Bluetooth – an Enabler for Personal Area Networking

Per Johansson
Ericsson Research

Manthos Kazantzidis, Rohit Kapoor, Mario Gerla
Computer Science Department, UCLA

Abstract

We find ourselves today often carrying numerous portable electronic devices, such as notebook computers, mobile phones, PDAs, digital cameras and mp3/MD/DVD players, used to help and entertain us in our professional life as well as in our private life. For the most part, these devices are used separately and their applications do not interact. Imagine, however, if they could interact directly and thus create a network where information may flow seamlessly between the devices – such a network of personal devices is often referred to as a *personal area network* (PAN). Moreover, access to the Internet via a (public) wireless LAN access point and/or via a 3G UMTS mobile phone would enable the PAN to be constantly on-line. The strongest candidate to provide with the cheap, short-range radio links necessary to enable such networks is the Bluetooth wireless technology. Seen from a networking perspective, a PAN will be expected to have participants, both of its "own" devices and "guest" devices from other PANs, continuously moving in and out of its coverage. To cope with this volatile nature of the network, the concept of *ad-hoc networking* may be applied to create a robust and flexible connectivity. A major technical step is taken when the Bluetooth *piconet* network architecture, a strict star topology, is extended into a *scatternet* architecture, where piconets are interconnected. A consequence of creating scatternet based PANs is that some nodes will form gateways between piconets and these gateways must be capable of time sharing their presence in each piconet they are members of. While the Bluetooth standard defines the gateway nodes, the actual mechanisms and algorithms that accomplish the inter-piconet scheduling (IPS) is left rather open. Given the lack of research literature in the subject, an overall architecture for handling scheduling in a scatternet is presented in this paper. A family of feasible IPS algorithms, referred to as *Rendezvous Point* algorithms, is also introduced and discussed.

1. Introduction

The number of mobile devices that we carry around these days, both for work and pleasure, seems to steadily increase. Our notebook computer, cellular phone and PDA have all become almost necessary tools not only for the business traveler but also for an increasing part of the general public. Even if these devices have became both more powerful and still smaller in size, we now often add a MD/mp3 player or even a DVD player to complete the electronic "wardrobe". For the most part, these devices are used separately and their applications do not interact, mainly because interconnection via wires becomes tedious and time consuming. This niche was also identified by a group of mobile phone vendors and notebook computer vendors, who initiated the development of the Bluetooth [9] wireless technology. Bluetooth uses a frequency-hopping scheme in the unlicensed Industrial Scientific-Medical (ISM) band at 2.4 GHz. In 1998 the Bluetooth Special Interest Group (SIG) was formed to host the work of the Bluetooth specification and the first version of the specification was released in 1999. The major goal with the Bluetooth wireless technology is to allow relatively cheap electronic devices to communicate directly in an ad-hoc fashion, which requires the price of the Bluetooth radio to be only a few dollars. Moreover, the Bluetooth equipped devices can also form networks where information may flow seamlessly between the applications hosted in the devices – such a network of personal devices is often referred to as a *personal area network* (PAN).

One example of an application within one PAN is the electronic business card of a new contact that automatically find its way across the PAN into the address register on a notebook computer and into the number register on a mobile phone. Communication between PANs would, for instance, enable participants at a meeting to share documents or presentations. Access to the Internet via a (public) wireless LAN (WLAN) access point and/or via a third-generation (3G) cellular phone would enable the devices in the PAN to be constantly on-line. For instance, commuters may have a public WLAN access point in a train to their notebook computers and when exiting the train their notebook computers could remain online via the 3G phone, while incoming e-mail could now be diverted to their PDAs through the PAN. Finally, as they enter their offices, the access could again, automatically, go through the notebook computer via a wireless LAN access to the corporate campus network.

In parallel with the development of the Bluetooth PAN technology the introduction of the 3G mobile networks is under way at a very rapid pace. The 3G systems, i.e. Universal Mobile Telecommunications System (UMTS) and cdma2000 will move the Internet into the mobile world in a useful way, moving on to higher data rates in the order of hundreds of kbps instead of tenths. This development will create a push for mobile terminals that can offer high resolution images and high quality audio, requiring more computing capacity – the mobile phone will evolve from a pure voice service terminal to a multimedia platform. However, the introduction of the PAN concept may very well widen the design space for 3G terminals. Since the PAN would enable a network of devices, the 3G terminal can be kept rather simple while, e.g., a PDA could offer the more powerful computing necessary for multimedia applications. The Bluetooth wireless link interconnecting the two devices would be capable to carry the multimedia stream. Moreover, the form factor of the 3G phone may in this scenario vary from a simple voice terminal to a full-fledged palm-top multimedia terminal. Thus, the Bluetooth ad hoc PAN scenario enables the user to personalize her personal media environment, by allowing a free distribution of applications and functions within the Bluetooth PAN – a wider variety of terminals may be accepted to access, and take advantage of, the 3G networks.

Bluetooth presents a number of technical challenges from a networking perspective such as ad-hoc network formation and scheduling of traffic between nodes simultaneously operating on different channels. Bluetooth operates inherently in an ad-hoc manner since it is not relying on any infrastructure via say an access point or base station. This is reflected in the way nodes are detected and how networks, for instance PANs, are created without, or with a minimum, of pre-configuration. The participants of a Bluetooth network is expected to be mobile and will move in and out of coverage and nodes may also join or leave the network rather frequently. For instance, a PAN may have both "own" devices and "guest" devices from other PANs. The characteristics of a Bluetooth PAN will in many cases be such that the concepts of ad-hoc networking [1] fit very well and could help to create a robust and flexible network connectivity. Furthermore, a major technical step is taken when the Bluetooth *piconet* network architecture, a strict star topology[1], is extended into a *scatternet* architecture, when piconets are interconnected. The combinations of network topologies will increase dramatically and methods to create robust, but still efficient, scatternets become crucial. To create a good scatternet connectivity, bandwidth and delay requirements, among others, must be considered. In addition, many mobile phones and other electronic devices already are or will soon be Bluetooth-enabled. Consequently, the ground for building more complex ad hoc personal area networks is being laid.

The continuation of the paper is organized as follows. In Section 2, the ad-hoc PAN, its usage cases and typical characteristics are discussed in more detail. Networking of the Bluetooth wireless system is described in Section 3, together with some general technical background on Bluetooth. In Section 4 a general architecture for scatternet scheduling is given together with some categories of inter-piconet scheduling (IPS) algorithms. Finally, some conclusive remarks are given in Section 5 to summarize the work.

2. The Ad-hoc PAN — a network extension

Before continuing into a more detailed technical description of Bluetooth networking, a more general discussion around applications is given to motivate the design of Bluetooth PANs in the first place. Obviously, just a single Bluetooth link by itself will simplify intercommunication between various

[1] In a piconet the central, controlling node is called master and other client nodes in the star are called slaves.

mobile devices (such as a cellular phone and a PDA) by eliminating cables. However, it is with the introduction of Bluetooth ad-hoc networking on the PAN level that the true potential of ad-hoc PANs can really be utilized. The ad-hoc nature of these networks should be emphasized, not only as isolated PANs but also as PANs interconnected with larger public network infrastructures, e.g., UMTS. It is likely that the first ad-hoc network application really moving into the commercial market segment will be in the form of PANs.

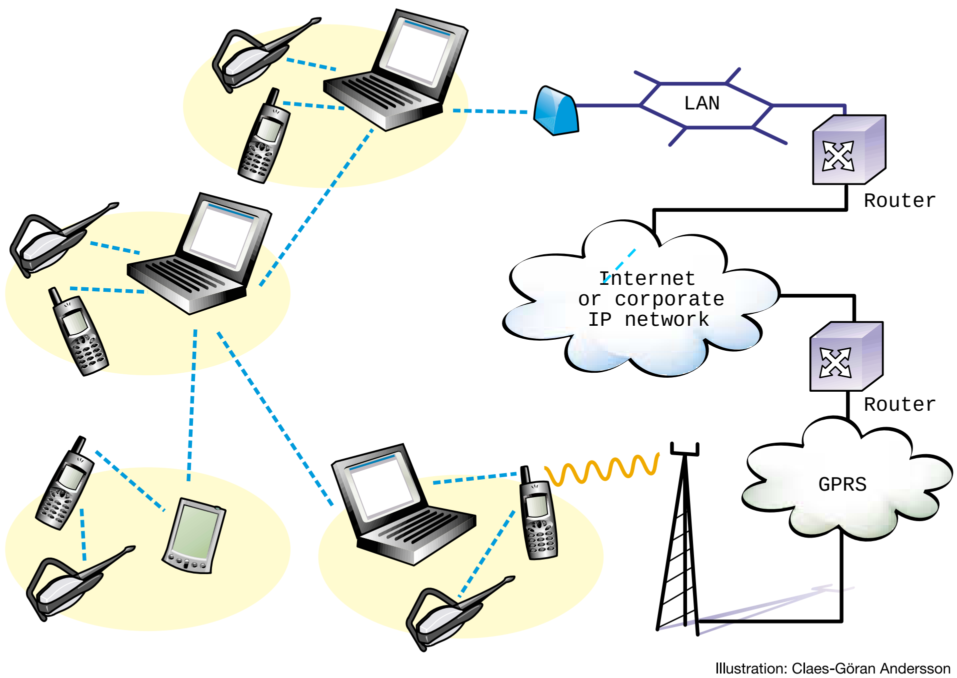

Figure 1   Four interconnected Bluetooth PANs, with both a WLAN and a GPRS access running simultaneously

Seen from the viewpoint of the traditional mobile network, a Bluetooth based PAN opens up a new way to extend the mobile network services into the user domain – out to other devices further out in the PAN. In terms of traffic load offered to the network, the aggregate traffic of the PAN would typically exceed that of the single mobile phone. In addition, if several Bluetooth PANs are interconnected this offered traffic will be further increased. For a mobile network operator this could be a very attractive way to get users to increase utilization of the new 3G networks as well as extending the user applications in the 3G phones. The high bandwidth consumer in these PANs would typically be a notebook computer or a high-end palmtop computer.

Figure 1 illustrates the case discussed above with a scenario of four Bluetooth PANs, consisting of notebook computers, cellular phones, headsets and PDAs. The PANs are interconnected via three notebook computers and one PDA. In addition, two of the PANs are connected to an IP backbone network, one via a LAN access point and the other via a General Packet Radio Service[2] (GPRS) phone, thus creating two possible ways to access the IP backbone network. In case the mobile operator also manages the WLAN access point, mechanisms could be deployed to divide the traffic between the wide area 3G access and the local area WLAN access to achieve an efficient utilization of the network resources. The PAN itself can also encompass several different access technologies—distributed among

---

[2] GPRS is often referred to as "2.5G" and is an intermediate step in the evolution from GSM (2G) to UMTS (3G).

its member devices—which exploit the ad hoc functionality in the PAN. For instance, a notebook computer could have a WLAN interface (such as IEEE 802.11 [16] or HiperLAN/2 [17]) that provides network access when the computer is used where a WLAN access is available. When no WLAN access are available the GPRS/UMTS phone would be used to access the IP backbone network. Thus, the PAN would benefit from the total aggregate of all access technologies residing in the devices of the PAN.

As the PAN concept matures, it will allow new devices and new access technologies to be incorporated into the PAN framework. It also eliminates the need to create hybrid devices, such as a combined PDA-mobile phone, because the PAN network will instead allow for wireless integration. In other words, it will not be necessary to trade off form for function.

2.1. PAN Performance Characteristics

An ad hoc Bluetooth PAN would be expected to operate in a network environment in which some or all the nodes are mobile, which contrasts to traditional wireline and wireless networks. In general, however, the same basic user requirements for connectivity and traffic delivery that apply to traditional networks will apply to ad hoc networks.

- Distributed operation: a node in the PAN cannot rely on a network in the background to support security and routing functions. Instead these functions must be designed so that they can operate efficiently under distributed conditions.
- Dynamic network topology: in general, the nodes will be mobile, which sooner or later will result in a varying network topology. Nonetheless, connectivity in the PAN should be maintained to allow applications and services to operate undisrupted. In particular, this will influence the design of routing protocols. Moreover, a user in the ad hoc network will also require access to a fixed network (such as the Internet via a 3G access network) even if nodes are moving around. This calls for mobility-management functions that allow network access for devices located several Bluetooth radio hops away from a network access point.
- Fluctuating link capacity: the effects of high bit-error rates are more profound in a multihop ad hoc network, such as the Bluetooth PAN, since the aggregate of all link errors is what affects a multihop path. In addition, more than one end-to-end path can use a given link, which if the link were to break, could disrupt several sessions during periods of high bit-error transmission rates. However the Bluetooth link layer uses both automatic repeat request (ARQ) and forward error correction (FEC) techniques to counter these problems.
- Low-power devices: in most cases, the PAN nodes will be battery-driven, which will make the power budget tight for all the power-consuming components in a device. This will affect, for instance, CPU processing, memory size/usage, signal processing, and transceiver output/input power.

3. Bluetooth networking

In order to implement ad hoc PANs based on the Bluetooth wireless technology [6-11], work is under way in the Bluetooth SIG to enhance the current Bluetooth functionality to provide an improved networking support. One key, necessary feature is a very good capability to carry IP efficiently to/from and within a Bluetooth PAN. This is called upon since the Bluetooth PANs will be connected to the Internet via either the 3G networks, or corporate/public WLANs, and will contain IP-enabled hosts. Generally speaking, a good capacity for carrying IP would give Bluetooth networks a wider and more open interface, which would most certainly boost the development of new applications for Bluetooth.

Before elaborating more on the technical aspects related to Bluetooth network functionality, a brief introduction to Bluetooth in general is given in the next subsection.

3.1. Bluetooth basics

Two or more Bluetooth units that share the same channel form a *piconet*. Figure 2 shows three examples of different piconet configurations. Within a piconet, a Bluetooth unit can play either of two roles: master or slave. Each piconet may only contain one master (and there must always be one) and up to seven active slaves. Any Bluetooth unit can become a master in a piconet.

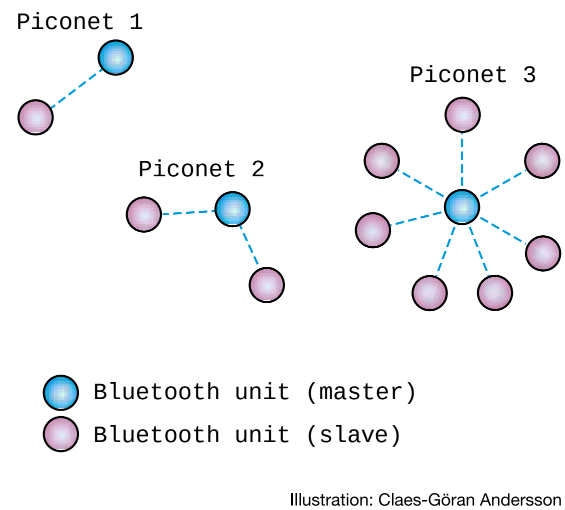

Figure 2 Three piconets.

Furthermore, two or more piconets can be interconnected, forming what is called a *scatternet*. In Figure 3 a scatternet consisting of twelve piconets is depicted. The connection point between two piconets consists of a Bluetooth unit that is a member of both piconets. A Bluetooth unit can simultaneously be a slave member of multiple piconets, but only a master in one. Moreover, because a Bluetooth unit can only transmit and receive data in one piconet at a time, its participation in multiple piconets has to be on a time division multiplex basis.

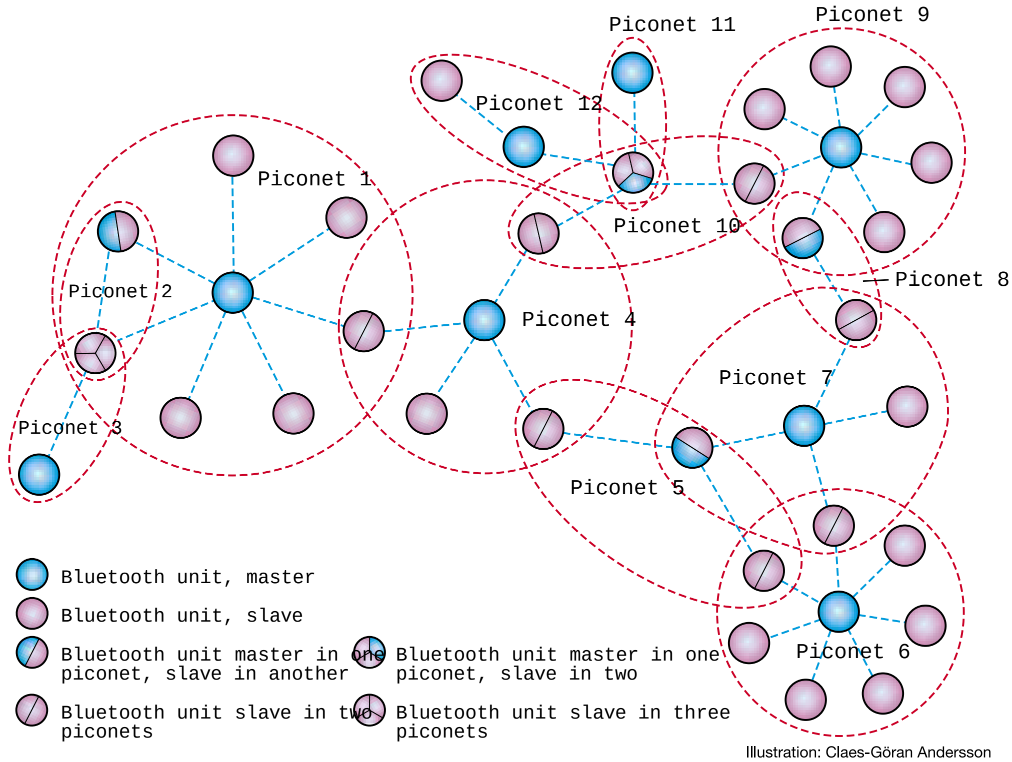

Figure 3 Twelve piconets interconnected into one single scatternet.

The Bluetooth system provides full duplex transmission based on slotted time-division duplex (TDD), where the duration of each slot is 0.625 ms. There is no direct transmission between slaves in a Bluetooth piconet, only from master to slave and vice versa.

Communication in a piconet is organized so that the master polls each slave according to a polling scheme. A slave is only allowed to transmit after having been polled by the master. The slave will start its transmission in the slave-to-master timeslot immediately after it has received a packet from the master. The master may or may not include data in the packet used to poll a slave. However, it is possible to send packets that cover multiple slots. These multislot packets may be either three or five slots long.

3.2. Scatternet-based PANs

In a Bluetooth PAN context the scatternet functionality is important to allow a flexible forming of ad hoc PANs. Even though a Bluetooth PAN often will be based on a single piconet, the possibility for a node in the PAN to be present in another piconet is essential to allow for combination of various Bluetooth usage cases. The nodes that are present in multiple piconets, i.e. inter-piconet nodes, may either have applications that are operating independently in the piconets, or function as a gateway between the piconets and forward traffic between them. The former, non-forwarding, case will typically occur when a notebook computer in the PAN has small accessories such as a mouse or joystick and at the same time needs to reach a Bluetooth LAN access point in the room. Then the notebook computer will in general be a master in the (single piconet) PAN and the accessories will be slaves. However the LAN access point will in most cases require to be the master of the nodes it communicates with. So if the notebook computer connects to the LAN access point it will remain as a master in its “own” PAN, while it will operate as a slave unit for the data access point, i.e. the notebook computer is an inter-piconet node. However, there is no need to forward traffic between the data access point and the accessories, thus not any need to consider networking in this case. Figure 4 shows a case where two PANs are connected to a Bluetooth access point, forming a scatternet of three piconets. In this scenario the inter-piconet nodes (the notebook computers in the PANs) are not forwarding traffic since the slave nodes in the PANs typically only need to exchange local information with its masters (the notebook computers in this case).

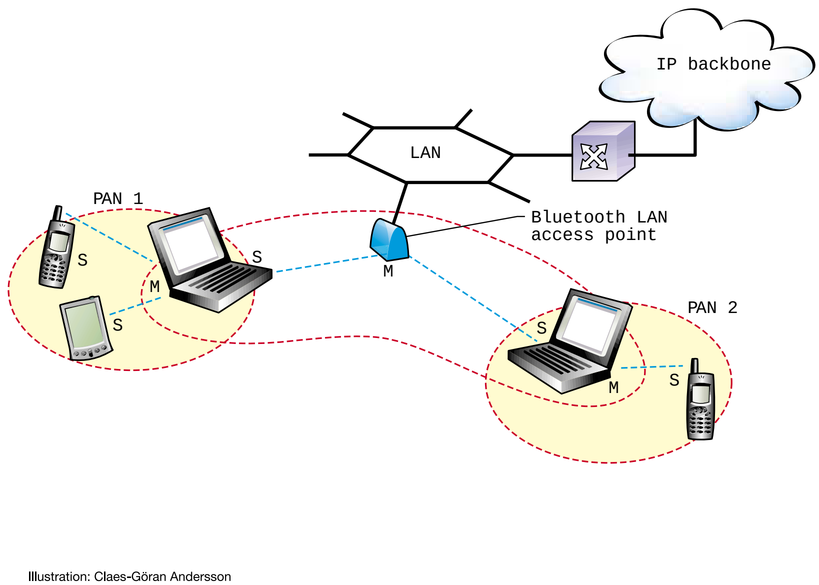

Figure 4 The notebook computer in PAN 1 operates as an inter-piconet node (master in PAN 1 and slave to the the Bluetooth access point). In this scenario packets are typically not forwarded by the inter-piconet node.

In the case when the inter-piconet nodes forward packet between piconets, the Bluetooth ad-hoc PANs pertain to the class of multihop ad hoc networks. These scatternet based PANs may be used when information needs to be spread widely among the PANs residing within reach through a “reasonable” number of radio hops. Here, the reasonable number of hops depends on the type of information in terms of required data rate and end-to-end delay. However, the number of possible hops depends heavily on the scatternet application. In a case where the scatternet is used to interconnect, say, a large network of Bluetooth equipped sensors, the flow of information may settle with very low data rates and high delays so the hop count could be in the orders of hundreds. On the other hand, a high quality interactive video application within a PAN may handle only one Bluetooth hop and cannot cross any inter-piconet gateway due to the imposed delay.

In Figure 5 a scenario in which a GPRS phone provides Internet access to all the notebook computers in the scatternet. This means that the notebook computer in PAN 1 operates as an inter-piconet gateway, forwarding packets directly between PAN 1 and the notebook computers. In this case the inter-piconet node needs to have network forwarding functions to be able to achieve the end-to-end path.

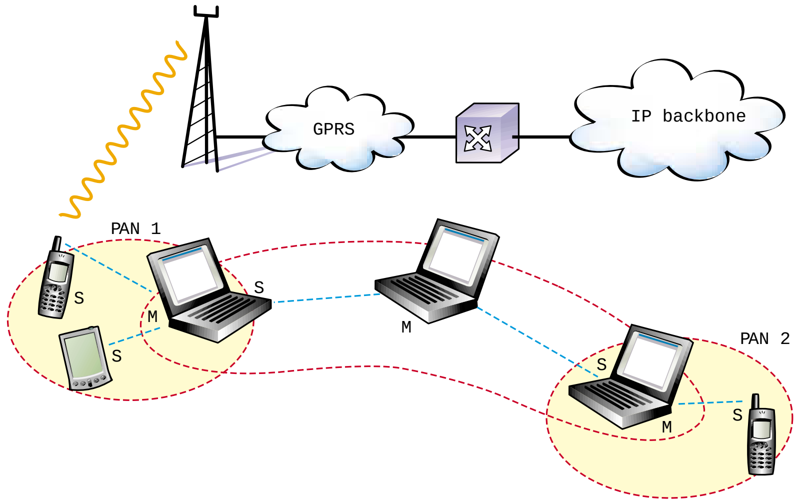

Illustration: Claes-Göran Andersson

Figure 5 The master of PAN 1 forward packets from the GPRS phone to the rest of the notebook computers so as to give internet access to the entire scatternet.

The scatternet functionality may also be used to improve the performance of a group of Bluetooth nodes that either already are part of a scatternet, or are part of separate piconets. The roles of the nodes in such a group may be rearranged to adapt to a new traffic distribution among the nodes by changing the allocation of masters, slaves and inter-piconet nodes. For instance, if two slave nodes need to communicate, it might be wiser to create a new piconet that solely contains these two nodes (see Figure 6). The nodes can still be part of their original piconets if traffic flows to or from them, or if they need to receive control information. Since the frequency-hopping spread-spectrum (FHSS) system combined with the fast ARQ makes Bluetooth very robust against interference, new piconets gain substantially more capacity than they lose as a result of increased interference between them.

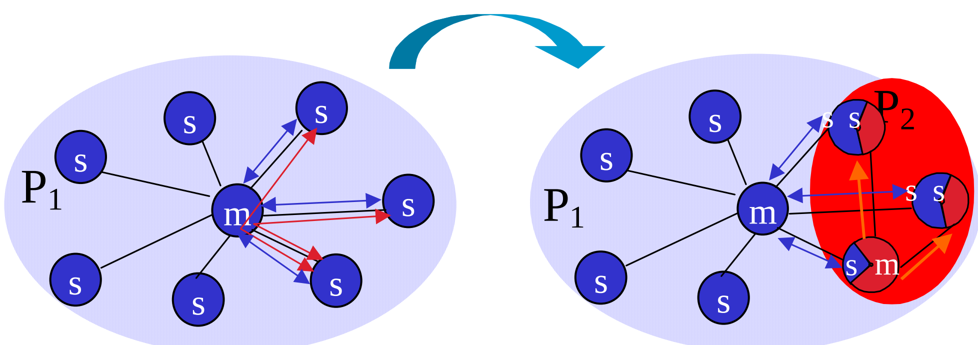

Figure 6 The new piconet (P2) adds more capacity for the slave-to-slave traffic in P2 and also for the traffic remaining in P1.

3.3. Packet forwarding in the scatternet

Packet forwarding—or routing—becomes necessary when packets must traverse multiple hops between the source and destination nodes. Given that IP will be commonplace in scatternet contexts, one might conclude that routing over the scatternet should be handled within the IP layer. However, there are good arguments for taking another course in this respect.

- The current IP dynamic host configuration protocols [13] (DHCP) and the emerging zero-configuration methods [13,14] (IETF Zero Configuration Networking Working Group, zeroconfig) rely on link layer connectivity. These protocols are typically used to attain a dynamic IP address for an IP host or to select a random IP address [14]; both operations that the devices of a PAN most likely will need to perform. Generally, the protocols will not work beyond an IP router, which means that they will not reach nodes located more than one Bluetooth hop away in an IP-routed scatternet. A scatternet that provides a broadcast segment-like connectivity would enable these protocols to work for Bluetooth-based IP hosts that are separated by multiple hops.

- To operate efficiently, the routing function should be joined with the function for forming scatternets. A routing function on the IP layer would thus need to be adapted to, or interact very closely with, the underlying Bluetooth layer, which violates the idea of keeping the IP layer as independent as possible from the link layer technology.

- Other, non-IP based, applications may use the scatternet functionality provided by the Bluetooth networking layer, i.e. the Bluetooth devices are not forced to host an IP layer to utilize the networking features of the scatternet.

In summary, the best way of providing networking among the nodes in a Bluetooth scatternet would be to perform the routing on a Bluetooth network layer residing below IP. This approach is also pursued within the Bluetooth SIG PAN working group, where a networking protocol, referred to as the Bluetooth Network Encapsulation Protocol (BNEP), is being developed to provide an Ethernet like interface to IP. In Figure 7 BNEP provides a broadcast segment across a scatternet, consisting of two piconets, which enables ordinary IP hosts to be interconnected in a Bluetooth ad hoc scatternet. The scenario in Figure 7 may be a case where several PANs are involved in an ad hoc meeting. Note that the piconets and PANs do not necessarily have to be mapped one to one, i.e. depending on the best topology a several PANs may be in one piconet or a PAN may consist of multiple piconets (a scatternet). This independence allows for efficient and flexible ad hoc PANs since rearrangements may be done without changing the PAN memberships.

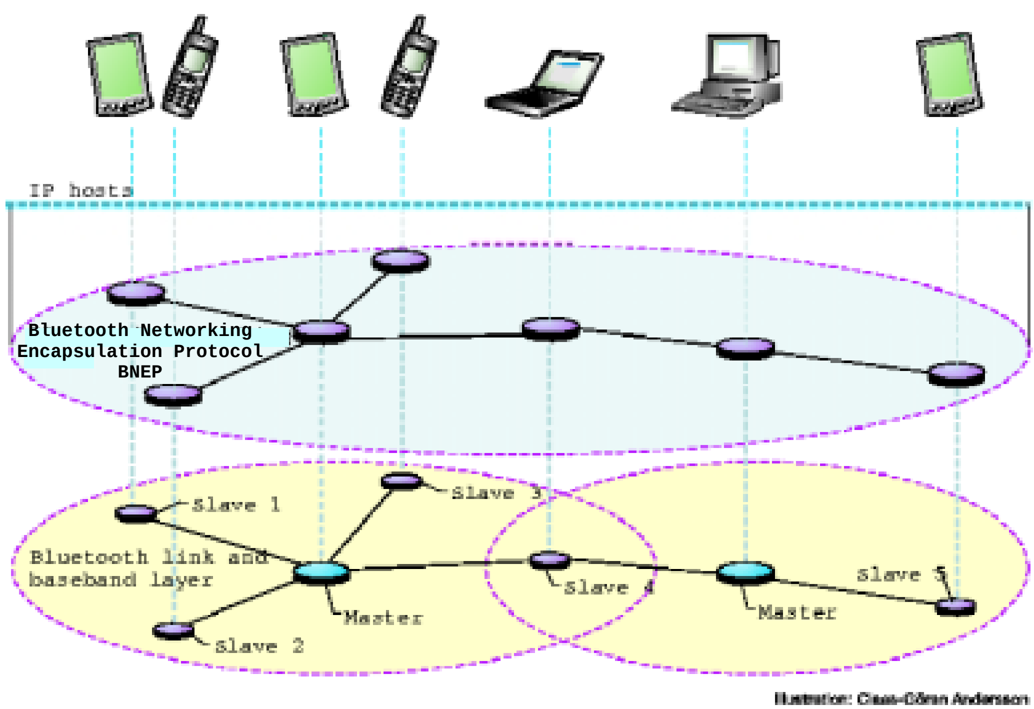

Figure 7 The Bluetooth Networking Encapsulation Protocol will offer a broadcast segment infrastructure to the IP layer, similar to Ethernet, potentially spanning over an entire scatternet.

In the first release of the BNEP, the focus is to provide the broadcast segment within one piconet only, but BNEP hosts the potential to offer a scatternet wide broadcast segment. From a protocol point of view, BNEP will be placed on top of the link layer protocol of Bluetooth, the L2CAP (Logical Link and Control Adaptation Protocol). L2CAP comprises functions for segmentation and reassembling between packets from higher layer protocols and Bluetooth baseband packets. Multiplexing of protocols apart from BNEP (e.g. OBEX, RFCOMM) over the Bluetooth radio link is also made possible in L2CAP. However, L2CAP is strictly link oriented and cannot forward packets beyond one link since there is no concept of network wide addressing within L2CAP. BNEP will add the networking functionality by utilizing the unique MAC address of each Bluetooth interface. The location of BNEP in the Bluetooth IP stack is depicted in Figure 8.

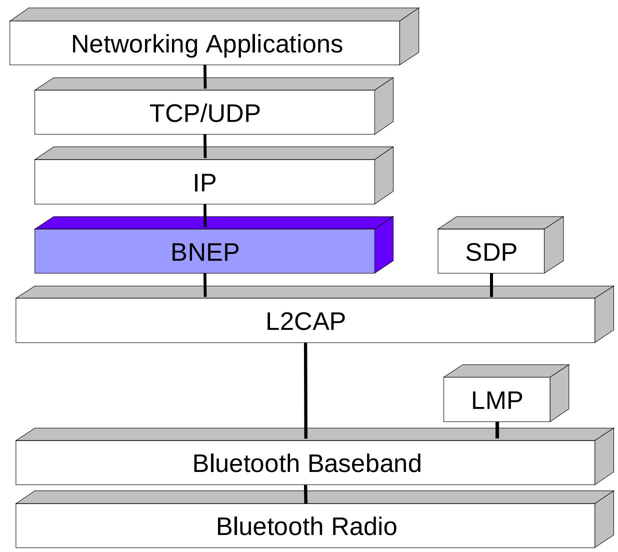

Figure 8. Location of BNEP in the Bluetooth IP stack.

For many IP applications BNEP will replace the point-to-point approach which currently is the only supported IP stack within Bluetooth, i.e. via PPP over RFCOMM. The point-to-point approach may still be used for point-to-point oriented applications such as dial-up via a Bluetooth enabled mobile phone. However, the introduction of a broadcast support for IP will enable easy interworking with protocols that are commonplace in the LAN/WLAN environment and also bring the same networking concepts into the Bluetooth PAN domain. Similar to Ethernet, BNEP may also be used by other protocols than IP, thus making the scatternet capabilities available for more than just IP based applications. Moreover, BNEP will be able to interact closely with the Bluetooth baseband functions during the establishment or teardown of piconets, which will facilitate the formation of efficient Bluetooth scatternets.

3.3.1. Scatternet Routing

Given the approach that IP will be serviced by a broadcast segment provided by BNEP, the issue arises for BNEP to forward packets between source and destination across a scatternet. However, based on its ad hoc nature the Bluetooth PAN, or network of such PANs, can be seen as a mobile ad hoc network (MANET). This means that results from research on ad hoc routing protocols may be used also for scatternet routing. Several proposed routing protocols for ad hoc IP networks, e.g. [3,4,5], have been submitted to the MANET working group [1] within the IETF. Even though these protocols are intended for IP routing the algorithms can in general be applied to any network where the task is to establish a path between any two nodes with unique addresses. In the case of a scatternet, the 48-bit Bluetooth Device

Address (BD_ADDR) may be used as the unique identifier of the scatternet nodes. Thus, an ad hoc routing protocol can use these in the same way as the IP addresses are used in the IP versions of the routing algorithms. The capability to aggregate the IP address to allow scalable IP routing protocols can in most cases never be used in a mobile ad hoc network since the nodes may move randomly relative to each other. That is, the IP address is purely used as an identifier of the node (or the interface of the node).

The inherent node mobility in an ad hoc network may cause a path, considered optimal at one time instant, to not work at all a few moments later. Moreover, the stochastic properties of the wireless channels and operating environment add to the uncertainty of path quality. Traditional routing protocols in the Internet are proactive in that they maintain routes to all nodes, including nodes to which no packets are being sent. They react to any change in the topology even if no traffic is affected by the change, which requires periodic control messages to maintain routes to every node in the network. The rate at which these control messages are sent must reflect the dynamics of the network in order to maintain valid routes (see e.g. [2]). Thus, scarce resources such as power and link bandwidth will be used more frequently for control traffic as node mobility increases. An alternative approach involves establishing reactive routes, which dictates that routes between nodes are determined solely when they are explicitly needed to route packets. This prevents the nodes from updating every possible route in the network, and instead allows them to focus either on routes that are being used, or on routes that are in the process of being set up.

Applied to the scatternet (e.g. as part of BNEP), a reactive ad hoc routing algorithm would also have to consider establishment of new Bluetooth connections in order to find efficient paths through the scatternet. Hence, forming and re-arrangements of the scatternet should be a part or closely integrated with the scatternet routing algorithm. If the routing had been purely IP based, this integration would be more difficult to achieve since the scatternet is then seen as a link layer and routing initiated link establishments would need to be indirectly triggered through a message exchange over the Bluetooth host controller interface (HCI). Otherwise, the IP implementation for Bluetooth scatternets need to be modified to contain Bluetooth specific elements and this is in general not a preferred approach to IP networking.

3.4. Scatternet forming

In order to have an efficient infrastructure for IP networking on Bluetooth, piconets and scatternets must be able to adapt to the connectivity, traffic distribution, and node mobility in the network. This is mainly achieved by setting up new piconets or terminating others, to attain the optimal scatternet topology. In this context, optimal refers to a scatternet that, for instance, yields minimum delay or maximum throughput. But it could also mean minimizing energy consumption in network nodes. Obviously the scatternet routing protocol should have an impact on how the scatternet is formed. If a reactive ad hoc routing approach is used, the search for new paths (routes) will be one important way to adapt the scatternet topology to the current traffic distribution. Consequently, the routing protocol needs to consider what the optimal formation objective is (e.g. best throughput per energy unit).

Apart from the (reactive) routing oriented scatternet forming operations, formation functions are also needed to establish a more generic connectivity in the scatternet in order to be able to find nodes searched for by the routing protocol. Typically these functions operate in the background to include new nodes or re-arrange old nodes in the scatternet and will use the Bluetooth connection establishment functions (INQUIRY and PAGE) on a periodic basis. In [18] one of the first known a simulation studies on scatternet forming aspects is presented. The study showed that the total number of Bluetooth links and the overhead imposed in conjunction with inter-piconet forwarding have a large impact on the system performance.

3.5. Intra- and Inter-piconet Scheduling

The master unit of a piconet controls the traffic within the piconet by means of polling the slaves according to any preferred algorithm, e.g. Round Robin, which determines how the bandwidth capacity will be distributed among the slave units. The polling of slaves within a piconet results in scheduling of the slaves in the master unit, which is referred to as intra-piconet scheduling (IRPS). The IRPS function in the master should assess the capacity needs of the units in the scatternet so as to ensure that capacity is shared fairly, or according to any other preferred capacity-sharing policy. In [12] an efficient intra-piconet scheduling algorithm denoted Fair Exhaustive Polling (FEP) was introduced and was shown to provide both high utilization and fair distribution of the piconet capacity among the participating slaves. The main contribution of the algorithm is that it focuses the capacity on the active nodes with maintained fairness, while still being very simple from an implementation perspective.

In a scatternet, at least one Bluetooth unit is member of more than one piconet. These inter-piconet nodes might have a slave role in numerous piconets but can have the master role in only one of them. Irrespective of the roles, the inter-piconet node must schedule its presence in all the piconets it is a member of, hence, an inter-piconet scheduling (IPS) algorithm is necessary in addition to the intra-scheduling algorithm. The main challenge for the inter-piconet scheduler is to schedule the presence of the inter-piconet node in its different piconets in such a way that the traffic can flow within and between the piconets as efficiently as possible. Given that the inter-piconet node is a single transceiver unit, only one of its entities (master or slaves) can be active at a time, resulting in that the node is blind in all other piconets except for the one it is active in for any given moment. This means for any data that needs to be exchanged with an inter-piconet node in a particular piconet it must be done while it is present in that piconet.

The master unit in a piconet always expects a slave unit to be present when it sends a data or poll packet to it. If the slave unit is not present, the master may choose to disconnect the slave after some predefined time-out period since it may be seen as a case of an erroneous radio channel (e.g. the slave unit is out of range or the level of interference is too high). However, in a scatternet, the slave may be an inter-piconet node and may be visiting another piconet when the master sends it a packet. Thus, the main issue with inter-piconet scheduling is to coordinate the simultaneous presence of an inter-piconet unit and a master (this node may be a slave if the inter-piconet unit itself has a master role).

Moreover, the time-slots of the piconet channels are, in general, not synchronized, which introduces a delay of one or two slots (0.625 ms each) each time a unit switches between piconets. This results in a piconet switching overhead that should be taken into account in the design of the IPS algorithm.

In addition, the combination of IPS and IRPS schedulers in the scatternet should be coordinated in order to give an efficient scheduling of the units of the scatternet. For instance, an inter-piconet node may be removed from the scheduling list of the intra-piconet scheduler for the period it is not present, but re-installed in the list for more frequent scheduling once it is back in the piconet. In Figure 9, a scatternet consisting of two piconets with one inter-piconet unit (double slave role) is depicted. The two time lines in the figure show how the time division multiplexing for the inter-piconet unit needs to take the phase shift into account when switching between the two piconets. When the inter-piconet unit is active in one of the piconets, the IRPS algorithm of that piconet determines the amount of polling received by it.

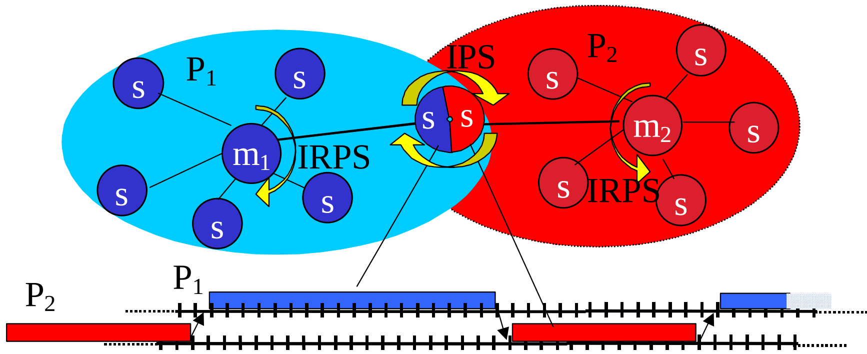

Figure 9 A scatternet with one inter-piconet unit that divides its time between the two piconets.

4. Functional Architecture for Scatternet Scheduling

In its current version, the Bluetooth specification contains very limited information on inter-piconet scheduling in general and does not mention how inter-piconet scheduling should fit into the overall Bluetooth architecture. This section proposes a generic architecture for scheduling functions within a scatternet that will allow a wide variety of different scheduling algorithms to co-exist. Figure 10 gives a proposed functional architecture of how an inter-piconet scheduling function (IPSF) would interact with other functions in a generic Bluetooth unit. A subset of these functions is active in a Bluetooth unit, depending on the role that the unit has in the scatternet.

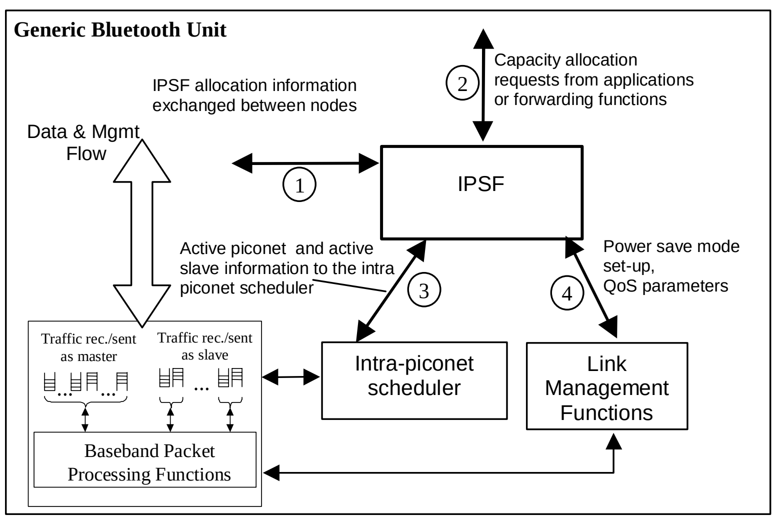

Figure 10. Functional architecture of the interpiconet scheduling function (IPSF) and its relation to other functions in a generic Bluetooth unit.

The IPSF will have the following relation to other functions in the Bluetooth unit (the list relates to the corresponding numbers in Figure 10):

1. The IPSF will need to exchange coordination information with a peer IPSF unit in some other node. This information may contain the points in time when a master and an inter-piconet device will meet and may also involve an information exchange for negotiation of these points.

2. The IPSF may receive information on requirements posed by applications, such as delay constraints etc. For an inter-piconet unit, the "duration of its visit" in a piconet and the "time elapsed" between

such visits are parameters that determine the capacity and delay to the piconet. Thus, the IPSF may try to negotiate these parameters in accordance with the application requirements.

3. The interaction of the IPSF with the intra-piconet scheduler is only active when the device is a master. The IPSF should inform the intra-piconet scheduler about the presence of inter-piconet slave devices, so that the intra-piconet scheduler may appropriately "poll" them.

4. The Link Manager describes functions to handle establishment, security and control of the radio link within a piconet. It also defines a set of power save modes (SNIFF, HOLD, PARK) that may be used by the IPSF to implement coordination between inter-piconet units. Moreover, new releases of the Link Manager may host specialized inter-piconet functions that the IPSF may invoke.

4.1. Rendezvous Point IPS Algorithms

The functional architecture for scatternet scheduling outlined in the previous section allows for a wide variety of ways to do inter-piconet scheduling (IPS). In this section, we discuss a family of algorithms that we refer to as "Rendezvous Point" algorithms. A "Rendezvous Point" is a slot at which a master and an inter-piconet unit have decided to meet, i.e., at this slot, the master has agreed to address a packet to the inter-piconet unit and the inter-piconet unit has agreed to listen to the master.

Two major issues that need to be addressed by a Rendezvous Point inter-piconet algorithm are:

- The rendezvous point (RP) issue: how do the master and slave decide on the RP and how strict is the commitment to this RP? In a strict algorithm, the master and slave units will always honor a RP. In fact, a wide variety of rules on honoring RPs may be defined. The distribution of RPs over time could be periodic, decided upon each visit, or spread out in a pseudo-random sequence (known by both nodes). The delay property of an inter-piconet algorithm will, to a great extent, depend on the time between mutually honored rendezvous points.

- The rendezvous window (RW) issue: given that the master and slave units are both present at a RP, how much data will they be able to exchange? This depends upon the duration of the period that the units are mutually present and how much of this time is used to send data. In a strict algorithm, a time window, called a rendezvous window, could be defined in which both master and gateway must be present and exchange data in every available slot. A relaxed algorithm would not require any of the master or slave units to stay any pre-determined time. The achieved throughput for the master slave pair will depend on the duration of these rendezvous windows.

Based on how the issues listed above are approached, categories of IPS algorithms of the RP family can be identified and below a list of some of these categories are given. Note that the list is not exhaustive, but includes the most feasible ones.

1. Honoring-Periodic Static-Window (HPSW): In this category, units always honor RPs. Also, RPs occur periodically, which gives a constant period between RPs for a particular pair of units, henceforth referred to as a superframe. The size of the RW for a piconet is static and remains the same throughout the duration of a connection.

2. Honoring-Periodic Dynamic-Window (HPDW): This algorithm category always honors RPs that are distributed periodically. The RW is dynamic in order to adapt to both topology changes as well as traffic dynamics. For instance, if the RW is defined to be the time between RPs, the RW size will change if a new inter-piconet node joins a piconet or an existing one leaves.

3. Honoring-Random Static-Window (HRSW): The RPs are always honored but unlike the previous schemes, they are spread out according to a pseudo-random pattern known both to the master and the inter-piconet node. Due to the random spreading, RWs may end up quite close to each other. The size

of the RW is in this case static and may give a limit on the number of piconets the inter-piconet node may take part in.

4. Master-Honoring Dynamic-Window (MHDW): In this category the master unit always honors the RPs but the (slave) inter-piconet unit may skip an RP in order to give priority to another piconet. The RPs may be distributed periodically as in HPSW or according to a pseudo-random sequence as in HRSW. Since the inter-piconet unit may not honor RPs, the RW size may change to adapt to new traffic conditions or topology changes.

Generally, the always honoring (HPSW, HPDW, HRSW) IPS categories will give a more strict scheduling that at on the one hand enables better traffic delay guarantees but on the other hand is less flexible to change and less adaptive to traffic. When an HPDW algorithm wants to change the RVW size, it may need to reallocate some RPs as well, which requires the exchange of control information causing a penalty in terms of overhead. The non-always honoring category (MHDW) will potentially result in less delay guarantees but will enable traffic adaptivity without requiring exchange of control information.

Simulation studies of the various IPS algorithms are ongoing within several research groups both from academia and industry, e.g. [19], but so far no results have been presented in any comparative form to position one algorithm in favor of another. However, several studies of IPS algorithms are expected to be published in the near future.

5. Conclusions

Ad hoc networks have mostly been used in the military sector, where being able to establish ad hoc communication is often a necessity. On the other hand, in the commercial sector, successful examples of ad hoc radio networks are few so far, if any. However, instead of large-scale networks, small-scale personal area networks are emerging in response to the introduction of short-range radio technologies, such as Bluetooth. Here, ease of use and flexibility are fueling the demand for ad hoc operation. In addition, a centralized network architecture would have serious problems trying to control all PAN devices. In particular, ad hoc Bluetooth networks—scatternets—will give rise to a whole new set of business and consumer applications for small, battery-driven user devices, such as mobile phones, PDAs, and notebook computers. The combination of wide-area IP connectivity via 3G (mobile phone) access, and personal area connectivity in the PAN presents new opportunities for the user on the go. End-to-end IP networking is a key component in this respect, providing the basis on which to develop applications for PAN products. Thus, the current development of IP support for Bluetooth networks in general and for 3G access of Bluetooth PAN in particular, is crucial.

The current work in the PAN WG of the Bluetooth SIG focuses on developing IP support based on a network layer on OSI level 2, creating a broadcast segment similar to Ethernet. This will enable a straightforward reuse of a number of IP related protocols for configuration and address resolution (such as DHCP, ARP etc.) typically suited for LAN access environments, but also for standalone or interconnected PANs. Using the Bluetooth PAN protocol will, seen from upper protocol levels, be very similar to connecting devices together on an Ethernet segment. As an interesting parallel, one may consider how mobile phones would support laptops or other mobile devices to access the phone, or the Internet through the phone, if the phone was equipped with a physical Ethernet adapter.

The use of scatternets in Bluetooth networks introduces inter-piconet gateways, which need to be scheduled between the piconets they are members of. The choice of inter-piconet scheduling (IPS) scheme is crucial for the overall performance of the scatternet since it controls the efficiency of Bluetooth packet forwarding between piconets and, moreover, it affects the performance within the piconets the gateways are residing in. An overall architecture for scheduling in scatternets was outlined herein to capture how the various Bluetooth functions will interact to achieve an efficient system. Moreover, a family of IPS algorithms, referred to as Rendezvous Point algorithms, was introduced as a feasible way to design IPS algorithms.

The Bluetooth wireless technology in PANs will most likely change the way we handle and access information in the near future. A similar development during the past ten years can be observed in the way the mobile phone has changed our behavior in terms of information vs. independence of location. The actual impact of Bluetooth in general, and its use in PANs in particular, are of course yet mostly speculations. However, it is the first realistic attempt in large scale to solve the last meter problem we often encounter – simply to get our personal devices to "talk" to each other without it being a hassle. Thus, this carried PAN, put in combination with the emerging 3rd generation mobile systems, provides with an intriguing entry for the Internet into more than one of our pockets.

6. REFERENCES

1 Mobile Ad hoc Networks (MANET). URL: http://www.ietf.org/html.charters/manet-charter.html. (2000-05-28). Work in progress.

2 P. Johansson, T. Larsson, N. Hedman, B. Mielczarek, M. Degermark. "Scenario-based Performance Analysis of Routing Protocols for Mobile Ad hoc Networks", In Proceedings of the Fifth Annual International Conference on Mobile Computing and Networking, August 1999.

3 C. E. Perkins, "Ad Hoc On Demand Distance Vector (AODV) Routing." Internet draft, draft-ietf-manet-aodv-02.txt, November 1998. Work in progress.

4 J. Broch, D. B. Johnsson, D. A. Maltz, "The Dynamic Source Routing Protocol for Mobile Ad hoc networks." Internet Draft, draft-ietf-manet-dsr-00.txt, March 1998. Work in progress.

5 C. E. Perkins, P. Bhagwat, "Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers." In Proceedings of the SIGCOM '94 Conference on Communications Architecture, protocols and Applications, pages 234-244, August 1994. A revised version of the paper is available from http://www.cs.umd.edu/projects/mcml/papers/Sigcomm94.ps. (1998-11-29).

6 J. Haartsen, "Bluetooth – The universal radio interface for ad hoc, wireless connectivity", Ericsson Review Vol. 75(1998):3, pp. 110-117.

7 J. Hartsen, M. Naghshineh, J. Inouye, O. J. Joeressen, W. Allen, "Bluetooth: Visions, goals, and architecture," ACM Mobile Computing and Communications Review, pp. 38-45, No. 2, Vol. 4, 1998.

8 J. Haartsen, "The Bluetooth Radio System", IEEE Personal Communications Volume: 7 1, February 2000 , Page(s): 28 –36.

9 The Bluetooth Special Interest Group, "Baseband Specification", http://www.bluetooth.com.

10 P. Johansson et al. "Short-Range Radio-Based Ad hoc Networking: Performance and Properties", ICC'99, Vancouver, 1999.

11 P. Bhagwat, "Bluetooth: Technology for Short-Range Wireless Apps", IEEE Internet Computing, vol 5, Number 3, May/June 2001. Available from http://www.computer.org/internet/ic2001/w3toc.htm

12 N. Johansson, U. Korner, P. Johansson, "Performance Evaluation of Scheduling Algorithms for Bluetooth", In Broadband Communications: Convergence of Network Technologies, Edited by Danny H. K. Tsang and Paul J. Kuhn, Kluwer Academic Publishers, pp 139-150, 2000.

13 R. Droms, "RFC 2131: Dynamic Host Configuration Protocol," March 1997, available from http://www.ietf.org/rfc/rfc2131.txt

14 IETF, Charter of the Zero Configuration Networking (zeroconf) WG, Available from http://www.ietf.org/html.charters/zeroconf-charter.html.

15 S. Cheshire, “Dynamic Configuration of IPv4 link-local addresses,” Internet Draft, 8th October 2000, Available from http://www.ietf.org/internet-drafts/draft-ietf-zeroconf-ipv4-linklocal-00.txt

16 IEEE Computer Society LAN MAN Standards Committee, “Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” IEEE Std 802.11-1997. The Institute of Electrical and Electronics Engineers, New York.

17 Khun-Jush, J., Malmgren, G., Schramm, P. and Torsner, J.: HIPERLAN type 2 for broadband wireless communication. Ericsson Review Vol. 77(2000):2, pp.108-119.

18 G. Miklos, A. Racz, Z. Turanyi, A. Valko, P. Johansson, “Performance Aspects of Bluetooth Scatternet Formation”, MobiHoc 2000, Boston MA, August 11, 2000.

19 WAM (Wireless Adaptive Mobility) Lab, http://www.cs.ucla.edu/NRL/wireless/, available June 18, 2001.