Internal Assessment Test 2 – Nov. 2017

Sub:	**Computer Networks**					Sub Code:	**15CS52**	Branch:	ISE
Date:	07/10/17	Duration:	90 min's	Max Marks:	50	Sem / Sec:	V / A,B		OBE

<u>Answer any FIVE FULL Questions</u>

		MARKS	CO	RBT
1 (a)	Label router architecture and explain the elements.	[5+5]	CO4	L1 & L4
2 (a)	Illustrate Distance Vector Routing with an example.	[4+6]	CO4	L3
3 (a)	Compare direct routing and indirect routing.	[10]	CO5	L4
4 (a)	Label 3G architecture and explain the elements.	[5+5]	CO5	L1 & L4
5 (a)	Explain Mobile IP in detail.	[10]	CO5	L4
6 (a)	Summarize broadcast routing.	[10]	CO4	L2
7 (a)	Demonstrate how to manage mobility in cellular networks.	[10]	CO5	L3

1 a. Label router architecture and explain the elements.

Router architecture (5)

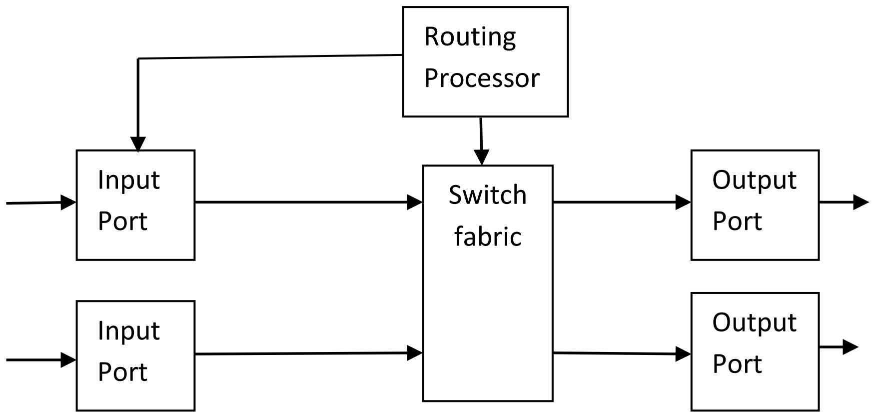

Components: (5)

- Input ports. An input port performs several key functions. It performs the physical layer function of terminating an incoming physical link at a router; this is shown in the leftmost box of the input port and the rightmost box of the output port in Figure. An input port also performs link-layer functions needed to interoperate with the link layer at the other side of the incoming link; this is represented by the middle boxes in the input and output ports. Perhaps most crucially, the lookup function is also performed at the input port; this will occur in the rightmost box of the input port. It is here that the forwarding table is consulted to determine the router output port to which an arriving packet will be forwarded via the switching fabric. Control packets (for example, packets carrying routing protocol information) are forwarded from an input port to the routing processor.
- Switching fabric: The switching fabric connects the router's input ports to its output ports. This switching fabric is completely contained within the router— a network inside of a network router.
- Output ports: Output port stores packets received from the switching fabric and transmit these packets on the outgoing link by performing the necessary link-layer and physical-layer functions. When a link is bidirectional (that is, carries traffic in both directions), an output port will typically be paired with the input port for that link on the same line card (a printed circuit board containing one or more input ports, which is connected to the switching fabric).
- Routing processor: The routing processor executes the routing protocols maintains routing tables and attached link state information, and computes the forwarding table for the router.

A router's input ports, output ports, and switching fabric together implement the forwarding function and are almost always implemented in hardware. These forwarding functions are sometimes collectively referred to as the router forwarding plane. Forwarding plane hardware can be implemented either using a router vendor's own hardware designs, or constructed using purchased merchant-silicon chips (e.g., as sold by companies such as Intel and Broadcom). While the forwarding plane operates at the nanosecond time scale, a router's control functions—executing the routing protocols, responding to attached links.

2 (a) Illustrate Distance Vector Routing with an example.

Whereas the LS algorithm is an algorithm using global information, the distance vector (DV) algorithm is iterative, asynchronous, and distributed. It is distributed in that each node receives some information from one or more of its directly attached neighbours, performs a calculation, and then distributes the results of its calculation back to its neighbours. It is iterative in that this process continues on until no more information is exchanged between neighbours. The algorithm is asynchronous in that it does not require all of the nodes to operate in lockstep with each other. Let $d_x(y)$ be the cost of the least-cost path from node x to node y. Then the least costs are related by the celebrated Bellman-Ford equation, namely,

$$d_x(y) = \min_v\{c(x,v) + d_v(y)\}, \quad (4.1)$$

ROUTING ALGORITHM (2)

```
Initialization:
for all destinations y in N:
D x(y) = c(x,y) /* if y is not a neighbor then c(x,y) = ∞ */
for each neighbor w
D w(y) = ? for all destinations y in N
 for each neighbor w
send distance vector Dx = [Dx(y): y in N] to w
loop
wait (until I see a link cost change to some neighbor w or
until I receive a distance vector from some neighbor w)
for each y in N:
Dx(y) = minv{c(x,v) + Dv(y)}
if Dx(y) changed for any destination y
send distance vector Dx = [Dx(y): y in N] to all neighbors
forever
```

Example: (8)

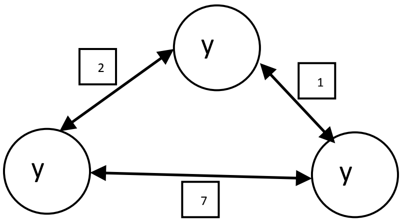

Node x table

Iteration 1:

	x	y	z
x	0	2	7
y	∞	∞	∞
z	∞	∞	∞

Iteration 2:

	x	y	z
x	0	2	3
y	2	0	1
z	7	1	0

Iteration 3:

	x	y	z
x	0	2	3
y	2	0	1
z	3	1	0

Node y table

Iteration 1:

	x	y	z
x	∞	∞	∞
y	2	0	1
z	∞	∞	∞

Iteration 2:

	x	y	z
x	0	2	7
y	2	0	1
z	7	1	0

Iteration 3:

	x	y	z

x	0	2	3
y	2	0	1
z	3	1	0

Node z table

Iteration 1:

	x	y	z
x	∞	∞	∞
y	∞	∞	∞
z	7	1	0

Iteration 2:

	x	y	z
x	0	2	7
y	2	0	1
z	3	1	0

Iteration 3:

	x	y	z
x	0	2	3
y	2	0	1
z	3	1	0

For example, node x computes

Dx(x) = 0 Dx(y) = min{c(x,y) + Dy(y), c(x,z) + Dz(y)} = min{2 + 0, 7 + 1} = 2 Dx(z) = min{c(x,y) + Dy(z), c(x,z) + Dz(z)} = min{2 + 1, 7 + 0} = 3

3 (a) Compare direct routing and indirect routing.

Indirect Routing (5)

Indirect Routing to a Mobile Node Let's first consider a correspondent that wants to send a datagram to a mobile node. In the indirect routing approach, the correspondent simply addresses the datagram to the mobile node's permanent address and sends the datagram into the network, blissfully unaware of whether the mobile node is resident in its home network or is visiting a foreign network; mobility is thus completely transparent to the correspondent. Such datagrams are first routed, as usual, to the mobile node's home network. The home agent have to lookout for arriving datagrams addressed to nodes whose home network is that of the home agent but that are currently resident in a foreign network. The home agent intercepts these datagrams and then forwards them to a mobile node in a two-step process. The datagram is first forwarded to the foreign agent, using the mobile node's COAand then forwarded from the foreign agent to the mobile node. The home agent will need to address the datagram using the mobile node's COA, so that the network layer will route the datagram to the foreign network. On the other hand, it is desirable to leave the correspondent's datagram intact, since the application receiving the datagram should be unaware that the datagram was forwarded via the home agent. Both goals can be satisfied by having the home agent encapsulate the correspondent's original complete datagram within a new (larger) datagram.

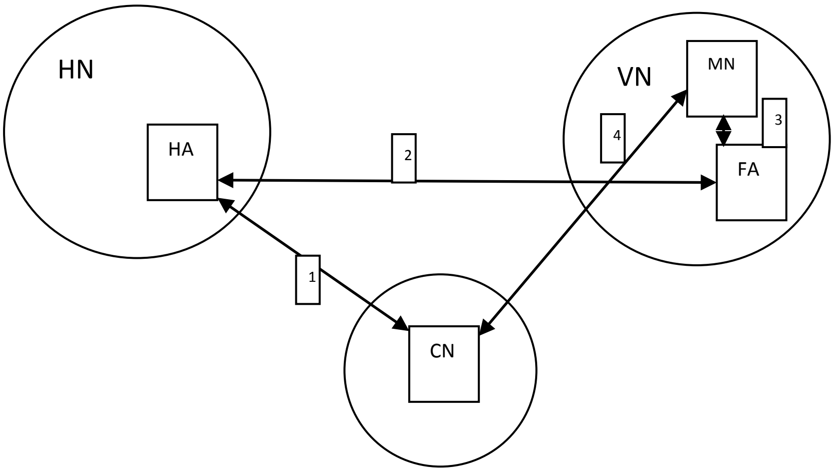

Indirect routing to a mobile node

HA-Home agent

HN-Home network: 128.119.40/24

VN-Visited network: 79.129.13/24

MN-Mobile node

Permanent address: 128.119.40.186 Foreign agent

Care-of address: 79.129.13.2

Figure 6.23 Indirect routing to a mobile node

The foreign agent, who "owns" the COA, will receive and decapsulate the datagram—that is, remove the correspondent's original datagram from within the larger encapsulating datagram and forward the original datagram to the mobile node.

Direct routing (5)

Direct Routing to a Mobile Node The indirect routing approach suffers from an inefficiency known as the triangle routing problem—datagrams addressed to the mobile node must be routed first to the home agent and then to the foreign network, even when a much more efficient route exists between the correspondent and the mobile node. In the worst case, imagine a mobile user who is visiting the foreign network of a colleague. The two are sitting side by side and exchanging data over the network. Datagrams from the correspondent (in this case the colleague of the visitor) are routed to the mobile user's home agent and then back again to the foreign network! Direct routing overcomes the inefficiency of triangle routing, but does so at the cost of additional complexity. In the direct routing approach, a correspondent agent in the correspondent's network first learns the COA of the mobile node. This can be done by having the correspondent agent query the home agent, assuming that (as in the case of indirect routing) the mobile node has an up-to-date value for its COA registered with its home agent. It is also possible for the correspondent itself to perform the function of the correspondent agent, just as a mobile node could perform the function of the foreign agent. The correspondent agent then tunnels datagrams directly to the mobile node's COA, in a manner analogous to the tunneling performed by the home agent.

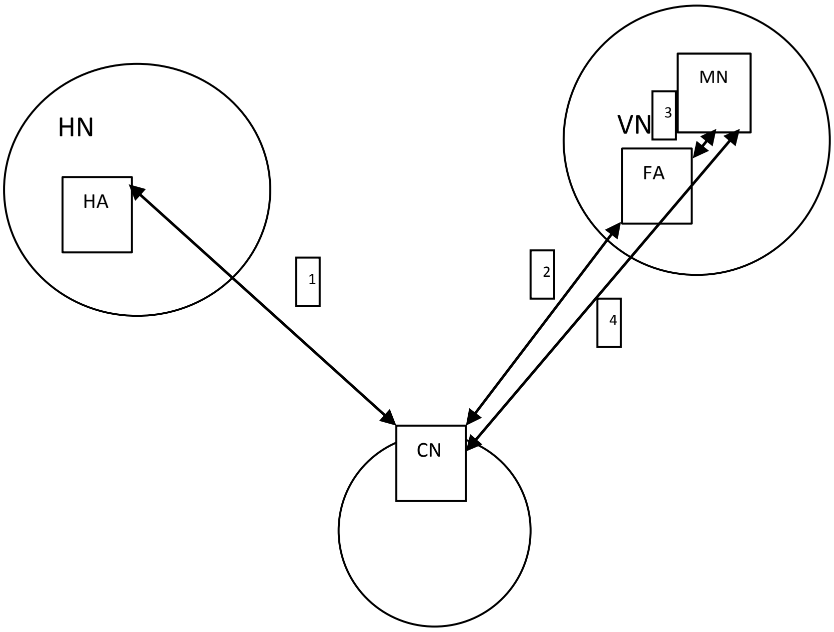

Direct routing to a mobile node

4 (a) Label 3G architecture and explain the elements.

Architecture (6)

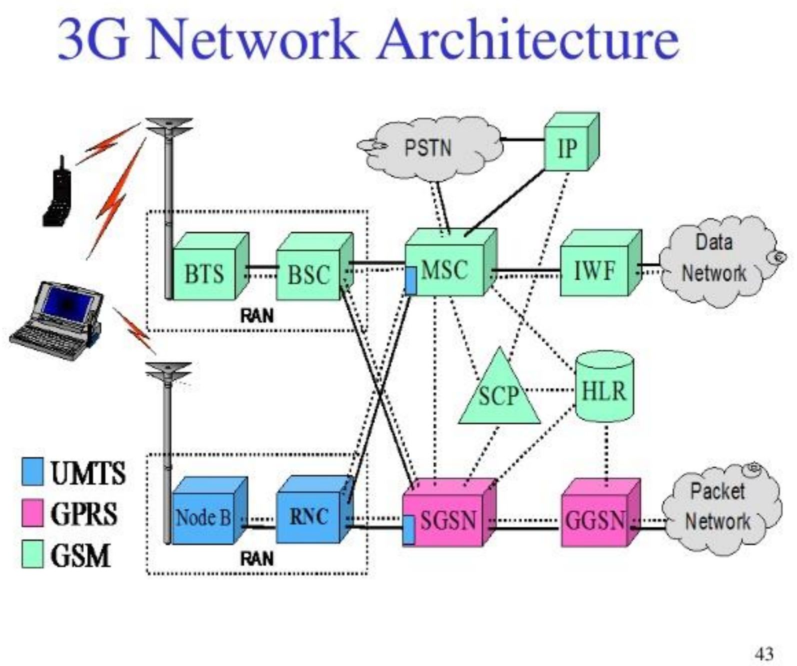

The UMTS 3G architecture is required to provide a greater level of performance to that of the original GSM network. However as many networks had migrated through the use of GPRS and EDGE, they already had the ability to carry data. Accordingly many of the elements required for the WCDMA / UMTS network architecture were seen as a migration. This considerably reduced the cost of implementing the UMTS network as many elements were in place or needed upgrading.

With one of the major aims of UMTS being to be able to carry data, the UMTS network architecture was designed to enable a considerable improvement in data performance over that provided for GSM.

3G UMTS network constituents (4)

The UMTS network architecture can be divided into three main elements:

1. ***User Equipment (UE):*** The User Equipment or UE is the name given to what was previous termed the mobile, or cellphone. The new name was chosen because the considerably greater functionality that the UE could have. It could also be anything between a mobile phone used for talking to a data terminal attached to a computer with no voice capability.
2. ***Radio Network Subsystem (RNS):*** The RNS also known as the UMTS Radio Access Network, UTRAN, is the equivalent of the previous Base Station Subsystem or BSS in GSM. It provides and manages the air interface for the overall network.
3. ***Core Network:*** The core network provides all the central processing and management for the system. It is the equivalent of the GSM Network Switching Subsystem or NSS.

The core network is then the overall entity that interfaces to external networks including the public phone network and other cellular telecommunications networks.

User Equipment, UE

The USER Equipment or UE is a major element of the overall 3G UMTS network architecture. It forms the final interface with the user. In view of the far greater number of applications and facilities that it can perform, the decision was made to call it a user equipment rather than a mobile. However it is essentially the handset (in the broadest terminology), although having access to much higher speed data communications, it can be much more versatile, containing many more applications. It consists of a variety of different elements including RF circuitry, processing, antenna, battery, etc.

There are a number of elements within the UE that can be described separately:

- ***UE RF circuitry:*** The RF areas handle all elements of the signal, both for the receiver and for the transmitter. One of the major challenges for the RF power amplifier was to reduce the power consumption. The form of modulation used for W-CDMA requires the use of a linear amplifier. These inherently take more current than non linear amplifiers which can be used for the form of modulation used on GSM. Accordingly to maintain battery life, measures were introduced into many of the designs to ensure the optimum efficiency.
- ***Baseband processing:*** The base-band signal processing consists mainly of digital circuitry. This is considerably more complicated than that used in phones for previous generations. Again this has been optimised to reduce the current consumption as far as possible.

- ***Battery:*** While current consumption has been minimised as far as possible within the circuitry of the phone, there has been an increase in current drain on the battery. With users expecting the same lifetime between charging batteries as experienced on the previous generation phones, this has necessitated the use of new and improved battery technology. Now Lithium Ion (Li-ion) batteries are used. These phones to remain small and relatively light while still retaining or even improving the overall life between charges.
- ***Universal Subscriber Identity Module, USIM:*** The UE also contains a SIM card, although in the case of UMTS it is termed a USIM (Universal Subscriber Identity Module). This is a more advanced version of the SIM card used in GSM and other systems, but embodies the same types of information. It contains the International Mobile Subscriber Identity number (IMSI) as well as the Mobile Station International ISDN Number (MSISDN). Other information that the USIM holds includes the preferred language to enable the correct language information to be displayed, especially when roaming, and a list of preferred and prohibited Public Land Mobile Networks (PLMN).

  The USIM also contains a short message storage area that allows messages to stay with the user even when the phone is changed. Similarly "phone book" numbers and call information of the numbers of incoming and outgoing calls are stored.

The UE can take a variety of forms, although the most common format is still a version of a "mobile phone" although having many data capabilities. Other broadband dongles are also being widely used.

3G UMTS Radio Network Subsystem

This is the section of the 3G UMTS / WCDMA network that interfaces to both the UE and the core network. The overall radio access network, i.e. collectively all the Radio Network Subsystem is known as the UTRAN UMTS Radio Access Network.

The radio network subsystem is also known as the UMTS Radio Access Network or UTRAN.

3G UMTS Core Network

The 3G UMTS core network architecture is a migration of that used for GSM with further elements overlaid to enable the additional functionality demanded by UMTS.

In view of the different ways in which data may be carried, the UMTS core network may be split into two different areas:

- ***Circuit switched elements:*** These elements are primarily based on the GSM network entities and carry data in a circuit switched manner, i.e. a permanent channel for the duration of the call.
- ***Packet switched elements:*** These network entities are designed to carry packet data. This enables much higher network usage as the capacity can be shared and data is carried as packets which are routed according to their destination.

Some network elements, particularly those that are associated with registration are shared by both domains and operate in the same way that they did with GSM.

Circuit switched elements
The circuit switched elements of the UMTS core network architecture include the following network entities:

- ***Mobile switching centre (MSC):*** This is essentially the same as that within GSM, and it manages the circuit switched calls under way.
- ***Gateway MSC (GMSC):*** This is effectively the interface to the external networks.

Packet switched elements
The packet switched elements of the 3G UMTS core network architecture include the following network entities:

- ***Serving GPRS Support Node (SGSN):*** As the name implies, this entity was first developed when GPRS was introduced, and its use has been carried over into the UMTS network architecture. The SGSN provides a number of functions within the UMTS network architecture.
  - <u>Mobility management</u> When a UE attaches to the Packet Switched domain of the UMTS Core Network, the SGSN generates MM information based on the mobile's current location.
  - <u>Session management:</u> The SGSN manages the data sessions providing the required quality of service and also managing what are termed the PDP (Packet data Protocol) contexts, i.e. the pipes over which the data is sent.
  - <u>Interaction with other areas of the network:</u> The SGSN is able to manage its elements within the network only by communicating with other areas of the network, e.g. MSC and other circuit switched areas.
  - <u>Billing:</u> The SGSN is also responsible billing. It achieves this by monitoring the flow of user data across the GPRS network. CDRs (Call Detail Records) are generated by the SGSN before being transferred to the charging entities (Charging Gateway Function, CGF).
- ***Gateway GPRS Support Node (GGSN):*** Like the SGSN, this entity was also first introduced into the GPRS network. The Gateway GPRS Support Node (GGSN) is the central element within the UMTS packet switched network. It handles inter-working between the UMTS packet switched network and external packet switched networks, and can be considered as a very sophisticated router. In operation, when the GGSN receives data addressed to a specific user, it checks if the user is active and then forwards the data to the SGSN serving the particular UE.

Shared elements
The shared elements of the 3G UMTS core network architecture include the following network entities:

- ***Home location register (HLR):*** This database contains all the administrative information about each subscriber along with their last known location. In this way, the UMTS network is able to route calls to the relevant RNC / Node B. When a user switches on their UE, it registers with the network and from this it is possible to determine which Node B it communicates with so that incoming calls can be routed appropriately. Even when the UE is not active (but switched on) it re-registers periodically to ensure that the network (HLR) is aware of its latest position with their current or last known location on the network.

- ***Equipment identity register (EIR):*** The EIR is the entity that decides whether a given UE equipment may be allowed onto the network. Each UE equipment has a number known as the International Mobile Equipment Identity. This number, as mentioned above, is installed in the equipment and is checked by the network during registration.
- ***Authentication centre (AuC) :*** The AuC is a protected database that contains the secret key also contained in the user's USIM card.

5. Explain Mobile IP in detail.

Agent Discovery (5)

A mobile node uses a method known as agent discovery to determine the following information:

- When the node has moved from one network to another
- Whether the network is the node's home or a foreign network
- What is the foreign agent care-of address offered by each foreign agent on that network

Mobility agents transmit **agent advertisements** to advertise their services on a network. In the absence of agent advertisements, a mobile node can solicit advertisements. This is known as **agent solicitation**.

Agent Advertisement

Mobile nodes use agent advertisements to determine their current point of attachment to the Internet or to an organization's network. An agent advertisement is an Internet Control Message Protocol (ICMP) router advertisement that has been extended to also carry a mobility agent advertisement extension.

A foreign agent can be too busy to serve additional mobile nodes. However, a foreign agent must continue to send agent advertisements. This way, mobile nodes that are already registered with it will know that they have not moved out of range of the foreign agent and that the foreign agent has not failed.

Also, a foreign agent that supports reverse tunnels must send it's advertisements with the reverse tunnel flag set on.

Agent Solicitation

Every mobile node should implement agent solicitation. The mobile node uses the same procedures, defaults, and constants for agent solicitation, as specified for ICMP router solicitation messages.

The rate at which a mobile node sends solicitations is limited by the mobile node. The mobile node can send three initial solicitations at a maximum rate of one per second while searching for an agent. After registering with an agent, the rate at which solicitations are sent is reduced, to limit the overhead on the local network.

Mobile IP Registration (5)

When the mobile node receives an agent advertisement, the mobile node registers through the foreign agent, even when the mobile node might be able to acquire its own co-located care-of address. This feature enables sites to restrict access to mobility services. Through agent advertisements, mobile nodes detect when they have moved from one subnet to another.

Mobile IP registration provides a flexible mechanism for mobile nodes to communicate their current reachability information to their home agent. The registration process enables mobile nodes to perform the following tasks:

- Request forwarding services when visiting a foreign network
- Inform their home agent of their current care-of address
- Renew a registration that is due to expire
- Deregister when they return home
- Request a reverse tunnel

Registration messages exchange information between a mobile node, a foreign agent, and the home agent. Registration creates or modifies a mobility binding at the home agent, associating the mobile node's home address with its care-of address for the specified lifetime.

The registration process also enables mobile nodes to:

- Register with multiple foreign agents
- Deregister specific care-of addresses while retaining other mobility bindings
- Discover the address of a home agent if the mobile node is not configured with this information

Mobile IP defines the following registration processes for a mobile node:

- If a mobile node is registering a foreign agent care-of address, the mobile node registers using that foreign agent.
- If a mobile node is using a co-located care-of address, and receives an agent advertisement from a foreign agent on the link on which it is using this care-of address, the mobile node registers using that foreign agent (or another foreign agent on this link).
- If a mobile node uses a co-located care-of address, the mobile node registers directly with its home agent.
- If a mobile node returns to its home network, the mobile node deregisters with its home agent.

These registration processes involve the exchange of registration requests and registration reply messages. When registering using a foreign agent, the registration process takes the following steps, which the subsequent illustration depicts:

1. The mobile node sends a registration request to the prospective foreign agent to begin the registration process.
2. The foreign agent processes the registration request and then relays it to the home agent.
3. The home agent sends a registration reply to the foreign agent to grant or deny the request.
4. The foreign agent processes the registration reply and then relays it to the mobile node to inform it of the disposition of its request.

Figure 1–7 Mobile IP Registration Process

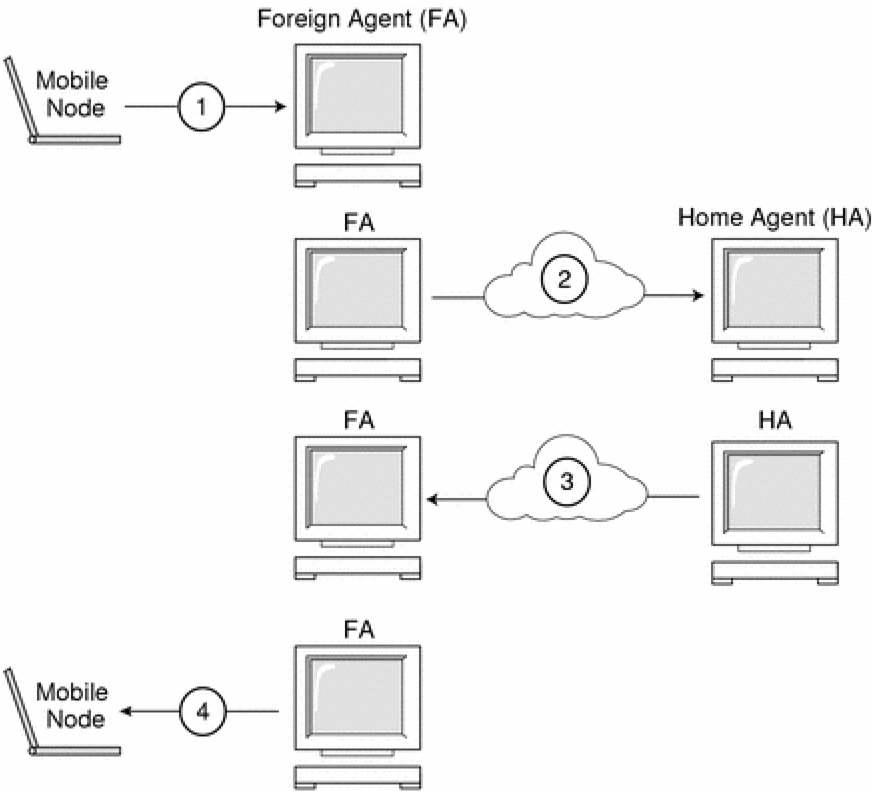

When the mobile node registers directly with its home agent, the registration process requires only the following steps:

- The mobile node sends a deregistration request to the home agent.
- The home agent sends a registration reply to the mobile node, granting or denying the request.

Also, a reverse tunnel might be required by either the foreign agent or the home agent. If the foreign agent supports reverse tunneling, the mobile node uses the registration process to request a reverse tunnel. The mobile node does this by setting the reverse tunnel flag on in the mobile node's registration request.

6. Summarize broadcast routing.

In unicast (i.e., point-to-point) communication, a single source node sends a packet to a single destination node. In broadcast routing, the network layer provides a service of delivering a packet sent from a source node to all other nodes in the network; multicast routing enables a single source node to send a copy of a packet to a subset of the other network nodes.

Broadcast Routing Algorithms (5)

Perhaps the most straightforward way to accomplish broadcast communication is for the sending node to send a separate copy of the packet to each destination. Given N destination nodes, the source node simply makes N copies of the packet, addresses each copy to a different destination, and then transmits the N copies to the N destinations using unicast routing. This N-way unicast approach to broadcasting is simple—no new network-layer routing protocol, packet-duplication, or forwarding functionality is needed. There are, however, several drawbacks to this approach. The first drawback is its inefficiency. If the source node is connected to the rest of the network via a single link, then N separate copies of the (same) packet will traverse this single link. It would clearly be more efficient to send only a single copy of a packet over this first hop and then have the node at the other end of the first hop make and forward any additional needed copies. That is, it would be more efficient for the network nodes themselves (rather than just the source node) to create duplicate copies of a packet.

We again model the network as a graph, G = (N,E), where N is a set of nodes and a collection E of edges, where each edge is a pair of nodes from N. We'll be a bit sloppy with our notation and use N to refer to both the set of nodes, as well as the cardinality (|N|) or size of that set when there is no confusion.

Uncontrolled Flooding The most obvious technique for achieving broadcast is a flooding approach in which the source node sends a copy of the packet to all of its neighbors. When a node receives a broadcast packet, it duplicates the packet and forwards it to all of its neighbors (except the neighbor from which it received the packet). Clearly, if the graph is connected, this scheme will eventually deliver a copy of the broadcast packet to all nodes in the graph. Controlled Flooding The key to avoiding a broadcast storm is for a node to judiciously choose when to flood a packet and (e.g., if it has already received and flooded an earlier copy of a packet) when not to flood a packet. In practice, this can be done in one of several ways. In sequence-number-controlled flooding, a source node puts its address (or other unique identifier) as well as a broadcast sequence number into a broadcast packet, then sends the packet to all of its neighbors. Each node maintains a list of the source address and sequence number of each broadcast packet it has already received, duplicated, and forwarded. When a node receives a broadcast packet, it first checks whether the packet is in this list. If so, the packet is dropped; if not, the
packet is duplicated and forwarded to all the node's neighbors (except the node from which the packet has just been received.

A second approach to controlled flooding is known as reverse path forwarding (RPF) also sometimes referred to as reverse path broadcast (RPB). The idea behind RPF is simple, yet elegant. When a router receives a broadcast packet with a given source address, it transmits the packet on all of its outgoing links (except the one on which it was received) only if the packet arrived on the link that is on its own shortest unicast path back to the source. Otherwise, the router simply discards the incoming packet without forwarding it on any of its outgoing links. Such a packet can be dropped because the router knows it either will receive or has already received a copy of this packet on the link that is on its own shortest path back to the sender. Note that RPF does not use unicast routing to actually deliver a packet to a destination, nor does it require that a router know the complete shortest path from itself to the source. RPF need only know the next neighbor on its unicast shortest path to the sender; it uses this neighbor's identity only to determine whether or not to flood a received broadcast packet.

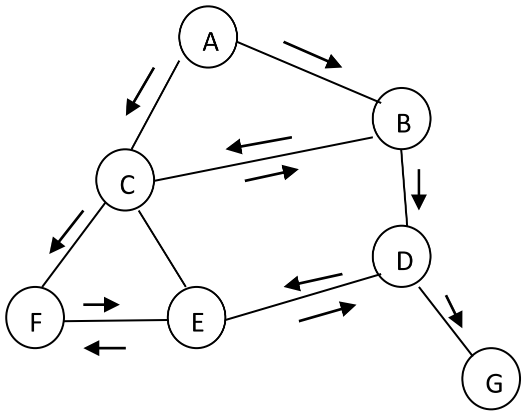

Reverse path forwarding (2)

Spanning-Tree Broadcast(3): While sequence-number-controlled flooding and RPF avoid broadcast storms, they do not completely avoid the transmission of redundant broadcast packets.
Spanning tree—a tree that contains each and every node in a graph. More formally, a spanning tree of a graph G = (N,E) is a graph=(N,E) such that E is a subset of E, G is connected, G contains no cycles, and G contains all the original nodes in G. If each link has an associated cost and the cost of a tree is the sum of the link costs, then a spanning tree whose cost is the minimum of all of the graph's spanning trees is called (not surprisingly) a minimum spanning tree. Thus, another approach to providing broadcast is for the network nodes to first construct a spanning tree. When a source node wants to send a broadcast packet, it sends the packet out on all of the incident links that belong to the spanning tree. A node receiving a broadcast packet then forwards the packet to all its neighbours in the Spanning tree (except the neighbour from which it received the packet).

7. Demonstrate how to manage mobility in cellular networks.

Like mobile IP, GSM adopts an indirect routing approach first routing the correspondent's call to the mobile user's home network and from there to the visited network. In GSM terminology, the mobile users's home network is referred to as the mobile user's home public land mobile network (home PLMN). The home network is the cellular provider with which the mobile user has a subscription (i.e., the provider that bills the user for monthly cellular service). The visited PLMN, is the network in which the mobile user is currently residing. As in the case of mobile IP, the responsibilities of the home and visited networks are quite different.

• The home network maintains a database known as the home location register (HLR) (3) , which contains the permanent cell phone number and subscriber profile information for each of its subscribers. Importantly, the HLR also contains information about the current locations of these subscribers. That is, if a mobile user is currently roaming in another provider's cellular network, the HLR contains enough information to obtain (via a process we'll describe shortly) an address in the visited network to which a call to the mobile user should be routed. As we'll see, a special switch in the home network, known as the Gateway Mobile services Switching Center (GMSC) is contacted by a correspondent when a call is placed to a mobile user. Again, in our quest to avoid an alphabet soup of acronyms, we'll refer to the GMSC here by a more descriptive term, home MSC.

• The visited network maintains a database known as the visitor location register (VLR) (3). The VLR contains an entry for each mobile user that is currently in the portion of the network served by the VLR. VLR entries thus come and go as mobile users enter and leave the network. AVLR is usually co-located with the mobile switching center (MSC) that coordinates the setup of a call to and from the visited network. In practice, a provider's cellular network will serve as a home network for its subscribers and as a visited network for mobile users whose subscription is with a different cellular provider.

1. The correspondent dials the mobile user's phone number. This number itself does not refer to a particular telephone line or location. The leading digits in the number are sufficient to globally identify the mobile's home network. The call is routed from the correspondent through the PSTN to the home MSC in the mobile's home network. This is the first leg of the call.

2. The home MSC receives the call and interrogates the HLR to determine the location of the mobile user. In the simplest case, the HLR returns the mobile
Mobile user station roaming number (MSRN), which we will refer to as the roaming number. Note that this number is different from the mobile's permanent phone number, which is associated with the mobile's home network. The roaming number is ephemeral: It is temporarily assigned to a mobile when it enters a visited network. The roaming number serves a role similar to that of the care-of address in mobile IPand, like the COA, is invisible to the correspondent and the mobile. If HLR does not have the roaming number, it returns the address of the VLR in the visited network. In this case (not shown in Figure 6.29), the home MSC will need to query the VLR to obtain the roaming number of the mobile node. But how does the HLR get the roaming number or the VLR address in the first place? What happens to these values when the mobile user moves to another visited network? We'll consider these important questions shortly. 3. Given the roaming number, the home MSC sets up the second leg of the call through the network to the MSC in the visited network. The call is completed, being routed from the correspondent to the home MSC, and from there to the visited MSC, and from there to the base station serving the mobile user.

An unresolved question in step 2 is how the HLR obtains information about the location of the mobile user. When a mobile telephone is switched on or enters a part of a visited network that is covered by a new VLR, the mobile must register with the visited network. This is done through the exchange of signaling messages between the mobile and the VLR. The visited VLR, in turn, sends a location update request message to the mobile's HLR. This message informs the HLR of either the roaming number at which the mobile can be contacted, or the address of the VLR (which can then later be queried to obtain the mobile number). As part of this exchange, the VLR also obtains subscriber information from the HLR about the mobile and determines what services (if any) should be accorded the mobile user by the visited network.

Placing a call to mobile user: Indirect Routing (4)