# IEEE P802.15
# Wireless Personal Area Networks

| | |
|---|---|
| Project | IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs) |
| Title | **Managing the performance of ad hoc mesh networks** |
| Date Submitted | [29, April 2004] |
| Source | [Francis daCosta]<br>[MeshDynamics]<br>[1299 Parkmoor Ave, San Jose, CA 95126]<br>Voice: [(408) 373-7700]<br>Fax: [(408) 516-8987]<br>E-mail: [fdacosta@meshdynamics.com] |
| Re: | [802.15 Study Group 5 mesh development.] |
| Abstract | [An approach based on the concepts of toll cost and hop cost are applied to dynamic modification of routing paths to meet application requirements for controlled latency and throughput.] |
| Purpose | [To present an approach for QoS in Ad hoc Mesh Networks.] |
| Notice | This document has been prepared to assist the IEEE P802.15. It is offered as a basis for discussion and is not binding on the contributing individual(s) or organization(s). The material in this document is subject to change in form and content after further study. The contributor(s) reserve(s) the right to add, amend or withdraw material contained herein. |

## CHALLENGES FOR 802.15 WPAN MESH

Figure 1. An ad-hoc Mesh Networking application for home networking

The needs for mesh networking in Personal Area Networks (PAN), are different than mesh networking in local area (LAN) 802.11 networks. In 802.11, the preferred application for mesh is with APs that form the nodes of the mesh and are stationary. Clients attach to the APs and are not part of the mesh network. This is referred to an infrastructure mesh and described in a companion white paper.[1]

Conversely, in 802.15 PANs, the clients are both consumers and nodes of an ad hoc mesh network. The "routers" are therefore mobile and this affects both the reliability and stability (?) of the network. There are also a host of 802.15 PAN issues related to beacon alignment, CTA[2], ACI[3], etc. that are more amenable to being solved with a mesh control layer than by other means. Thus, the need for a mesh control layer extends beyond the ability to provide coverage and range for 802.11 applications, and into the domain of network performance and control.

Additionally, 802.15 PANs focus more on multimedia applications that require enhanced QoS, as in control of the latency and throughput of the transmission more so than is needed in 802.11 data only networks. These concerns require a QoS aware approach to mesh routing.

Finally, WPAN devices may be battery operated. It may be more economical to use an intermediary to send packets to a destination node, even if that destination node is within range of the sending node. The mesh routing must support proactive energy management.

## DISTRIBUTED CONTROL LAYER

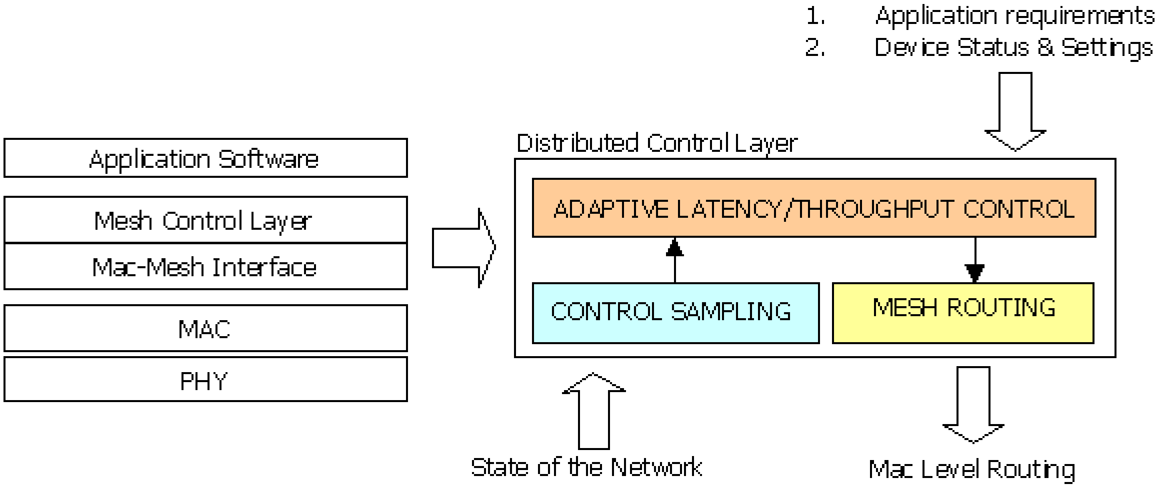

Figure 2: A distributed control layer addresses mediation issues for mesh networking

Managing the dynamics of wireless mesh networks requires a distributed control system approach - to judge network conditions correctly and adjust both the traffic flow and the overall topology of the network in order to meet the differing application requirements of multiple clients being serviced by the same AP. A distributed software control layer approach

---

[1] www.meshdynamics.com/Publications/MDWMANOVERVIEW.pdf
[2] Channel Time Allocation
[3] Adjacent Channel Interference

addresses dynamic control of the network topology ensuring consistent performance despite:

- Client mobility resulting in interference between Beacons
- Proactive power management resulting in changes in routing
- Changes in application requirements for latency/throughput

## ASSURING QUALITY OF SERVICE

In an ad-hoc mesh network, each node of the network manages its connectivity graph related to the routing path to every other node in the network. This implies that if there are 10 nodes in the network, each of the 10 nodes maintains a routing table for 9 other nodes. Figure 4 shows the routing paths selected for two nodes of the network. The boxes are the elements of the mesh, and the solid blue lines show how the highlighted device (the one with the heavy border) will connect to any other device in the mesh.

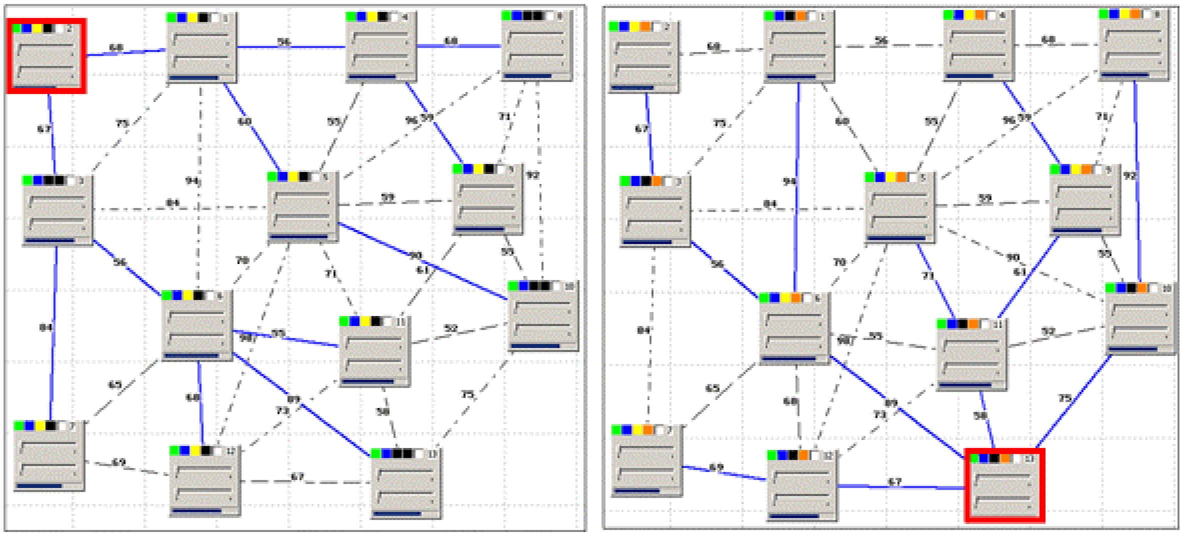

Figure 4: The routing paths that two different nodes will take to connect to other nodes

Since all notes are mobile, these routing paths are constantly changing for every node in the network. As the number of nodes in the network increases, so does the amount of routing information – by the square of the number of nodes.

To limit the amount of information that has to travel across all nodes of the network, many mesh routing algorithms focus on local information e.g. link state or distance vector. The selection of the “best” node to use in a multiple hop routing path is made based on which node closest to the sender has stronger signal strength or is fewer hops away.

While effective in smaller, stationary networks, this approach results in sub-optimal routing in large and/or dynamic networks since all nodes are making local decisions and no node has the “big” picture. This quick and dirty approach does not assure Quality of Service (QoS) for mission critical data.

A global approach is more QoS aware but requires that all nodes share information with all other nodes. A compact and efficient means of updating routing tables now becomes essential as does the speed with which the information is broadcast over the network. Changes are propagated through the network to reach a node far away, through a series of hops. Until the update reaches the node, it is basing decisions on obsolete data.

To address this, we have developed patent pending techniques we refer to as **heartbeats**. Heartbeats, sent periodically by every node, provide information needed by other nodes to make decisions related to routing. In additions to link state and distance vector, these heartbeats also include toll cost and hop cost information, and beacon alignment data. These are described below.

## TOLL COST AND HOP COST

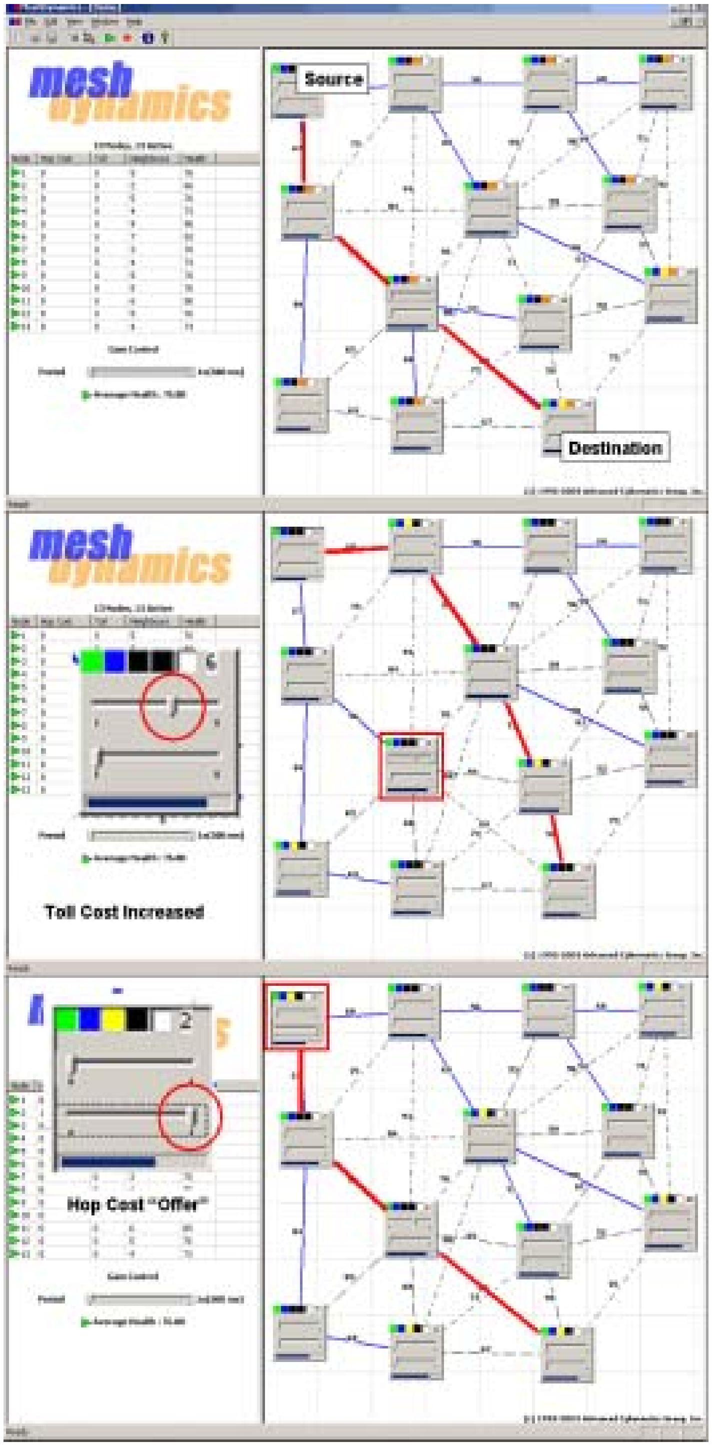

Figure 5: Toll Cost/Hop cost control for low latency QoS.

The heartbeat based mesh routing technique addresses the complexity of dynamic environments by modeling them as free market enterprises with consumers and providers of connectivity. To route a packet, a node must enlist the services of intermediaries that must cooperate to ensure that the packet reaches its destination in a timely manner. Quality of Service is driven by cooperation between intermediaries.

In some cases, a node may be used like a hub by its neighbors to relay packets. If that node needs to reduce that traffic load, for example to provide better service to the packets that it is generating and needing to send (recall from Figure 1 that each of these nodes is an end user device, as well as a participant in the mesh), that node – by virtue of the demand to use it – can raise its "price" to route data. We refer to this as the **toll cost** of using that node – as in the toll paid to cross a bridge.

With increasing toll cost, nodes with higher priority traffic get preferred treatment, since they are willing to pay for a lower latency path (path with fewer hops) – they are therefore willing to pay a higher hop cost.

With this free market approach, the system automatically adjusts its routing paths and maintains QoS despite local congestions (high toll cost). Low latency traffic moves along one set of paths, based on the ability of nodes to pay the **hop cost** for nodes charging a toll cost under congested node conditions. Lower priority traffic will move by less direct routes, thereby containing potential congestion at popular nodes.

Figure 5 (left) shows how the dynamics of this free market exchange ensure a low latency path for high priority traffic. The image on the top shows the default configuration for traffic flowing from the source node to the destination node, along the highlighted path, and traversing two intermediary nodes.

The image in the middle shows what happens when one of these intermediary nodes (highlighted) increases its toll cost. Traffic from the source to the destination will no longer use that node as an intermediary, finding a new, lower cost route, instead.

The image on the bottom shows what happens if the source node increases its hop cost, indicating a willingness to pay the higher toll cost. Traffic resumes its original routing path despite local congestions.

## PROACTIVE POWER MANAGEMENT

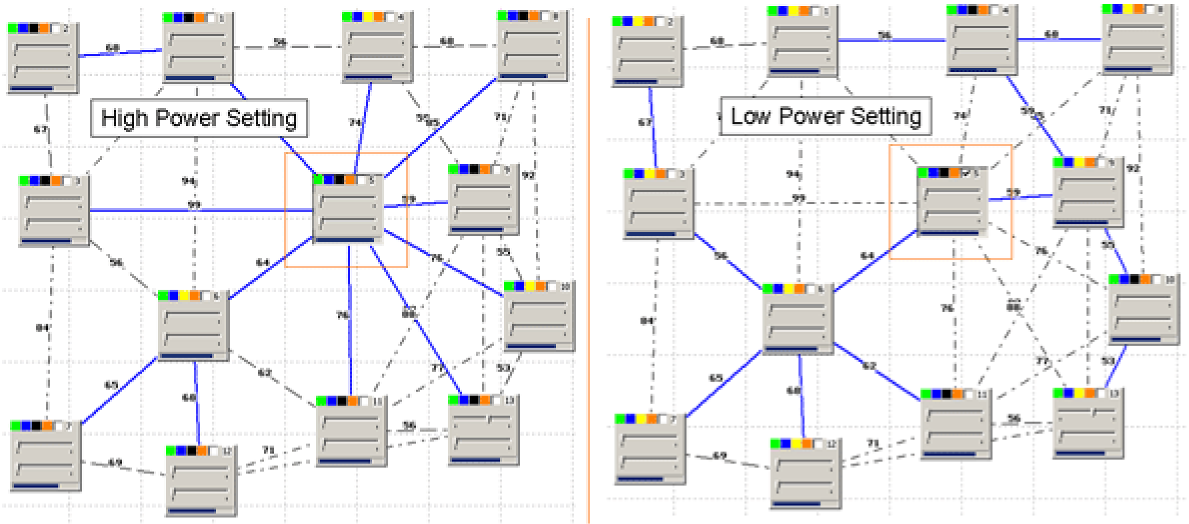

Demonstration of Proactive Mesh Routing for UWB devices in a mesh network, as power level decreases.

© Advanced Cybernetics Group, Inc. 1992-2004. All rights reserved.

Figure 6: Device being used as a hub on left broadcasts a low power setting on right

A key component of the adaptive control layer is to change mesh routing to reduce the power usage for battery operated devices. In Figure 6, the highlighted device in the image on the left is being used to provide routing services to other devices on the network. This limits its ability to sleep. By changing it mode to a low power setting, fewer devices use it.

In effect, if an alternate path exists that provides equivalent quality of service (in terms of latency and throughput) as the path using the battery operated device, then other devices will shift proactively to those routing paths. If the battery operated device wishes to go into sleep mode, it sends a sleep mode message causing devices to communicate with it only to send packets intended for it. as the destination node.

## BEACON ALIGNMENT ISSUES

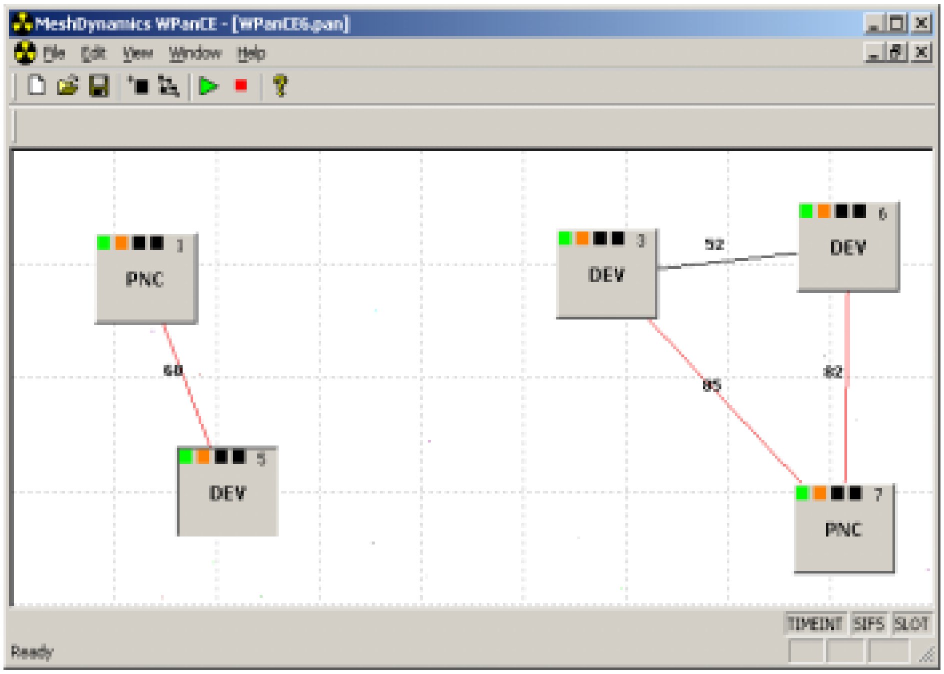

Fig. 7: Two Piconets operating - but not interfering - in the same area.

Fig. 7 shows two Pico Net Controllers (PNC) operating simultaneously in a Wireless Personal Area Network (PAN). In their current configuration, they are not interfering with each other's transmissions. However, since PAN devices support mobility, one could move into the airspace of the other. Even if it is temporary, this affects the transmission quality of both piconets. Problems associated with Simultaneously Operating Piconets (SOP) are endemic to home networking applications where both mobility and the ability to provide isochronous transmissions are core requirements.

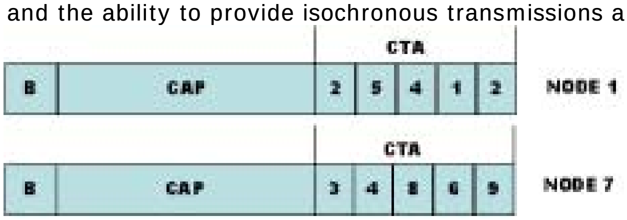

Fig. 8: Beacons are in interference with each other.

PAN devices use a beacon to ensure isochronous transmissions. In the IEEE 802.15.3 standard, as depicted in Fig. 8, CTA time slots are channel time allocation slots for regular transmissions of latency sensitive information such as video streams over a multihop network. A beacon synchronization pulse assures synchronicity. The PNC sends a beacon pulse marked B in Fig. 8. A device cannot effectively communicate if it loses the beacon synchronization pulse because of radio interference from devices in other Piconets. Such is the case when two piconets are sending beacons at the same time. Accordingly, there exists a need to coordinate between PNCs and their devices ensuring that beacons are sent at times when there is no interference from other devices.

Solutions to the Beacon interference problems are complicated by the fact that interference can occur at anytime- during the beaconing period (B), the Contention Access Period (CAP) or the CTA period. In the CAP, collision avoidance rules of CSMA will ensure eventual transmission – but with delay. However, interference in either the beaconing period or CTA period results in faulty transmissions. For high quality video, this is unacceptable.

## LOGICAL PICONET APPROACH

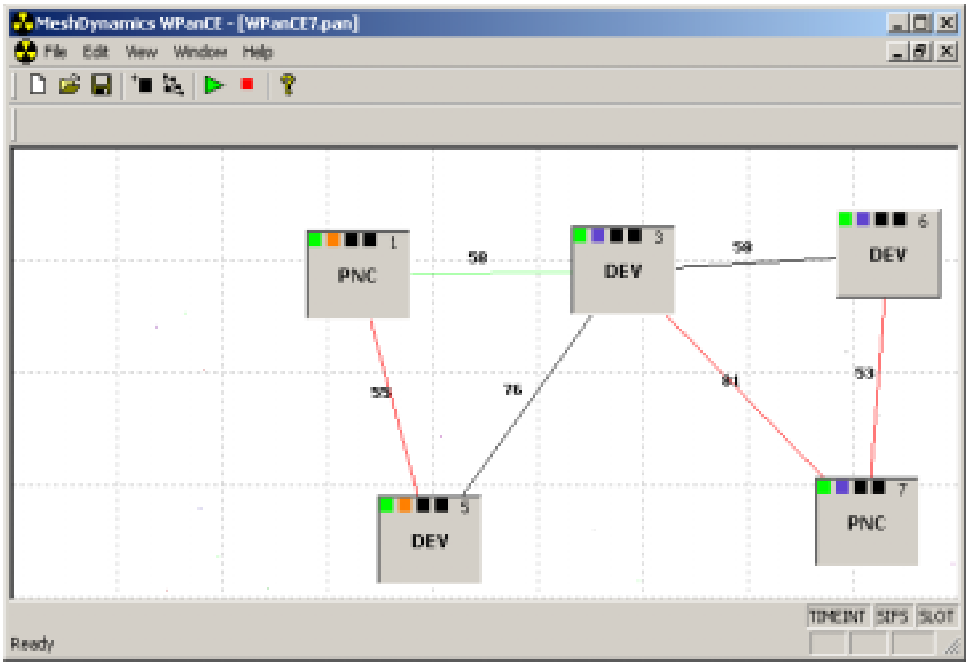

Fig. 9: A logical piconet is a grouping of piconets that support interference-free coexistence.

Fig. 9 depicts a change from Fig 9, when one piconet has moved closer to the other. Potential interference is avoided by aligning the beacon transmissions and channel time allocation (see Fig 10) ,which engenders interference free coexistence.

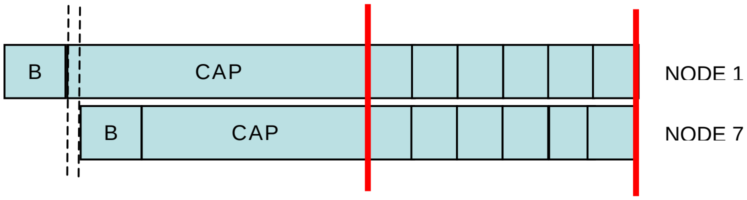

Fig. 10: Beaconing period staggered and CTA periods aligned to avoid interference.

The logical piconet is formed by mediation between the piconets that agree to support coexistence.
To address issues of mobility and requirements for isochronous transmission in Simultaneous Operating Piconets in a robust and scalable manner, we have developed and implemented software for managing beacon alignment issues in dynamic

environments. Our software only approach is layered on the MAC- no changes need to be made to the existing IEEE standard. Other driving factors that influenced our approach to beacon alignment and channel time allocation include:

***Robust Operation*** In addition to addressing beacon alignment and judicious channel time allocation issues, the system must be demonstrably robust, stable and scalable. Mobility is essential in personal area networks: both PNC and devices are moving and causing interference. Beacons could be affecting transmissions in a) beacon time slots b) CAP and c) CTA time slots. In all cases, the system must recover from changes to network topology swiftly and in a stable manner.

Interference effects may also be caused by external events e.g. the opening and shutting of doors changes the network topology suddenly. Algorithms that are proactive but not stable will generate unacceptable oscillatory behavior.

***Hidden Node Problem*** Fig. 11 is a snap shot from our simulation depicting a 4 node SOP problem. PNC nodes 1 and 5 cannot hear each other but device 2 hears both. Interference from the beacons sent by PNC nodes 1 and 5 can jam transmissions to Device 2 – but neither PNC is aware of the other.

The PNCs therefore need to be made aware of each other's existence for both beacon alignment and coordination of channel time allocations. In our approach, the intermediary Device 2 performs this function on hearing both PNCs transmit a "heartbeat". This is a retransmission of ASIE based information included in PNC beacons intended for use by other PNCs. The intermediary device also ensures that the handshaking between the two PNCs is performed in a stable manner.

***Mesh Networking*** The first applications of personal area networking, such as USB wire replacements, may not require mesh networking, but there is general consensus in the industry that mesh networking is needed to address the range limitations of UWB and in order to manage complex interactions between multiple piconets. Mesh networking is also desirable where transmit power control is required, especially for battery operated devices. With a mesh in place, devices may route, at lower power, through nearby devices, as opposed to using more power to reach devices directly.

## DYNAMIC BEACON ALIGNMENT

Referring to Fig. 11, Node 1 and Node 4 do not hear each other, but cannot transmit without restriction because Node 3 is in hearing distance from both of them. Conversely, Nodes 4 and 5 can transmit simultaneously as they do not have any common node in their "reachable" list of neighbors.

One approach to determining when beacons can be transmitted simultaneously is to apply set theory in determining if there is a NULL set of common reachable nodes. For example, Nodes 4 and 5 have no common nodes in their reachable list. Hence they can transmit at the same time as shown. We have implemented one such approach.

When PNC Nodes do have nodes in common in their reachable list then simultaneous transmission is not permissible, and beacon transmission must be staggered as shown in Fig. 11. Node 7 sends its beacon a short time after the beacon from Node 5 has concluded.

Set theory may aid us in determining which beacons need to be staggered. However, how do PNC nodes know that they need to stagger their beacons when they cannot hear each other? Consider the case for PNC node 7. It cannot hear any of the other PNC nodes. How can it determine what a safe beacon alignment slot should be? One means available is to help PNC node 7 align its beacon is through a re-transmission of beacon information by a device node (Node 6 in Fig. 11) that hears Node 7's beacon and apprises it of its surroundings. As such, this device acts as an intermediary or repeater node on behalf of the PNC node it belongs to. We refer to this re-transmission as a "heartbeat".

Heartbeats are thus re-transmissions sent by devices while PNCs sent beacons. Both heartbeats and beacons use ASIE information elements to relay information needed to perform beacon alignment. Appendix A shows the protocol used.

While devices generally send heartbeats on a periodic basis they also send a heart beat if they hear a request for information on its surroundings (e.g. "Where am I? Who are my neighbors?") When it hears such a request the device will relay the last beacon information it received from its PNC.

The beacon information, contained in ASIE information elements, informs the new PNC of other PNCs that the device's PNC knows of and it provides information on the current ordering of the PNC beacons. For example, Device 6 will relay the following information to Node 7:

- I am Device 6 and I belong to PNC node 1. I can also hear PNC Node 5.
- PNC node 1 belongs to a logical piconet- a family of aligned piconets.
- There are 3 piconets currently in this family and two beacon slots
- The beacon slot order of PNC node 1 is 1 of 2; That of PNC Nodes 4 and 5 is 2 of 2. (Node 7 has not yet joined).

Based in this information, Node 7 has limited options. It cannot take a used beacon slot such as 1 or 2 because device 6 can hear both node 1 and node 5 that use beacon slots 1 and 2 respectively. It must therefore request a new beacon slot, which when granted, places its beacon alignment position as shown in Fig. 11. For more information on our approach to dynamic beacon alignment please see: www.meshdynamics.com/Publications/MDPBEACONALIGNMENT.pdf

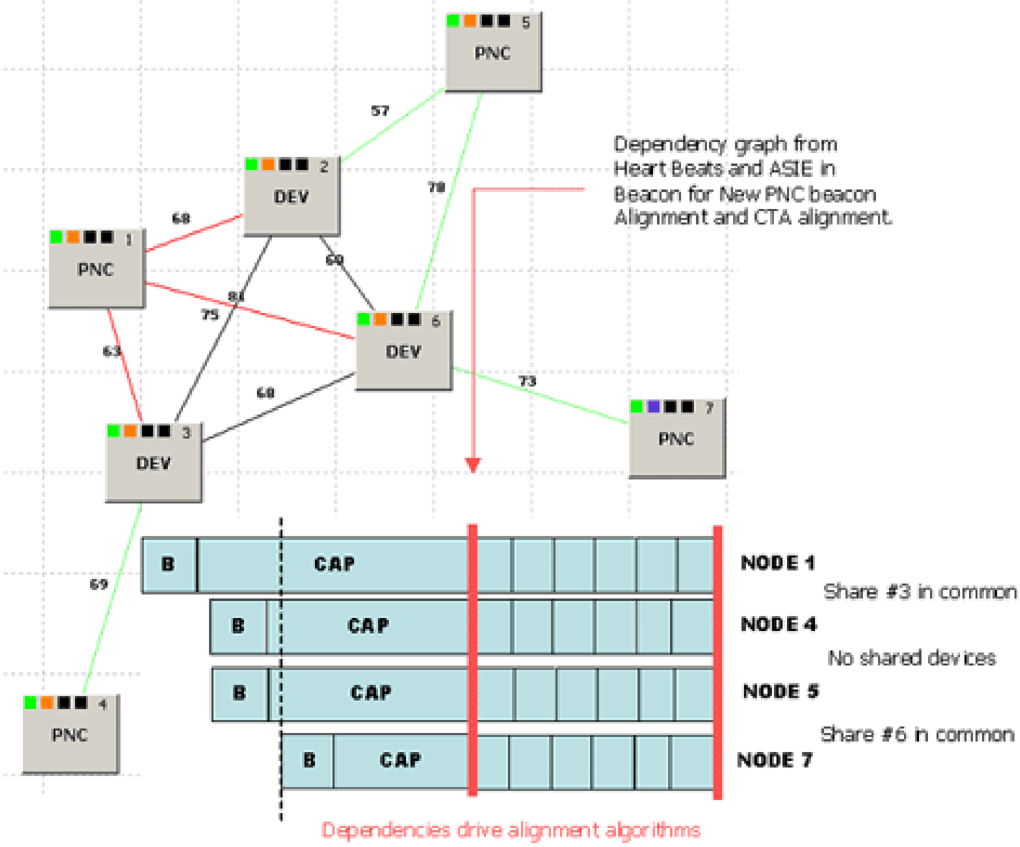

Fig. 11: Dynamic Beacon Alignment in Simultaneously Operating Piconets

## DYNAMIC CTA PERIOD ALLOCATION

Figs. 12 and 13 show two strategies for CAP alignment. Both strategies make the secondary PNC (node 7) begin its super-frame SIFS time units after the completion of the primary PNC's beacon. The SIFS wait ensures that Node 7 will get access to the medium before other devices - as they would normally wait for BIFS time units.

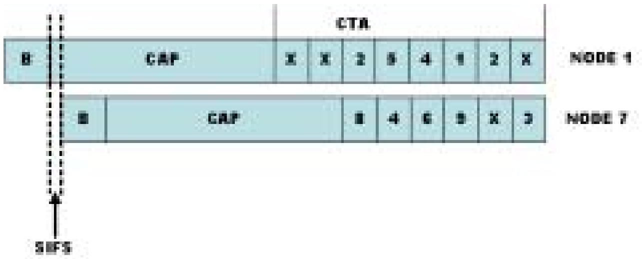

Fig. 12. Staggered but equal CAP for nodes 1,7

In Fig. 12, the CAP duration for both Nodes 1 and Node 7 is unchanged. Since the CAP of Node 7 interferes with CTA period of Node 1, some time slots may need to be marked private or be added to the CAP for node 1. After the completion of Node 7's CAP, both Node 1 and Node 7 begin their CTA Period.

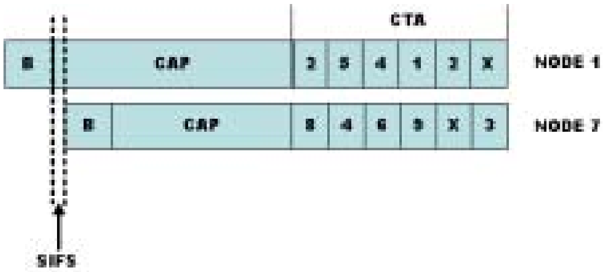

Fig. 13. Alignment of CTA Period for Nodes 1,7

In Fig. 13, the CAP duration for Node 7 is reduced so that its CAP end is aligned with Node 1's CAP end, after which both nodes begin their CTA periods. By the same token, Node 1 could have also increased its CAP duration so that its end is aligned with Node 7's CAP end. In this case Node 1 does not need to mark its first two CTA slots as private.

The two methods for CAP alignment depicted above are two approaches to allocating time slots ensuring that devices can talk with each other in specific time slots without stepping on each others toes. Both strategies need to be supported, based on the situation.

For example, if the CAP is not being used or there are many devices requiring the CTA allocations, this would favor reducing the CAP over overlapping CAP and CTA - which results in two slots in Node 1 becoming unusable. Conversely, if applications require more CAP than CTA, this drives the algorithms towards favoring Fig. 12 over Fig. 13.

The point being made is that the system must be adaptive: brute force algorithms are inappropriate. Consider the case where there are five Piconets in the same vicinity. If the choice is to align the CTA periods always it will result in progressive deterioration of CAP bandwidth for each PNC. The converse – as shown in Fig. 12, will result in reduced bandwidth for the CTA period. Clearly, neither approach is a one size fits all.

A patent pending approach developed and implemented by us provides configurable parameters for managing CAP and CTA periods, based on application requirements. The CTA allocation system is driven by these parameters to make time slot assignments optimally.

## FOR MORE INFORMATION

For more information or to arrange for live demonstrations: www.meshdynamics.com/contact_us.html