Bluetooth

Overview

- Cheap devices that reaches up to 3 Mbit/s.
- Security simple to not consume too much power.
- Users ISM frequencies: coexistence with 802.11 that implies interferences (also with microwave ovens); interferences are always present for communications within 4 meters.
- Possibility of exploiting cognitive networks (devices able to configure their operative frequency).
- FHSS with 79 channels of 1 MHz each one.
- Architecture composed of *piconets* (1 master, up to 7 slaves):
  - a master can be master only in one piconet each time;
  - slaves could belong to more piconets each time (*bridge nodes*).
- The *scatternet* is the network composed of piconets.

Technology

The documentation comprises:

- specifications that describes how the technology works;
- profiles that describes how the technology is used.

Specification

The radio interface can use 3 power classes:

- class 1: long range (up to 100 m), not used so much because is in competition with WLANs;
- class 2: range up to 10 m, largely used;
- class 3: short range (up to 10 cm), used for body area network.

Moreover, several operational states are provided (sniff, hold, park) that can reduce significantly the transmission power, allowing to save energy. Devices are composed of a single chip and, to keep simply the protocol, time division duplex is exploited to separate traffic (master/slave).

Protocol

The protocol stack is very well divided into:

- hardware: functions implemented in the chip;
- software: functions implemented on the top of the chip; several features are provided and they can be used accordingly to the application.

Lower levels are:

- RF: physical layer;
- Baseband: link layer;
- Link manager: helps the link layer.

RF

The radio channel uses FHSS (Frequency Hopping Spread Spectrum) to minimize the power consumption with 79 channels of 1 MHz each one: indeed, the larger is the number of channels, the more the signal is spread; this increases robustness to jamming and, therefore, leads to achieve better performances.

The piconet is a star topology in which the master is the center: communications among slaves are not allowed, only master-slave are permitted. Within a piconet, all nodes are synchronized with the master by adding an offset to their clock; each piconet uses its own frequency hopping sequence derived from the master's address: this avoids interferences among piconets (the probability that two piconets use the same frequency decreases as the number of channels increases). This is the reason why a node can not be master in two different piconets.

To deal with this problems, nodes within a piconet can use AFH (Adaptive Frequency Hopping): they measure parameters as transmission power and based on that classify channels. Bad channels are removed and the master can not more use them; the drawback is that, in this way, the benefit of frequency hopping is reduced.

Baseband

Piconet formation happens in two phases:

- *inquiry*: a node discover neighbors by acquiring information on their id and clock;
- *page*: allows to establish links.

The node that performs both phases become the master. Within a piconet there could be up to 7 slaves and up to 256 slaves in park mode: they are not active, so they can not exchange traffic.

Addresses are composed of:

. 3 bits for active nodes;

. 8 bits for parked slaves.

The channel is temporally organized in slots (each one is 625 μs long) and a packet is always transmitted on a single radio channel (of course the throughput depends on the packet duration). In each slot there is some idle time called *turnaround time*: it allows transceivers to switch mode (receiving/-transmitting).

The time division technique works as follow:

. during *even* slots, the master access the channel and poll slaves;

. during *odd* slots, slaves answer to master.

To support different kind of traffic, packets are classified as:

. Sco : real time traffic;

. eSco: used for transparent traffic, streaming for example;

. Acl : used for delay tolerant traffic; it has low priority with respect to the previous two kinds.

A master can support up to 3 Sco links: for this kind of packets, links are always symmetric while eScolinks could be asymmetric; if they are symmetric, master and slaves have the same resources in terms of number of slots.

Within a piconet, there is always ,for each slave, an Acl link established with the master used to exchange, in addition to possible data traffic, signalling: instead, Sco and eScohave first to reserve resources to be active.

If a slave does not decode the header, during an Acl connection, in the next slot it has to keep silent: this behaviour does not happens in Sco and eScobecause they are a kind of *circuits*.

Packets are:

Voice	Transparent traffic	Data
HV1	EV3	x-DH1
HV2	EV4	x-DH3
HV3	EV5	x-DH5

Number placed on the right correspond to different FEC techniques; headers, to be understood, are always sent at the basic rate.

HV packets achieve only 64 kbit/s: FEC techniques ensure reliability; for example:

- HV1 has 1 bit over 3 that carry information;
- HV2 has 2 bits over 3 that carry information;
- HV3 has all bits that carry information.

With ACL packets, if the transmission happens at the best modulation, 3-DH5, the maximum achievable throughput is 2.1 Mbit/s. Indeed, the x of x-DH, means that DH packets are sent with different modulations.

Inquiry This phase allows to recognize a communication although the link is not yet established. It is an asymmetric procedure in which listener and sender use the same hopping sequence, but the sender hops faster than the listener: this is done to increase the probability that they will meet on the same frequency soon. Indeed, senders transmit every 1.25 ms while listeners listen for a given frequency for 10 ms before hop to another one.

When they meet, the listener receives the id of the sender: it has to answer through a FHS packet that contains information on its address and local clock. The answer is not mandatory.

As a consequence of the procedure, only the sender gets information on other nodes, while nodes in inquiring scan state do not know anything about other.

Paging During inquiring phase, packets are sent in broadcast, while in paging they are sent directly to the destination: this allows to establish only one link.

In this phase, the sender has already some knowledge about the clock of the listener, but it estimate anyway the listener's clock because it is possible that some deviations occur.

When a node receive a paging message, it answers with an id packet; the pager replies to that packet through an FHS packet (this allows the listener to know the id and the clock of the pager) and, at the end, the listener acknowledge the pager with an id packet. After this procedure, the pager becomes the master.

A master can poll up to 7 neighbors; since masters have a much higher load to deal with, their energy consumption is high. It is possible that masters and slaves exchange their role with a specific procedure.

Scenarios

- A slave that belong to more piconets can communicate each time only in one of them because it need to change frequency hopping sequence.
- A node can be master in one piconet and slave in another one: if it is slave, it has to be inactive in the other piconet, therefore, in that moment, the traffic exchanged there is null.

Communication states

A node can be in one of the following operational states:

- standby: if it is not connected, but it listens to the channel;
- page/inquiry scan;
- connected:
  - active when receives and transmits packets;
  - sniff;
  - hold;
  - park.

Sniff During the period in which it is in *sniff* mode, a node is still synchronized with the master and maintains its MAC address. The sniff mode forecast that slaves listen to the channel to consume less energy and become active only for a small amount of time: this basic period is repeated until it reaches the maximum number for times that the slave agreed with the master. Performances degradation is not too high using this mode.

Hold Also for the hold state, slaves and master have to agree on the duration of the period, but, in this state, the node does not work at all. Due to this fact, when the hold period finishes, the slave has to listen for the master to re-synchronized its clock offset.

During this period, the node maintains its MAC address and the hold state could also be used by bridge nodes to switch frequency hopping sequence.

Park The park mode is the state that allows to save more energy: either master or slave can require that a node is put in park mode, but if the request is done by the master, only one slave at a time could be force to go in that state.

During a park period, the node wakes up sometimes to listen to the master's beacon because, the master can also force the slave to wake up; it is possible too, that through the beacon, the slave asks the master to change status.

Power control

Devices belonging to class 1 do not apply power control that, in practise, is always provided by class 2 devices: it is a numerical value, discrete, that represents the level of power received by a device; if it is above or below the *golden window* ($6 - 20\,\mathrm{dB}$), the current level is, respectively, decreased or increased.