

**IEEE P802.15
Wireless Personal Area Networks**

Project	IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs)	
Title	Berlin to San Antonio conference call minutes	
Date Submitted	[19 October, 2004]	
Source	[James P. K. Gilb] [Appairnt Technologies] [16990 Via Tazon, Suite #125 San Diego, CA 92127]	Voice: [858-485-6401] Fax: [858-485-6406] E-mail: [last name at ieee dot org]
Re:	[802.15.3-2003.pdf]	
Abstract	[The minutes of the 802.15.3b conference calls from Berlin (September 2004) to San Antonio (November 2004).]	
Purpose	[]	
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.	
Release	The contributor acknowledges and accepts that this contribution becomes the property of IEEE and may be made publicly available by P802.15.	

1. October 19, 2004, 8 am PDT

TG3b call minutes

Attendees:

Sanjeev Sharma - Freescale

John Sarallo, James Gilb, Mike Rudnick, Jim Allen, Mark Schrader - Appairant

John Barr, Charlie Mellone - Motorola

Julian Hall - Artimi

Dave Thompson - Agere

Jay Bain - Consultant

Al Heberling - Pulse~LINK

Agenda:

- Review agenda for San Antonio
- Status of text for first draft
- Review any new items for the draft
- AOB

San Antonio Agenda:

Reviewed TG3b agenda for San Antonio. Okay for now, will adjust as necessary Monday in San Antonio.

All submission for presentation are due to the reflector by start of TG3b on Monday.

Status of first draft text:

Update on current drafting status on 384/r2. Integration into a draft on-going. 0384r2 has been posted, has the latest work on comment resolution

The MLME proposal, 0401r2 has been posted with changes marked. The TG will discuss unresolved MLME/MAC SAP primitives Tuesday afternoon in San Antonio.

New items for the draft:

Discussion on how to handle QoS requests with more information than just time required. Will the DME translate QoS requirements (need to know how to express) based on channel characteristics and available time in superframe, or will the DME just pass through an applications time request. If the application does this, is their sufficient information available for the application to make the right allocation? Need to be able to handle changes in channel characteristics that may reduce ability to handle requested QoS.

Team led by J. Gilb including A. Heberling, K. Odman, J. Sarallo, J. Hall, M. Rudnick, and M. Schrader will address and report back on the next call. One of the deliverables will be a trial QoS architecture mapped onto the new MAC architecture.

John S. wants to make sure that we provide what is necessary to support requirements for 1394.

M. Rudnick, J. Bain, and M. Schrader will come up with recommendation on how to handle large async data transfer requests when the MAC doesn't have the entire block queued but the application knows a lot of data is coming.

J. Barr asked about the polling/token discussion and was told the current recommendation is in 384/r2.

J. Allen is waiting for feedback from Roger Marks, Bob Heile and Andy Ickowicz regarding the proposed conformance PAR.

Next conference call will be on Oct. 26, 8:00 am for 1 hour at the same time. Primary topic is QoS Abstract Architecture Review.

M. Rudnick will circulate an architecture for QoS for discussion within the next two days for discussion at the next conference call.

The group began discussion on a straw-person QoS architecture.

Meeting adjourned at 8:55 am PDT.

2. Tuesday, October 26, 2004

Attendees: John Barr (JB), Allen Heberling (AH), Jay Bain, Julian Hall, Knut Odman (KO), Mark Schrader (MS), James Gilb (JPKG), Mike Rudnick (MR), Bill Shvodian (WMS), Sanjeev Sharma, John Sarallo, Charlie Mellone, Peter Johansson (PJ), Jim Allen

Agenda:

- QoS Architecture Discussion
- AOB

Discussion on 15-04-0588-00

MR: Presented 15-04-588-00

JB: Where does the QoS manager reside, in your opinion?

MR: Most likely in the MAC due to the complexity of the SAP required to support it at higher layers.

AH: Do you have use cases for some of the protocols out there? IP is best effort, IETF suggests RSVP and RTP, have you looked at the requirements from their perspective.

MR: Not yet.

AH: This would help to provide us an idea of the services required. The other use case would be what 1394 or USB requires.

MR: This is a compromise between what the higher level wants and what 802.15.3 provides.

AH: Did you include a set of definitions for QoS?

MR: Yes, delivery of real-time data.

AH: IEEE Communications magazine had an article that defined various levels of QoS.

MR: I haven't seen that, there are many different definitions.

PJ: We (1394) define QoS as bounded latency

KO: If we do it based on the source, it will be based on statistics. It could be either in the MAC or outside of it. Wherever you put it will determine the location of the QoS manager. Can it be exposed to a larger group?

MR: QoS can be pretty complex, especially if you include other parts of the network.

JPKG: I will bring up the QoS architecture at the 802 architecture group meeting on Sunday before the meeting in San Antonio.

PJ: The QoS contract shouldn't depend on the arrival of later streams. Is first-come, first-serve a characteristics of this proposal.

MR: This is a management policy, first-come, first-serve is one technique. This document does not presuppose a specific management policy.

PJ: It is essential to be addressed for uniform operating environments. Prioritized QoS is simply a fancier best effort scheme.

JPKG: This is a separate topic from this call.

JS: The key part of the proposal is that this is a wireless environment and so a level of QoS that was agreed to can no longer be maintained. The MAC informs the upper layer and allows the upper layer to determine what it will do. The upper layer may be able to change the throughput requirements to match the changes in the environment.

PJ: Contemporary applications cannot change the data rate. It is good to inform upper layers, but currently they can't do much about it.

MR: That is true now, possibly because current applications haven't had to deal with this. Some of the internet streaming is beginning to deal with this.

PJ: If we only allow applications to change data rates that will limit us.

JS: The proposal doesn't limit the application to changing the data rate, but it does allow it.

JH: The proposed model allows the upper layers to be able make decisions. It provides the flexibility for the upper layers to try to fulfill the requirements.

AH: It would be good to work through some use cases, e.g., 1394 when the bandwidth is degrading. Another would be IP based video streaming.

JG: Can you write one up?

AH: Yes.

KO: If we put something into the MAC, it becomes our problem.

JS: Wasn't QoS part of the MAC?

PJ: Aren't the boxes in the figure done by the 1394 PAL?

CM: While some of it needs to be done in the MAC, some needs to be done in the application layer.

MR: Maybe a better approach is to distribute the work, move some to the higher layers.

CM: Can we get some of the key parameters passed to the higher layers. Perhaps this can be put into the MAC later.

MR: What would you suggest moving up?

CM: RSSI, power level, lost frames, data rate. Some of these apply to mesh routing as well.

PJ: 1394 would leave data rate as a minimum for the MAC.

MS: Ultimately, QoS is determined by the human being and hence by the application. When the system asks for a certain amount of throughput or latency, it is because that is exactly what they want. To the degree that the MAC can smooth out any wrinkles, it should be allowed to do what it can to deliver the requirements.

MS: Need to add a QoS controller element to the application in Figure 1. This doesn't eliminate QoS control at the MAC level. There is a piece of it that can be done at the MAC level. You don't want fine control at the application, but you don't want the application to not be able to control.

PJ: Can we assume that QoS control resides at the application and then determine what the application needs to know to manage the QoS?

MR: In the list of WPAN parameters, should channel selection should reside in the MAC?

PJ: Yes.

CM: For multi-hop networks, the channel select is part of the algorithm.

SS: What to we move to the higher layers?

CM: All of them unless they have timing issues. How often can you control those?

MR: We have two extremes, one the MAC fully controls the QoS, at the other, all of these parameters are controlled at some higher layer, below the application. At some point where the applications all talk, probably at the PAL.

AH: IETF is trying to define QoS for IP, but there currently is not agreement on the method. A WPAN that supports 802.15.3 at the edge of the internet, Diffserv with RSVP might be an appropriate solution. Other than the 1394 adaptation layer that has been defined, has anybody defined an adaptation layer?

PJ: In RFC 2734, they put in the hooks even though the internet doesn't have the capability to use them.

AH: We can define a detailed use case for 1394 for the services that need to be provided by the MAC. Do we have an idea of what is required for IP QoS methods?

JS: We are trying to support multimedia applications as indicated in our PAR.

KO: IP doesn't have end-to-end QoS because not all of the nodes support this.

WMS: Figure 4 is new, his understanding this was controlled in the DME.

Summary:

- Need some use case, AH has volunteered to do one, need other volunteers.
- AH will submit a reference for another definition of QoS
- Two visions, QoS controller is all above MAC, other is is most is in the MAC.

Next meeting: 8 am PST, November 1, 2004

3. November 1, 2004, 8:00 am PST

Attendees: John Barr, James Gilb, Julian Hall (JH), Jay Bain, John Sarallo, Charlie Mellone (CM), Mike Rudnick (MR), Colleen McGinn (CMcG), Bill Shvodian (WMS), Mark Schrader, Jim Allen, Dave Thompson (DT)

Agenda:

- QoS Architecture Discussion

– AOB

MR presented 15-04-0591-00-003b

CMcG: How does the higher layer interpret the statistics that depend on lower level artifacts.

JH: Some issues should be common across streams, e.g., interference scenarios. We can separate stream specific QoS vs. QoS parameters that apply to all streams. The MAC would do what per stream, e.g., max retries, lifetime, it is told, but would take on responsibility, such as avoiding interference. The existing MAC-ISOCH-DATA.confirm gives enough information on performance of a stream because it allows you to indirectly determine the performance.

Main QoS issues.

- Are you meeting latency deadline.
- PAL knows when a frame is sent and when it is actually sent
- Possibly include a transmit timestamp.
- Need to add a lifetime for a frame to expire and max retries.

CMcG: What about an upper limit to max retries to prevent a single application from blocking traffic.

JH: We shouldn't allow a DEV to be irresponsible.

MR: What about changing the ordering of frames in the async queues?

WMS: Only suggest management and data frames have separate queues internally to prevent blocking.

JPKG: Does the MAC manage the latency over the air?

JH: This is independent on how the information is passed back.

CMcG: So how does the higher layer handle a problem?

JH: It would ask for more time.

CMcG: What if you specify latency, jitter and throughput, what is wrong with that?

JH: I am in favor of requesting it in time, because of the requirements of the upper layers. For example, a VBR stream might allocate two CTAs, one for the average and one for additional traffic for the variable portion.

CMcG: What if you have multiple DEVs in the piconet, each running a different PAL?

JS: How does the upper layer know what time to ask for? Does it have the superframe duration, CTAP size, rate, etc.

JH: It would need to know something from the MAC.

CMcG: It would need to know the rate set and be able to change as the distance changes.

JH: Exposing the raw service allows you to control it anyway you like.

JS: The 1394 PAL used the older interface and made assumptions about the latency, which turned out to be wrong. The upper layers would need to know about the MAC's inner workings.

CMcG: How can we meet needs if we have to report all the information up to the higher layers.

CM: Worried about writing the interface that will work for all streams. We can work on getting the data back in a timely manner.

CMcG: Getting all of this information to higher layers is more difficult to than specifying it in the MAC.

CM: Take on action to write a case of how we do QoS. Where is the QoS? Don't worry about it now.

MR: QoS is end-to-end, but we are working on 802.15.3 right now.

WMS: Still thinks that 802.11 approach of a management plane on the side handles this, but the upper SAP is undefined. I haven't seen demand for the 1394 PAL as written.

CM: The current 1394 PAL doesn't work for them either.

DT: A standard interface would be better

JB: A standard that works is Motorola's goal. The current PAL doesn't work. The question is: can you interoperate?

CM: If you are discussing end-to-end QoS, it is application dependent rather than a standard.

WMS: The SME handles QoS for 802.11e. The QAP handles requests for allocations,

CMcG: Focus on core parameters.

Next call on Tuesday, 8 am PST. JPKG to host, discuss incoming document.

Call adjourned at 9:06.