1

Prologue: Wireless Personal Area Networks

1.1 Wireless Ad Hoc Networks

Wireless ad hoc networks are a category of wireless networks that utilize multi-hop relaying of packets yet are capable of operating without any infrastructure support (Perkins 2001; Ram Murthy and Manoj 2004; Toh 2002). Such networks are formed by a number of devices, possibly heterogeneous, with wireless communication capabilities that connect and disconnect at will. In addition, some or all of those devices may be mobile and are thus able to change their location frequently; ad hoc networks with mobile nodes are often referred to as mobile ad hoc networks, or MANETs. Even without mobility, nodes can join and/or leave an ad hoc network at will, and such networks need to possess self-organizing capability in terms of media access, routing, and other networking functions. Ad hoc networking includes such diverse applications as mobile, collaborative, and distributed computing; mobile access to the Internet; wireless mesh networks; military applications; emergency response networks; and others.

The design and deployment of those networks present a number of challenges which do not exist, or exist in rather different forms, in traditional wired networks:

- Self-organization, since individual nodes in an ad hoc networks must be able to attach to, and detach from, such networks at will, and without any fixed infrastructure. Protocols that can support and facilitate the tasks of topology construction, re-configuration, and maintenance, as well as routing, traffic monitoring and admission control, are needed.

- Scalability of the network refers to its ability to retain certain performance parameters regardless of large changes in the number of nodes deployed in that network. This is highly dependent on the amount of overhead at various layers (physical, medium access control, networking/routing, transport) of the network protocol stack.

Wireless Personal Area Networks Jelena Mišić and Vojislav B. Mišić

- Delay is the critical parameter in certain types of applications, e.g., in military applications such as battlefield communications or detection and monitoring of troop movement, or in health care applications where patients with serious and urgent medical conditions must be continuously monitored for important health variables via ECG, EEG, or other probes. Low delays can be achieved by bandwidth reservation, scheduling, or through some kind of admission control; the last two mechanisms require the presence of a controller or coordinator to monitor and prevent network congestion.

- Throughput is the most important performance target in a number of collaborative, distributed computing applications and in mobile access to the Internet, which might include significant amounts of multimedia traffic. At the PHY (physical) layer level, throughput may be impaired by packet errors caused by noise and interference. At the MAC (Medium Access Control) level, throughput may be impaired by collisions, if a contention-based medium access mechanism is used, or by unfairness, if bandwidth reservation- or scheduling-based access mechanism is used. (Detailed descriptions of these mechanisms can be found below.) Cross-layer optimization that accounts for those effects – preferably, all of them – may be needed in order to achieve high throughput.

- Packet and data losses. Loss of information is not tolerated in ad hoc networks, and active measures to restore reliability of data transfers must be undertaken both at the MAC and at the upper layers.

- Fairness among different nodes, applications, and/or users is also of importance.

- Power management is important when nodes operate on battery power, although facilities to recharge the batteries may be readily available at home or in the office.

- Finally, all maintenance tasks in ad hoc networks should be automated, or (at worst) simple enough to be undertaken by non-specialist human operators such as owners of laptop computers and PDAs.

1.2 Design Goals for the MAC Protocol

The medium access control (MAC) protocol is that part of the overall network functionality that deals with problems of achieving efficient, fair, and dependable access to the medium shared by a number of different devices (Stallings 2002). The role of the MAC protocol is particularly important in wireless networks which differ from their wired counterparts in many aspects. The most important among those differences stem from the very nature of the wireless communication medium, where two devices need not be explicitly connected in order to be able to communicate–instead, it merely suffices that they are within the radio transmission range of each other.

For example, when two or more packets are simultaneously received, the receiver may encounter problems. At best, the unwanted packets are treated as noise which impairs the reception of the packet intended to be received but can be filtered out. At worst, the correct packet may be damaged beyond repair and the receiver may be unable to make any

sense out of it; this condition is referred to as a collision. Collisions waste both network bandwidth and power resources of individual devices, transmitters and receivers alike, and active measures should be taken to reduce the likelihood of their occurrence.

Common approaches for collision minimization in wired networks include detection and avoidance. Collision detection is widely used in wired networks, where it involves the simple act of listening while transmitting. However, this is not feasible in wireless communication, where few devices are equipped with the required capability (Stallings 2002). Furthermore, packet collisions in wireless networks may occur in scenarios that cannot occur in wired ones, such as the so-called hidden and exposed terminal problems (Ram Murthy and Manoj 2004).

Since collision detection is not available, MAC protocols for wireless networks must rely on collision avoidance techniques, which include explicit scheduling, bandwidth reservation, and listening to the medium before attempting to transmit a packet. This last procedure is commonly known as clear channel assessment (IEEE 2003a, 2006; O'Hara and Petrick 1999), although other terms may be occasionally encountered as well.

Obviously, MAC protocols in wireless networks face both traditional challenges encountered in wired networks and new ones that stem from the use of the wireless communication medium. According to Ram Murthy and Manoj (2004), the most important features of MACs in ad hoc wireless networks can be summarized as follows:

- The operation of the protocol should be distributed, preferably without a dedicated central controller. If the use of such a controller cannot be avoided, the role should be only temporary, and devices with appropriate capabilities must be allowed to undertake it for a certain period of time.
- The protocol should be scalable to large networks.
- The available bandwidth must be utilized efficiently, including the minimization of packet collisions and minimization of the overhead needed to monitor and control network operation. In particular, the protocol should minimize the effects of hidden and exposed node problems.
- The protocol should ensure fair bandwidth allocation to all the nodes. Preferably, the fairness mechanism should take into account the current level of congestion in the network.
- The MAC protocol should incorporate power management policy, or policies, so as to minimize the power consumption of the node and of the entire network.
- The protocol should provide quality of service (QoS) support for real-time traffic wherever possible. Real-time, in this context, implies data traffic with prescribed performance bounds; these may include throughput, delay, delay jitter, and/or other performance indicators.

Two additional issues deserve to be mentioned. First issue is time synchronization among the nodes, which is required for the purpose of bandwidth reservation and allocation. Time synchronization is usually achieved by having one of the nodes periodically broadcast some sort of synchronization signal (the beacon) which is then used by other nodes. While the use

of periodic beacon transmissions facilitates the process of placing the reservation requests and subsequent broadcasting of reservation allocations, it requires that some node is capable of, and willing to, act as the central controller – somewhat contrary to the distributed, self-organizing character of an ad hoc network. In particular, additional provisions must be made to replace the controller node when it departs from the network or experiences a failure; this is part of the self-healing property of ad hoc networks described above. Furthermore, the use of beacons consumes the bandwidth and affects the scalability of the MAC algorithm.

The second issue is related to the interference from neighbouring nodes. As this interference is harmful, steps have to be taken to reduce it, most often through appropriate multiplexing techniques. According to Stallings (2002), multiplexing techniques are available in the following domains:

- in the frequency domain (FDMA), wherein different frequency bands are allocated to different devices or subnetworks;
- in the code domain (CDMA), wherein different devices use different code sequences;
- in the time domain (TDMA), wherein different devices transmit at different times; and/or
- in the space domain, where the range and scope of transmissions are controlled through the use of transmitter power control and directional antennas, respectively.

Strictly speaking, all these techniques belong to the PHY layer; while the MAC layer is completely oblivious to the first two techniques, it can utilize the latter two (multiplexing in time and space domain), or even integrate them to a certain extent. (For example, time multiplexing is a close relative of scheduling.) This cross-layer integration and optimization allow the MAC protocol to better address the requirements outlined above. We note that such integration is not too common in ad hoc networks, where the MAC layer is more likely to cooperate with the network and, possibly, transport layers above it, than with the PHY layer below; however, MAC protocols exist that make use of it (Ram Murthy and Manoj 2004).

1.3 Classification of MAC Protocols for Ad Hoc Networks

Before we present some of the important MAC protocols for wireless ad hoc networks, we will give a brief overview of some among the possible criteria for classifying those protocols; the reader will thus be able to grasp main features of different MAC protocols and identify the important similarities as well as differences among them.

Mechanism for accessing the medium. Probably the most intuitive among the classification criteria is the manner of accessing the medium, which comes in three main flavours:

- Contention-based protocols are those in which a potential sender node must compete with all others in order to gain access to the medium and transmit its data.
- Bandwidth reservation-based protocols have provisions for requesting and obtaining bandwidth (or time) allocations by individual senders.

- Finally, scheduling-based protocols, in which the transmissions of individual senders are scheduled according to some predefined policy which aims to achieve one or more of the objectives outlined above, such as the maximization of throughput, fairness, flow priority, or QoS support.

Note that the third option requires the presence of an entity which is responsible for implementing the aforementioned policy. In most cases, this requirement translates into the requirement for a permanent or temporary central controller. Note also that the policy to be pursued should be adaptive, depending on the traffic and/or other conditions in the network. The presence of a central controller is sometimes needed in protocols that use the second option as well.

Quite a few among the existing MAC protocols offer more than one of those mechanisms. This may be accomplished by slicing the available time into intervals of fixed or variable size, referred to as cycles or superframes (IEEE 2003a, 2006; O'Hara and Petrick 1999), and assigning certain portions of those intervals to different categories of access from the list above. For example, the IEEE 802.11 Point Coordinator Function (PCF) uses superframes in which the first part is reserved for (optional) contention-free access, while the second part is used for contention-based access (ANSI/IEEE 1999; O'Hara and Petrick 1999). A similar approach is adopted in the IEEE 802.15.4 protocol in its beacon enabled, slotted CSMA-CA mode (IEEE 2006), except that the contention access period precedes the contention-free period in the superframe. More details on the structure of the superframe are presented in the next chapter.

On the other hand, some MAC protocols offer optional features which modify the manner in which the protocol operates, and effectively introduce a different mechanism for medium access control. For example, the IEEE 802.11 Distributed Coordinator Function (DCF) utilizes pure contention-based access in its default form, but allows bandwidth reservation on a per-packet basis through the optional RTS/CTS handshake (ANSI/IEEE 1999).

Alternative classifications on the basis of medium access mechanism. An alternative classification criterion could be devised by assuming that contention-based access will always be present, and then using the presence or absence of the latter two access mechanisms as the basis for classification. This approach results in the common (and marginally more practical) classification into pure contention-based MACs, contention-based MACs with reservation mechanisms, and contention-based MACs with scheduling mechanisms (Ram Murthy and Manoj 2004). A variant of this approach distinguishes between contention- or random access-based protocols, scheduling or partitioning ones, and polling-based ones. Yet even these classifications are neither unambiguous, as the presence of optional features outlined above leads to the same protocol being attached to more than one category, nor comprehensive, as some of the existing protocols cannot be attached to any single category (Ram Murthy and Manoj 2004); on account of these shortcomings, it is listed as an alternative only.

Mechanism used for bandwidth reservation and its scope. These two criteria applies only to MAC protocols that employ some form of bandwidth reservation, and thus actually represent sub-classifications within the previous one based on the mechanism used to access

the medium. With respect to the mechanism used for bandwidth reservation, we can distinguish between the protocols that use some kind of handshake, e.g., RTS/CTS, and those that use out-of-band signalling, most notably the Busy Tone approach which is an extension of the familiar concept from the traditional telephony systems.

With respect to the scope of bandwidth reservation, we can distinguish between the protocols which request bandwidth for a specified time (i.e., for a single packet or for a group of consecutive packets, commonly referred to as a burst) and those that request bandwidth allocation for an unspecified time. In both cases, time can be measured in absolute units or in data packets. In the former case, bandwidth allocation is valid for the transmission of a specified number of packets only, while in the latter, it has to be explicitly revoked by some central authority, or perhaps waived by the requester itself.

Another scheme based on the concept related to bandwidth reservation is the family of the so-called multi-channel MAC protocols. Namely, most communication technologies use only one channel out of several available in the given frequency band. Multi-channel MACs exploit this feature to employ channel hopping in order to improve bandwidth utilization and/or reduce congestion.

Presence and scope of synchronization. The presence or absence of time synchronization among the nodes in the network is another criterion that can be used to classify MAC protocols for wireless ad hoc networks. Synchronization, if present, may be required to extend to all the nodes in the network (global synchronization); alternatively, it may apply to just a handful of nodes which are physically close to one another (local synchronization). In the former case, a central controller may be needed to initiate and broadcast the necessary synchronization information.

Synchronization is most often required in protocols that use scheduling or bandwidth reservation, as basic synchronization intervals serve to apportion the available bandwidth to appropriate sender nodes. However, bandwidth reservation and allocation can be accomplished in an asynchronous manner, in particular when reservation is requested on a per-packet basis, while synchronous protocols can be used even with pure contention-based access. For example, the IEEE 802.15.4 protocol in its beacon enabled, slotted CSMA-CA mode without guaranteed time slots uses pure contention-based access, yet all transmissions must be synchronized to the beacon frames periodically sent by the network coordinator (IEEE 2006).

Synchronization is one of the most important factors that may affect scalability of the network. As the size of the network grows, synchronization becomes more difficult and more costly to establish and maintain. In particular, protocols which rely on global synchronization will suffer the most degradation; for example, it has been shown that the construction and maintenance of a globally optimal schedule in a multi-level Bluetooth network (a scatternet) is an NP-complete problem (Johansson et al. 2001).

Presence of a controller and its permanence. Another possible classification criterion is the presence and permanence of a central network controller or coordinator. While wireless ad hoc networks, by default, should be able to function without a permanent or dedicated central controller, quite a few protocols rely on certain monitoring and control functions that can only be provided by a local or global controller. This is the case with several of the MAC protocols that use bandwidth reservation, as well as with all of the MAC

protocols which use scheduling. In fact, even some pure contention-based protocols rely on the presence of a controller for administrative tasks such as time synchronization and sometimes even node admission.

Again, the presence of a controller affects the scalability of the network, as the amount of work the controller has to do – most of which is administrative and control overhead – must grow with the number of nodes. Hierarchical decomposition or layering is often used to reduce this overhead, but it leads to additional problems regarding synchronization and delays.

Interdependence of the classification criteria. As can be seen, not all of the classification criteria outlined above are entirely independent of each other; rather, they exhibit a certain overlap or redundancy. Still, they are useful in the study of MAC protocols, as they tend to highlight different aspects of their design and operation.

1.4 Contention-Based MAC Protocols

We will now look at two contention-based MAC protocols: the basic CSMA protocol and the IEEE 802.11 DCF. They are interesting because the 802.15.4 protocol uses a variant of CSMA which is rather similar to those two. While many other protocols exist, contention-based, polling-based, and those that use bandwidth reservation, multiple channels, out-of-band signalling, and directional antennas, they are beyond the scope of the present work.

1.4.1 Basic CSMA

Many MAC protocols are derived from the basic Carrier Sense Multiple Access (CSMA) mechanism (Bertsekas and Gallager 1991). CSMA is a pure distributed protocol without centralized control, which operates as follows. The node that wants to transmit a packet first performs the clear channel assessment procedure, i.e., it listens to the medium, for a prescribed time. If the medium is found to be clear (or idle) during that time, the node can transmit its packet. Otherwise, i.e., if another transmission is in progress, the node backs off – i.e., waits for a certain time before undertaking the same procedure again.

Different MAC algorithms use different ways to calculate the time they need to listen to the channel during the clear channel assessment procedure and to calculate the time to wait (i.e., the duration of the backoff period) before the next transmission attempt.

It is possible that the transmissions from two or more nodes overlap in time, which results in a collision and loss of all packets involved. If lossless communication is desired, collisions must be detected so that the lost packets can be retransmitted. Since a collision can be detected only at the receiver side, some form of acknowledgment from the receiver may be needed; some MAC protocols provide this facility, while others leave it to some of the upper layers – most likely, the transport layer. The former approach is more efficient in terms of reaction time, whereas the latter allows for much simpler implementation of the MAC protocol used.

In the basic CSMA protocol, carrier sensing is performed only at the sending node. Therefore, the hidden terminal problem is still present. Moreover, the exposed terminal problem leads to deferred transmissions and thus reduces bandwidth utilization.

1.4.2 IEEE 802.11 MAC

The IEEE 802.11 protocol (O'Hara and Petrick 1999) is, strictly speaking, intended for wireless local area networks (LANs), rather than wireless ad hoc networks. However, it is interesting to examine it in some detail, mainly on account of its ubiquity, and because it uses most of the main concepts which are reused in many MAC protocols for ad hoc networks. The protocol covers the functional areas of access control, reliable data delivery, and security; in the following we will focus on the first two areas, as the last one (security) is beyond the scope of this chapter.

Reliable transfer is achieved through the use of special acknowledgment (ACK) packets or frames, sent by the destination node upon successfully receiving a data packet. Medium access is regulated in two ways, the first of which is a distributed contention-based mechanism known as Distributed Coordination Function (DCF), which does not require a centralized controller. The DCF, based on the CSMA protocol described above, operates as follows. The node that wants to transmit a packet first performs the clear channel assessment procedure, i.e., it listens to the medium, for a time equal to Interframe Space (IFS). If the medium is found to be clear (or idle) during that time, the node can transmit its packet immediately; otherwise, i.e., if another transmission is in progress, the node waits for another IFS period. If the medium remain idle during that period, the node backs off for a random interval and again senses the medium. During that time (referred to as the backoff window or contention window), if the medium becomes busy, the backoff counter is halted; it resumes when the medium becomes idle again. When the backoff counter expires and the medium is found to be idle, the node can transmit the packet.

A possible scenario in which this procedure is applied is shown in Figure 1.1. There are several points worth mentioning. First, the backoff interval is chosen as a random number from a predefined range. After each collision, the range is doubled in order to reduce the likelihood of a repeated collision. After each successful transmission, the range is reset to its initial value, which is typically small. This approach is known as binary exponential backoff, or BEB (Stallings 2002). In this manner, the protocol ensures a certain level of load smoothing in case of frequent collisions caused by heavy traffic.

Second, in order to enhance reliability and avoid the hidden/exposed terminal problems to a certain extent, the RTS/CTS handshake – well known from wired communications – may optionally be used. In this case, the node that wants to send a data packet first sends a Request To Send (RTS) packet to the designated receiver which, if ready, responds with a Clear To Send (CTS) packet. Both RTS and CTS packets contain information about

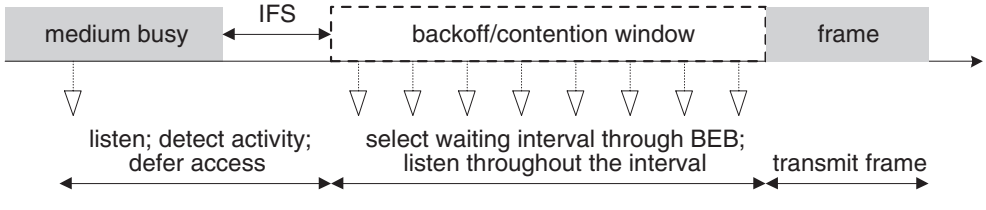

Figure 1.1 Basic access method in IEEE 802.11 DCF.

the duration of the forthcoming transmission, including the optional acknowledgment. Once the sender receives the CTS packet, it may begin actual data transmission, which may optionally be followed by an ACK packet. The RTS/CTS handshake constitutes a simple form of bandwidth reservation on a per-packet or per-group basis, as will be explained below.

Reliability of transmission is enhanced because the RTS and CTS packets are generally much shorter than data packets; if they collide, the time waste is not high – but the risk that subsequent data packets will experience a collision is substantially reduced. The hidden terminal problem is avoided because other nodes within the transmission range of the receiver, upon hearing the CTS packet, become aware of a forthcoming data transmission and defer their transmission for the time interval specified. On the other hand, a transmission from an exposed terminal may prevent the sender from initiating the RTS/CTS handshake. However, once the sender receives a proper CTS packet, it can assume that the receiver is not affected by the interfering transmission and can, thus, proceed with the data packet transmission.

Third, in order to ensure the proper functioning of the protocol, three different IFS intervals are used: a short IFS (SIFS), a medium duration Point Coordination Function IFS (PIFS), and a long duration one, referred to as Distributed Coordination Function IFS (DIFS). The existence of several IFS intervals of different duration actually serves to implement different priority levels for different types of access. The DIFS interval is used for ordinary asynchronous traffic, while the SIFS interval, being the shortest, is used in the following cases:

- When the receiver sends an ACK packet upon successful reception of a data packet; in this manner, ACK packets are safe from collisions since regular data packets wait longer.

- When the sender wants to send another data packet upon receiving an ACK packet for a previous one. In this manner, a burst of packets (commonly obtained by segmenting a longer packet from the upper layers) can be delivered quickly and with little risk from collision. However, such transmissions can result in unfairness, since there is limit on the duration of the burst that can be transmitted.

- When the node sends a CTS packet upon receiving a RTS packet from a prospective sender; again, the use of the SIFS interval minimizes the risk that the CTS packet will experience a collision.

The PIFS interval is used in an alternative access method known as the Point Coordination Function (PCF), which is implemented on top of DCF. The PCF requires the presence of a central point coordinator, hence the name. The point coordinator defines an interval known as superframe. In the first part of the superframe, the coordinator issues polls to all nodes configured for polling. The polls are sent using the regular CSMA algorithm outlined above. When a poll packet is sent, the polled node may respond using the SIFS interval. If the coordinator receives the response, it issues another poll but using the PIFS interval. The polling continues in round-robin fashion (i.e., one node at a time), until all the nodes are polled. Then, the point coordinator remains idle until the end of the superframe, which allows for DCF-style contention-based access by all other nodes.The duration of the superframe is fixed, but an ongoing transmission may force the coordinator to defer the

beginning of a polling cycle; in this case the useful duration of the superframe will be reduced.

While the IEEE 802.11 DCF is able to deal with asynchronous traffic, the presence of synchronous traffic with specified (and reasonably stable) throughput over prolonged periods of time is well served by its PCF counterpart. Still, the PCF functionality is designated as an optional facility in the 802.11 standard (ANSI/IEEE 1999), and it is rarely used in practice.

1.5 New Kinds of Ad Hoc Networks

Recently, new families of wireless ad hoc networks for specialized applications have emerged, most notably sensor and personal area networks.

Wireless sensor networks, or WSNs, are aimed at monitoring environmental phenomena (e.g., temperature, humidity, light but also the presence of a specific object or movements of persons and objects) in a given physical space. Such networks find increasing use in areas as diverse as military applications, object surveillance, structural health monitoring, and agriculture and forestry, among others.

Wireless Personal Area Networks, or WPANs, are intended to provide advanced capabilities such as cable replacement, interconnection of various electronic devices, monitoring of physical parameters on the human body, and the like, all within a person's workspace. Different application areas for WPANs have widely differing requirements in terms of data rate, power consumption, and quality of service, such networks are typically classified into the following three classes:

- High data rate WPANs are needed for real-time and multimedia applications. Such applications are supported through the IEEE 802.15.3 standard (IEEE 2003a), with the maximum data rate of 55 Mbps (megabits per second).
- Medium data rate networks for cable replacement and consumer devices. This was the original use of WPANs, as envisioned in the IEEE 802.15.1 (Bluetooth) communications standard. The original Bluetooth specification (Bluetooth SIG 2003; IEEE 2002) allowed raw data rates of up to 1 Mbps, but recent improvements allow data rates of up to 3 Mps (Bluetooth SIG 2004; IEEE 2005).
- Finally, low data rate WPANs are intended for use in wireless sensor networks and other similar application scenarios. A typical example of a LR-WPAN is the 802.15.4 standard (IEEE 2003b, 2006), which allows data rates of up to 250 kbps (kilobits per second).

In this book, we will focus on the performance of WPANs that utilize the 802.15.4 standard in its various configurations.

1.6 Sensor Networks

Sensor networks are a class of wireless networks intended for monitoring environmental phenomena in a given physical space; such networks find increasing use in areas as diverse

as military applications, object surveillance, structural health monitoring, and agriculture and forestry, among others. Monitoring may be continuous, with a prescribed data rate which may change over time; it may also be triggered by an explicit demand from a controlling node or a specific event in the environment. Environmental phenomena to be monitored include simple physical variables such as temperature, humidity, light, pressure, pH value, and the like; but other phenomena such as the presence or absence of a specific object (say, an inventory item with a RFID tag), or movements of persons and objects (e.g., cars) can be monitored as well. The spaces to be monitored include rooms, hallways, foyers, homes, backyards, streets, larger buildings and structures (e.g., bridges), but also open spaces such as fields or forests. Sensor nodes can be deployed in large numbers, from tens through hundreds to even thousands. Sensor networks are often expected to operate autonomously, with little or no human intervention, for prolonged periods of time. Sensor nodes are seldom mobile, and even when mobility is present, not all of the nodes are equipped with appropriate capabilities. Given such a diverse set of applications and requirements, it should come as no surprise that the constraints which guide the design and deployment of wireless sensor networks differ, sometimes substantially, from those that hold in wireless ad hoc networks (Achir and Ouvry 2004; Sohrabi et al. 2000). Let us now discuss those constraints in more detail.

Energy efficiency. Probably the most important difference is due to the fact that sensor nodes typically operate on limited battery power, which means that the maximization of network lifetime (and, consequently, minimization of power consumption) is a sine qua non for sensor networks. On the contrary, power consumption is seldom the critical requirement for ad hoc networks.

According to Jones et al. (2001), the constraint of minimal energy consumption translates into two distinct, yet closely related design requirements:

1. The communication efficiency has to be maximized through the design of simple yet flexible and effective communication protocols and functions.

2. Those protocols and functions have to be implemented by small chips with limited computational and memory resources.

Simultaneous achievement of these objectives necessitates some kind of cross-layer protocol optimization in which the MAC layer would use the information obtained from, and control the operation of the PHY layer. At the same time, optimal operation of the upper, network and transport layers requires the knowledge of appropriate information from both the PHY and MAC layers. Again, such tight integration is not too common in ad hoc networks.

An important consequence of the requirement for energy efficiency is the limited transmission range of most sensor node radio subsystems; few real devices have a transmission range of more than 100 meters (300 feet), and ranges of 10 meters (30 feet) and even less are not uncommon.

Protocol efficiency. Regarding communication protocols, the main sources of inefficiency are packet collisions, but also overly complex handshake protocols, receiving packets destined for other nodes, and idle listening to the medium (Ye et al. 2004). Actual power consumption of sensor nodes, often called motes, depends mostly on the radio subsystem

and its operating mode. In most (but not all) cases, transmitting and receiving use about the same amount of energy, depending on the power level used for transmission. However, most savings can be made by putting the node to sleep, when power consumption drops by one to two orders of magnitude, depending on the hardware (Jung and Vaidya 2005; van Dam and Langendoen 2003).

Use of redundant sensors. Since nodes are small and cheap to produce and the network lifetime needs to be maximized, it is often feasible to deploy the sensors in a given physical space in much larger numbers than necessary to obtain the desired rate of information flow. If redundant sensors are used, they can be periodically sent to sleep in order to minimize their duty cycle, which extends the lifetime of individual sensors and of the entire network and reduces or eliminates the need for operator intervention, thus reducing the operational cost of the network (Akan and Akyildiz 2005). The use of redundant sensors has profound implications on the design of MAC protocols, as will be seen below.

Node specialization. Another important distinction is related to the role of individual nodes. An ad hoc network allows its nodes to choose the specific role, or roles, they would like to play – i.e., data source, destination, or intermediate router – at any given time. In most cases, a node is free to switch to a different role, or roles, whenever it finds appropriate or is instructed to do so by the specific application currently executing on it. On the contrary, nodes in a sensor network have specific roles that do not change often, or never change at all. Most of the nodes act as sensing nodes, some act as intermediaries which route the traffic and (possibly) perform some administrative duties, and a small number of nodes (sometimes only a single node) act as the network sink (or sinks) toward which all the sensed data ultimately flows (Akyildiz et al. 2002). A group of sensor nodes under the control of an intermediary is sometimes referred to as a sub-network or cluster, while the intermediary itself is known as cluster head. We note that the number of intermediate levels interposed between the sensing nodes and the network sink(s) depends on a number of variables such as the size of the network, the size of the physical space which the network has to monitor, the transmission range of individual nodes, and (to some extent) the actual MAC protocol used.

Traffic characteristics. The traffic in sensor networks is rather asymmetric, as the bulk of it flows from the sensing nodes toward the network sink (this is often referred to as the uplink direction). The traffic in the opposite direction is generally much smaller and consists of control information and, possibly, queries issued by the network sink on behalf of the corresponding sensing application (Intanagonwiwat et al. 2003). Furthermore, traffic patterns in sensor networks are rather different than in ad hoc networks. For example, temperature or humidity monitoring might require periodic or nearly periodic transmissions – in essence, synchronous traffic with low data rate – while object surveillance and other event-driven sensing applications exhibit low average traffic volume and random bursts with considerably higher peak rates.

Furthermore, data packets are often much smaller in sensor networks. Original data from sensing nodes typically consists of only a few data values reported by appropriate sensors. Intermediate nodes may choose to aggregate those values in order to improve

energy efficiency and reduce bandwidth and energy consumption; data aggregation is more common in networks with a larger number of hierarchical levels. At the same time, the number of sensor nodes and their spatial density may be very large, depending on the size of the space to be monitored and the requirements of the sensing application.

Quality of Service requirements. Delay considerations are of crucial importance in certain classes of applications, for example, in military applications such as battlefield communications and detection and monitoring of troop movement, or in health care applications where patients in special care units must be monitored for important health variables (via ECG or EEG) due to a serious and urgent medical condition. Maintaining prescribed delay bounds in a network of resource-constrained nodes with limited transmission range is a complex issue. Low delays can be achieved either by bandwidth reservation, as utilized in variations of the TDMA approach, or by some kind of admission control that will prevent network congestion, if the CSMA approach is used.

At the same time, the requirement for maximum throughput is relaxed due to the following. First, the exact value of the throughput requirement is usually prescribed by the sensing application, unlike general networks where the goal is to obtain as much throughput as possible. Second, energy efficiency dictates the use of protocols that incorporate power control, which will strive to keep the nodes inactive for as long as possible (Akan and Akyildiz 2005). In order to obtain the desired throughput, it suffices to adjust the mean number of active nodes.

Even packet losses can be catered to in this manner, since we don't care whether a given packet from a given node will reach the network sink – as long as the sink receives sufficient number of packets from other nodes. Any packet loss can be compensated for (in the long term) by varying the mean number of active nodes. In a certain sense, fairness is not needed at the node and packet level as long as it is maintained at the cluster level (Callaway, Jr. 2004). On the contrary, fairness at the node/packet level is important in ad hoc networks.

Differences from ad hoc networks. The requirements outlined above lead to a number of important differences between sensor networks and ad hoc networks, most notably the following:

- Power efficiency and lifetime maximization are the foremost requirements for sensor networks.
- Self-organization is important in both ad hoc and sensor networks. In the former case, this is due to dynamicity and node mobility, which cause frequent topology changes and make self-organization more difficult; in the latter, this is mostly caused by sensor nodes exhausting their battery power (i.e., dying), although mobile sensors are used in some applications.
- Throughput maximization is often required in ad hoc networks but is not too common in sensor networks.
- Delay minimization is typically assigned much higher priority in sensor networks than in their ad hoc siblings.

- The use of redundant sensors allows for a certain level of fault tolerance; on the contrary, packet losses are intolerable in ad hoc networks.
- Scalability is an important issue due to the potentially large number of sensors; scalability is also important in ad hoc networks, but it is limited by the available bandwidth and the desired throughput.
- Nodes in ad hoc networks are often mobile, while most sensor networks have no mobile nodes.

In more than one sense, wireless ad hoc networks are a class of networks with flexible topology but without infrastructure, that should cater to all kinds of networking tasks. On the other hand, sensor networks are highly specialized networks that perform a rather restricted set of tasks under severe computational and communication restrictions.