to transmit this particular packet to so and so output port number, can you establish a link between this input port and that output port and the switch fabric will come and do that link establishment and the actions that we want to take that is the routing decision or the forwarding decision that you are doing.

In addition to that, we are doing the traffic management, and the actions of the traffic management are also done at the input port and also at the output port. So, what this picture is showing is what actions are done, what kind of traffic management components are sitting at the input line card, and what kind of operations are done at the output port. So, you can see that, here is the notation P stands for policing, B stands for buffer management, S stands for scheduling, and R stands for shaping.

So, buffer management, policing, and scheduling happen at the input port, and buffer management, scheduling, and shaping happen at the output port. So, policing, buffer management, and scheduling happen at the input port I repeat, and on the output port, you have the buffer management, scheduling, and shaping actions happening. So, remember there are two queues one is called as the input queue, where the arriving packets are put, this is at the input line card or the input port number. And once the FIB consultation is done and you decide on which port number you want to forward this packet to, so, there is a queue at the output port as well, and the packet comes and sits inside this output port.

Now, you can decide in what order you want to pick up the packets from the output queue itself. So, let us say there are some series of packets that have arrived on the input port number 1 there are some other packets that have arrived on the input port number 2 and all of them want to go to take the exit route of port number 3 and the next hop whatever is connected to the port number 3, and then you have a mixture of the packets coming from the port number 1 and another set of the series of packets coming from the input port number 2, they are all sitting inside the queue that is available at the output port number 3 and now maybe the P1 has come from port number 1, P2 has come from the port number 2, again P3 has come from port number 2, and P4 has come from the port number 1.

Now you can decide so, whether I want to transmit packets that are on the input port number 2 with a priority or not that is why the scheduling operation is required. So, in a nutshell, policing, buffer management, and scheduling happen at the input port, and the buffering, scheduling, and shaping operations happen at the output port when you want to transmit the packets on the output line.

(Refer Slide Time: 18:34)

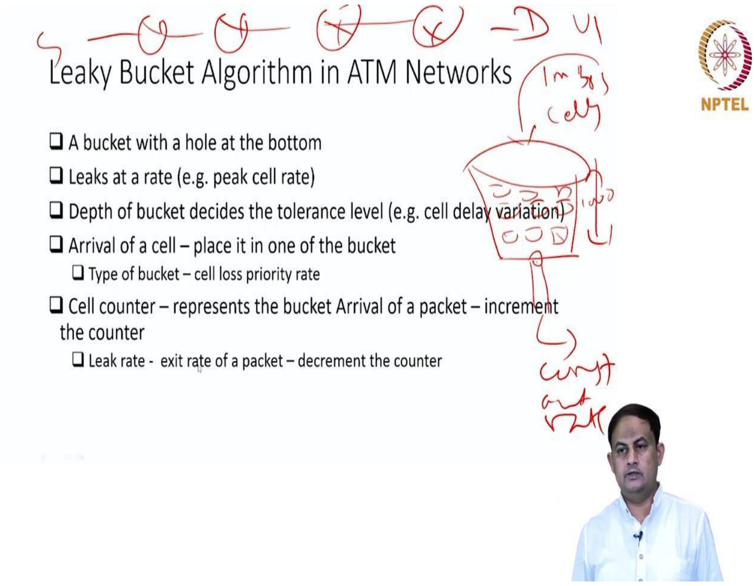

Leaky Bucket Algorithm in ATM Networks

- ❑ A bucket with a hole at the bottom
- ❑ Leaks at a rate (e.g. peak cell rate)
- ❑ Depth of bucket decides the tolerance level (e.g. cell delay variation)
- ❑ Arrival of a cell – place it in one of the bucket
  - ❑ Type of bucket – cell loss priority rate
- ❑ Cell counter – represents the bucket Arrival of a packet – increment the counter
  - ❑ Leak rate - exit rate of a packet – decrement the counter

Now, the question is how do we exactly do both policing and shaping? So, policing we want to do, we said that anything in excess of the agreed-upon rate could be either dropped or marked, and while shaping, you might put them inside the buffer and delay the transmission. Now, the question is how exactly this is done, and how do I determine at what rate the end user is transmitting?

So if it is transmitting at more than a certain rate, what exactly want to do so in practice, this is done both shaping and policing are done with an algorithm called as the Leaky bucket algorithm. So the conceptual idea is the following, you can assume that there is a bucket which has got a finite capacity, and the packets which are arriving on the input line card are put inside this bucket. And this bucket has got a hole at the bottom, and the hole is of a fixed size, and the packets are exiting from that hole at a constant rate.

So that is how you bring the notion of the constant rate. So let us visualize how the Leaky Bucket algorithm would work. Here is a bucket, and the packets are coming to this bucket at a certain rate. And let us say this user U1 is transmitting to this router, and the bucket is available at the router. And so the user U1 is supposed to transmit at 1 Mbps rate, and if he is transmitting at more than 1 Mbps rate, those packets are buffered inside this bucket. So, maybe the packet P1, P2, P3 all of them are put inside this bucket, and they take turns.

So, depending upon the size of the hole available at the bucket, they are exiting at a constant rate. So, such an algorithm is called the Leaky Bucket algorithm. So, the bucket leaks at a constant rate. This algorithm was primarily designed keeping the ATM networks in mind.

So, ATM networks have got the packets, whatever we are calling them, as they are technically in the ATM networks called as the cells, cells arrive at a certain rate, and they are put inside this bucket, and the cells are exiting from this bucket at a constant rate. So, one of the nice things about the ATM network the cells is all the cells are up to equal size. So, unlike the IP packet, an IP packet can carry a variable amount of data inside it, but in the ATM network, the cell has got a fixed sized content inside it. So, every time something exits the network a fixed cell size cell is exiting the bucket.

So, exiting is called as the bucket leaks at a certain rate. So, we call it the peak cell rate when the bucket is full, something is given by the end user so that is exiting at this particular peak rate. If something is not available, then nothing is actually flowing through this bucket. So, in a way, the bucket itself has got a finite capacity and the capacity can vary. So it can accommodate maybe 10 cells, 100 cells, 1000 cells.

So, the bucket size dictate the term something called as the tolerance level. So, this tolerance level has something to do with the end-to-end delay. So let us say the sources and destination are 10 hops away, and at every router in between, you are having such a bucket and the packets are actually here is your destination, here is your source. And the cells are now put inside this bucket at every router.

Now, bucket has got a capacity and larger the capacity that you have, larger the number of cells get accumulated inside the router or this bucket, when the transmission rate is exceeding the available capacity. So, what it means is more the capacity you have, more the number of cells getting accumulated, more the time it takes for them to exit the bucket. So the source is transmitting the particular cell, and it is expecting it to be delivered at a certain time. And by virtue of putting more capacity, more sized buffer, more sized bucket inside the routers, what we need is the cells to wait in each of these routers for a longer duration of time. If the source is constantly pumping the packets, these cells will spend more time inside these routers or inside the bucket, then I may not be able to meet the end-to-end delay that we agreed upon.

So that is why it is a little counterintuitive, the larger the capacity of the bucket, the larger the time they spend and then larger the end to end delay, so, that, we may not be able to afford. So, in a nutshell, what I want to achieve, what is the end-to-end delay I want to bring in, and how much is the buffer capacity that I want to include inside this leaky bucket is the question, I can actually adjust that accordingly.

So, as I said, as the packeta arrive, they are put inside this bucket, and then they are exiting. So once the bucket is full, then the packets are lost, or the cells are lost. If let us say, the 1000 cells you can accommodate inside this buffer. And when 1000 cells are sitting inside this bucket, the 1001 cell that is coming would be automatically dropped. So I can bring some kind of the definition while dropping whether the newly arrived packets are dropped or something else which is already sitting inside the bucket itself is removed and then made room for the newly arrived packet that is the question.

So, that is something to do with a scheduling operation. So, that differentiation you can bring, you can assign a notion of the priority while deciding which packet to drop from the bucket. So, every time something comes inside the bucket, there is a counter kept that counter increments, that will indicate how many cells are actually buffered, and every time a cell exits from the bucket from the bottom, there is the counter is actually decremented.

So, if let us say the counter can go up to 1000. So, if the counter ceiling value is 1000 and the minimum value is 0, so, when one cell exits, you may decrement by 1, 999 and then a new packet can come and sit inside the bucket. So, that is how the Leaky bucket algorithm actually works.

So, rememberthis is actually working operating when the ATM network cells have got a fixed sized capacity, and every packet is of the same size. So, the notion of the number capacity defined in terms of the number of the cells is actually going to work.

(Refer Slide Time: 26:26)

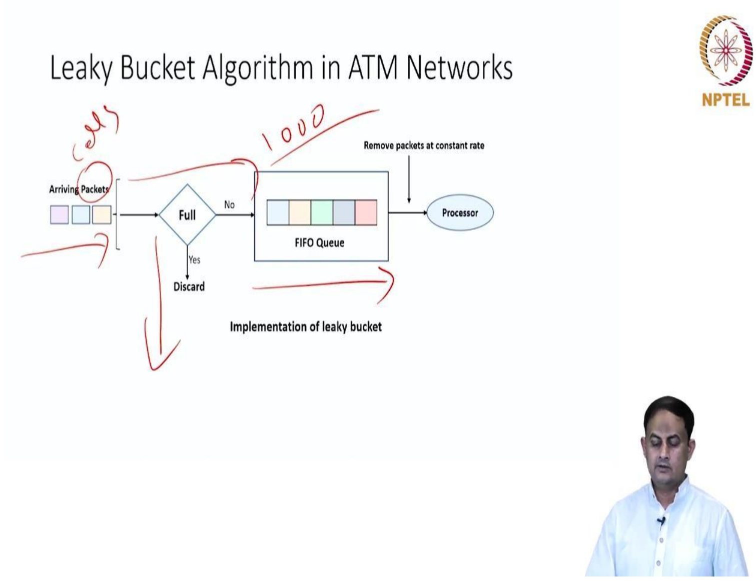

Leaky Bucket Algorithm in ATM Networks

Implementation of leaky bucket

But when you use the IP network, this is not the case. So, here is a picture which is actually showing how exactly the Leaky Bucket algorithm works in ATM networks. So the cells are arriving from the left-hand side here, what is shown as packets is technically the cells, and there is a queue here, and this is where the buffer or the bucket is.

So this bucket is now operating in the form of the FIFO queue where the first cell that has come into the bucket is the first one that is exiting the bucket. And every time a cell arrives, you ask this question whether this bucket is full or it has the capacity. If it can accommodate 1000 cells, are there 1000 cells already queued up, or do you have less than 1000 cells? That is the question that you ask.

If you have 1000 cells already queued up, then you discard the packet right here. If there is a capacity available, then you put it inside the queue. So once you put the cells inside this queue, then you can do a decision, subsequent processing, where to forward it to and how many of them to forward it and all that decision can be subsequently taken. So that is how the Leaky Bucket algorithm in the ATM network works with the FIFO queue. So FIFO queue is a very simple queue, but that is not the only type of queue you might have. There are other types of queues. We will come back and discuss that subsequently

.

(Refer Slide Time: 27:59)

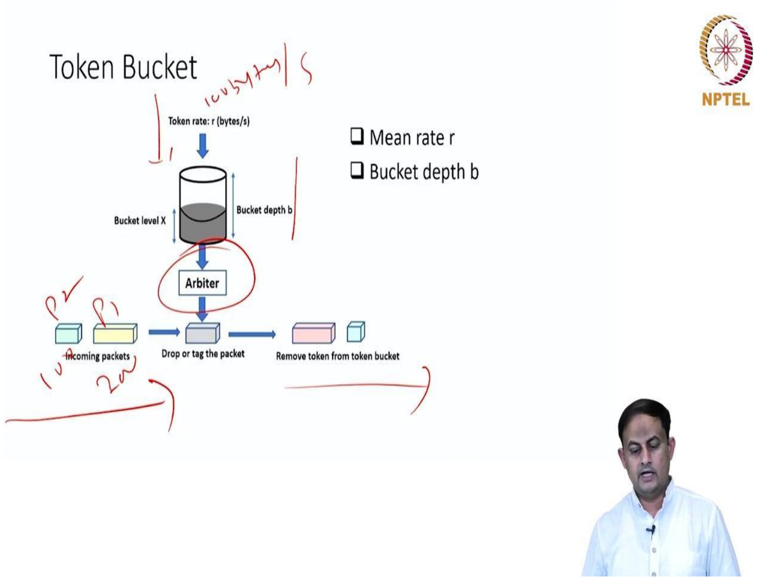

Token Bucket

❑ Mean rate r

❑ Bucket depth b

So that is the Leaky Bucket algorithm with respect to the ATM networks. And the concept of the Leaky Bucket algorithm in the IP network takes a little bit change. So the companion version in the IP network is called as the Token bucket. So, the reason why this change was required is because of the variable-sized packets inside the IP network, so here is a picture, the incoming packet; let me call this packet as P1 this packet is P2; P1 might have, let us say 200 bytes of the data and the P2 might have 100 bytes of the data, the size of the packets inside the IP network is variable.

So it can have 0 bytes, or at the maximum, it can have 65,535 bytes of data. The size can actually vary between these two, between 0 to 65,535. So now, how exactly do I define at what rate do I want to transmit? Because now the counter on a packet number is not going to work out, I need to define something else. So the token bucket, the way it works is at a constant rate, I want to transmit certain traffic at a certain rate, and I am going to generate something called as the token per unit of time.

So this has got a fixed-sized quantity. So let us say I am going to generate 100 Bytes worth of tokens per second. So that is the rate at which the tokens are generated and put inside the bucket, and these tokens are getting accumulated inside the bucket. Remember in the Leaky Bucket algorithm the cells are the ones that are getting accumulated inside bucket, but in the token bucket, the tokens are the ones that are getting accumulated inside the bucket. So, at any point of time, there is certain number of tokens available inside the bucket, and the packets are arriving on the input link.

So, you match what is the size of the incoming packet, you take that and then how many bytes I need to transmit if the tokens worth that capacity is available inside the bucket, and then you will be able to transmit this packet, otherwise you need to drop that particular right there. So, if in this case, P1 is carrying 200 bytes worth of data. So, if the tokens worth 200 bytes are available in the bucket, then you will be able to transmit it otherwise, the packet P1 is actually dropped, or you might put it inside a queue and wait till 200 byte worth of the tokens are getting accumulated inside the bucket.

So, the one shown here, the Arbiter is the one that is actually checking or comparing what is the size of the data that is arriving and how many bytes worth of the tokens are available inside the bucket. So, that is a change, otherwise, the rate at which you want to transmit since the tokens are arriving at a constant rate, you are actually guaranteeing that the exit rate is constant, confirming to what is the token generation rate.

So, if your arrival rate is smaller than the rate at which the tokens are getting generated, then none of the packets are actually technically queued up or dropped inside the network or the routers. So, that is how this algorithm works. So, the Leaky bucket algorithm or the token bucket algorithm is used in both the policing operation and also in the shaping operation. So both policing when you want to curtail or reduce the rate of the transmission, or when you want to receive the traffic from some user, you do not want to receive the excess amount of traffic, then both times the Leaky bucket algorithm can be used that is precisely for policing and the shaping operations.