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.