Solutions for Networks Qualifying Examination – Fall 1998

1. Fundamentally, bridges operate at the data-link layer and routers operate at the network layer. This means that they look at different kinds of network traffic. Bridges inspect data-link frames (e.g., Ethernet packets), while routers inspect network-layer packets (e.g., IP packets).

   A bridge connects two or more local area networks, and makes those networks appears as one local area network to the hosts on all of the connected networks. It does this by transparently forwarding data-link frames to the other networks. It does not simply forward all frames, which would cause unnecessary traffic on the networks that did not contain the destination host. Rather, a bridge learns which hosts are on which networks by inspecting the data-link frame for the sending host address, and recording which network the frame came from. The bridge then forwards a frame only to the network which it knows the host is on. If it does not yet know where the destination host is, the bridge will flood the frame to all networks, but a bridge rapidly learns where all hosts are, and is quite efficient. This automatic configuration means that a bridge is "plug-and-play" and requires virtually no setup.

   Bridges tend to be used to connect geographically close local area networks, although they are sometimes used with a dedicated long- distance connection to connect distant LANs. Bridging technology is most appropriate for connecting networks totaling hundreds of hosts.

   A router, on the other hand, operates at the network layer. It connects two or more local area networks or sets of bridged local area networks. A router might also connect a local network to other routers, not other networks. The other routers then may be connected to networks or still more routers. The routers not connected to any networks with hosts on them, but just connected to other routers, form a backbone transmission network for the host networks to communicate on.

   Routers are not transparent to the hosts—the hosts need to know about their closest routers. A host specifically sends network packets destined for far-away hosts to a selected router, and relies on the router to forward the packets on their way towards the destination. Routers decide where to send packets to based on either statically configured routes or dynamically learned routes. The static routes are configured by a person, but dynamic routes are learned by communicating with other routers on the internetwork. The routers share information about what networks (e.g., groups of network addresses) are connected to them, as well as what routers they are connected to.

   Hosts are also statically or dynamically configured with routing information, though static routing is predominant. Internetworks with routers can handle thousands, and even millions of hosts, as the Internet demonstrates.

2. The two main classes of routing algorithms are distance vector and link state. In a compact sense, distance vector algorithms communicate global information locally, and link state algorithms communicate local information globally.

Distance vector methods communicate only with neighboring routers; hence, they only send information locally. What they send is all knowledge about network reachability—a "vector" of distances from themselves to all other known routers. If a neighboring router sees a distance to a third router that is less than it previously knew, it updates its own distance vector to indicate that that third router is reached most easily by forwarding the packet on to the neighbor that sent it the new distance. There is no difficult computation by a router—all it needs to do is to compare distances and update its own table. In this way, the reachability of the network propagates to each router a step at a time. This method has problems with pathological cycles in the network, where a router updates its distance on a route that ends up going through itself again, and loss of connection can take a long time to detect (the count to infinity problem).

With link state methods, on the other hand, each router only sends local information—the state of each link it has with its neighbors, but it sends this information to all other known routers. In this way the slow propagation of network changes is avoided, unlike distance vector. Each router, then, receives information on all of the links, rather than whole pre-computed distances across the network. This means that the router must take the graph of the network, and compute its own optimal paths to each point in the network. The router uses a shortest-path algorithm to do this, in which it computes shortest paths through the graph to each other node. It is these paths that it uses when send packets on through the network. With link state methods, changes in the network are seen throughout the network immediately, and the problems with distance vector are avoided. More computational power is needed at each router, however.

3. Because bridges simply inspect the source address of a frame to decide which network the source host is on, and forward frames blindly based on that, multiple local area networks connected by bridges cannot, in general, be connected in cycles. A cycles would cause an "infinite loop" of forwarding to occur. Thus, at this level, bridges offer no capacity for redundancy, because multiple paths from one network to another are not handled.

   To provide redundancy, bridges can in fact connect networks in cycles. They communicate with each other and agree to operate as a spanning tree so that communication cycles are avoided. They do this by all running the same spanning tree algorithm, selecting a subset of all the bridges that still connects all networks. Those that do not make up the spanning tree take themselves off-line and do not actively operate. If a bridge or network fails, the live bridges recompute a new spanning tree (without operator intervention), allowing an inactive bridge to become active, thus continuing to provide complete connectivity for the networks. During normal operation, however, these inactive bridges and connections are, in a sense, wasted resources, since their bandwidth is not used.

   Routers, on the other hand, since they send packets along specific routes, are inherently able to handle redundant network paths and cycles in connectivity. Their route calculation algorithms, such as distance vector and link-state, calculate (heuristically) the most efficient paths through the internetwork. If a connection or a router fails, the live routers will recompute new paths through the network, since the paths using that failed connection or router are now not efficient (they don't work at all!). In this way, routers can even handle network congestion, not just completely failed connections. In addition, no potential routes are taken "off-line", as with bridges, but all are used to provide efficient usage of all available resources.