Like FDDI, RPR consists of two counterrotating optical fiber rings; unlike FDDI, it takes advantage of the bandwidth of both rings during normal operation. And unlike the previously described rings, an RPR frame is removed from the ring by the receiving node instead of leaving it to be removed by the sender, thereby freeing up some of the bandwidth on the ring in what is called *spatial reuse*.

Most strikingly, RPR does not use tokens. Instead, RPR uses a technique called *buffer insertion*. In a buffer insertion ring, a node can transmit its own frames whenever it has no other frames to forward. If a frame arrives while the node is transmitting its own frame, then the node temporarily buffers that frame. One of the major challenges for buffer insertion rings is how to avoid starvation and enforce QoS guarantees, since in its simplest form a buffer insertion ring could allow a station to hog the link indefinitely. RPR addresses this issue with fairly sophisticated QoS and fairness mechanisms.

RPR supports three QoS classes: class A provides low latency and low jitter (e.g., for phone calls), class B provides predictable latency and jitter (e.g., for prerecorded multimedia), and class C provides a best-effort transport.

To meet the resiliency goals, RPR uses two mechanisms to recover from the failure of a link or node. The first, *wrapping*, is similar to the approach described above for FDDI. The second, *steering*, is more sophisticated: nodes adjacent to the failure notify the other nodes, which are then able to direct packets in the correct (unbroken) direction around the ring toward any given destination, even if that is the "long" way around the ring—assuming the destination is not the node that just failed, of course.

A final interesting aspect of RPR is that it was designed to run over previously defined physical layers, including

Where Are They Now

The Future of Rings

The history of rings has seen them compete against Ethernet and ultimately lose on several occasions. 802.5 eventually lost out to 10-Mbit Ethernet for a variety of reasons, not least of which being the development of switched Ethernet, a topic we will discuss in the next chapter. FDDI was proposed as the faster alternative to Ethernet, but then Ethernet got faster too, and without the need for costly fiber optics, and FDDI never really caught on. The one ring technology that is still seeing some significant deployment is RPR, primarily in metropolitan area networks (MANs), although it seems likely that "metro Ethernet" will eventually come to dominate here just as Ethernet has done in LANs. There is, however, at least one reason RPR has had some success in MANs, which is the fact that rings are something of a natural fit for this kind of network, in a way that they are not in the LAN. Whereas it is cheap enough in a LAN

SONET and the physical layer specified for Ethernet. This saved the designers the time and effort of developing their own physical layer specs and hardware—a good example of the value of layered architectures.

2.8 Wireless

Wireless technologies differ in a variety of dimensions, most notably in how much bandwidth they provide and how far apart communicating nodes can be. Other important differences include which part of the electromagnetic spectrum they use (including whether it requires a license) and how much power they consume (important for mobile nodes). In this section we discuss four prominent wireless technologies: Bluetooth, Wi-Fi (more formally known as 802.11), WiMAX (802.16), and third-generation or 3G cellular wireless. In the following sections we present them in order from shortest range to longest. Table 2.6 gives an overview of these technologies and how they relate to each other.

The most widely used wireless links today are usually asymmetric, that is, the two endpoints are usually different kinds of nodes. One endpoint, sometimes described as the *base station*, usually has no mobility, but has a wired (or at least high bandwidth) connection to the Internet or other networks as in Figure 2.38. The node at the other end of the link—shown here as a "client node"—is often mobile, and relies on its link to the base station for all its communication with other nodes.

Observe that in Figure 2.38 we have used a wavy pair of lines to represent the wireless "link" abstraction provided between two devices (e.g., between a base station and one of its client nodes). One of the interesting aspects of wireless

Where Are They Now

to string cables in a hub-and-spoke manner from a central switch to each workstation, a ring actually provides a very cost-effective way to interconnect nodes in a MAN, where the cost of obtaining rights-of-way and laying fiber can be significant. The resiliency of a ring is also attractive in this environment—the fact that you have both a "clockwise" and an "counterclockwise" path between any two points ensures that a single fiber cut won't cut off a customer. RPR was also developed with some fairness mechanisms that ensure that a node's location on the ring doesn't put it at an unfair advantage or disadvantage to another node in another location when it comes to getting access to the bandwidth—this is harder to achieve with Ethernet. Thus, while there is certainly plenty of momentum behind Ethernet in the MAN, it is probably too soon to predict the demise of RPR in this environment.