Back to admission control

- A simple example of admission control in a scenario with a single link served with a FIFO scheduler is shown in the figure
- *N* statistically identical flows share the link
- The statistical SLA is (*d*, *p*)
- Delay increases as the number of flows grows
- There exists a maximum number of admissible flows, N_{max}, such that the SLA is fulfilled
- If $N>N_{max}$, the SLA is violated

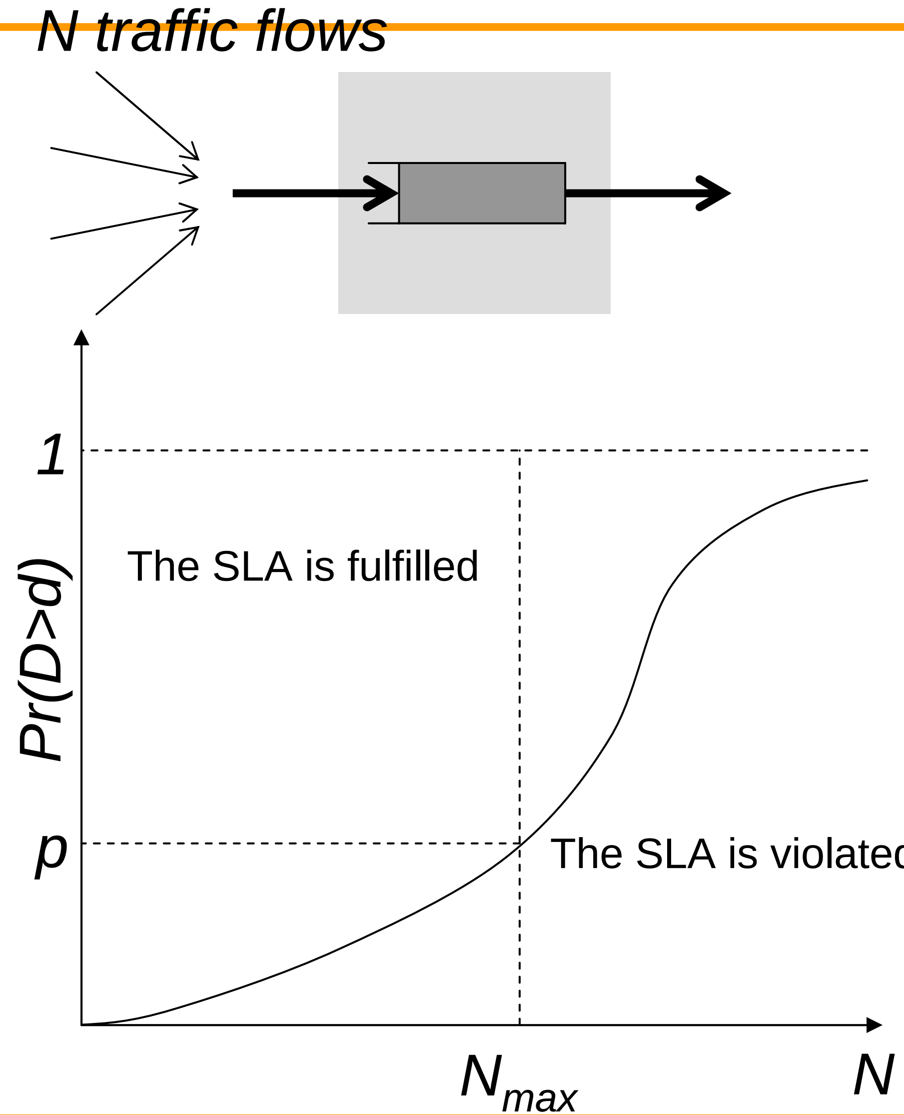

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Back to admission control

- For a single-class scheduler, the value of N_{max} defines the ***admission region*** of the scheduler
- For a single-class scheduler, the admission region is uni-dimensional and defines the set of admissible values of N, in $[0, N_{max}]$
- In the figure, the admission region is highlited in greed and, in order to meet the SLA, the red region should be avoided, that is, requests that would lead the system into the red region should be blocked
- The amplitude of the admission region depends on the link's capacity and on the statistical properties of traffic

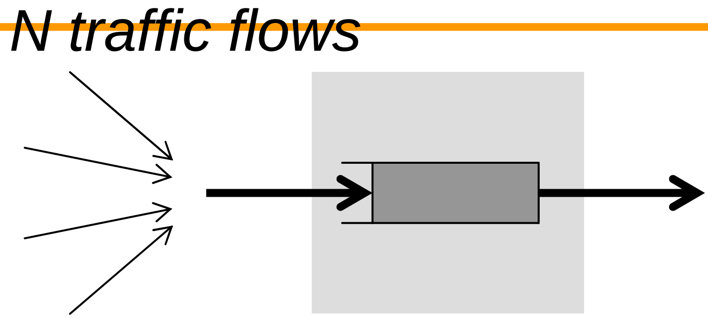

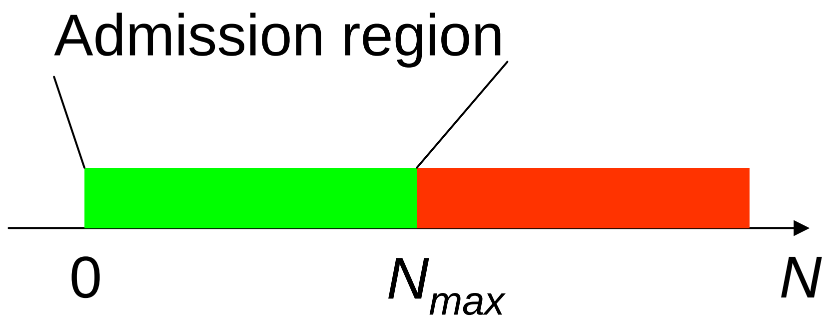

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Back to admission control

- For a two-class scheduler, the admission region is clearly bi-dimensional
- The admission region is the set of all pairs (N_1, N_2) such that the SLA (d_1, p_1) of class-1 flows and the SLA (d_2, p_2) of class-2 flows is fulfilled
- The admission region (in green in the figure) is delimited by the admission line $N_2(N_1)$
- The admission line is the curve representing the maximum number of class-2 flows that can be admitted, given N_1 admitted flows in class 1

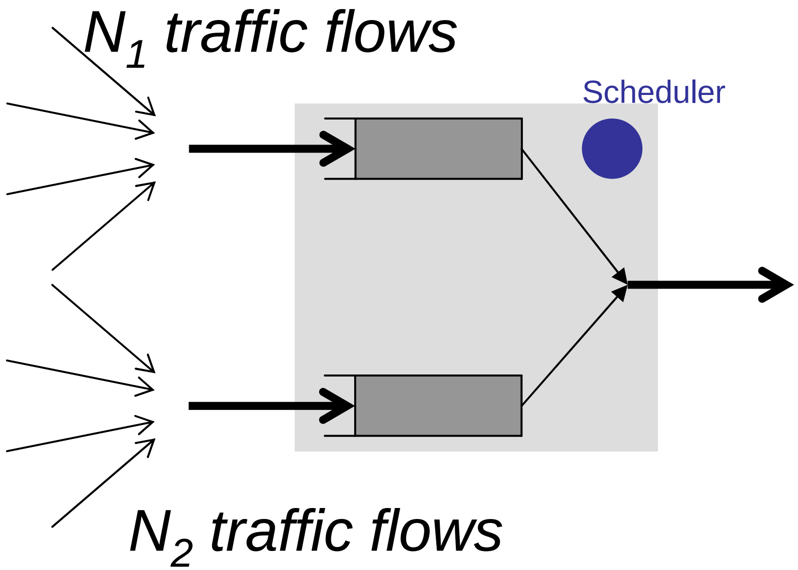

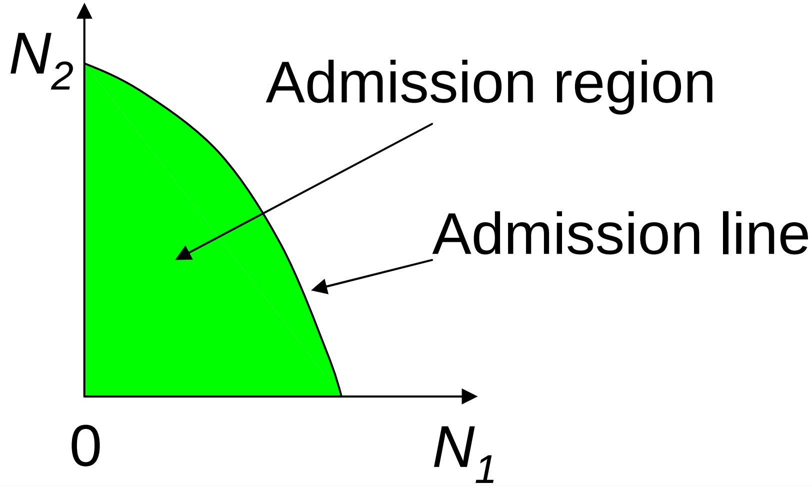

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Back to admission control

- In a three-class scheduler, the admission region is three-dimensional
- There are N_1, N_2 and N_3 flows in class 1, 2 and 3, respectively
- The SLAs of the three classes are (d_1, p_1), (d_2, p_2) and (d_3, p_3)
- The admission region is the set of triplets (N_1, N_2, N_3) such that the three SLAs are simultaneously met
- The extension to more than three classes is intuitive, even if calculations to determined the shape of the admission region are increasingly complex

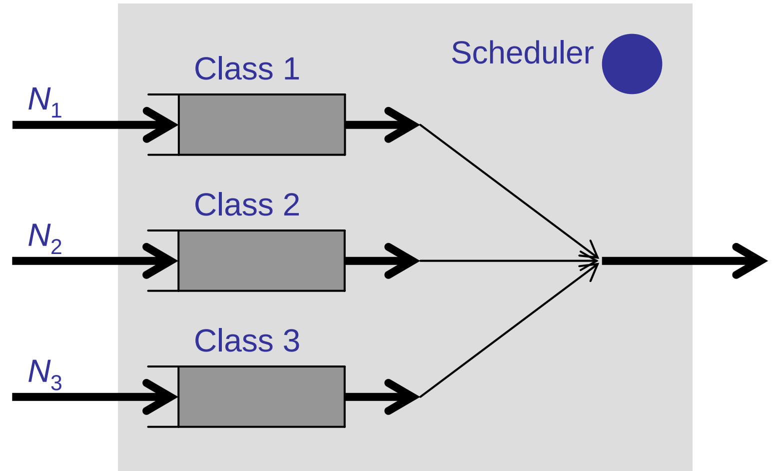

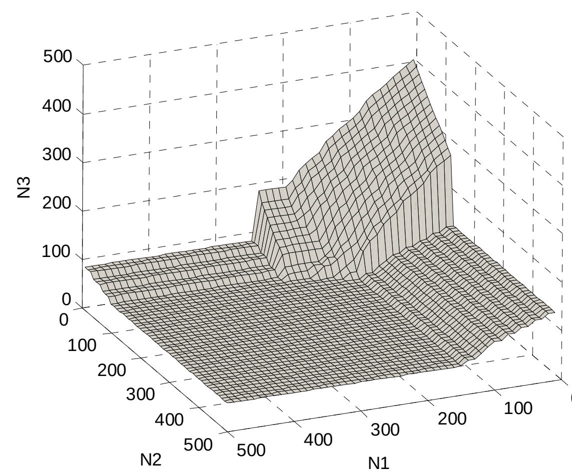

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Packet classification

- In network architectures providing differentiated QoS, packet classification is a basic step to deliver differentiated service level agreements
- In the IP Differentiated Services Architectures, these different service levels are referred to as Per Hop Behaviors
- Nodes treat differently packets, according to their Per Hop Behavior (PHB)
- Clearly, nodes must be able to determine the PHB to which a packet belongs, otherwise, the associated SLA cannot be implemented
- Packets carry their PHB code in a specific field
- However, this specific field must be filled at the network ingress by an edge node
- This operation is referred to as **classification**

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Packet classification

- Classification is based on keys and rules
- The classification key, in general, is a set of *N* bits in a packet header, concentrated in one field or distributed in multiple fields
- With a key composed of *N* bits, 2^N different classes of packets can be differentiated
- The act of classifying a packet involves matching the classification key against a set of classification rules
- The result of classification is to assign a treatment behavior to each packet
- Clearly, all packets within the same treatment behavior receive the same service by a node

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Packet classification

- The simplest classification schemes are called single-field: they take into account a single field of the packet header as classification key
- Multi-field classification considers multiple fileds of the packet header
- Multi-field classification is more complex than single-field classification and problems of computing capacity may arise, as classification must operate at wire speed
- However, multi-filed classification is more powerful than single-field, as it provides a greater amount of context to the router's subsequent stages

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Packet classification

- Typically, single-field classification uses the Type of Service field of the packet header
- However, this field is currently being replaced by the Diferentiated Services field
- The field is the same byte in the packet header
- However, the allocation of bits and their semantics is different
- Currently, six bits of the Differentiated services field are used to code the packet's class, two bits are reserved
- Multifield classification, by considering other fields such as Protocol, and transport protocol ports, can identify the application to which the packet belongs

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Packet classification

- The problem with multi-field classification is that some fields are long, for example, the IP addresses
- For example, a flow is univocally identified by the fields:
  - Source Address
  - Destination Address
  - Source Port
  - Destination Port
  - Protocol
- Which sum up to 104 bits (a large number!)
- A significant computation is required to examine these fields

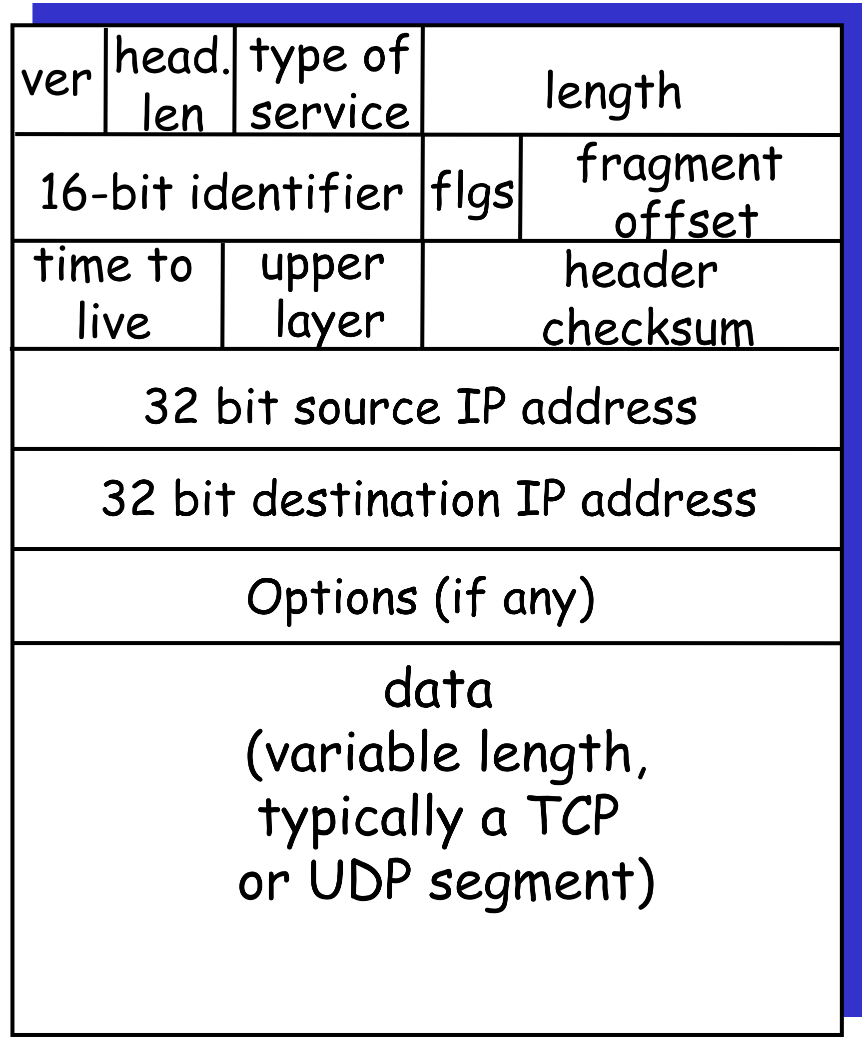

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Packet classification

- With IPv6, the problem is even more serious, as addresses are 128 bit long
- An address + port numbers classification requires 288 bits
- Moreover, the IPv6 extension headers may require a considerable amount of parsing before finding the transport port numbers
- However, the flow label field (20 bits) can be useful to allow a simple single-field classification

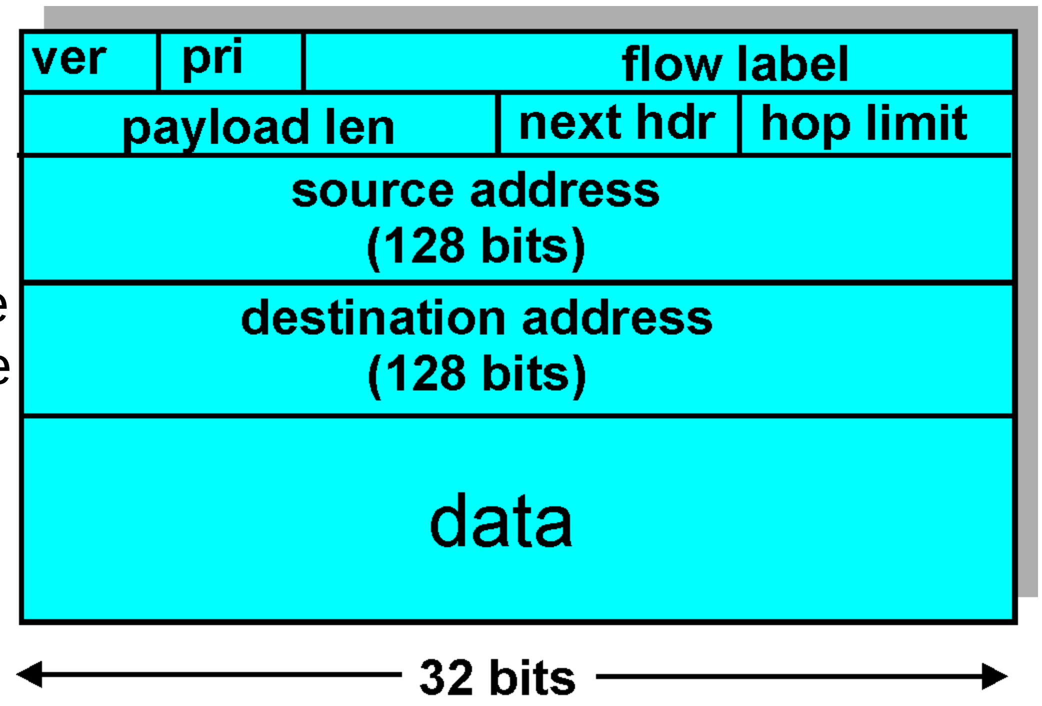

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Classification rules

- A MF classifier must be able to deal with having multiple rules that match the same packet
- Matching can be
  - Exact: a field must be equal to a specific value, for example, the source IP source address must be equal to 192.10.10.1
  - Ranging: a field must be in a specified interval of values, for example, the IP source address must be in the range [192.10.10.1, 192.10.10.254]
- Rules can be composed with logical operators:
  - The IP source address must be in the range [192.10.10.1, 192.10.10.254] **AND** the protocol must be TCP **AND** the destination port must be 25
- Any logical expression composing field values can be instantiated
- However, complex rules consume a large amount of CPU

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Classification rules

- The output of classification is the class to which the packet belongs
- Usually the packet's class is written into the Differentiated services field of the IP packet
- For example, a possible service class is the traditional Best-Effort
- This class, in the Differentiated services architecture, has code field 000000 in the Differentiated services field of the IP packet
- For example, the network administrator might assign to all HTTP packets coming from a given subnet the Best-Effort service class
- In this case, the rule would be
  - **IF** SA in [192.10.10.1, 192.10.10.254] **AND** protocol = TCP **and** DP = 80 **THEN** class = Best-Effort

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Line rate processing

- Each input port of a node has a classifier examining incoming packets
- If a packet arrives when another packet is being classified, it is stored in a FIFO queue
- Thus, a backlog of packets to be classified builds up in this queue
- Packets accumulate delay while waiting for classification
- For packets needing a stringent delay requirement, this delay may disrupt the SLA

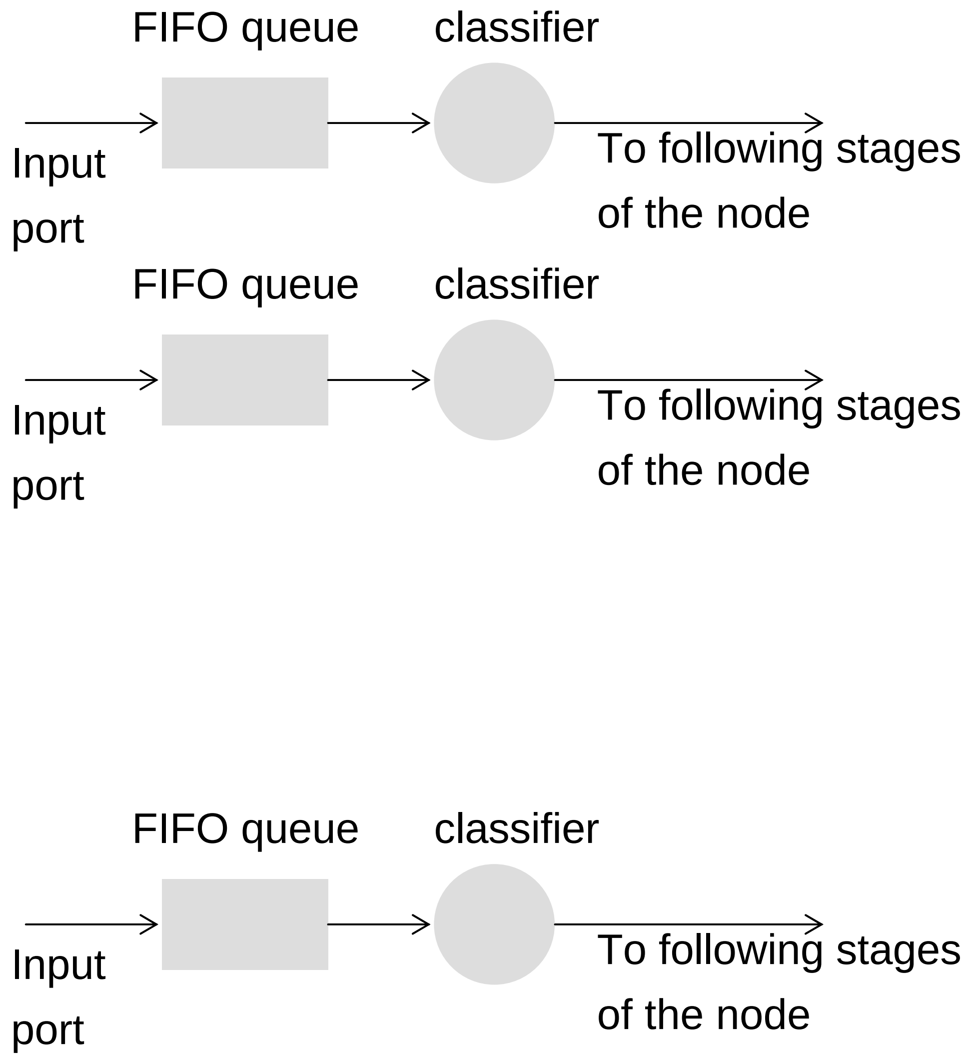

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Line rate processing

- Moreover, it is not possible to give precedence to packets with stringent delay requirement
- This is because the service requirements of unclassified packets are clearly unknown
- Thus, the only way to make classification transparent for delay-critical packets is to ensure that the classifier has enough CPU to deal fastly with the incoming stream of packets, in the worst case hypothesis

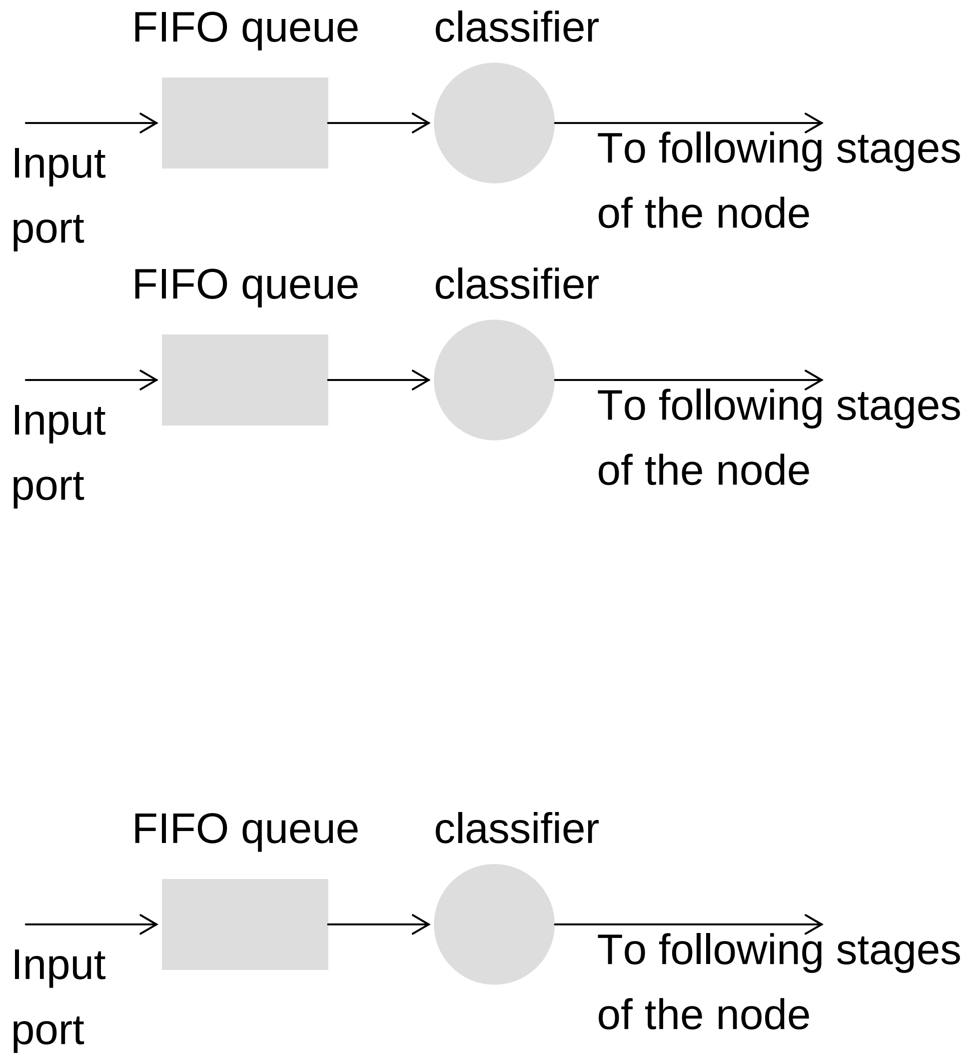

Copyright © 2009 paolo Giacomazzi. All rights reserved.

Line rate processing

- The worst case hypothesis is a stream of short packets at wire speed with long multi-field classification rules, carried out on long fields
- Classification is a complex function, with high computational requirements, and the hardware issues are rather complex

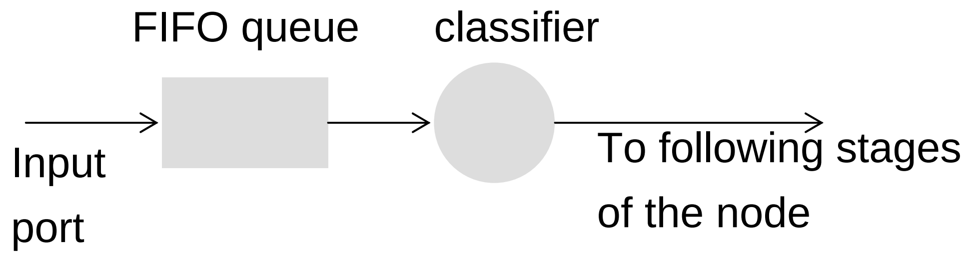

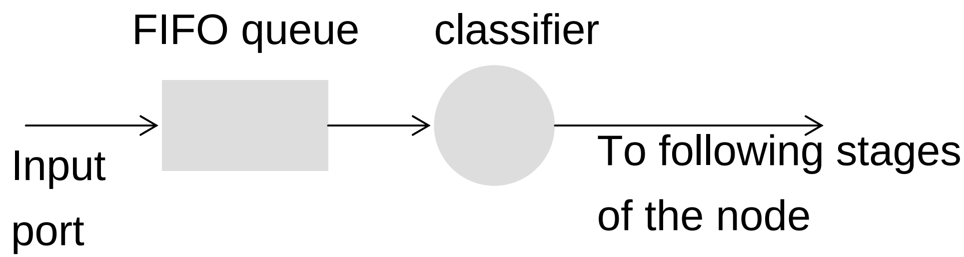

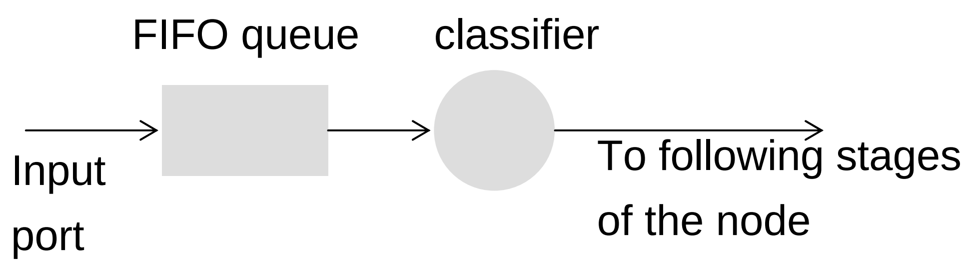

Copyright © 2009 paolo Giacomazzi. All rights reserved.