Dear NMRA Board -

Re: Proposed S/RP9.x.1 and Changes

It is the opinion of the signatories below that the original S9.x.1 and TN9.x.1 should be accepted as submitted. The submitted draft S9.x.1 and TN9.x1 are the result of extended work by experienced people, including professional engineers with relevant experience and model railroad manufacturers. They represent a significant new model railroading capability, while at the same time using conservative design to ensure reliable performance.

The undersigned recommend that the changes proposed by Di Voss be rejected. We do not believe that the majority of model railroad users will be reliably able to achieve the performance attributed to NMRAnet with those changes, because they ignore other limiting conditions such as node spacing, capacitance, and reflection. Ignoring real engineering issues is not the way to create a standard that results in reliable product networks. The changes also unnecessarily limit potential implementations and uses of the bus, which will dissuade manufacturers and slow adoption of NMRAnet.

<u>Background</u>

As you are aware, the NMRA is developing a bus system, NMRAnet, to provide controls for a DCC or DC layout, and which will carry bidirectional information between track-side equipment. To this end, a NMRAnet committee was constituted, that brought together several groups of modellers whom had been developing similar projects. Unfortunately, the end result was several proposals for consideration for NMRAnet.

As a way forward, at the recent Milwaukee Convention, Di Voss constituted an informal committee under the guidance of Karl Kobel, with membership from the above proposals, and other interested parties including manufacturers. This committee was tasked with developing a standard for the physical aspects of one possible transport, namely the Control Access Network (CAN), which is an international standard bus (ISO 11898-1,2) used extensively in both vehicles and lab and factory automation. The CAN standard completely specified the bus except for options in the area of speed, connectors, and a few physical parameters. The committee subsequently debated and considered all proposals, came to a consensus decision, produced two documents: the S9.x.1 standard; and its associated recommended practice document RP9.x.1, and these were duly submitted by the chair Karl Kobel to the NMRA board via Di Voss.

Don Voss, Di Voss's brother, was a member of the committee. Don took part in the discussion, proposing details and discussing others. He suggested several changes that he felt would enhance the capabilities of the proposal. The group discussed these at considerable length, several were included in the draft, but some of them were rejected as not being appropriate for an NMRA Standard. Several of these were Don's ideas for enhanced performance as a result of constrained implementation, which several members of the committee thought were much better suited for a follow-on Recommended Practice.

After submission of the standard, Di Voss made changes to the document to include the rejected proposals of his brother. We have grave concerns regarding both the technical merits, the effect on the standard, and the process with which these changes took place. We document these below in greater detail.

Introduction

The following sections can be quite detailed and technical. We suggest you read them in light of an analogy: consider the process involved in setting a standard for a curve on a road. This would take into account the radius of the curve, the road surface and the expected traffic, among other things. Such a standard would be set to ensure the safety of all the traffic, so that it can be reliably travelled by the intended traffic. Developing a standard for NMRAnet is similar, in that it must take into account the physical characteristics of the bus, in order to specify one that successfully carries its traffic.

Such standards must balance that which is technically possible with that which is practical, reliable and safe. In the analogy, it would balance the ability to manoeuvre the curve in a racing car in perfect conditions with the need for the safety of ordinary cars and trucks. Further, it would not be acceptable for any of the traffic to fail to negotiate the curve. The NMRAnet group has taken great care and worked hard to get this balance right. The original unmodified proposal is careful to set conservative expectations while not being too proscriptive. We were also careful to make use of existing standards, so as to not reinvent or invalidate them. All submissions were debated and given due consideration. The members of this group included two electrical engineers, one with extensive experience in bus systems and CAN in particular, a PhD in physics with extensive practical CAN experience, two members with computer science backgrounds and electronics experience, and a model railroad equipment manufacturer with extensive experience designing, building, selling and supporting LocoNet bus-resident equipment.

Mandate of the Committee:

The NMRAnet committee was charged by Di Voss with specific mandates, which included:

> 6. Maximum utilization of existing *globally standardized communication technologies* should be used to develop this bus. *Multiple forms of communication transport shall co-exist.*
> 8. Develop a bus that *minimizes the amount of certification testing* required of NMRA. Certification of this bus should not require special testing equipment beyond a computer and appropriate connections.
> 9. Develop a set of NMRA Standards and Recommended Practices that *fully define the characteristics of the bus* which can be used by manufacturers to develop products.

The Voss changes are contrary to the above mandates, as expounded below, and should be rejected.

Bus Length

The Voss proposed changes to the standard include a graph of bus length vs number of nodes. This graph is based on two formulae from two manufacturers' data-sheets. We agree that the calculation indicates conditions beyond which the bus is likely to not perform properly. We disagree with the Voss changes, because we believe that they overstate the performance because they ignore other technical considerations that may limit performance. This is detrimental is several ways:

- It documents the limits beyond which the bus will certainly fail, instead of the more conservative approach taken in the original proposal of documenting limits within which the bus will certainly work,
- It will require additional testing by manufacturers; and
- It will require sophisticated equipment in order develop a bus that approaches the limits given.

The analogy here is to setting the speed limit on a particular curve at 60MPH, because not even a race car can go around it faster than that. Larger vehicles or inattentive drivers might not be able to get around it at 30MPH, and <u>following</u> a 60MPH speed limit sign would give a rude shock. Similarly, CAN networks are limited by more than just the two DC conditions considered in the two formulae. There are AC issues, node-spacing issues, and noise margin concerns that indicate the proposed limits are not likely to be reached.

Contrary the Voss approach the original S/TP-9.x.1 presents three simple rules:

- Total cable length of 1000 ft / 300m; stub cables count double their length in that total; nodes count as 20 ft / 6m in that total;
- Never less than 1 ft / 30cm of cable between nodes, nor between a stub connection and a node;
- No more than 50 nodes connected;

which are based on extensive experience of CAN networks, and takes into account extensive testing of configurations using higher bus-speed and better wire and connectors. In the analogy, this is similar to extensive testing a tighter corner, and limiting the speed to that curves maximum safe speed.

<u>Industry Standards</u>

The group was mandated to use existing industry standards where they exist. The Voss change violate this mandate because:

- It restricts the global CAN standard, such that much CAN-conforming equipment will fail to conform to the Voss changes;
- As such these changes fail the mandate that the bus be based on *global standards;*

The Voss proposal seeks to extend the CAN standard by insisting on using hardware with better characteristics than the CAN standard requires. This means that industry standard equipment, such as CAN bridges, may not be conform to the Voss changes and must be separately certified. Further, many existing hobby, test and demo boards will not conform, forcing model railroaders to either purchase their equipment from a small number of vendors, or risk confusion about whether a collection of units will perform properly.

<u>Bus Termination</u>

The Voss changes mandate a specific high-performance termination for the bus. This may in some cases be beneficial, in that it may reduce sensitivity to common-mode bus noise. However, this change is also more restrictive, in that not all commercially available devices have terminators with the required characteristics. This type of termination was not included in the International Standards Organization's specification of CAN because in some cases it will introduce other problems, such as ground offsets.

In the analogy, this would be equivalent to mandating studded tires, which are beneficial in ice conditions, but be detrimental in dry conditions. The committee recommends that a solution should not be mandated, but rather that manufacturers be able to choose an appropriate specific solution.

<u>Power Distribution</u>

The original drafts provided a mechanism for distributing limited power via the bus. This allows the easy implementation of small power conservative nodes, such a fast clocks, LED drivers, bridges and gateways. It minimizes the need for separately supplying power to small nodes. It also makes it simple and inexpensive to provide isolated nodes for when needed, for example to connect to a computer or other equipment. Labelling will make clear the limits of the power distribution, so that model railroaders can easily see whether they need to add additional power. In addition, this can be automated by including power requirements as a reportable item via the bus.

The Voss change would disallow all power distribution. This unnecessarily severely limits the ability of manufacturers to build small nodes, nodes that need to be isolated, etc.

<u>Summary</u>

The committee strongly recommends the rejection of the a posteriori Voss changes, in order to ensure a bus that has conservative features that ensure user satisfaction. If the Board feels that any of the Voss changes have merit, then protocol would suggest that these changes be sent back to the committee for reconsideration or for inclusion in an new Recommended Practice. At the very least they would require a serious rewrite of the original S/RP9.x.1 as they introduce dead paragraphs and other inconsistencies into the text.

Signed-
John Day PhD, Silicon Railway
Bob Jacobsen PhD
David Harris PhD MD, Silicon Railway
Alex Shepherd
Dick Bronson, RR-CirKits, Inc.

Appendix

Number of nodes and differential input resistance

In this case the proposed change mandates the choice of CAN transceivers that have a differential input resistance greater than a specific value, which value is greater than the CAN bus specifications require. This precludes developers from using parts which are otherwise totally in compliance with CAN bus. It seems the only motivations for this change is to allow larger numbers of nodes on the bus. But then even the calculation of maximum node numbers in the supporting TN is not truly valid. It is valid only for the case where an idealized system is constructed with no ground bounce or offsets, no external noise and minimal capacitive loading with no or minimal stub connections. Highly unrealistic.

If we revert to the original standard, using the CAN bus standard Rdiff range, then we would be more inclusive in terms of the certified CAN bus parts that could be used. We could also easily evolve a standard for marking nodes with "UNIT LOAD" which would be based on the Rdiff of the transceiver. A simple web based calculator could then be devised which would allow a user to enter the load values for the nodes he has and see if they approach any sort of limit which might impede performance.

But it must also be said that striving to build a network with over one hundred nodes is also flawed. Disposition of nodes along the length of the bus is an important issue, and one which also limits the number of nodes. However reality will likely dictate that a single NMRAnet network sector will naturally tend to limit itself to a value closer to 32 nodes, based on considerations that are part of the development of the higher layers of the NMRAnet standard.

In the analogy, this is equivalent to calculating, say, tire grip based on tire composition, tread, road materials, and the curve radius, and calculating the maximum speed at which the corner can be negotiated. While this gives a theoretical maximum speed, it does not take into account other more limiting conditions, such as the weather, the centre of gravity of the vehicle etc. A road speed standard needs to use the most limiting estimates to ensure safety. Similarly, the NMRAnet standard must specify limits that ensures a bus that guarantees success. The Voss proposals do not do this, rather than determining limits were success is assured, they determine limits beyond which failure is assured.

Termination of the bus:

The proposed change mandates a specific termination connection and only one possible value of terminating resistance. One issue here is that because we use CAN bus and make no apology for it, users and developers will seek to utilize other products with CAN bus interfaces. Whilst many may be compliant with the Voss modified standard, many will not. And herein lies the issue, the original was drafted with respect being accorded to the already extant standards, applications notes and supporting material. The changes tend to cherry pick, looking at what might be good practice, in some instances, and mandating it for the entire NMRAnet structure.

Within the professional community using CAN bus there are a number of additional termination methods which can improve the performance even more. But the Voss changes now preclude those being applied within NMRAnet should they be accepted.

Power distribution

Of all three fundamental Voss changes, the loss of power on the NMRAnet cable is possibly the least discussed, but also the most sinister. Much of the concentration of discussion on power in the NMRAnet cable centred around using the cable to distribute operational power to the nodes, power to actually run the node and power other external devices such as turn-out motors or whatever. This issue is perplexing because the current rating of the proposed connectors is not great. So on those grounds it might seem wise to remove power wires from the cable.

In fact it is a disaster. Good practice, and not even the best practice, dictates that the CAN interfaces on every node should be at "a COMMON UNITPOTENTIAL VALUE". Which can be achieved by having separate CAN bus power which can be used to power galvanically (electrically) isolated CAN interfaces on the node. The operational power distribution being separate will allow for larger losses and thus more widely varying ground potentials.

The original standard, without the Voss changes, does not mandate the use of isolation and a CAN bus power circuit used for the bus alone, but it does allow for a developer to implement it. In the case of the changed version the possibility of implementing best practice is precluded. This alone will tend to limit the number of nodes on the NMRAnet segment much more severely than the reverting the change in (2) above would.

Use of application notes as authority

Manufacturers write application notes in order to facilitate and encourage the use of their products. Often, sadly, they also attempt to cast their owners products in a better light than something equally good from a competition. One of us (JD) has a close acquaintance and one time colleague who now works as a field applications engineer for a company that makes CAN transceivers. It occurs that in their labs they have a small CAN network, "of maybe 10 or 20 nodes where we do all our testing." he said. That's it, that's all they have! And what of the legendary equations that tell us how many CAN transceivers we can connect? I was told "Well, we assume that (company x) did the work, so I think that since then everybody has just taken the formula and used it without re-validating it."

"Our notes attempt to give our clients a basis for comparison and some rough rules to start from."

But the same discussion with the relevant applications people in two other companies yielded basically the same information.

The authors do not set out to write an authoritative reference, they are not given the time or the funds to do so. When asked why we couldn't get SPICE models for the CAN transceivers the standard answer is that after the silicon is produced for prototypes there is about a week allowed for the production of data-sheet, initial applications info and a model if it is absolutely necessary. For CAN transceivers it must be deemed unnecessary.

Accepting the formulae and graphs in the application notes as being anything other than a starting point is risky, if not outright dangerous. An engineer adopting that position would find himself severely challenged when he faced a hazard and risk analysis, or a critical design review. Yet people within the NMRAnet community would have us accept the information in one or two application notes as though they were tablets of stone, when in fact they are more like the beach sand, able to hold us up for now, but don’t stand in one place too long and don't look back!

Accepting these missives as authoritative or normative is at least as risky as accepting any of the assertions implied by changes 1, 2 or 3 above.