

This file contains

- 1. The AMBSI Requirements review meeting report; and*
- 2. the original call for review; and*
- 3. the comment list with replies.*

ALMA Review – AMBSI Requirements, Review Meeting Report 2000-Feb-01, 15:30UT (08:30 MDT, 16:30 CET)

Present:

L. D’Addario (Presenting Author)
B. Glendenning (Chair)
B. Gustafsson
W. Koski
K. Morita
G.H. Tan

1. Is the purpose and scope of this document clear and correct? (Tan)

Tan introduced this subject by noting that this document should define the boundaries into which the AMBSI should fit. The AMB is well specified – what is needed in this document is information towards the device side of the interface. The scope should be as deep as possible.

D’Addario remarked that this document should have been produced earlier, that there had been essentially no project wide discussion of these concepts even though two interfaces have been designed. Glendenning disagreed, noting that essentially all the requirements in this document have been in various versions of the AMB specification, which has been widely disseminated. Glendenning agreed that the extra exposition provided by this document is useful.

D’Addario also mentioned that the device side is not so detailed because most devices have not been designed yet. We want to give as much freedom to the interface designers as possible, constrained by these requirements.

A particular issue that Tan noted had to do with data rates. D’Addario pointed out that these are defined implicitly in the AMB specification, and moreover that devices should be able to keep up with the full bus speed. Slow devices are an interesting case that should be discussed at a design review. Glendenning pointed out that transactions per second (related to the 150us response time specification) is probably more relevant (harder) than the bits per second specification. We agreed that these items should be brought out more explicitly in this document.

2. Is it necessary to specify here how many standard interfaces will be produced, or is that a design decision that follows later? (D’Addario)

D’Addario said that the document should revert to his original language which states that one interface is preferred but more allowed, and in particular if there is more than the reasons for that should be documented and policies on this topic published.

Glendenning stated that he will issue such a policy, and that he’s asked Koski and Brooks to produce a requirements matrix showing the compliance (and any deviations) from the requirements in this document to their two designs. This can be discussed at the review of those designs.

3. Is it necessary to specify here the number of I/O lines for each interface style, or is this a design

parameter (subject to tradeoffs) as the draft document says? (Koski)

Koski pointed out that the document suggest exactly 32 I/O lines, whereas his current design in fact only has 31. Where does this requirement come from? D'Addario suggested, with no dissent, going back to leaving the number of lines as a design parameter, discussed in the design tradeoffs. This was accepted.

4. The fact that the interface must provide a means of generating a bus-unique 11b node address was not covered in the draft document, nor is any suggested method given in the Bus Specification. (D'Addario)

This was accidentally omitted both from the draft under review and previous drafts. There's a requirement that each node on the bus must respond only to its own node address (11 bits). The node address must be settable after fabrication – it can't be hardwired into the device. Additionally, it is desirable that some of the bits, 4 suggested, be available to the device or backplane to be set. That is, at least 4 bits should be determined off the interface board via some mechanism. This is particularly useful for devices which will have multiple copies on the same bus. This would allow, e.g., devices to be swapped and have this fact reflected in the node address.

Glendenning commented that the requirement is only that the node addresses be unique, and in fact as the interface defined (7) bits might be identical, it increases the chance that duplicate node numbers might be accidentally created (e.g., if an interface is swapped to another device). Discovering the slot number could, e.g., be supported via a monitor point.

Nobody supported this viewpoint and D'Addario's wording on this topic will be adopted.

5. Is low cost a valid requirement? If so, what is a reasonable target? (Tan)

Tan stated that cost is a valid requirement. According to his back of the envelope calculations an interface board like this should cost <\$100, including labor.

Glendenning pointed out that a difference of \$100 is a fractional cost difference of only 0.1% on average per node (\$3M of electronics per antenna, 30 nodes). Moreover all the interfaces combined at an antenna will be less expensive than the ABM and general networking equipment at that antenna. Tan pointed out that in absolute dollars saving \$100 per node is ~\$200k for the array. Glendenning pointed out that time is also a precious resource to be expended wisely, and significant NRE at this time to save \$100 per node might not be prudent, and argued that the requirement is only that the cost should be reasonable, and this can be checked at AMBSI review time.

D'Addario pointed out that in small quantities Brooks board costs \$180, and it would probably be significantly cheaper in quantity.

As a compromise, the committee agreed on a \$200 price target with mention of the desirability of lower cost.

6. For the serial data connection style, is it better to require synchronous or asynchronous communication? (Gustafsson)

Gustafsson reported that he hadn't entirely appreciated the nature of the serial communication – typically expected to be microcontroller to microcontroller. With this in mind, he agrees that synchronous communication is appropriate.

D'Addario commented that it would be desirable if a design supported asynchronous serial I/O (e.g. to support COTS RS232 devices). However if we have such devices (none are known), they can probably be catered to by connection to a serial port on the real-time computer.

7. Document numbering and format (Glendenning)

Glendenning reported that the computer division document template has not yet been issued, so there is no

formal need for it to be used. He suggests that in addition to the project-wide standard, the document be numbered with the new-style computing division number (e.g., Computer Memo#7 becomes ALMA-SW-0007). The new template should be used if doing so would be easy and not cause any delays in issuance.

D'Addario commented that while it's fine for a particular division to have its own secondary numbering system, as these are project wide documents the project wide standard described in ALMA memo #293 should be used. Do to some unfortunate telephone problems these remarks were cut off prematurely – however the same subject is described in D'Addario's formal comments on the document.

8. Summary and Actions (Glendenning)

The document is accepted once updated for the replies to comments and following actions.

- D'Addario and Brooks to insert more explicit reference to speed (both bits per second and transactions per second) into this document.
- Glendenning to issue policy statement on proliferation of AMBSI designs.
- Brooks and Koski to produce requirements compliance matrix for their AMBSI designs against these requirements in advance of any AMBSI review meeting.
- D'Addario and Brooks to also have a number in the form ALMA-SW-xxxx (where xxxx is the computer division memo number), and to consider using the computer division document template if that wouldn't delay document issuance.

=====

SOFTWARE DOCUMENTATION REVIEW PROCEDURE - REVIEW CALL

=====

Date : 2001-01-19

Document 1 : ALMA Monitor and Control Bus Interface Specification
Author(s) : Mick Brooks, Larray D'Addario

Document 2 : ALMA Monitor and Control Bus Standard Interface Requirements
Author(s) : Larry D'Addario, Mick Brooks

Invited reviewers:

L. D'Addario
B. Glendenning (Chair)
B. Gustafsson
W. Koski
K. Morita
G. Raffi
G.H. Tan

=====

The above documents may be found at the following URL, in the DRAFT table:
<http://www.alma.nrao.edu/development/computing/docs/memos/index.html>

Please send your comments on the first document directly to the author,
mbrooks@nrao.edu. As this is a revision of computing memo #7, only comments on
changes, as well as errors or clarifications, are requested.

Please send your comments on the second document directly to the author,
ldaddari@nrao.edu. This is a new document and all aspects are available for
comment.

All are invited to send comments, not just reviewers. Reviewers are *expected*
to send comments. All comments, with disposition, will be recorded in a
report.

Written comments by: 2001-01-25

Review meeting: By phone, approximately 2001-01-31.
Exact time and date to be negotiated amongst the reviewers

It is a convenience if the comments are recorded in the following format:

p.11 4.3.2 I disagree that it is best to...

Where this comment refers to text starting on page 11, in section 4.3.2.

_____oOo_____

RESPONSES TO COMMENTS ON
"ALMA Monitor and Control Bus: Standard Interface Requirements"
Draft dated 2001-01-17

Comments were received from all assigned reviewers:

Brian Glendenning
Wayne Koski
Larry D'Addario
Birger Gustafsson
Gie Han Tan
Gianni Raffi
Koh-Ichiro Morita

The comments are reproduced below in the order of their receipt. Each individual comment is followed by a response from me, introduced by '>' and indented.

--Larry D'Addario

=====
From: "Brian Glendenning" <bglenden@nrao.edu>

Mostly minor comments:

p.5 s2.1 "up at" -> "at"

> I think the writer really meant "up to", but there are additional problems with this sentence as discussed later.

p.5 s2.2 para 2 sentence 1. "which"->"that", "many"->"many more"

> I think "which" is correct; will check. OK on "many more".

p.5 s2.3 The error could come back on the same line via some defined protocol. That is, two lines seems to me to be a design choice rather than a requirement.

> The wording suggests that full-duplex communication is required, but you are right that half-duplex would be possible as well. Will re-write to allow this.

p.7 s? At one time it was described as desirable for designers to be able to directly lay the circuit out on their own boards, as opposed to using a daughtercard. That usage is not listed here - is it not desirable?

> It's true that this concept was discussed. But it creates many difficulties in achieving the standardization described in sec 1, and it does not support that idea that the AMBSI should be centrally mass-produced at low cost. Whether it's feasible to provide AMBSI design data in such a way that it can be integrated into a larger board is not obvious; it depends on the design details, and the design should not be constrained by a requirement that this be done. Nevertheless, the possibility of doing it deserves mention in this document, so I'll try to find a way to insert it.

p.7 s3.1 Should AT style module be referenced for people at, e.g. ESO (or give absolute dimensions)?

> Yes. Unfortunately, I have been unable to find an accessible reference. In lieu of that, I'll insert some data on the dimensions.

p.7 s4 Formally this section does not provide requirements, so I'd be tempted to elide it, but I don't feel very strongly about this.

> Perhaps, but the discussion helps to make it clear that some requirements are soft and that compromises are possible. It is here that the possibility of >1 AMBSI is discussed. Moreover, some actual requirements are embedded, such as the minimum-useful pins for parallel, addressed mode.

=====
From: "Wayne Koski" <wkoski@nrao.edu>

Page 5 section 2.1

The set requirement of 32 I/O lines. I see this similar to the table on page 8 where you have a minimum/maximun. I think that this requirement should also have a minimum/maximum amount. Notice that the minimum Address/Data/Handshake combined doesn't add up to 32.

> I'm afraid I don't understand the point of this comment. The statement about 32 pins to which you refer is part of the latched/sampled connection style (sec 2.1), whereas the table in sec 4 refers to the parallel, addressed style (sec 2.2). The claim that "up to at least 32 pins" are needed for the latched/sampled style is unjustified and should be deleted, as I said in my own comments (below).

=====
From: "Larry D'Addario" <ldaddari@tuc.nrao.edu>

Perhaps it seems strange that I should be writing critical comments on a document for which I seem to be the principal author, so a brief explanation may be in order. Although most of the words in this document are mine, I wrote them only as a rough draft and sent them only to a few colleagues for comments. I received none in six weeks. The present version has is now out for formal review, yet I did not see it until it was distributed. I would like to have had a chance to polish up the language here and there, as the following comments will show. Also, in the time since I wrote the rough draft I have thought of some points that should be added. And it appears that the distributed version has a few changes from my draft; I was not consulted about these, and I disagree with some of them. Although they affect rather few of the words, they are significant. Thus, my role as author should be taken with a grain of salt, and I now take up the role of reviewer instead.

> I guess it's obvious that I agree with my own comments. Whether they result in changes to the document depends on decisions by the management, which I presume to mean the heads of Computing and Systems. Perhaps there will be a consensus after the review meeting. Anyway, the following comments are submitted for the

review committee's consideration without "responses."

Throughout: Spacing after headings is inconsistent; needs minor cleaning up. Sec 3,4,5 headings are all-caps, but 1,2 are not.

Comments were requested by page number, but the pages are not numbered. However, the table of contents shows that the main text begins on page 4, so I've used that numbering in the following.

p.4 s1: The first and last paragraphs discuss the AMBSI in the plural, contrary to my draft. It has not been established that more than one is needed, and no such thing will be established in this document. The possibility of needing more than one is discussed in Sec 4, and it makes no sense to mention it until then. But even there no conclusion is reached. (See also later comments.)

p.4 s1: As an official document, more explanation of its purpose and scope is needed. Add sentence as follows:

"Any proposed AMBSI design should be shown to meet these requirements, and if it does not do so fully then the exceptions should be fully explained."

p.4 s2.1: Make secs 2.1,2.2,2.3 into subsections 2.1.1,2.1.2,2.1.3 and make new sec 2.1 as follows. This produces a clearer organization, and it provides more explanation of the point of these three subsections.

"2.1 Device Connection Styles

"To accommodate a wide range of device complexity, the AMBSI should provide several styles of sending digital command data and retrieving digital monitor data to/from the device. It is expected that the following three styles are sufficient to satisfy the needs of all devices needed for ALMA."

p.5 s2.1 (new sec 2.1.1): The second half of the sentence beginning, "The number of bits supported..." should read "this is a design parameter" not "one of the ... 32 pins." The language was changed from my draft. First, it once again wrongly implies that more than one AMBSI is required. Second, the phrase "up to at least" is nonsense. The tradeoff in selecting the number of pins is covered in sec 4, and there is no point in prejudging that discussion here. If we eventually come to the conclusion that 32 pins are needed, then it must be fully justified rather than pulled out of the blue as this language does.

p.5 s2.1 (new sec 2.1.1): Last sentence is poorly worded (my fault). Try: "The mapping between each device connection and a specific data bit within a CAN bus message having a specific ID depends on firmware running in the AMBSI."

p.5 s2.2 (new sec 2.1.2), par 2, 1st sentence: Delete parenthetical reference to serial interface.

p.5 s2.2 (new sec 2.1.2), par 2, 1st sentence, 2nd half: Word left out (my fault). It should read: "or devices which must support many more

command and/or monitor bits than can be handled by a practical number of separate connections...".

p.5 s2.3 (new 2.1.3), par 1, middle: Delete sentence "Alternatively...updating the cached value periodically." This was added to my draft, and I believe that it would be a bad design. The device must be a slave, and it should never tell the AMBSI what to do (just as no interface on the bus may tell the bus master what to do). The problem of slow devices that cannot supply monitor data fast enough to meet the stringent AMB spec exists for *all* connection styles, not just this one. If it is to be addressed in this document, then a thorough discussion in the more general context is needed. My opinion is that we should not address it here; it is a matter for detailed design of any proposed AMBSI. We imply here that the timing on the device side must be fast enough to meet the bus spec. If an AMBSI design includes provisions for handling slow devices, perhaps via a cache, then that is a feature which goes beyond the requirements.

p.6 s2.3 (new 2.1.3), last sentence: Change parenthetical "item 2" to "section 2.1.2".

p.6 s2.4: This now becomes sec 2.2, in view of the above changes.

p.6 s2.4 (new sec 2.2), par 2: The following sentences from my draft were deleted, and they should be restored (slightly reworded from draft):

"Nevertheless, if it is impractical to include analog monitoring in the AMBSI without large size or cost penalties, then a separate add-on 'analog module' should be provided. It would connect to the device side of the AMBSI and would provide the ADC, MUX, and control logic needed to digitize several analog signals from the device. The analog module should also be a standard item, to be used in each device that needs it, without its having to be separately designed for each case."

p.6, s2.5 (reset sig): This now becomes sec 2.3, in view of the above changes.

p.6: Add new sec 2.4 after new 2.3 (old 2.5), as follows.

"2.4 Bus Node Addressing

"As explained in [1], each interface on a single bus must have a unique 11-bit node address. Since the AMBSI will be mass produced, some means of setting the node address of each unit must be provided. It is unlikely that the all addresses allowed by the 11-bit number will be needed, so it is possible to provide a smaller number of settable bits. Some of these should be in non-volatile memory on the board, such as DIP switches. Electronic NV memory may be used, but then a simple means of setting the address bits under field conditions (away from the laboratory) must be provided.

"At least four bits of the address must be determined by off-board wiring. This is to allow the AMBSI to be embedded in a module of which several copies are used on one bus (e.g., at an antenna); the modules are identical but must have different addresses, according to where they are installed in the system. This is accomplished by wiring jumpers to

the AMBSI's address pins through the module's back-plane connector. The largest number of identical modules on one bus within ALMA is currently 8, which could be handled by 3 bits; 4 bits are specified to allow some headroom."

p.6 s2.6 (device-specific): This now becomes sec 2.7, in view of the above changes.

p.6 s2.6 (new sec 2.7), par 1, 1st sentence: Change "...code in that processor that performs..." to "...code in that controller to perform..." (better wording).

p.7 s2.6 (new sec 2.7), par 2: In 3rd and 4th sentences, change "will" to "must" (more appropriate to a requirements document). In 5th sentence, one important adverb and an additional clause were deleted from my draft; the sentence should read:

"These rules may significantly limit the functionality that a device designer can place in the AMBSI; complex devices may need a separate microcontroller to achieve sufficient functionality."

p.7 s3.1: I agree with Brian's comment that "AT style module" needs more explanation. I do not know of a suitable reference, so we should give the inside dimensions of a 1-wide module here.

p.7 s3.4: The phrase "but unlikely" was added to the wording in my draft. It should be deleted. I don't believe that it's true, but even if it is, it's inappropriate here.

p.8 s4, par 2, sentence 3: The original wording said, "Ideally, there should be only one version of the hardware ..." and this is a correct statement of the ideal. No justification has been given for more. The original also said, "But two or perhaps three different hardware versions could be considered if necessary." This sentence was deleted. The original wording should be restored. It allows for multiple versions with sufficient justification (to be provided outside this document) but not otherwise.

General: This document will have served its purpose only when those responsible for the AMBSI design (the Computing Division) produce decisions and policies that implement the requirements. For example, it could be decided that the cost goal of \$200 per unit is difficult and that the "much lower cost" should not be pursued. It could be decided to produce more than one version of the AMBSI, with clear justification and clear separation of the roles of the versions. So, once the requirements given here are agreed upon (presumably upon completion of the present review), we need an official follow-up document responding to them. There already exist two draft design documents for AMBSIs, although they have not yet been presented for review. As written, these are not responses to the requirements. Something in between the requirements and the designs is needed to tie it all together.

=====
From: Birger Gustafsson <bgustafs@eso.org>

Comments on ALMA Monitor and Control Bus: Standard Interface
Requirements,
Draft dated 2001-01-17

General. Why are you not using the new document standard for sw documents? OK this is not a pure software document, but I think you should use this anyhow (I guess this ememo was written by Mick who should know this).

> Actually, it was mostly written by me and I am not familiar with your "document standard." I have no objection to its use iff it has no conflict with our needs for project-wide documents. This is now *primarily* an ALMA project document. If you wish to make it *secondarily* a computing division document as well, and if you can do so without constraining the rest of the project, then that's fine.

You are now using a document number, is this now a standard? When and how was this defined?

> The document numbering system is explained in ALMA Memo 323. Since 2000-Sep, it has been an official ALMA standard for all documents produced in the US, and it is expected to be adopted by the joint project.

There is no page numbering.

> Will be fixed.

p.4 s1 In the first sentence you refer to ALMA08001Nx0001, but that document has the number ALMA08001.0001.

> The citation here is correct. The spec document has not yet been issued except as a computing memo. When it's issued, it will bear the correct number.

p.4 S2 I think it is also important that device designers review this document to see if the requirements are suitable for them. I am working on software and have not much idea what is really required.

> Agreed. Several reviewers (Wayne and Gie Han and myself) are attempting to represent the device designers.

p.5 s2.3 Why are you requiring a synchronous serial interface. I would imagine that asynchronous interface s are much more common and easier to implement.

I fwe for example want to hook up some commercial equipment with serial interface they normally have asynchronous serial interface with ASCII protocol.

> As explained in the 2nd par of this section, this style is intended to support direct communication with a microprocessor (or other fast hardware) in the device. In this context, synchronous communication is much simpler, and nearly always much faster, than asynchronous. As explained in sec 1, the AMBSI is intended for embedding in specialized devices that we build, not for support of commercial equipment. Including commercial equipment support via RS232 or

IEEE488 was considered, but rejected. Such interfaces are largely incompatible with the AMB protocol, have different semantics for each commercial product, and would therefore be extremely difficult to support in a *standard* interface. Nevertheless, it sometimes happens that an asynchronous serial port is useful; if physically present in the AMBSI, it could be supported by device-specific firmware as described in sec 2.6. Thus, having an async port might be considered an option (especially if it comes for free) but not a requirement.

p.7 s 3.1 What is AT style module? There are many other abbreviations in the text, like CPLD in section 2.3) which are not explained.

> Agreed. These should all be spelled out.

p.8 s5 ref 1. Apart from the doc number (see my first comment) this document is named "ALMA Monitor and Control Bus Interface Specification"

> You are right. We will make this reference consistent with the actual title when the bus specification is finally issued. IMHO, the word "interface" in that title is not appropriate.

=====

From: Gie Han Tan <ghtan@eso.org>

Regarding your document I have the following comments:

General: Use the ALMA software documentation standard, as was recently reviewed, for this document. Change document number accordingly to what is described in this standard.

> See response to Gustafsson's comment on document standards. This is not a software document, but rather a project-wide one. To the extent that the software document numbering conflicts with project-wide numbering, the latter must prevail. Any group may number and index documents however it likes for its internal purposes.

p.4 s1: Document number ALMA08001Nx0001 is not compliant to what has recently been decided on for software documents. See ALMA software documentation standard for details.

> That document is not a software document either, but rather a project-wide specification. The remarks above apply.

p.4 s2: Requirements are too vague to be useful for defining an interface design. At least data rates, bandwidth, signal levels, etc. need to be added. Exact values to be determined in consultation with sub-system designers.

> The requirements are not intended to define a design, but rather to set parameters within which any acceptable design must fit. It is expected that detailed designs will be proposed in response to these requirements. Each such proposed design should also be reviewed.

p.4 s2: The need for the unique 64 bit serial number, as is mentioned

in ALMA computing memo #7, should be added.

> Perhaps. But that document is referenced, and it's emphasized in sec 1 that *all* interfaces on the bus (the AMBSI as well as any non-standard ones) must comply fully with the AMB specification, including this point. This present document is therefore concerned with the *device* side of the AMBSI, not the bus side.

p.7 s3: Section on susceptibility to external EMI, e.g. conducted/radiated/ESD, should be included.

> Possibly, but in practice EMI *susceptibility* is never a problem for digital devices in radiotelescopes. If you care to propose text for such a section, it might be included. EMI *emissions* from the AMBSI might be a more significant issue, as is discussed in sec 3.3.

p.7 s3.4: Cost objective for production of interface design should be less than \$100,-. For the given functionality this should be feasible with a clever design.

> I agree, but it appears that some designers do not. We can discuss this at the review meeting. I propose to delete the phrase "but unlikely."

=====

From: "Gianni Raffi" <graffi@eso.org>

I have no comments on (the latest version of) these documents.

I see purely a formal problem with document numbering. the AMBSI document uses the new system (approved ?) and refers therefore to an AMB unexisting document.

> The document numbering system is explained in ALMA Memo 323. Since 2000-Sep, it has been an official ALMA standard for all documents produced in the US, and it is expected to be adopted by the joint project. The referenced document exists in draft form as Computing Memo No. 7; after review (now in progress), it will be issued as a project-wide specification and numbered as indicated.

=====

From: "Morita, Koh-Ichiro" <morita@nro.nao.ac.jp>

General:

For me, this document is a little bit confusing. It to clarify the ideal interface for ALMA Control and Monitor, although two version of AMBSI are now developing? If so, I want to have a clear explanation about that.

> Development of proposed standard interfaces has occurred without much consideration of underlying principles or requirements. This

document is an attempt to fix that. It would have been better if we had written and agreed upon the requirements much earlier. To this I can only say, "better late than never."

p.4 s1

Clear description of the purpose and outline of this document should be needed.

> I agree. See my own comment on this section, above.

p.4 s2.1 and p.5 s2.2

Could you show us device examples for this connection type? To discuss the number of data bit needs such a list.

> If the descriptions are not clear, more explanation and perhaps some diagrams could be added.

p.7 s3.1

What is AT style?

> This refers to the Australia Telescope packaging for electronic modules, consisting of plug-in metal boxes of several sizes. This system has been adopted for many ALMA devices. However, it is not widely known within the project, so we must include more information here. I will add a brief description and dimensions of the smallest box.

p.8 s4 par 2. sentence 3

"Ideally there should be no more than two versions ..."

Why two versions? There is no explanation of the justification in this document.

> I agree. See my own comments above.

p.8 s4 par 2

If we want to discuss the number of versions of the AMBSI, we should estimate the number of devices of each connection type.

> In principle you are right, but this is difficult to do since most devices have not yet been designed. In addition, we have a "chicken and egg problem": If AMBSIs are mass-produced, then devices designed later will adapt to the connection styles provided. This requirements document represents a judgment, based on experience, that these three connection styles will (a) each be useful in ALMA for some devices, and (b) be sufficient to support all reasonable device designs.