University of Uppsala

Department of Computer Systems (DoCS)

Final Examination

November 3, 2000

IT3 (Datorsystem II : (Del 1) Datakomunikation)

Data Communication and Networks

9:00 - 14:00
Polacksbacken

INSTRUCTIONS TO CANDIDATES

- This is a FIVE (5) hour examination
- Answer all questions in English or Swedish
- Dictionaries are Permitted
- The exam has 10 questions on 5 pages
- Marks total 85, 55 points are required for a pass grade.

Question 1

Use the weighted network graph shown in the following diagram to answer the following questions related to network layer routing protocols for heterogenous networks. Assume that each point in the graph represents a local sub-network and that an heirarchical routing approach is being used.

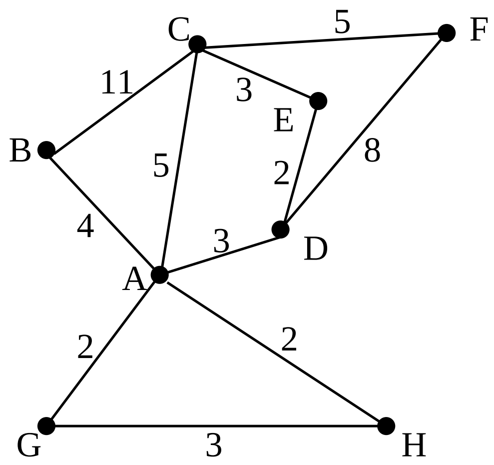

a. Assume that a distance vector routing protocol is being used to maintain the routing table information and that the network graph shown above has been in a stable configuration and running for some time.

Describe the process that B will follow to compute its routing vector when the next routing table update takes place. What data will it receive? Show the data that will be received and how that data is used to compute a new routing vector at node B.

[6]

b. If node A fails in the network shown, what is the effect on the routing information stored at B over the next few routing table update cycles? Show the changes for 3 update cycles after the failure of A. Comment on what is happening, is there a problem?

[4]

c. Describe the basic algorithm used by a link state routing protocol to discover routing information. How is this protocol different from that used in distance vector routing? Discuss why link state routing (OSPF) replaced distance vector routing (RIP) as the popular Internet router table maintainance algorthm.

[4]

SOLUTION

a. *The network begins in a stable state, and that means that we can assume that all the routing vectors have converged on the true state of the network. The routing vectors received at B come from A and C, and the following tables show how B uses this information to create its new vector.*

Routing Vector at B

Dst	Cost	Next
A	4	A
B	-	-
C	9	A
D	7	A
E	9	A
F	14	A
G	6	A
H	6	A

Routing Vector A

Cost
-
4
5
3
5
10
2
2

Routing Vector C

Cost
5
11
-
5
3
5
7
7

Cost= min(cost to A + cost A to Dst, cost to C + cost C to Dst)

Next= winning cost node in above equation

b. *When A fails the nodes G and H form a sub-net that is disconnected from the rest of the network, but this will not be discovered immediately since the routing vectors stored in the network will count up to a high cost for reaching G and H relatively slowly. We can see this in the vector stored at B, which changes in response to the data received from the other nodes in the manner shown in the following figure. To produce the figure below I have assumed synchronous information updates and that the information in the tables is that obtained from only C in each cycle.*

c. *Link state routing collects information about its transmission metric to all its immediate neighbours and stores that data into a packet that it broadcasts globally to all other routers in the network. Upon receiving these state packets the other routers calculate their best paths to all the reachable destinations. Thus link state router table construction has access to global information about the state of the graph.*

In distance vector routing a router collects routing knowledge from its neighbours, but has no idea about the actual reachability of any destination. Since data collection is local, graph loops can create a situation where each node tells others that it can reach a disconnected host (at a slightly increased cost).

Link state routing's global data collection policy solves the count to infinity problem that occurs when using distance vector routing to compute the tables using only transit cost vector transmission. The weakness with the distance vector approach is that it has no knowledge of the underlying structure of the routes that are being used to reach the hosts. Link state routing explicity provides reachability information about all the hosts and this solves the problem of not understanding the graph topology when creating routing vectors.

Question 2

Name two data transmission services that are provided by the TCP/IP transport layer for the use of "Applications"? Why are both types of service provided? Describe situations in which an application might choose to use one or other of these services.

[6]

SOLUTION

Two services are UDP and TCP. The first provides a connectionless datagram service, the other a reliable connection based service. Many services prefer not to have to deal with the mechanics of maintaining a connection, and such applications can take advantage of the reliable connection service TCP provides. In generating reliable service model TCP is a complex and expensive protocol. This means that the header overhead is greater per packet, and that the cost associated with providing high reliability is incurred irrespective of the needs of the application.

Many applications that do not need to have a highly reliable connection (many client server applications) elect to replace TCP with a UPD based protocol, and implement the subset of protocol features that they require as a part of the application functionality.

Even if an application uses a high level protocol such as TCP it cannot be completely insulated from protocol failure. Not all data corruption will be detected, and not all data will necessarily arrive. Thus, applications need to have some fallback position and knowledge of failure modes in order to remain robust even when using reliable underlying transport protocols such as TCP.

Question 3

Discuss the reasons for using a stack based approach to designing and implementing communication protocols. What is gained by using the "stack" approach? Name and discuss three (3) key objectives addressed in the design and specification of the OSI 7-layer protocol model?

Contrast the objectives of the OSI design committee with the objectives of the designers of TCP/IP. How are the objectives similar and how are they different?

[7]

SOLUTION

The stack based approach limits the complexity of the protocol by limiting the number of concerns that are dealt with at each level. Consequently, high level applications can define powerful abstract protocols with which to communicate, and rely on a few lower level services which are used to implement the more abstract facilities provided in the higer layer.

Three key aspects of the design were,

- *Separation of function to provide layers with well defined interfaces, and minimal data transfer over those interfaces.*
- *To provide an abstract model that could be used to specify and design future communication protocols. to also provide an international foundation for such designs that was agreed on by the ISO as an appropriate standard.*
- *Represent the emerging development of protocols, and to provide a standard within which protocols can be compared and classified.*

The OSI protocol committee set out to develop a generic and comprehensive specification about how protocols for data communication should be designed and implemented. They divided the areas of concern for protocols and attempted to group these concerns into activity levels. This attempt was largely successful, however there are limitations, since the resulting 7 layers include two (presentation and session) that have limited function and are still almost empty in terms of protocols and services.

TCP/IP is a description and formalization of the structures and protocols in the most widely used internetworking suit of protocols so far in use. The TCP/IP set of protocols and their related facilities are used throughout the Internet, and represent the evolution of protocols and services first proposed for the initial ARPA-net in the 1960's.

Question 4

What do you think the term "Quality of Service" (QoS) means in terms of a network connection. Describe 4 factors or metrics that might be of interest when negotiating QoS for a connection?

Frequently QoS parameters vary depending on the type of traffic flow over the connection. Write a paragraph discussing the factors that are important for three different types of traffic with very different requirements. You should say what factors are important and what factors are not important, and give reasons why you believe that factors are important or not important.

[5]

SOLUTION

Quality of Service (QoS) is the ability of a network infrastructure to grant, and guarantee resources to applications. If this can be implemented in a nice way, that opens up the possibility to share the network beween traffic in many ways, and to trade off one traffic flow against another.

It also provides mechanisms for deciding whether to allow applications access to the network. In this context we can use the knowledge of the current utilization of resources, and the requested QoS to decide whether the application requesting a network service can be

- *given what it needs to function in a reasonable way.*
- *what impact granting a new traffic request will have on other users*
- *how to handle traffic when resources are momentarily exceeded by demand for those resources.*
- *decide what "degradations" of quality are acceptable to applications whose requests for resources can no longer be sustained by the network.*

I will describe the requirements for the following three traffic types.

Streaming video: *requires constant bandwidth and rate, low variation in delay, and a minimum bandwidth allocation in order to provide a reasonable service to the customer. This type of application is sensitive to timely arrival of data, and is not sensitive to some loss on the line, or some corruption in packets, so long as the minimum rate and maximum corruption level requirements are met.*

FTP: *might require zero packet loss on the end to end connnection and a guaranteed average bandwidth over the life of the connection, in order to ensure that all the data is delivered to the destination in a reasonable time. FTP does not require constant rate, or care about resent packets, or internal loss rates, so long as all the data is delivered eventually.*

telnet: *might require low end to end delay, and zero less on the connection, however it is not sensitive to constant rate, or variation in delay that might not be noticed by a person.*

Question 5

Describe multicast. What is a multicast communication?

Using the following picture of a network show using two versions of this picture how a message is distributed from the SRC to the destination set D1..D4 using broadcast and multicast.

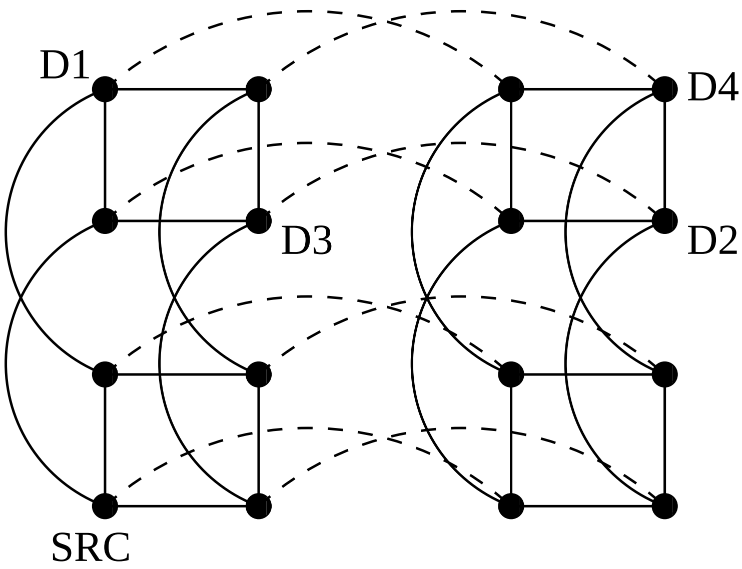

Using these pictures to motivate your answer comment on why using multicast can help to improve network performance?

[8]

SOLUTION

Multicast is a method for delivering a single packet to multiple destinations. The set of destinations can include all nodes other than the source, or a subset of the nodes in the network. Using the diagram one can demonstrate the links that need to be covered to deliver a packet to the final destinations. In the following illustrations it can be seen that multi-cast employs a single packet which is split at key routing points to deliver the data packet to all destinations. This is achieved using a spanning tree, and results in a considerable reduction in network resource utilisation.

Flood based broadcast on the other hand, addresses independent packets to each possible destination. This results in clogging on key paths as traffic waits to traverse the link, or in the use of a greater number of network paths than are required. In addition, due to the queueing delays, broadcast traffic will generally take longer to arrive at all the required destinations.

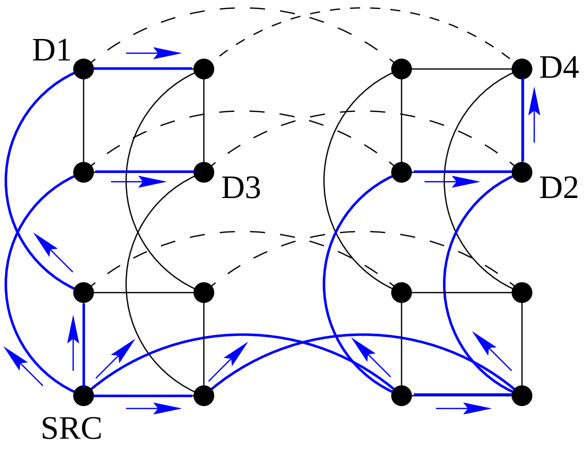

4 pkts delivered using 13 links in 4 cycles

a: Broadcast

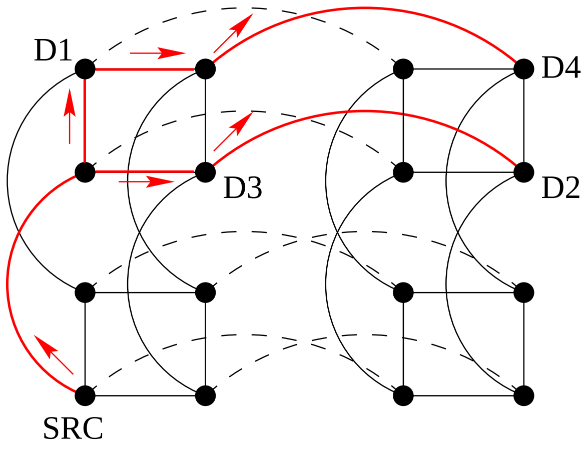

4 pkts delivered using 6 links in 4 cycles

b: Multi-cast

Question 6

What is "network capacity"? Define this quantity using a simple formula. Write a paragraph that explains why the capacity of a network is relevant to the design of sliding window protocols.

[6]

SOLUTION

Network capacity is defined by the following equation, and relates to the quantity of data that can be inserted into a link/connection before an ack is received.

$$Network_capacity = Delay_{end-to-end} * Bandwidth$$

The capacity of a connection defines the number of bytes that the sender can insert into a network before the first ack from the receiver is expected. This quantity should be taken into account when selecting a window size because the window used for "sliding window" should be able to ensure that buffering is available to reconstruct the order of all bytes of data currently in transit, irrespective of the order in which they arrive. Sliding window does this by keeping a record in both the sending and receiving sides of the connection of which data is in transit and what data is received and to be sent.

A firewall can be constructed using two routers, one on either side of an application gateway. Draw a picture that shows the configuration of a firewall and describe how each of the routers and the application gateway function to protect a local network, and the hosts on that network.

SOLUTION

A classic firewall depiction is found in the following figure.

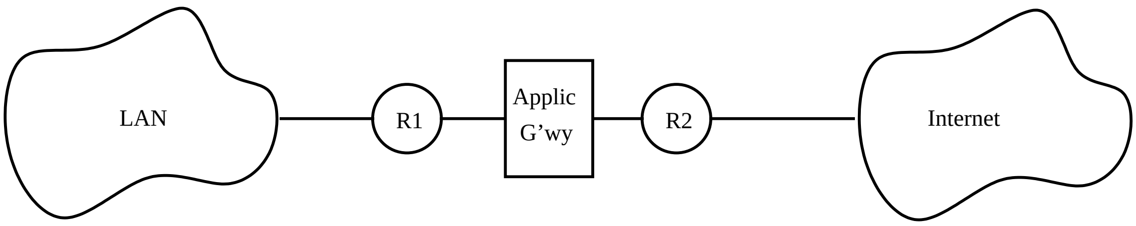

A firewall acts to deny access to traffic from outside a local network access to protected resources in the local environment. In general terms a firewall can be considered to be composed of an application gateway (G) and two routers (R1, and R2). The application gateway can be engaged in keyword monitoring of traffic flowing from and to the local network. It can also examine file types being transfered to look for types of traffic that are considered a security risk.

The function of router inspection of traffic is to look at the flows of packets from inside the network to the rest of the networking world (R1) and from the outside world into the local environment (R2). Incomming traffic can be inspected for packet type, and excluded (often UDP traffic is excluded since it is harder to verify the source and purpose of UDP traffic), or filtered on the basis of allowing connections to well defined services on defined hosts and port numbers.

[2]

Name two factors that can compromise firewall security and indicate how they affect the security of the internal network.

SOLUTION

Two factors that can have an influence on security of firewalls are

wireless: *since traffic can be snooped using a simple receiver and the traffic is freely available (doesn't have to pass through the firewall to be received). Solutions include frequency shifting and encryption of traffic, as well as validation of the mobile devices.*

dialup: *external available dialup often gives access to the internal network, and people do not think of this as being a vulnerability. However, modem numbers are quickly distributed, and remain in the hands of ex-employees (who may be motivated to make attacks).*

[4]

Question 7

Explain the problems with estimating RTT for networks. What is this value used for? Give an example of a situation where standard techniques for calculating RTT estimates are inaccurate and explain why.

[5]

SOLUTION

RTT is the so called "round trip time" that predicts the time taken for an ack to be received for a given data packet/frame. The RTT value associated with a network connection is dynamic and depends on the load in the network at the time, and on hardware and environmental factors that may affect the transmission medium (eg. wireless during a thunderstorm).

*The RTT value is important in computer networking because it lets protocol elements set timers which are used to estimate if a packet has been discarded or lost in the network. In TCP connections for example when an ACK is not received in 2*RTT it is usually considered dropped or lost, and will be retransmitted.*

Calculating RTT is a difficult proposition, since the estimation of RTT is generated by starting timers associated with sequence numbers and then recording the delay to the ACK. Each instantaneous RTT estimate is denoted RTT_sample and is integrated into a weighted average of RTT called RTT_est using a variety of formulae. This process however is prone to error in the situation where a large percentage of packets time out and are retransmitted. If timers are reset on retransmission it is not possible to distinguish between the RTT_est's that are real, and those that are too small due to a re-transmission (see the diagram).

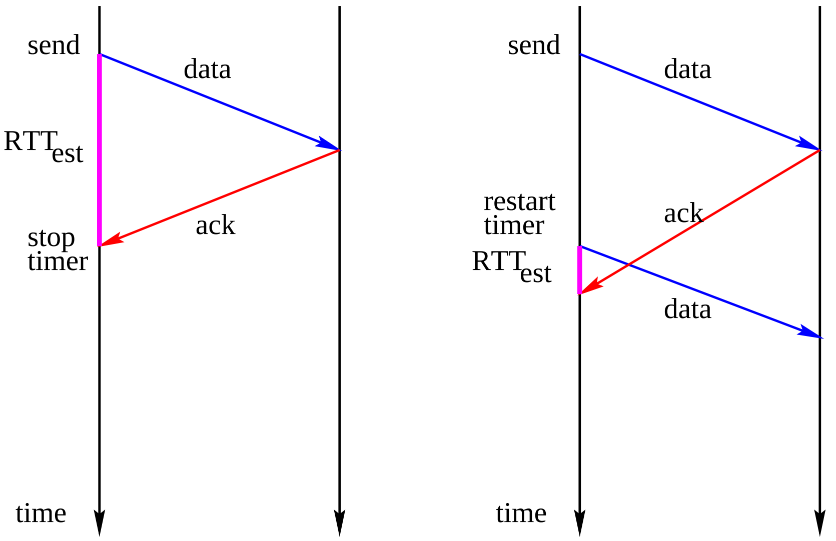

Correct RTT calculation

Incorrect RTT calculation

Question 8

Describe the idea of a sliding window protocol for frame transmission? Explain using diagrams what the sender and receiver window represent. In this context describe the operation of the "stop and wait" protocol. How does this protocol work, and how does it recover from data loss on the link?

[5]

SOLUTION

Sliding window keeps track of the progress of packets over the network between the sender and receiver using a buffer space and a record of what is assumed to be in transit. Communication between the sending and receiving agents is used to synchronise the advance of the window over the data stream, permitting the sender to send data only when the receiver (or the network) is capable of supporting the transmission.

Buffer allocations and sequence numbers are used to store the packets that arrive out of order, and to request copies of packets that do not arrive. This allows the receiver to assemble a complete byte stream at the receiving end of the connection. The receiver also informs the sender (via ack messages) that data has been received, and the availability of buffer space or network capacity allowing the sender to insert more traffic into the network.

Stop and Wait is a simple protocol that uses single bit sequence number to preserve order and to distinguish between new and resent packets. Since Stop and Wait sets the window size to a single packet, this means that the sender sends a packet, and waits for the ACK before it can send another. This results in sending another data packet which will be either a retransmission, if the packet was lost, or the next data packet in the byte stream.

Describe a situation when "stop and wait" results in very bad utilization of network resources. What can be done to improve performance? What extensions to the protocol are required to implement the solution you propose?

[5]

SOLUTION

This results in poor utilization of the link when we have high bandwidth and long delay, since it would be possible to insert many packets onto the network in the time that it takes to receive the ACK.

Consider a situation where the delay is 8ms and the link capacity is 5 packets, then we can insert a packet into the network every 1ms (up to a total of 5 packets) before an ACK is received. In this situation stop and wait obtains 20% utilisation on the link, leading to underutilisation of resources.

To address this situation we can extend Stop and Wait to become a windowed protocol, which in turn requires us to add some things to the protocol definition in order for it to operate correctly. Sequence numbers need to be provided that are appropriate to the maximum window size being considered, as well as methods for nogotiating window buffer allocation in the sender and receiver and allocating and managing buffers via acknowledgements and packet send operations.

Question 9

Wireless communication is becoming increasingly popular for local area networks. Discuss in a few paragraphs the practical and technical implications for

- security,
- wide area connectivity (coverage),
- traffic routing to hosts that are moving between cells in a wireless communication environment.

[9]

SOLUTION

security: *Security in wireless environments is hard to enforce, especially if there is a requirement for high mobility. In a local environment one can rely on encryption and on access control on the basis of known devices with known ID's (such as MAC addresses). However, in a wider mobile context security and trust building, especially in ad hoc and active networks is a difficult problem to solve.*

connectivity: *Providing connectivity is a two edged sword for mobile internet over wireless. Wireless access requires that base stations be installed with required coverage, and wider coverage at lower bandwidth and higher delay can be achieved using satellites. However, despite the advantages of wireless in terms of cost involved in laying cables and providing wired infrastructure, there are still large costs involved in providing wireless coverage of remote and low population density areas. There are also environmental and health concerns that may limit the development and deployment of wireless communication in the future.*

routing: *Routing for wireless mobile computing devices is likely to be managed in a similar manner to that used to handle current mobile telephone traffic. The existing cellular network will doubtless be enhanced and features that permit mobile IP devices will be added. These may be based on the existing* foreign and *local agent proposals, or on some other techniques adapted from cellular telephony.*

Question 10

Comment on the use of classes for network routing in the Internet protocol. What classes of networks are defined and how are they allocated. What problems are emerging in the internet as a result of this approach to routing.

What does "sub-netting" do to assist with the issues you have identified in the first part of this question? Illustrate the idea of sub-netting using a diagram and showing IP numbers and other necessary information.

[9]

SOLUTION

This is a basic outline of the solution.

Basic IP address structure is based on the association of an IP class address with each existing physical network. This was intended to match an hierarchical model in which large service providers and organisations provided services to a multitude of smaller private and corporate networks. In this model three classes of network address were proposed, namely A (providing 2^21 *hosts numbers), B (providing* 2^14 *host numbers) and C (providing* 2^8 *host numbers).*

Due to a number of factors the A class has proven unpopular, and too large for realistic corporations. The C class addresses are unacceptably limiting for most applicants, since the possibility of a maximim netowrk size of 255 hosts provides for very limited expansion. This has lead to early depletion of the B class addresses, and significant address waste. It appears that many B class networks contain between 2000 and 3000 hosts. This problem is exacerbated by the association between network class and physical network being envisaged as 1 to 1. There are few single physical nets with 64K hosts.

Two methods are currently provided to resolve this situation. The first is CIDR, which allows sets of C class networks to be aggregated.