ITU - Telecommunication Standardization Sector Temporary Document **2012-11-4A-024**

STUDY GROUP 15 Original: English

Chengdu, China, 5 – 9 November 2012

Question: 4a/15

SOURCE[1]: KPN, TNO

TITLE: G.fast: KPN requirements for G.fast

ABSTRACT

This contribution revisits the KPN requirements for G.fast. It refines earlier requirements, brings additional arguments for some of them and identifies a few new requirements. This contribution does not address requirements that are already agreed (e.g. on bitrate/reach performance), but concentrates on those requirements from KPN that are currently missing or are not yet agreed upon.

1. Introduction

So far, a variety of requirements has been identified by different operators, and many of these requirements have been agreed or are gaining acceptance. Most of them are applicable for KPN as well. There are, however, a few requirements missing or not yet seen by ITU as essential while they are very relevant for the Dutch situation.

Contribution [1] contained a list of KPN requirements. The current contribution refines and updates these requirements, focussing on those that have not yet been agreed. This contribution also provides additional arguments for some of these requirements, in response to questions that were raised by Q4a during the discussion of [1].

It is recognised that not all the requirements listed in this contribution lead to a simple, clear requirement on the G.fast transceiver itself. In fact, some of the requirements are aimed more at system level specification. Nevertheless, many of the requirements (e.g. the size of the coordinated group) may have implications for the transceiver, which is why it is considered useful for the sake of awareness to have these requirements as items on the G.fast issue list.

This contribution also explores some different deployment scenarios for G.fast, related to the case study of G.fast deployment in Amsterdam [2]. Exploring various deployment scenarios for G.fast can help in building a common understanding of the system functionalities that G.fast receivers need to support.

2. Role of G.fast in the Netherlands

For the Netherlands, G.fast is envisaged to provide the telcos with the means to deliver many hundreds of Mbps to customers, while avoiding the costs and the complexities of deploying full Fiber-to-the-Home. It should provide Telcos with a cost-effective answer to competitive solutions from cable operators.

[1]

Contact: KPN, Dennis Francois Email: Dennis.Francois@kpn.com
KPN, Bob Hoefnagel Email: Bob.Hoefnagel@kpn.com
TNO, Rob van den Brink Email: Rob.vandenBrink@tno.nl
TNO, Bas van den Heuvel Email: Bas.vandenHeuvel@tno.nl

G.fast is to be used in a "Hybrid Fiber to the Home" scheme, where the last 200m of existing copper wiring is reused. It is estimated that G.fast will be particularly useful for deployment in apartment buildings (FttB) and in other areas with a high population density. In the accompanying contribution [2], the deployment case of Amsterdam is discussed. That case study shows that there are large opportunities for the deployment of G.fast, even already from existing manipulation points (cabinets).

It is our opinion that hybrid FttH with G.fast ***should not be positioned as an interim solution*** towards full FttH, but should offer a long term solution. If hybrid FttH with G.fast can offer aggregate bitrates up to 500-1000Mb/s, and to a significant number of homes, then there may not be a need to replace it by Full FttH for a very long time.

3. Support of legacy functionality

If G.fast cannot offer aggregate bitrates up to 500-1000Mb/s to a significant number of homes, then Hybrid FttH with G.fast ***may not be the right alternative*** for full FttH. Therefore KPN is concerned with all requirements causing a significant limitation of the attainable bitrates. Depending on the details, demanding support of legacy DSL solutions on a G.fast platform may lead to such limitations.

That having said, migration from legacy situations is also important to KPN. Features of G.fast that facilitate easy migration are welcomed, as long as they do not lead to significant limitation of the attainable bitrates with G.fast.

It would be "nice to have" if G.fast has a special mode where it can emulate vectored and non-vectored VDSL2. However it will be an undesired feature if such requirement has a big penalty for the G.fast performance when disabled. In such a case, KPN prefers to leave such requirements on interworking with legacy equipment out of the G.fast specification. Such a functionality can also be implemented by equipment having G.fast and VDSL2 chips on board.

Regarding implementations of the CPE, KPN envisages a need for CPEs that support both VDSL2 and G.fast.

4. G.fast should be possible in conjunction with bonding

The availability of a double wire pair, organized in a quad geometry, is quite common in the Dutch access Network. About 58% of the links to about 7 million homes in the Netherlands are implemented as a double wire pair between street cabinets and customer premises. They have been pre-installed in the past as spare wire pair, for offering a second (analog) telephony line when needed. As a result, most of these extra wire pairs are unused and thus available as a bonded second wire pair for DSL equipment. This 58% refers only to those wire pairs that are fully pre-installed, but this number increases to about 95% after digging in the ground and making new splices. Deployment of VDSL2 solutions with bonding has already been started in the Netherlands, and KPN expects to deploy G.fast mainly in cases where bonding is feasible.

Bonding can (roughly) double the bit rate that can be delivered to a customer. This is such a major performance gain, that KPN cannot afford to miss a bonding feature for G.fast, especially in the view (as described above) that G.fast is going to be a long-term solution for a substantial part of the access network. Hence, KPN has a clear requirement for bonded G.fast. In the bonding vision of KPN, this requires a single CPE that connects to two wire pairs.

This contribution proposes that Q4a agrees to a goal to specify G.fast in conjunction with bonding. It is an open issue whether:

- "G.fast should support bonding" (implying that the G.fast standard shall support a 'native' bonding option as part of the transceiver, similar to SHDSL), or
- "Bonding should support G.fast" (i.e. 998.2 style, implying that changes may be needed to the G.998.2 standard).

Other options than the two described above could also be considered, e.g. the possibility to implement bonding at a higher protocol layer, and/or to pull back the bonding/reassembly functionality from the DSLAM to a higher point in the network.

At this point in time, KPN has no preference on how bonding should be specified for G.fast as long as it can be implemented with the same stability as without bonding (there may be complications caused by the retransmission scheme), and as long as the implementation loss of bonding is limited. Regarding the bonding loss, we think it is too early to specify an exact requirement for this, but a value of 5% appears as a reasonable target.

We propose to agree to the following (new) goal in the issues list:

2.4.x	Open	Should Q4a specify a method for using G.fast in conjunction with bonding ?	2012-11-4A-024

5. G.fast shall facilitate a coordination group size of up to 100

In [1] it was argued that a module size of up to 48 is too restricted for certain deployment scenarios, and that a maximum module size of up to 100 is more appropriate. The term "module" implies coordination (both TDD and vectoring) between all the ports of the module. Although expressed in terms of module size, the real issue is that a maximum of 48 for the size of the coordinated group is too small for a significant number of deployment scenarios in the Netherlands.

In the accompanying contribution [2], the deployment case of Amsterdam is studied. That analysis shows:

- There are big opportunities for G.fast to be deployed from existing manipulation points (cabinets),
- The customers within reach of G.fast from those cabinets are typically connected by 100 wire pair cables
- In a substantial number of situations high numbers of ports in the same 100 wire pair cable need to be connected, thus requiring coordination of all the connected lines.

This would require either a single (coordinated) module with up to 100 ports, or it would require the functionality to coordinate e.g. two 48-port modules (as proposed in e.g. open issue 6.5.8).

More generally, it must be realised that the economies of scale are going to be very important to build the business-case for G.fast. There will be a ***big*** business-economic pressure to deploy G.fast from the locations that maximise the number of connectable customers per location. For the Dutch situation this often means that G.fast needs to able to be used in highly filled 100 wire pair cables.

In view of the above, we propose to agree to the following existing issue in the issues list:

2.1.2.4	Open	Should the G.fast specification not preclude the implementation of FTU-O module sizes of up to 100 ports, where each wire pair counts for one port and where a bonded double wire pair counts for two ports?	2012-06-4A-039 2012-11-4A-024

In the above issue, the term "module" is meant to imply coordination (both TDD and vectoring) between all the ports of the module. To make this explicit, it could be considered to agree to the following alternative (new) issue:

2.1.2.x	Open	Should the G.fast specification support a coordination group size of up to 100 ports ?	2012-11-4A-024

6. Different classes of Distribution Point Units

The recently initiated Broadband Forum project "Fiber To The distribution point (FTTdp)" [3] distinguishes two classes of Distribution Point Unit (DPU) that are (for the time-being) denoted as "small" and "big". The requirements for these nodes could differ substantially.

- A "small" DPU could be typically deployed from a pole, or from another point in the network (e.g. a buried, sealed unit) where it is serving a relatively small number of ports (say, up to 16). A typical example is the BT deployment scenario from the top of poles. Such nodes require reverse powering, zero-touch OAM and other typical requirements discussed so far for FttDP solutions. In particular, power dissipation limits will apply to these nodes.

- A "big" DPU would be more suited for deployment from an existing street cabinet or man-hole, and thus resemble the familiar concept of a DSLAM. A typical example are the Amsterdam street cabinets described in [2], or possibly the Swisscom deployment from man-holes. Such a DPU may need to support coordination groups of up to 100 ports (as argued in the section above), which may seem a heavy requirement for G.fast. On the other hand, various other requirements (power dissipation limits, reverse powering or even zero touch OAM) may possibly be relaxed to some extent for such equipment.

These two classes of DPUs can be considered as two extremes, and intermediate implementations may be needed, e.g. a buried, sealed (water-tight) unit that must service 48 ports, but with most other requirements similar to those of a "small" DPU.

The distinction between various types of DPU is useful to understand that the system requirements may differ depending on the type of deployment. This does not imply that different classes of G.fast transceivers are needed: ideally, a single G.fast transceiver specification should be capable to support all possible classes of DPU. To capture these considerations, it is proposed to agree to the following new issue:

2.7.x	Open	Should different classes of DPUs be distinguished, that may have different sets of requirements related to e.g. coordination group size, power dissipation or reverse powering?	2012-11-4A-024

7. G.fast should have features to limit the QoE impact of disruptions

The early experiences with vectored VDSL2 revealed that the long training times of vectored VDSL2 are a real problem from a Quality-of-Experience perspective:

- A single retrain of a vectored VDSL2 link can lead to minutes of service unavailability and hence to a dissatisfied customer.
- The reset of a line card can sometimes leads to tens of minutes of service unavailability for all the connected users.

The vectored G.fast solution should avoid these issues. A possible remedy could be to have G.fast feature a short training period (i.e. the time from handshake to reaching showtime). An alternative solution could be to have a quick "warm start" feature in G.fast that allows the system to re-enter showtime quickly after a disruption, e.g. by "skipping" some steps from the full training sequence.

It is proposed to discuss these issues and reach agreement on (one or more of) the following new issues in the Qos/QoE section of the issue list. The timeframes mentioned in these issues serve as ball-park targets for further discussion.

2.10.x	Open	Should G.fast specify a "warm start" feature that allows the link to re-enter showtime within 10 seconds after a disruption?	2012-11-4A-024
2.10.y	Open	Should the initialisation time of a vectored G.fast link be less than 60 seconds ?	2012-11-4A-024
2.10.z	Open	Should a reset of a G.fast DPU result in no more than 3 minutes of unavailability for the connected customers?	2012-11-4A-024

8. References

[1] KPN, TNO, *"G.fast: KPN requirements for G.fast"*, Contribution ITU-T SG15/Q4a 2012-06-4A-039, Antwerp, Belgium, June 2012.

[2] TNO, *"G.fast: Is G.fast deployable in Amsterdam?"*, Contribution ITU-T SG15/Q4a 2012-11-4A-023, Chengdu, China, November 2012.

[3] BBF FAN & E2E, *"Fiber To The distribution point (FTTdp)"*, Broadband Forum New Project Initiation form, bbf2012.1046.03

1. Introduction | 2. Role of G.fast in the Netherlands | 3. Support of legacy functionality | 4. G.fast should be possible in conjunction with bonding | 5. G.fast shall facilitate a coordination group size of up to 100 | 6. Different classes of Distribution Point Units | 7. G.fast should have features to limit the QoE impact of disruptions | 8. References | 9. Summary

9. Summary

This paper should be presented under the G.fast agenda item, and concentrates on requirements from KPN that are not yet covered or agreed by the issues list. The list below reproduces the issues discussed in this contribution. It is proposed to agree to these issues.

2.4.x	Open	Should Q4a specify a method for using G.fast in conjunction with bonding ?	2012-11-4A-024
2.1.2.4	Open	Should the G.fast specification not preclude the implementation of FTU-O module sizes of up to 100 ports, where each wire pair counts for one port and where a bonded double wire pair counts for two ports?	2012-06-4A-039 2012-11-4A-024
2.1.2.x	Open	Should the G.fast specification support a coordination group size of up to 100 ports ?	2012-11-4A-024
2.7.x	Open	Should different classes of DPUs be distinguished, that may have different sets of requirements related to e.g. coordination group size, power dissipation or reverse powering?	2012-11-4A-024
2.10.x	Open	Should G.fast specify a "warm start" feature that allows the link to re-enter showtime within 10 seconds after a disruption?	2012-11-4A-024
2.10.y	Open	Should the initialisation time of a vectored G.fast link be less than 60 seconds ?	2012-11-4A-024
2.10.z	Open	Should a reset of a G.fast DPU result in no more than 3 minutes of unavailability for the connected customers?	2012-11-4A-024