Title:

Mr

Forename:

Larry

Surname:

Taylor

Representing:

Organisation

Organisation (if applicable):

DTC (UK) Ltd

Additional comments:

DTC (UK) Ltd is a specialist Consulting company that has been involved in the regulatory and standardization aspects of the 870-876/915-921 MHz spectrum re-designation for use by SRDs. DTC welcomes this opportunity to provide those interested in this spectrum with some insights into the work behind the draft Harmonised Standard EN 303 204 which may be relevant to considerations of licencing and deployment of Network Relay Points.

Question 1: Do you have any evidence to inform Ofcom?s view on the density of higher duty cycle (up to 10%) NRPs deployments, whether this is likely to exceed 10 NRPs/ km2 and the total number of higher duty cycle NRPs that might be deployed?:

Although DTC cannot offer any evidence from practical deployments, the contributions to the SRDoc and Regulatory discussion were based on considerations of commercial viability of the eventual regulations for the frequency band.
An NRP is a relatively expensive device compared to an end SRD device, and the applications targeted are highly cost sensitive. Hence the end device cost has to be very low, with extremely low operating costs, in order for the total network offering to be competitive against other technological approaches (e.g. cellular modems, narrowband low data rate modems etc.).
Operators of these Rural/Metropolitan Networks must maximize the number of end devices per NRP in order to minimize the deployment and operating costs of the total network.
Only in dense urban deployments, where heavy shadowing or severe path losses limit the number of reachable end devices, would DTC expect to see a higher density of NRPs. In this case, of course, the number of end devices per NRP would be reduced and a corresponding reduction in the traffic load via the NRP would be seen, reducing pro rata the interference potential of such NRP deployments.

Question 2: Do you have a view on how intra-network interference caused by NRPs deployed in large numbers within a network will be managed?:

In order to understand the potential interference from NRPs to other ‘intra’ Network Based SRDs using the spectrum, it is necessary to understand what an NRP is.
ECC Report 189 includes the definition:
Network Relay Points - Device deployed by organisations, such as smart utilities, municipal, industrial, transport, logistics or other metropolitan/rural area network operators, to support wider area operations. NRPs provide connectivity for one or more otherwise isolated network devices by forwarding traffic in both directions between the network and the isolated device(s). Such devices will be limited in their deployment and will not be operated by the general public/consumers.
EN 303 204 defines NRP as:a class of device intended to provide network infrastructure to support communications between devices and an external communications network or service and in its Scope includes: Network Relay Points which are specific fixed Network Based SRDs supporting interconnection of a network of SRDs with an external network or service.
These definitions make it clear that an NRP is an infrastructure device i.e. its communications load arises from transferring information between a network of SRDs and an external network or service. Apart from traffic such as (relatively infrequent) keep-alive packets, NRPs should not initiate any significant traffic load.
SRD traffic in operated networks, such as for Smart Cities and Utility Networks, is not likely to be continuous traffic. The operating duty cycle limit for a device is 2.5% to allow construction of MESH Networks where traffic from outlying devices may be repeated by MESH Network neighbours on the path to the destination. (The repeating of traffic for outlying devices adds to the load generated by the intermediate MESH device and leads to a 2.5% requirement.) The destinations will be the NRPs which act as the interconnection point for the exchange of information between the network of SRDs and the external network to which to Smart City or Utility Applications are connected (control centres, back office applications etc.). The need for higher duty cycles for the NRPs is to allow support of a large population of SRDs for each interconnection point to the external networks since multiple MESH sub-networks can connect to the same NRP.
DTC agrees that the operation of such a topology will be essentially self-limiting. If SRDs are not able to access the medium to send information towards an NRP, the NRP will not have response information for the SRD concerned.
The worst case would most likely be for information towards the SRDs failing to be delivered and requiring re-transmission attempts. Several factors will then intervene to limit the eventual traffic – including both higher protocol layer (e.g. timeouts) and application layer effects (message re-petitions etc.) which will have timescales longer than the device medium access cycles and hence spread the re-transmission over suitably long periods.
If a potential operator abused the notion of NRP to try to deploy a dense network of NRPs directly supporting traffic to end devices, the NRP/device ratio would be small leading to unlikely commercial viability of such a network topology. In such cases, existing cellular modems for M2M applications would likely be highly competitive and the existing network infrastructure of such cellular systems would likely defeat any attempts to deploy new networks based on SRD technology.

Question 3: Do you have any evidence that networks may fail if the aggregate density of higher duty cycle NRPs reaches or exceeds 10 NRPs/ km2?:

No.
Again, the commercial factors are considered by DTC to be dominant and the cases where more than the recommended 10/km2 would be where the propagation environment was too poor to support a large population of end devices and consequently the traffic loads in higher density NRP deployments would be expected to be lower.

Question 4: Do you have any views on whether exchanging NRP deployment information between licensees and developing and using an industry-managed code of practice would be practical and sufficient to manage the risk of some networks failing?:

DTC believes that it will not be necessary to require inter-network coordination of NRP devices. Information from large scale deployments of metering networks in the US and other locations indicate that very large populations of end devices (of the order 1000s) per NRP are possible since the application traffic load is very light for these M2M/IoT applications. Where application event rate is higher (e.g. more frequent meter readings), aggregation of multiple readings per transmitted packet can keep the spectrum load low.
DTC believes that it will be more beneficial to ensure the maximum effectiveness of the politeness protocols which will follow initial deployments. The data carrying capacity of the spectrum, coupled with spectrum–reuse inherent in SRD systems design, will permit a very large number of low activity devices to be operating concurrently without coordination.
More than NRP coordination, DTC believes benefits will be obtained by good uniform distribution of end devices over the available spectrum, coupled with adaptive behavior to avoid unsuitable or heavily loaded channels in favour of less loaded spectrum. This behavior should be inherent in all devices, end SRD devices as well as infrastructure devices, obviating the need for coordination between operators.

Question 5: Do you think CCA as defined by ETSI will be an effective protocol for (a) managing interference between networks? (b) managing interference to short range devices using the 870-876 MHz band?:

Yes.
EN 303 204 requires NRPs to implement CCA and Adaptivity. CCA prevents transmissions beginning when the channel is already in use and Adaptivity requires the NRP to defer to such existing transmissions on a given channel or to change operation to a different channel, reapplying the CCA/Adaptivity access behavior.
There is a vast accumulation of evidence in favour of Listen Before Talk as a fair spectrum sharing technique, both theoretically (dating back to the seminal work of Kleinrock, Tobagi & others in the late 1970’s) and practical with WiFi/802.11 being the dominant deployed technology.
Although limitations and shortcomings certainly exist, channel sensing before transmission coupled with adaptivity is a highly effective basis for multiple access in shared spectrum. However, DTC believes that CCA alone is only a partial solution to the polite spectrum access goal. One difficulty with SRD regulation and associated Harmonised Standards is the requirement for technology neutrality. Harmonised Standards for licence exempt devices cannot be defined assuming any particular design choices. Manufacturers must be free to innovate and find better ways to use the scarce spectral resource.
This means, for example, that a specific channelisation scheme cannot be mandated, or a

specific modulation scheme or bandwidth. The Harmonised Standard can place limits where compatible with the regulations. For example, REC 70-03 Annex 2 Band C recommends a maximum channel spacing of 200kHz and this is reflected in the EN 303 204 Harmonised Standard. Furthermore, the limits defined in the Harmonised Standard take into account the limits built into the recommendations to give a framework for equipment design.
For CCA, the issue is how long should be the channel sensing time in order to make a reasonable estimate of channel occupancy. This depends on the bandwidth of the receiver, the bandwidth of the transmitted signal, the number of transitions in the modulated transmitted signal, the signal strength received etc. etc. The value chosen, 160us, is a compromise on these parameters and is aligned with product related standards expected to be important for this spectrum.
EN 303 204 deliberately separated the channel sensing from the random access mechanism i.e. separated the well-known Listen Before Talk into its 2 parts. The choice of random access algorithm is highly dependent on the characteristics of the network (especially the number of simultaneously active devices and the duration of each transmission) and is therefore best left to the implementation.
CCA does not address the inherent limits of channel sensing, particularly the susceptibility to ‘hidden terminal’ effects and ‘near-far’ effects. Nor does it allow use of the spectrum when a transmission could be successful because of relative position of devices.
However, CCA is the basis on which more polite behavior should be built. Its purpose is to avoid knowingly transmitting when that transmission would destroy an already in place use of the spectrum. Polite Spectrum Access should minimize the wastage of the spectral resource.
Consider, for example, Short Signalling Transmissions. If the spectrum is sensed free, a very short probe signal can be sent to a potential destination device with a low probability of interference. If received, the destination device can quickly send a short response signal if the spectrum is sensed free in its neighbourhood. Although there is no guarantee that neighbouring devices will detect or be able to interpret these signals, the very fact the destination is available already prevents the source device wasting spectrum by transmitting data to a device that is not available. If the destination correctly receives the transmission, a short confirming acknowledgement can avoid unnecessary re-transmission of the same data, again preventing waste of the spectral resource.
The probe/response transaction can be a good technique to limit the effects of hidden terminals. The acknowledgement of received transmission certainly avoids unnecessary repeating of transmissions. An almost ideal case would be the use of a probe/response/data transmission/acknowledgement transaction as a general polite spectrum access policy.
DTC does not know of any shared spectrum access mechanism to guarantee interference avoidance with unlike systems, but believes the kind of techniques permitted (and mandated for NRPs) by EN 303 204 can significantly increase the available capacity of the spectrum and hence should be strongly encouraged and perhaps mandated for all devices at some future time.

Question 6: Do you have a view on the costs and benefits of adding effective mitigation protocols such as Clear Channel Assessment to higher duty cycle NRPs?:

Yes.
The choice of CCA duration took into account contributions from leading semiconductor vendors on the technical aspects of the detection process. These contributions in turn took

account of practical device implementations to ensure that parts available to product manufacturers would be able to meet the requirements of the Harmonised Standard.
There should be no device cost added in order to support the CCA timing except for any firmware/software development to implement a vendor’s choice of multiple access mechanism. In most cases, such a mechanism is already defined by a product standard and therefore will typically be available in standard libraries of firmware/software.

Question 8: Do you agree with our proposals to authorise spectrum for NRPs using non-exclusive network licences available on demand?:

Yes.
NRPs should be clearly labelled by vendors to avoid abuse of the increased duty cycle limit for non-infrastructure inter-network connection purposes.

Question 9: Do you agree with the proposed licence conditions for higher duty cycle NPRs?:

Yes.
However, DTC does not believe spectrum trading would be applicable in this frequency range.