The above example was simplified in that we assumed that all the channels can be used simultaneously in the given network, and we implicitly assumed single-hop flows. In general, as we have seen previously, the routes used for the flows may need to include multiple hops in order to maximize the performance. Also, the number of interfaces, say m, on a given nodes may be smaller than the number of channels c, making it impossible for a single host to use all the channels simultaneously. Thus, the capacity of the wireless network in this case depends on both the number of channels c, and the number of channels a single host may use simultaneously (say, m).

Suppose that each host is equipped with just one interface. Then a given host can only use one channel at any given time. But if there are a sufficiently large number of nodes in its "neighborhood" then they can use the remaining channels. In other words, so long as we have a sufficiently dense network, the nodes can *collectively* use the available channels, even if a single node cannot use them all at the same time. However, when we have a large number of channels, we cannot use the channels as effectively. In particular, when we have too many channels (that is, large c), we run into an *interface bottleneck*. With n nodes in the network, each with m interfaces, we cannot use more than $nm/2$ channels at any given time, leaving $c - nm/2$ channels idle. Thus, at some point, increasing the number of channels simply amounts to throwing away more spectrum.

The interface bottleneck may also affect routing decisions, as we now illustrate. Suppose that we have two flows in the network A to B, and C to D. Also suppose that we want to use 2-hop routes for these flows, with two alternative intermediate nodes E and F. Each host is equipped with 1 interface, and the available spectrum is divided into 2 channels. Also assume that, for all practical purposes, nodes E and F are exactly at the same location, and experience identical channel conditions. Now consider the two alternative routings shown in Figure 9.4; in both cases, flow A-to-B is routed via E. However, in Figure 9.4(a), flow C-to-D is routed via E, but via F in Figure 9.4(b). The numbers shown on the links represent the channel to be used on the corresponding link. Thus, different channels are being used for the two flows. Although the channels used in both figures are identical, the performance in figure (b) is likely to be better. In figure (a), although node E can use two channels, it only has one interface, so it can only use one channel at any given time, limited by the interface constraint. In figure (b) on the other hand, since nodes E and F together have two interfaces, they can collectively use the two channels simultaneously. To facilitate the performance improvement, as in figure (b), the routing protocol must prefer routes that balance the load across the nodes, even if the nodes are in each other's vicinity.

In designing practical protocols, other issues also need to be addressed, particularly the channel switching delay (that is, the delay incurred when a wireless interface switches its channel).

9.7 Impact of Infrastructure on Capacity

In our analysis of wireless capacity so far, we have assumed that only the wireless channel must be used to transmit information. What if we have a *hybrid network* in which some of the nodes had access to a wired network? Suppose that some of the hosts act as *access points*, and the access points are connected to each other by a separate high-speed wired network that is not the performance bottleneck. Thus, in the hybrid network, the information can be transmitted on the wireless channel or the wired network connecting the access point. The wired network has the benefit of being able to transmit the information between potentially distant access points, without having to use the wireless resources. Thus, intuition suggests that the availability of the access points should help improve performance. This intuition is indeed accurate, as capacity results for infrastructure-based networks have shown.

9.8 Capacity with Mobility

In our previous discussion, we have implicitly assumed that the hosts are fixed. If the hosts were to be mobile, can we improve performance? A hint at the answer comes from the above discussion of hybrid networks. In hybrid networks, we can improve performance by utilizing the wired network when appropriate: these transmissions can move data across the network without using the wireless channel. Can we exploit mobility to the same end? For instance, if the source and destination nodes periodically come close to each other, we could simply wait for this to occur, and the have the source transmit the information to the destination (instead of using a multi-hop wireless route as in our prior discussion). It has been shown that under certain assumptions about host mobility, it is indeed possible to improve capacity by exploiting mobility, although the above strategy of waiting until the destination comes close to the sender does not yield the optimal performance.

We briefly discuss a strategy to exploit mobility informally, omitting some important mathematical details. Let us consider a "random" mobility model such that that each node comes close to each of the other nodes equally frequently. Suppose that the time is divided in slots, such that a slot duration is sufficient to transmit one packet from any host to its nearest neighbor. In each slot, a source host transmits one packet to its nearest neighbor. If this neighbor is the intended destination of the packet, then no more forwarding of the packet is needed. If, however, the recipient of the packet is not the intended destination, then the recipient waits until it finds the destination as its nearest neighbor, and then transmits the packet to the destination. Thus, each packet is forwarded over at most two hops (no packet traverse more than two hops). This strategy results in all nodes "carrying" roughly $1/n$ fraction of the packets for each flow. Each such node will find each destination node as its closest neighbor $1/n$ fraction of the time, allowing enough time for the packets to be delivered to their destinations. One important detail that we have ignored is that for each transmission, one node must be a receiver. To address this issue, in each time slot, we can randomly choose half the nodes as transmitters, and have them pick the nearest node among the remaining half of the nodes as the receiver.

While the above approach has the benefit of improving per-flow throughput capacity of the wireless network to $\Theta(1)$, it has the disadvantage of significantly increasing the delay in delivery of the information.

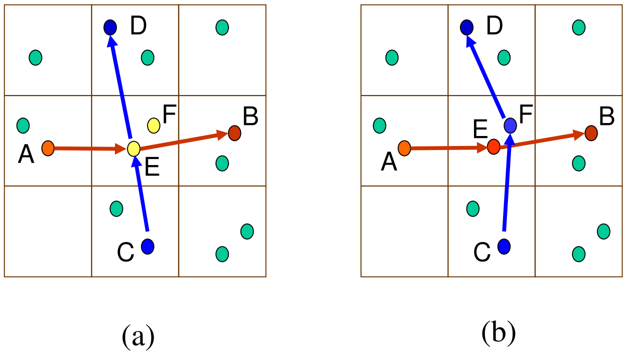

Figure 9.4 Alleviating interface bottleneck