Power management

Power management is a feature of some electrical appliances, especially copiers, computers and computer peripherals such as monitors and printers, that turns off the power or switches the system to a low-power state when inactive. In computing this is known as PC power management and is built around a standard called ACPI. This supersedes APM. All recent (consumer) computers have ACPI support.

Power Management Techniques

The previous section discussed WLANs and WPANs and the various standards that exist for them. The differences between each type of network were introduced with an emphasis put on their requirements for performing power management that each of them have. This section discusses the various power management techniques used by these standards for reducing the power consumed in each type of network. Many of the techniques introduced in this section do not appear in any of these standards, but are used in common practice to reduce the power of devices in both WLANs and WPANs. These techniques exist from the application layer all the way down to the physical layer of a traditional networking protocol stack. Techniques specific to a particular type of network are annotated as appropriate.

Application Layer

At the application layer a number of different techniques can be used to reduce the power consumed by a wireless device. A technique known as load partitioning allows an application to have all of its power intensive computation performed at its base station rather than locally. The wireless device simply sends the request for the computation to be performed, and then waits for the result. Another technique uses proxies in order to inform an application to changes in battery power. Applications use this information to limit their functionality and only provide their most essential features. This technique might be used to suppress certain "unnecessary" visual effects that accompany a process. While these techniques may be adapted to work with any application that wishes to support them, a number of techniques also exist for specific classes of applications.

Some applications are so common that it is worth exploring techniques that specifically deal with reducing the power consumed while running them. Two of the most common such applications include database operations and video processing. For database systems, techniques are explored that are able to reduce the power consumed during data retrieval, indexing, as well as querying operations. In all three cases, energy is conserved by reducing the number of transmissions needed to perform these operations. For video processing applications, energy can be conserved using compression techniques to reduce the number of bits transmitted over the wireless medium. Since performing the compression itself may consume a lot of power, however, other techniques that allow the video quality to become slightly degraded have been explored in order to reduce the power even further.

Transport Layer

The various techniques used to conserve energy at the transport layer all try to reduce the number of retransmissions necessary due to packet losses from a faulty wireless link. In a traditional (wired) network, packet losses are used to signify congestion and require backoff mechanisms to account for this. In a wireless network, however, losses can occur sporadically and should not immediately be interpreted as the onset of congestion. The TCP-Probing and Wave and Wait Protocols have been developed with this knowledge in mind. They are meant as replacements for traditional TCP, and are able to guarantee end-to-end data delivery with high throughput and low power consumption.

Network Layer

Power management techniques existing at the network layer are concerned with performing power efficient routing through a multi-hop network. They are typically either backbone based, topology control based, or a hybrid of them both. In a backbone based protocol (sometimes also referred to as Charge Based Clustering), some nodes are chosen to remain active at all times (backbone nodes), while others are allowed to sleep periodically. The backbone nodes are used to establish a path between all source and destination nodes in the network. Any node in the network must therefore be within one hop of at least one backbone node, including backbone nodes themselves. Energy savings are achieved by allowing non-backbone nodes to sleep periodically, as well as by periodically changing which nodes in fact make up the backbone.

Fig. 3 shows how packets would be routed from node 3 to node 4 and from node 1 to node 2 using the backbone that has been established. Black nodes signify backbone nodes, while numbered nodes signify non-backbone nodes. Solid lines indicate paths along which a packet may travel, while dashed ones show paths that will not be followed. Given this backbone structure, packets traveling from node 3 to node 4 will have to travel through 4 different backbone nodes before reaching their destination. If node 5 had been chosen as a backbone node as well, packets would only have had to traverse through 2.

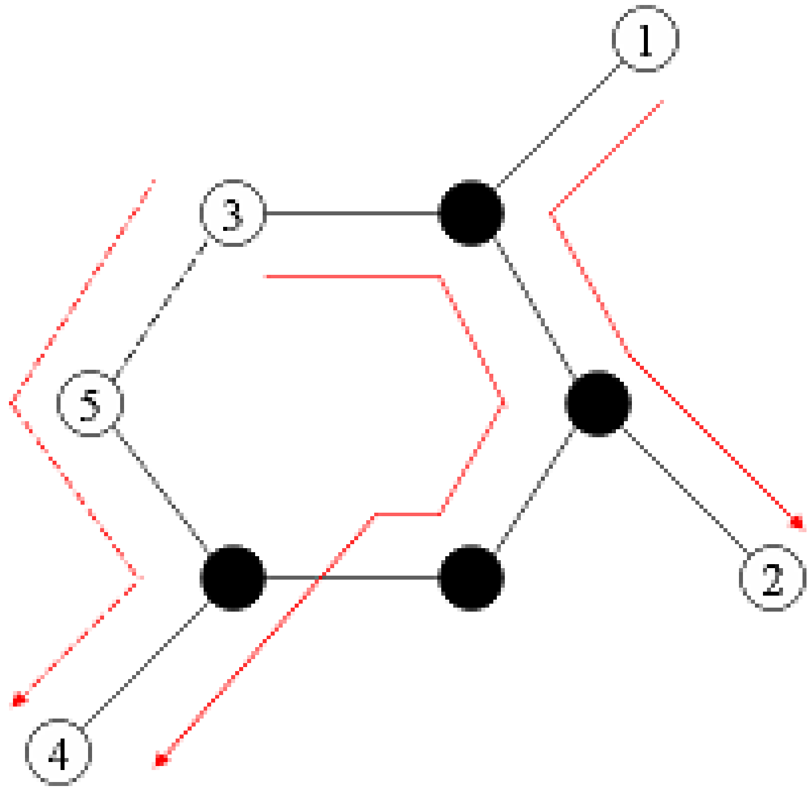

Fig. : Backbone based routing

Topology based routing protocols achieve energy savings in a different way. Their goal is to reduce the transmission power of all nodes in a network such that the network remains connected, but all nodes operate with the lowest transmission power possible. In a homogeneous network, this means that the transmission powers of all nodes are adjusted so that they are just within range of their nearest one-hop neighbor. In heterogeneous networks (i.e. networks with nodes of different type, power limitations, etc.) the transmission powers may be adjusted according to the needs of that network. A summary of the different types of topology based protocols that exist can be seen in Fig. 4.

As seen in the figure, certain location based topology control protocols attempt to use the topology of the network to provide the most energy efficient communication path possible. These protocols produce a sort of "Localized Power-Aware Routing" mechanism for the network. In some cases, providing this path means taking a larger number of hops through the network than would otherwise be taken when transmitting directly from one node to another. While this may seem counterintuitive at first, it makes sense if the amount of energy expended in transmitting to a node very far away is significantly greater than the energy expended when transmitting between a large number of nodes that are within closer range of one another. The rational behind the other topology based protocols found in Fig. 4.

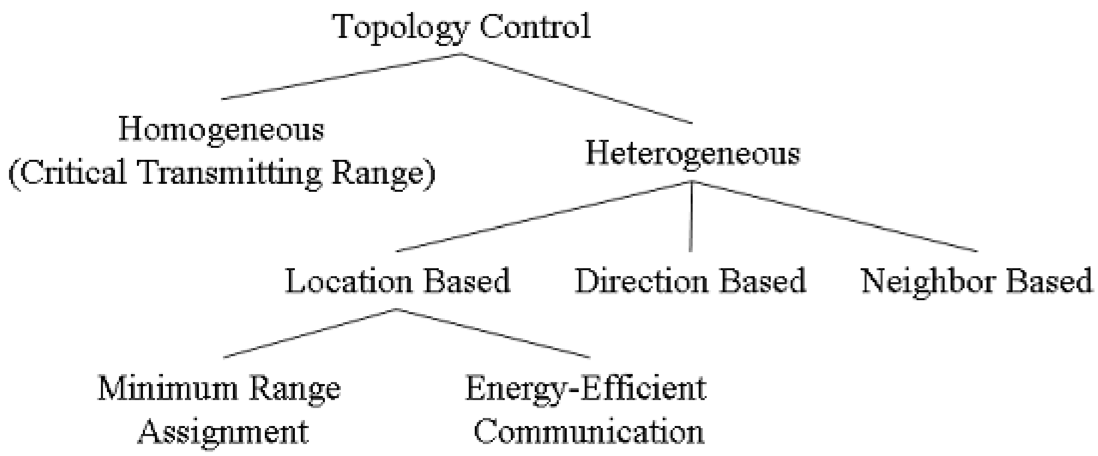

Fig. : Topology based routing protocols

Transmission power control schemes are combined with backbone based ones to produce a hybrid of them both. Using hybrid based protocols, the benefits of both backbone based and topology based routing protocols can be achieved simultaneously.

Data Link Layer

The two most common techniques used to conserve energy at the link layer involve reducing the transmission overhead during the Automatic Repeat Request (ARQ) and Forward Error Correction (FEC) schemes. Both of these schemes are used to reduce the number of packet errors at a receiving node. By enabling ARQ, a router is able to automatically request the retransmission of a packet directly from its source without first requiring the receiver node to detect that a packet error has occurred. Results have shown that sometimes it is more energy efficient to transmit at a lower transmission power and have to send multiple ARQs than to send at a high transmission power and achieve better throughput. Integrating the use of FEC codes to reduce the number of retransmissions necessary at the lower transmission power can result in even more energy savings.

Other power management techniques existing at the link layer are based on some sort of packet scheduling protocol. By scheduling multiple packet transmission to occur back to back (i.e. in a burst), it may be possible to reduce the overhead associated with sending each packet individually. Preamble bytes only need to be sent for the first packet in order to announce it presence on the radio channel, and all subsequent packets essentially "piggyback" this announcement. Packet scheduling algorithms may also reduce the number of retransmissions necessary if a packet is only scheduled to be sent during a time when its destination is known to be able to receive packets. By reducing the number of retransmissions necessary, the overall power consumption is consequently reduced as well.

MAC Layer

Power saving techniques existing at the MAC layer consist primarily of sleep scheduling protocols. The basic principle behind all sleep scheduling protocols is that lots of power is wasted listening on the radio channel while there is nothing there to receive. Sleep schedulers are used to duty cycle a radio between its on and off power states in order to reduce the effects of this idle listening. They are used to wake up a radio whenever it expects to transmit or receive packets and sleep otherwise. Other power saving techniques at this layer include battery aware MAC protocols (BAMAC) in which the decision of who should send next is based on the battery level of all surrounding nodes in the network. Battery level information is piggy-backed on each packet that is transmitted, and individual nodes base their decisions for sending on this information.

Sleep scheduling protocols can be broken up into two categories: synchronous and asynchronous Synchronous sleep scheduling policies rely on clock synchronization between nodes all nodes in a network. As seen in Fig. 5., senders and receivers are aware of when each other should be on and only send to one another during those time periods. They go to sleep otherwise.

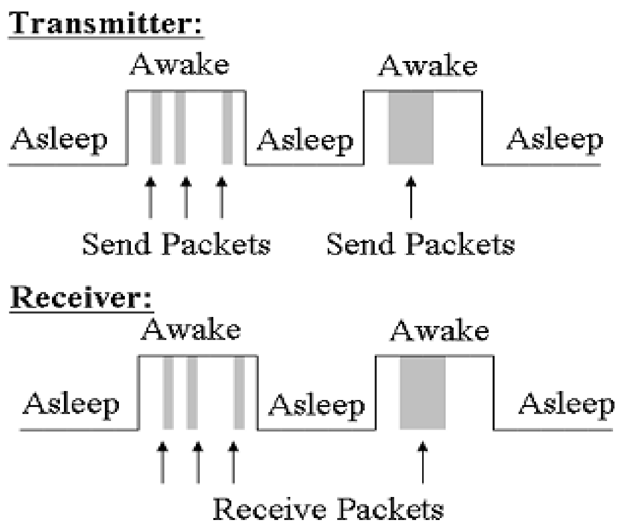

Fig. : Synchronous sleep scheduler

Asynchronous sleep scheduling, on the other hand, does not rely on any clock synchronization between nodes whatsoever. Nodes can send and receive packets whenever they please, according to the MAC protocol in use. Fig. 6 shows how two nodes running asynchronous sleep schedulers

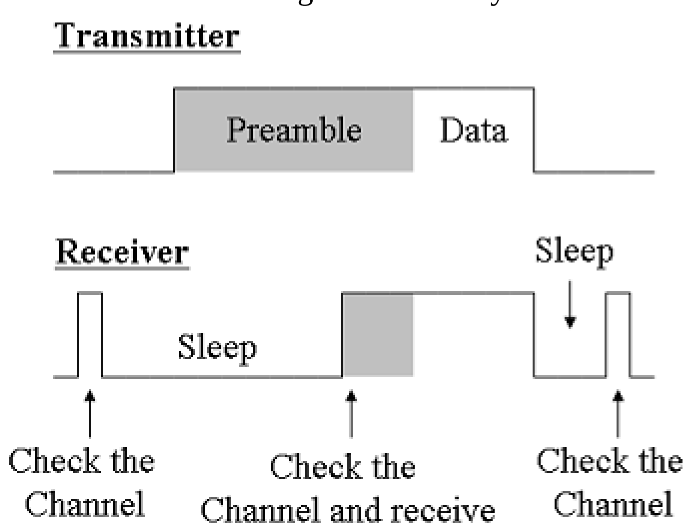

Fig. 6: Asynchronous sleep scheduler

Nodes wake up and go to sleep periodically in the same way they do for synchronous sleep scheduling. Since there is no time synchronization, however, there must be a way to ensure that receiving nodes are awake to hear the transmissions coming in from other nodes. Normally preamble bytes are sent by a packet in order to synchronize the starting point of the incoming data stream between the transmitter and receiver. With asynchronous sleep scheduling, a significant number of extra preamble bytes are sent per packet in order to guarantee that a receiver has the chance to synchronize to it at some point. In the worst case, a packet will begin transmitting just as its receiver goes to sleep, and preamble bytes will have to be sent for a time equal to the receiver's sleep interval (plus a little more to allow for proper synchronization once it wakes up). Once the receiver wakes up, it synchronizes to these preamble bytes and remains on until it receives the packet.

It doesn't make sense to have a hybrid sleep scheduling protocol based on each of the two techniques. The energy savings achieved using each of them varies from system to system and application to application. One technique is not "better" than the other in this sense, so efforts are being made to define exactly when each type should be used.

Physical Layer

At the physical layer, techniques can be used to not only preserve energy, but also generate it. Proper hardware design techniques allow one to decrease the level of parasitic leak currents in an electronic device to almost nothing. These smaller leakage currents ultimately result in longer lifetimes for these devices, as less energy is wasted while idle. Variable clock CPUs, CPU voltage scaling, flash memory, and disk spin down techniques can also be used to further reduce the power consumed at the physical layer. A technique known as Remote Access Switch (RAS) can be used to wake up a receiver only when it has data destined for it. A low power radio circuit is run to detect a certain type of activity on the channel. Only when this activity is detected does the circuit wake up the rest of the system for reception of a packet. A transmitter has to know what type of activity needs to be sent on the channel to wake up each of its receivers.

Energy harvesting techniques allow a device to actually gather energy from its surrounding environment. Ambient energy is all around in the form of vibration, strain, inertial forces, heat, light, wind, magnetic forces, etc. Energy harvesting techniques allow one to harness this energy and either convert it directly into usable electric current or store it for later use within an electrical system.