Unresolved negative comments and rebuttal report from SB3

Background

After the IEEE 802 LMSC July 2001 plenary the IEEE P802.15 WG for WPANs worked with the IEEE-SA Balloting Center to conduct two (2) Sponsor Ballots during the late summer and early fall of 2001. The P802.15.1/D0.9.2 initial 30-day ballot opened 27Jul01 and closed 25Aug01. Six (6) Balloters provided a total of 95 new comments; three (3) of which were DISAPPROVING Balloters. Additionally, this initial ballot met the 75% returned ballot requirement and the 75% affirmation requirement was met too. By virtue of meeting these requirements, the ballot is considered to have passed.

A Ballot Review Committee (BRC) was formed and led by the Project Chair. There were sixty-five (65) Editorial comments received (includes Balloters and Coordinators) and of these sixty-two (62) were accepted and three (3) were rejected. There were thirty (30) Technical comments received and of these eleven (11) were accepted and nineteen (19) were unresolved.

All balloting group members, observers, and coordinating groups have been advised of substantive changes made with respect to P802.15.1/D0.9.2 the balloted draft standard (in response to comments, in resolving negative votes, or for other reasons) and have received copies of all unresolved negative votes with reasons from the negative voter and the rebuttal, and have been advised that they have an opportunity to change their votes. The edits were applied and draft standard P802.15.1/D1.0.1 was produced; with both change bar and clean versions.

The P802.15.1/D1.0.1 10-day recirculation ballot opened 2Oct01 and closed 11Oct01. No new DISAPPROVING Balloters were introduced and the only comments received were from coordinators submitting approving votes and/or that the "IEEE P802.15.1/D1.0.1 meets all aspects of IEEE editorial coordination." Additionally, this first recirc ballot met the 75% returned ballot requirement and the 75% affirmation requirement was met too.

Based on the results of the first recirculation the 802.15 WG for WPANs submitted to the IEEE-SA Standards Board Review Committee (RevCom) their application for approval of 802.15.1 on 26Oct01. However, on 7Dec01 the Project 802.15.1 received a disapproval notification from the IEEE-SA RevCom.

During and after the Dec01 IEEE-SA Standards Board Committee Meetings the Project 802.15.1 conducted discussions with our RevCom Mentor, Sponsor, etc. and we discerned that six (6) (SB1 Comments #1, #2, #9, #26, #27, and #28) were problematic and needed *proper technical rebuttals.* A Ballot Review Committee (BRC) was formed and new proper technical rebuttals were prepared for the recirculation package. D1.0.1 was balloted unchanged.

Second recirculation ballot summary

The P802.15.1/D1.0.1 10-day second recirculation ballot opened 7Jan02 and closed 17Jan02. No new DISAPPROVING Balloters were introduced; actually the consensus increased when a DISAPPROVING Balloter changed to a yes with comments. The second recirc ballot met the 75% returned ballot requirement and the 75% affirmation requirement was met too. There were two (2) Editorial comments received (includes Balloters and Coordinators) and both were accepted. There were six (6) Technical comments received and of these **only two (2) were valid i.e., Comment #26 and #27** based on the ballot instructions. All six (6) Technical comments received were originally received during the Sponsor Ballot #1 (SB1) and the BRC, once again, thanks the voter for taking the time to further elaborate his vote on these comments from SB1, however, all six (6) Technical comments remain unresolved.

Based on the results of the second recirculation, 27/2/2 or 93% affirmation, the 802.15 WG for WPANs has resubmitted to the IEEE-SA Standards Board Review Committee (RevCom) their application for approval of 802.15.1 on 1Feb02. The change bars below reflect the new incoming commentary and the subsequent BRC technical rebuttals.

Unresolved negative comments and rebuttal report legend:
COMMENT TYPE: T/technical E/editorial, COMMENT STATUS: D/dispatched A/accepted R/rejected, SORT ORDER: Comment #, Clause, Page, Line, RESPONSE STATUS: O/open W/written C/closed U/unsatisfied Z/withdrawn

Note: "r" in front of the comment ID number indicates a revision from the first and/or second ballots; therefore superceding the prior entry.
--
Ian Gifford, Vice Chair, IEEE 802.15 Working Group for WPANs
giffordi@ieee.org

The 6 unresolved "T"echnical comments follow:

COMMENT #: r11

IEEE #: 06810238
NAME: Fischer, Michael
E-MAIL: mfischer@choicemicro.com
PHONE: +1-210-614-4096
FAX: +1-210-614-8192
CO/ORG: Intersil Corp.
PAGE: 55
LINE: 52
CLAUSE: 8.4.5.2

TYPE OF COMMENT: T

COMMENT: "Given the stated reason for not using payload flow bit "stop" in AUX1 packets (lack of payload CRC in that packet format), a stronger enforcement of this non-use is appropriate."

SUGGESTED REMEDY: 18Jan02 - Comment #11 is attempting to cause specification of the proper handling of an identified ERROR CASE. The draft standard recommends not using payload flow bit "stop" in AUX1 packets BECAUSE OF THE LACK OF PAYLOAD CRC in that packet format. Because data errors DO occur (especially on wireless media), there WILL be cases where a station receives an AUX1 packet with "stop", DUE TO AN UNDETECTED ERROR (which is undetectable because there is no CRC), even though such a packet "should not be used." My suggested remedy of specifying that should such a packet be received the receiver ignore that bit would prevent an incorrect receiver state transition from occurring pursuant to this erroneous reception. To reject this comment because it is "totally different" than the cited statement is absurd. The cited recommendation is explicitly identified as being recommended due to lack of a CRC, which lead this reviewer to consider the possible consequences of the undetected error in this control bit. In fact, I see NO negative consequences of accepting this comment, because to have the receiver ignore a control that "should not be used" ought not to break any existing Bluetooth functionality, whereas having the receiver act on an erroneous stop bit might impede the intended communication with the sender that followed the spec and "did not use" this control bit.

"It would be far safer to specify that the flow bit="stop" is ignored in received AUX1 packets instead of suggesting that this bit "should not be used."

RESPONSE: "REJECT.
30Jan02 REJECT.
The BRC thanks once again the voter for taking the time to further elaborate his vote on this comment from SB1. The BRC would like to again provide the original SB1 #11 response: "The 802.15.1 (draft) standard says that "...AUX1 packets should not be used with payload flow bit of "stop"..." On the other hand the comment suggest that the payload flow control bit shall be ignored when received in AUX1 packets. What the standard says and what the comments suggest are totally different. The standard recommends not using AUX1 when payload flow control equals 0, the comment proposes ignoring payload flow control when AUX1 packets are sent."

Again, the BRC respectfully rejects the voters comments.

"REJECT.
The 802.15.1 (draft) standard says that "...AUX1 packets should not be used with payload flow bit of "stop"..." On the other hand the comment suggest that the payload flow control bit shall be ignored when received in AUX1

packets. What the standard says and what the comments suggest are totally different. The standard recommends not using AUX1 when payload flow control equals 0, the comment proposes ignoring payload flow control when AUX1 packets are sent.

COMMENT STATUS: R
RESPONSE STATUS: U

COMMENT #: r14

IEEE #: 06810238
NAME: Fischer, Michael
E-MAIL: mfischer@choicemicro.com
PHONE: +1-210-614-4096
FAX: +1-210-614-8192
CO/ORG: Intersil Corp
PAGE: 63
LINE: 26-29
CLAUSE: 8.5.4

TYPE OF COMMENT: T

COMMENT: "Using the HEC and CRC to perform part of the address matching function means that some of the HEC and/or CRC errors detected in received packets will be due to UAP mismatch instead of errors occurring during transfer over the wireless medium."

SUGGESTED REMEDY: 18Jan02 - Comment #14 was rejected based on the reasoning used to reject comment #13. However, the very reason I made 2 separate comments on this single aspect of the draft was because I was identifying 2 different issues, and the issue in comment #14 is actually supported by the reasoning stated as the basis for rejecting comment #13. The general issue concerns the potential side-effects of initializing the HEC and CRC registers with subsets of the station address, rather than with a constant value as is done in most other 802 standards. Comment #13 concerned the possibility of an erroneous packet to a different address being accepted as valid; and was rebutted with arguments as to why such an occurrence was exceedingly unlikely. (Note: Do not misconstrue this as indicating my agreement with the rebuttal to comment #13 -- I do agree that the "false positive" case is unlikely, but there has still been nothing presented that shows it to be impossible nor quantifies that its likelihood is at least as low as that achieved by the error detection used on other 802 networks.) ON THE OTHER HAND, Comment #14 concerned the case of a mis-addressed packet being rejected due to HEC or CRC error -- which is the expected case, and indeed the apparent reason that the CRC register was being initialized with address information in the first place. The requested remedy was to alert the reader of the standard to the fact that, unlike all other 802 standards, the count of CRC errors at a Bluetooth station cannot be assumed to be equal to the number of packets that incurred errors during transfer (hence being indicative of link quality) because some of these errors on a Bluetooth station might be due to non-erroneous packets addressed to other stations. THEREFORE, IF the qualitative argument used in the rebuttal to comment #13 IS valid, then the EXACT situation to which the remedy given in comment #14 pertains is WHAT WILL NORMALLY OCCUR pursuant to receipt of a mis-addressed, possibly non-erroneous packet. It is inappropriate to reject comment #14, which concerns the EXPECTED case of proper detection of a HEC or CRC error due to mis-addressing, based on an argument as to the improbability of the UNCOMMON case of the "false positive" where an erroneous and mis-addressed packet might be accepted as apparently valid.

Please add a statement that makes it clear that in 802.15.1, unlike other 802 MAC/PHY protocols, HEC and/or CRC errors can and will occur on packets that are received without errors, and therefore HEC and/or CRC error counts in 802.15.1 must not be used in assessment of communication link reliability.

RESPONSE: "REJECT.
30Jan02 REJECT.
The BRC thanks once again the voter for taking the time to further elaborate his vote on this comment from SB1. The BRC would like to again provide the original SB1 #13 response: "The actual calculation is not required as part of a standard. That is left for the user, but the user must understand that this calculation is not solely base on the LAP, HEC and CRC. Therefore this is not an issue. In order to get a packet to this level (i.e. baseband), a packet must have passed at least two other error checking mechanism and a frequency hopping sequence, plus the Master/Slave procedures. One is the access code that uses the LAP ((24-bit Bluetooth) (MAC)) and the second is the 1/3 FEC. In the data transfer phase a frequency hopping sequence is used, thus if the intended (or unintended) device is not using the same hopping sequence, then the likely hood of a packet being received by an unintended device is based on the chance that the device happened to have used the same hopping frequency for this time slot. This is the case of two piconets. However, since the piconets are not synchronized, even if the frequency hop matches for both piconets the time slots would not be aligned, so the packet would not be received by the unintended device. Also the procedures defining the Master/Slave relationship for transmission also provides error checking of received packets, since a slave is not to transmit a packet unless it receives a indication from the master that the slave is allowed to transmit. The Master can then match the received address to the indicated address."

Again, the BRC respectfully rejects the voters comments.

"REJECT.
If this statement is needed based on the previous comment [SB1 #13], then this statement is not relevant, since the previous comment was not deemed to be correct and therefore does not support this inclusion.

COMMENT STATUS: R
RESPONSE STATUS: U

COMMENT #: r22

IEEE #: 06810238
NAME: Fischer, Michael
E-MAIL: mfischer@choicemicro.com
PHONE: +1-210-614-4096
FAX: +1-210-614-8192
CO/ORG: Intersil Corp
PAGE: 482
LINE: 23-30
CLAUSE: B.1.2.2.1

TYPE OF COMMENT: T

COMMENT: "I disagree strongly with the statement "These {Bluetooth clock, buffers, flow control, ARQ, SEQN} are important, but not so from a behavior point of view in the overall understanding of things.

... Our plan is to implement these last, if then. {emphasis mine} As a minimum, if this statement is true there needs to be a much clearer explanation of what behavior this SDL model is attempting to define and to promote the understanding of. After all, the primary place within the scope of 802.15.1 where there is an exposed interface BETWEEN devices (which is where the interoperability created by this standard must exist) is the wireless medium -- and a substantial part of the behavior that can be observed on the wireless medium involves packet exchanges that use one or more of flow control, ARQ (retry), and SEQN (duplicate filtering), and all such exchanges occur under control of the Bluetooth clock. It appears to me that the behavior of these items is of critical importance in "the overall understanding of things."

SUGGESTED REMEDY: 18Jan02 - I also feel it is appropriate to respond for the record regarding the rebuttal to comment #22. While I disagree with the reasoning behind the rejection of my comment, the subject matter thereof is informative, and therefore less critical than many of my other comments. The reason for this response on #22 is that the rebuttal statement makes some statements about the SDL language which are erroneous and others which are misleading. Since I was the first person to make large-scale use SDL in an 802 standard (in 1997 for Annex C of 802.11), and remain a strong advocate of SDL as the best formal description tool for specifying network protocols, I do not like to see an incomplete specification blamed on the language, rather than on the decisions of the people writing the description (or the people directing those writing the description). There are portions of 802.11 Annex C that are incomplete because of exactly these kinds of decisions regarding what was and was not appropriate and/or allowable to specify in the standards document -- and I, as well as most of my colleagues who were heavily involved in definition of the 802.11 MAC, now regard this intentional incompleteness to have been a mistake. (A new resource on the subject of how to use SDL in a protocol specification, available since 1999, although not back when the 802.11 SDL was being written, are several documents from ETSI, including EG 201 015-Validation methodology for standards using SDL; EG 201 383-Use of SDL in ETSI deliverables; and EG 202 106-Guidelines for the use of formal SDL as a descriptive tool.) My current feeling is that if 802.15.1 wants to repeat (albeit on a smaller scale) a mistake made by 802.11, rather than to learn from it, that is their prerogative (especially considering that the description in question appears in an informative annex). However, they should take responsibility for this decision rather than attempting to blame the SDL language!

The rebuttal to comment #22 starts with the incorrect statement that “The SDL language has a built in channel and buffering system (First in First Out).” In fact, the SDL process input queue is very much NOT a FIFO. While the input queue can be used as a simple FIFO, it is also possible for messages of a specified type to be removed

from arbitrary locations in the queue using the "priority input" symbol, to be retained at their relative positions in the queue while messages of other types pass ahead of them using the "save" symbol, and in some cases to even be dropped from the queue without ever being processed (although these auto-discard cases should not be used in SDL written to describe a standard). The fact that each SDL process has a built-in input queue is a capability that may be used if the semantics of this queue match the system being described. The existence of the process input queue is not a restriction against describing systems which do not match these semantics. In the particular case discussed in the rebuttal to comment #22, it is completely possible to use SDL to show that portion of buffer manipulation behavior needed for accurate description of the externally observable aspects of flow control, retransmission (ARQ), and packet sequence control. One technique is to use sorts (SDL data types) that are sufficiently abstract so as to be "implementation neutral," as was done for analogous buffering descriptions (e.g. retransmission, power save, and duplicate filtering) in 802.11 Annex C. Another, possibly better, approach is to "hide" the buffer processing details using external procedures, as is recommended by ETSI in EG 202 106.

The rebuttal goes on to state "The Bluetooth Clock is not standardized either." however, clause 8.10.3 on pages 80-81 of the draft purports to do just that, and at least 5 other sections of clause 8 describe normative behaviors which make explicit reference to the Bluetooth Clock. The rebuttal then states "The SDL language does not provide a clock" -- which is true (of SDL as well as all of the other ISO formal languages, and of most non-formal computer languages other than HDLs like Verilog), but this is irrelevant. What SDL does provide (unlike most other languages) is explicit access to TIME as a primitive data type, and a very flexible TIMER function to generate events after controllable amounts of time have elapsed. (Actually, SDL provides both absolute time, called time, and relative time, called duration, as two, distinct, data types.) The rebuttal goes on "so again a clock behavior would need to be contructed using other features of the SDL." Which is true, but this construct requires far less than 1 page of SDL! All that is needed to model operation of the Bluetooth Clock is to define a subtype of INTEGER limited to the value range 0:(2^28)-1, to declare a variable as an instance of this data type to hold the value of the Bluetooth Clock, declare a timer with a static period of 312.5 microseconds (3.2kHz), and define a state transition which increments the clock variable clock by +1 on each occurrence of this timer event.

The rebuttal goes on "Finally the ARQ is use to confirm delivery or to retransmitt a lost packet, that is needed to be stored in a buffer/channel not a built-in SDL feature." Precisely because of the use of ARQ, which involves inter-station messages which traverse the wireless medium, and the receipt of which by the initiator of a data packet determines when it is possible to reclaim the transmit buffer holding that packet, a description of the ARQ process is appropriate to include in a formal description of the protocol. To say that is absence is due to the lack of built-in features in SDL is absurd! If the only protocol features described are those matching built-in features of the language it is pointless to attempt formal description, as any network protocol of sufficient complexity to need such a description is likely to use contructs absent from any formal language. One key advantage of SDL for protocol description is that it is one of the most extensible of the formal languages, permitting definition of custom data types, operators, procedures, process types, block types, and system types. SDL also has excellent facilities for hierarchical decomposition of structural entities, permitting details to be hidden in higher levels of the description, but defined precisely within the model to permit simulation and validation. I believe the actual reason for the missing material is stated near the end of the rebuttal as "Significant amounts of work would be needed to implement these "real" features in SDL, while the behavior gain would be minimal." The rebuttal text prior to this statement is mostly incorrect or misleading, and should be removed. The amount of work involved, and the committee's evaluation of the benefits of doing that work, are the reason for the omissions. I would consider such a statement, absent the attempts to blame the SDL language, as a valid basis for rejecting my comment, although I disagree with the committee as to the value of the "behavior gain" that would result from describing these portions of Bluetooth operation."

I recommend that the mention of of unimportance, and the "... if then" be deleted, and that the listed items be

included in the complete version SDL model. I agree that it is appropriate to add these items to the model near the end of its development, as doing so will reduce the effort required for testing of higher-level functions within the SDL model.

RESPONSE: "REJECT.
30Jan02 REJECT.
The BRC thanks once again the voter for taking the time to further elaborate his vote on this comment from SB1. The BRC would like to point out that an SDL is not meant to be an implementation. Yes, with all the extensions and SDL modeling tools one can be constructed. However, it would be easier to include the code of an implementation, if that is what one wants. A standard, as far as we understand it, describes behavior and contains coding rules and formats.

As for the ETSI guidelines, Project 802.15.1 is not part of ETSI, IEEE has not accepted ETSI standards for SDL development, so what is the relevance. ITU-T defined the SDL, and that is all that is stated. The SDL is not stated as following ETSI standards. Specifically, we stated the SDLs were conformant to SDL-88 and SDL-92, and so that is what we followed.

Additionally, here is a quote from ITU-T Z.100 3/93 (SDL-92) from section 2.5.1 Channel "A channel represents a transportation route for signals"..."Signals are presented at the destination endpoint of a channel in the same order they have been presented at its origin point. If two or more signals are presented simultaneously to the channel, they are arbitrary ordered A channel with delay may delay the signals conveyed by the channel. That means that a First-In-First-Out (FIFO) delaying queue is associated with each direction in a channel."

Again, the BRC respectfully rejects the voters comments.

"REJECT.
These three items are not standardizable things, especially the buffers. The SDL language has a built in channel and buffering system (First in First Out). Since real buffers require manipulation (i.e. not FIFO), the built-in features of SDL cannot be used. Therefore a buffering and channel scheme must be created using other features of the SDLs. However, since the standard can not require an implementation scheme on buffer management, then neither can the SDLs. The Bluetooth Clock is not standardized either. The SDL language does not provide a clock. So again a clock behavior would need to be contructed using other features of the SDL. Finally the ARQ is use to confirm delivery or to retransmitt a lost packet, that is needed to be stored in a buffer/channel not a built-in SDL feature. See 8.8 on Transmit and Receive procedures. Significant amounts of work would be needed to implement these "real" features in SDL, while the behavior gain would be minimal. An SDL model does not necessarily implement every aspect of a system. Therefore it is important to note what is not modeled.

COMMENT STATUS: R
RESPONSE STATUS: U

COMMENT #: r25

IEEE #: 06810238
NAME: Fischer, Michael
E-MAIL: mfischer@choicemicro.com
PHONE: +1-210-614-4096
FAX: +1-210-614-8192
CO/ORG: Intersil Corp
PAGE: 437 and 440
LINE: 15-29 on p.437, 34-35 on p.440
CLAUSE: 12.1 and 12.2.1

TYPE OF COMMENT: T

COMMENT: "Since L2CAP supports connections, it is unclear why LLC type 2 connection oriented service service is not supported along with LLC type 1 connectionless service."

SUGGESTED REMEDY: 18Jan02 - Comment #25 is rejected because "There is nothing prohibiting LLC type 2 from L2CAP." I have reason to believe that the rebuttal statement is what the committee believes to be true, but IF SO there is an error in clause 12.1, page 439, line 18, where it says "... shall support Type 1 operation only." In an 802 standard, the interpretation of a statement of the form "shall support --- only" IS to prohibit alternative mechanisms for that particular function. If it is the INTENT of 802.15.1 that nothing prohibit the use of LLC type 2, then, as a minimum, the word "only" needs to be deleted from the draft at the location identified above. As stated in my comment, I can readily accept that LLC type 2 is supported, but to reject the comment while wording contradictory to the basis for rejection appears in the draft seems to be wrong.

"Either allow LLC type 2 or provide an explanation. It may be acceptable, even appropriate, to omit support for LLC type 2, but as a minimum there should be mention of why connection oriented service is omitted when the L2CAP protocol is clearly capable of providing such a service. Also, if LLC type 2 is not going to be supported, remove the inconsistency between 12.1, which states that only LLC type 1 is supported, and 12.2.1, which mentions LSDUs generated internally by LLC as part of type 2 service."

RESPONSE: "REJECT.
30Jan02 REJECT.
The BRC thanks once again the voter for taking the time to further elaborate his vote on this comment from SB1. The BRC would like to point out, use of the Transmission_Status parameter is not used to determine or trigger any specific behavior based on the value returned. As per ISO/IEC 8802-2:1998(E) 2.3.2.3.2 Semantics of the service primitive: "The types of status that can be associated with this primitive are dependent on the particular implementation as well as the type of MAC sublayer that is used (e.g., "excessive collisions" may be a status returned by a CSMA/CD MAC sublayer)." Based on this statement we (TG1) decided that there was no need to define any values, since the Transmission_Status parameter is implementation dependent by definition, nor did we define specific values for the Bluetooth (MAC). Thus we neither restricted nor needlessly defined values for the implementation.

The SAP clause is necessary to show which primitives are defined for use for the LLC type 1 between the LLC and MAC, in this case Bluetooth. Since there are a number of architecture views for Bluetooth, there are a number of ways to map those primitives to Bluetooth triggers.

Again, the BRC respectfully rejects the voters comments.

"REJECT.
There is nothing prohibiting LLC type 2 from L2CAP."

COMMENT STATUS: R
RESPONSE STATUS: U

COMMENT #: r26

IEEE #: 06810238
NAME: Fischer, Michael
E-MAIL: mfischer@choicemicro.com
PHONE: +1-210-614-4096
FAX: +1-210-614-8192
CO/ORG: Intersil Corp
PAGE: 440-441
LINE: 49-54 and 1-47
CLAUSE: 12.2.2

TYPE OF COMMENT: T

COMMENT: "This clause appears to have been copied from 8802-2, clause 2.3.2.2, which defines MA-UNITDATA.indication from the LLC side of the MAC SAP. Much of this text is inappropriate when defining the MAC side of the MAC SAP (for example, line 22 on page 441)."

SUGGESTED REMEDY: 18Jan02 - The statements rejecting comments #26 and #27 are particularly troubling. I infer from these statements that the error of including generic wording from 8802-2, clauses 2.3.2.x as is, rather than modifying this text to specify the meanings of the parameter values actually used at the 802.15.1 MAC SAP, was made by the Bluetooth SIG, and was just passed through into the 802.15.1 draft by the 802.15 Working Group. However, the question before the sponsor ballot group members is whether the draft subject to ballot should be approved as an IEEE 802 standard -- and I cannot conceive of a reason that (as one, isolated example) the NORMATIVE SPECIFICATION of the Transmission Status parameter of the MA-UNITDATA-STATUS.indication primitive for 802.15.1 (whose MAC uses NEITHER CSMA NOR COLLISION DETECTION) should include the statement "(e.g., "excessive collisions" may be a status returned by a CSMA/CD MAC sublayer entity)." while containing NO statements regarding any status values that actually pertain to the MAC and/or PHY specified in the foregoing clauses of the draft. I was aware of the copyright license between the Bluetooth SIG and IEEE, and I made these comments precisely because the corrections required seemed to be well within the scope of the restriction to "Make such limited changes to the licensed portion of the Bluetooth Specification as the Licensee determines are required for the Derivative Work." (Or, to state this another way, the recitation of generic SAP definitions, copied almost verbatim from 8802-2, is bewildering when it appears in a specification published by the Bluetooth SIG, but is flat out WRONG when it appears in a specification published by IEEE LMSC.) THEREFORE, if Bluetooth politics prevent the BRC from fixing this problem IN CLAUSE 12, the proper solution is not to reject the comment but rather to find somewhere else within the 802.15.1 document where the necessary information can be made available to the readers. (One approach might be to create a new normative annex which "contains clarifications and additional details of the IEEE 802 interfaces introduced in clause 12.")

"Please modify this clause to be a definition of the MA-UNITDATA.indication primitive and associated parameter values that will actually be generated by 802.15.1 MAC entities.

RESPONSE: "REJECT.
30Jan02 REJECT.
The BRC thanks once again the voter for taking the time to further elaborate his vote on this comment from SB1 and SB2. It appears that voter may have missed our response in SB3 that replaced our previous ones for this comment. Admittedly it may have been easy to miss our last response as it was dwarfed by our responses in the previous sponsor ballots; the previous responses were repeated for reference purposes. We have provided a

URL to an SB3 unresolved comment report in ASCII (read non change bar), which may assist the voter: http://ieee802.org/15/pub/SB3/SB3-unresolved-comments.txt or for anything related to Std. 802.15.1/D1.0.1 Second Recirculation please refer to:
http://ieee802.org/15/pub/SB3/SB3.html

As our response in SB3 states, the BRC and the 802.15.1 editors believe that the clause in question (subclause 12.2), that provides general introductory and overview material regarding the SAPs from subclause 2.3 of the IEEE Std. 802.2, is appropriate for our draft standard. Subclause 12.2 states (as subclause 2.3 of 802.2 also does) that it represents work in progress [undertaken within project 802.2]. It is also stated (likewise in 802.2) that when the work in progress is completed, then a reference to the completed work will be provided. In view of this, focusing on the specific examples that the voter uses, e.g., the reference to the "excessive collisions" value of the Transmission_Status parameter in MA-UNITDATA-STATUS.indication() should be considered as an example (informative) rather than a normative statement; note the wording in the draft standard: "... (e.g., 'excessive collisions' may be a status returned by a CSMA/CD MAC sublayer entity)...".

Following the introductory overview material, the draft standard presents in subclause 12.3, the SAP and primitives identified in this draft standard through which control and data information can be passed and received from the MAC and PHY of a device participating in an 802.15.1 wireless personal area network.

Additionally, the BRC would like to point out, use of the Transmission_Status parameter is not used to determine or trigger any specific behavior based on the value returned. As per ISO/IEC 8802-2:1998(E) 2.3.2.3.2 Semantics of the service primitive: "The types of status that can be associated with this primitive are dependent on the particular implementation as well as the type of MAC sublayer that is used (e.g., "excessive collisions" may be a status returned by a CSMA/CD MAC sublayer)." Based on this statement we (TG1) decided that there was no need to define any values, since the Transmission_Status parameter is implementation dependent by definition, nor did we define specific values for the Bluetooth (MAC). Thus we neither restricted nor needlessly defined values for the implementation.

The SAP clause is necessary to show which primitives are defined for use for the LLC type 1 between the LLC and MAC, in this case Bluetooth. Since there are a number of architecture views for Bluetooth, there are a number of ways to map those primitives to Bluetooth triggers.

It is thus the belief of the BRC that clause 12 is appropriate and consistent with this draft and the BRC will respectfully reject the voters comments.

31Dec01 REJECT.
We, the 802.15.1 Ballot Review Committee (BRC), thank the voter very much for taking the time to comment on the draft standard. The BRC acknowledges that it did not communicate its position properly to the sponsor voter in earlier opportunities. The BRC has re-reviewed the voter's comments and it does not share the voter's view. We believe that the use of this section from the IEEE Std 802.2 is valid and consistent with this draft standard.

~~REJECT.~~
~~The BRC understands the comment but based on our understanding of the Bluetooth Specification we decline the suggested remedy. We are open to further discussion but based on the signed agreement (see note below) we refer the commenter to the two (2) Bluetooth document references to Clause 2:~~

~~2.4.5. Bluetooth Personal Area Networking Profile~~
~~Bluetooth Special Interest Group, "Bluetooth Personal Area Networking Profile Revision 0.95a", June 26, 2001. [PAN-Profile.pdf]~~

~~2.4.6 Bluetooth Network Encapsulation Protocol (BNEP) Specification~~
~~Bluetooth Special Interest Group, "Bluetooth Network Encapsulation Protocol (BNEP) Specification Revision 0.95a", June 12, 2001. [BNEP.pdf]~~

~~Note: Bluetooth documents are available from the IEEE website: http://ieee802.org/15/Bluetooth/~~

~~Note: License Agreement - The signed Bluetooth SIG - IEEE Copyright License Agreement to publish the Derivative Work states: "Make such limited changes to the licensed portion of the Bluetooth Specification as the Licensee determines are required for the Derivative Work." The Ballot Review Committee (BRC) considers the suggested remedy a misapplication of the license agreement and would therefore constitute an infringement and nullify the contract.~~

COMMENT STATUS: R
RESPONSE STATUS: U

COMMENT #: r27

IEEE #: 06810238
NAME: Fischer, Michael
E-MAIL: mfischer@choicemicro.com
PHONE: +1-210-614-4096
FAX: +1-210-614-8192
CO/ORG: Intersil Corp
PAGE: 441-442
LINE: 49-54 and 1-38
CLAUSE: 12.2.3

TYPE OF COMMENT: T

COMMENT: "This clause appears to have been copied from 8802-2, clause 2.3.2.3, which defines MA-UNITDATA-STATUS.indication from the LLC side of the MAC SAP. Much of this text is inappropriate when defining the MAC side of the MAC SAP (a glaring example is the discussion of an "excessive collisions" status value on line 19 of page 442)."

SUGGESTED REMEDY: 18Jan02 - The statements rejecting comments #26 and #27 are particularly troubling. I infer from these statements that the error of including generic wording from 8802-2, clauses 2.3.2.x as is, rather than modifying this text to specify the meanings of the parameter values actually used at the 802.15.1 MAC SAP, was made by the Bluetooth SIG, and was just passed through into the 802.15.1 draft by the 802.15 Working Group. However, the question before the sponsor ballot group members is whether the draft subject to ballot should be approved as an IEEE 802 standard -- and I cannot conceive of a reason that (as one, isolated example) the NORMATIVE SPECIFICATION of the Transmission_Status parameter of the MA-UNITDATA-STATUS.indication primitive for 802.15.1 (whose MAC uses NEITHER CSMA NOR COLLISION DETECTION) should include the statement "(e.g., "excessive collisions" may be a status returned by a CSMA/CD MAC sublayer entity)." while containing NO statements regarding any status values that actually pertain to the MAC and/or PHY specified in the foregoing clauses of the draft. I was aware of the copyright license between the Bluetooth SIG and IEEE, and I made these comments precisely because the corrections required seemed to be well within the scope of the restriction to "Make such limited changes to the licensed portion of the Bluetooth Specification as the Licensee determines are required for the Derivative Work." (Or, to state this another way, the recitation of generic SAP definitions, copied almost verbatim from 8802-2, is bewildering when it appears in a specification published by the Bluetooth SIG, but is flat out WRONG when it appears in a specification published by IEEE LMSC.) THEREFORE, if Bluetooth politics prevent the BRC from fixing this problem IN CLAUSE 12, the proper solution is not to reject the comment but rather to find somewhere else within the 802.15.1 document where the necessary information can be made available to the readers. (One approach might be to create a new normative annex which "contains clarifications and additional details of the IEEE 802 interfaces introduced in clause 12.")

"Please modify this clause to be a definition of the MA-UNITDATA-STATUS.indication primitive and associated parameter values that will actually be generated by 802.15.1 MAC entities.

RESPONSE: "REJECT.
30Jan02 REJECT.
The BRC thanks once again the voter for taking the time to further elaborate his vote on this comment from SB1 and SB2. It appears that voter may have missed our response in SB3 that replaced our previous ones for this comment. Admittedly it may have been easy to miss our last response as it was dwarfed by our responses in the

previous sponsor ballots; the previous responses were repeated for reference purposes. We have provided a URL to an SB3 unresolved comment report in ASCII (read non change bar), which may assist the voter: http://ieee802.org/15/pub/SB3/SB3-unresolved-comments.txt or for anything related to Std. 802.15.1/D1.0.1 Second Recirculation please refer to: http://ieee802.org/15/pub/SB3/SB3.html

As our response in SB3 states, the BRC and the 802.15.1 editors believe that the clause in question (subclause 12.2), that provides general introductory and overview material regarding the SAPs from subclause 2.3 of the IEEE Std. 802.2, is appropriate for our draft standard. Subclause 12.2 states (as subclause 2.3 of 802.2 also does) that it represents work in progress [undertaken within project 802.2]. It is also stated (likewise in 802.2) that when the work in progress is completed, then a reference to the completed work will be provided. In view of this, focusing on the specific examples that the voter uses, e.g., the reference to the "excessive collisions" value of the Transmission_Status parameter in MA-UNITDATA-STATUS.indication() should be considered as an example (informative) rather than a normative statement; note the wording in the draft standard: "... (e.g., 'excessive collisions' may be a status returned by a CSMA/CD MAC sublayer entity)...".

Following the introductory overview material, the draft standard presents in subclause 12.3, the SAP and primitives identified in this draft standard through which control and data information can be passed and received from the MAC and PHY of a device participating in an 802.15.1 wireless personal area network.

Additionally, the BRC would like to point out, use of the Transmission_Status parameter is not used to determine or trigger any specific behavior based on the value returned. As per ISO/IEC 8802-2:1998(E) 2.3.2.3.2 Semantics of the service primitive: "The types of status that can be associated with this primitive are dependent on the particular implementation as well as the type of MAC sublayer that is used (e.g., "excessive collisions" may be a status returned by a CSMA/CD MAC sublayer)." Based on this statement we (TG1) decided that there was no need to define any values, since the Transmission_Status parameter is implementation dependent by definition, nor did we define specific values for the Bluetooth (MAC). Thus we neither restricted nor needlessly defined values for the implementation.

The SAP clause is necessary to show which primitives are defined for use for the LLC type 1 between the LLC and MAC, in this case Bluetooth. Since there are a number of architecture views for Bluetooth, there are a number of ways to map those primitives to Bluetooth triggers.

It is thus the belief of the BRC that clause 12 is appropriate and consistent with this draft and the BRC will respectfully reject the voters comments.

31Dec01 REJECT.
We, the 802.15.1 Ballot Review Committee (BRC), thank the voter very much for taking the time to comment on the draft standard. The BRC acknowledges that it did not communicate its position properly to the sponsor voter in earlier opportunities. The BRC has re-reviewed the voter's comments and it does not share the voter's view. We believe that the use of this section from the IEEE Std 802.2 is valid and consistent with this draft standard.

~~"REJECT.~~
~~See comment resolution SB1 #26.~~

COMMENT STATUS: R
RESPONSE STATUS: U

Note: This Unresolved negative comments and rebuttal report is an excerpt from the posted IEEE 802.15 document -01/420r12 contribution. More info: http://ieee802.org/15/pub/SB3/SB3.html

-EOF-