

Network Working Group
Internet Draft
Updates: 3471, 6205 (if approved)
Intended Status: Standards Track
Expires: 10 March 2016

A. Farrel
D. King
Old Dog Consulting
Y. Li
Nanjing University
F. Zhang
Huawei Technologies

10 September 2015

Generalized Labels for the Flexi-Grid in
Lambda Switch Capable (LSC) Label Switching Routers

draft-ietf-ccamp-flexigrid-lambda-label-05.txt

Abstract

GMPLS supports the description of optical switching by identifying entries in fixed lists of switchable wavelengths (called grids) through the encoding of lambda labels. Work within the ITU-T Study Group 15 has defined a finer granularity grid, and the facility to flexibly select different widths of spectrum from the grid. This document defines a new GMPLS lambda label format to support this flexi-grid.

This document updates RFC 3471 and RFC 6205 by introducing a new label format.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Conventions Used in This Document	4
2. Overview of Flexi-Grid	4
2.1. Composite Labels	4
3. Fixed Grid Lambda Label Encoding	5
4. Flexi-Grid Label Format and Values	5
4.1 Flexi-Grid Label Encoding	5
4.2. Considerations of Bandwidth	7
4.3. Composite Labels	7
5. Manageability and Backward Compatibility Considerations	8
5.1. Control Plane Backward Compatibility	9
5.2. Manageability Considerations	9
6. Implementation Status	10
6.1. Centre Tecnologic de Telecomunicacions de Catalunya (CTTC)	10
7. Security Considerations	11
8. IANA Considerations	12
8.1. Grid Subregistry	12
8.2. DWDM Channel Spacing Subregistry	12
9. Acknowledgments	12
10. References	13
10.1. Normative References	13
10.2. Informative References	13
Appendix A. Flexi-Grid Example	15
Authors' Addresses	16
Contributors' Addresses	16

1. Introduction

As described in [RFC3945], GMPLS extends MPLS from supporting only Packet Switch Capable (PSC) interfaces and switching, to also support four new classes of interfaces and switching that include Lambda Switch Capable (LSC).

A functional description of the extensions to MPLS signaling needed to support this new class of interface and switching is provided in [RFC3471].

Section 3.2.1.1 of [RFC3471] states that wavelength labels "only have significance between two neighbors": global wavelength semantics are not considered. [RFC6205] defines a standard lambda label format that has a global semantic and which is compliant with both the Dense Wavelength Division Multiplexing (DWDM) grid [G.694.1] and the Coarse Wavelength Division Multiplexing (CWDM) grid [G.694.2]. The terms DWDM and CWDM are defined in [G.671].

A flexible grid network selects its data channels as arbitrarily assigned pieces of the spectrum. Mixed bitrate transmission systems can allocate their channels with different spectral bandwidths so that the channels can be optimized for the bandwidth requirements of the particular bit rate and modulation scheme of the individual channels. This technique is regarded as a promising way to improve the network utilization efficiency and fundamentally reduce the cost of the core network.

The "flexi-grid" has been developed within the ITU-T Study Group 15 to allow selection and switching of pieces of the optical spectrum chosen flexibly from a fine granularity grid of wavelengths with variable spectral bandwidth [G.694.1].

[RFC3471] defines several basic label types including the lambda label. [RFC3471] states that wavelength labels "only have significance between two neighbors" (Section 3.2.1.1); global wavelength semantics are not considered. In order to facilitate interoperability in a network composed of LSC equipment, [RFC6205] defines a standard lambda label format and is designated an update of RFC 3471.

This document continues the theme of defining global semantics for the wavelength label by adding support for the flexi-grid. Thus, this document updates [RFC6205] and [RFC3471].

This document relies on [G.694.1] for the definition of the optical data plane and does not make any updates to the work of the ITU-T.

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Overview of Flexi-Grid

[G.694.1] defines DWDM fixed grids. The latest version of that document extends the DWDM fixed grids to add support for flexible grids. The basis of the work is to allow a data channel to be formed from an abstract grid anchored at 193.1 THz and selected on a channel spacing of 6.25 GHz with a variable slot width measured in units of 12.5 GHz. Individual allocations may be made on this basis from anywhere in the spectrum, subject to allocations not overlapping.

[G.694.1] provides clear guidance on the support of flexible grid by implementations in Section 2 of Appendix I:

The flexible DWDM grid defined in clause 7 has a nominal central frequency granularity of 6.25 GHz and a slot width granularity of 12.5 GHz. However, devices or applications that make use of the flexible grid may not have to be capable of supporting every possible slot width or position. In other words, applications may be defined where only a subset of the possible slot widths and positions are required to be supported.

For example, an application could be defined where the nominal central frequency granularity is 12.5 GHz (by only requiring values of n that are even) and that only requires slot widths as a multiple of 25 GHz (by only requiring values of m that are even).

Some additional background on the use of GMPLS for flexible grids can be found in [FLEXFWRK].

2.1. Composite Labels

It is possible to construct an end-to-end connection as a composite of more than one flexi-grid slot. The mechanism used in GMPLS is similar to that used to support inverse multiplexing familiar in time-division multiplexing (TDM) and optical transport networks (OTN). The slots in the set could potentially be contiguous or non-contiguous (only as allowed by the definitions of the data plane) and could be signaled as a single LSP or constructed from a group of LSPs. For more details, refer to Section 4.3.

How the signal is carried across such groups of channels is out of scope for this document.

3. Fixed Grid Lambda Label Encoding

[RFC6205] defines an encoding for a global semantic for a DWDM label based on four fields:

- Grid: used to select which grid the lambda is selected from. Values defined in [RFC6205] identify DWDM [G.694.1] and CWDM [G.694.2].
- C.S. (Channel Spacing): used to indicate the channel spacing. [RFC6205] defines values to represent spacing of 100, 50, 25 and 12.5 GHz.
- Identifier: a local-scoped integer used to distinguish different lasers (in one node) when they can transmit the same frequency lambda.
- n: a two's-complement integer to take a positive, negative, or zero value. This value is used to compute the frequency as defined in [RFC6205] and based on [G.694.1]. The use of n is repeated here for ease of reading the rest of this document: in case of discrepancy, the definition in [RFC6205] is normative.

$$\text{Frequency (THz)} = 193.1 \text{ THz} + n * \text{frequency granularity (THz)}$$

where the nominal central frequency granularity for the flexible grid is 0.00625 THz

4. Flexi-Grid Label Format and Values

4.1 Flexi-Grid Label Encoding

This document defines a generalized label encoding for use in flexi-grid systems. As with the other GMPLS lambda label formats defined in [RFC3471] and [RFC6205], the use of this label format is known a priori. That is, since the interpretation of all lambda labels is determined hop-by-hop, the use of this label format requires that all nodes on the path expect to use this label format.

For convenience, however, the label format is modeled on the fixed grid label defined in [RFC6205] and briefly described in Section 3.

Figure 1 shows the format of the Flexi-Grid Label. It is a 64 bit label.

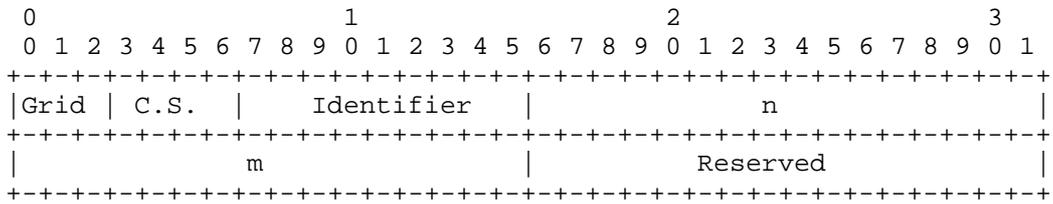


Figure 1 : The Flexi-Grid Label Encoding

This document defines a new Grid value to supplement those in [RFC6205]:

Grid	Value
ITU-T Flex	3

Within the fixed grid network, the C.S. value is used to represent the channel spacing, as the spacing between adjacent channels is constant. For the flexible grid situation, this field is used to represent the nominal central frequency granularity.

This document defines a new C.S. value to supplement those in [RFC6205]:

C.S.(GHz)	Value
6.25	5

The meaning of the Identifier field is maintained from [RFC6205] (see also Section 3).

The meaning of n is maintained from [RFC6205] (see also Section 3).

The m field is used to identify the slot width according to the formula given in [G.694.1] as follows. It is a 16 bit integer value encoded in line format.

$$\text{Slot Width (GHz)} = 12.5 \text{ GHz} * m$$

The Reserved field MUST be set to zero on transmission and SHOULD be ignored on receipt.

An implementation that wishes to use the flexi-grid label encoding

MUST follow the procedures of [RFC3473] and of [RFC3471] as updated by [RFC6205]. It MUST set Grid to 3 and C.S. to 5. It MUST set Identifier to indicate the local identifier of the laser in use as described in [RFC6205]. It MUST also set n according to the formula in Section 3 (inherited unchanged from [RFC6205]). Finally, the implementation MUST set m as described in the formula stated above.

4.2. Considerations of Bandwidth

There is some overlap between the concepts of bandwidth and label in many GMPLS-based systems where a label indicates a physical switching resource. This overlap is increased in a flexi-grid system where a label value indicates the slot width and so affects the bandwidth supported by an LSP. Thus the 'm' parameter is both a property of the label (i.e., it helps define exactly what is switched) and of the bandwidth.

In GMPLS signaling [RFC3473], bandwidth is requested in the TSpec object and confirmed in the Flowspec object. The 'm' parameter that is a parameter of the GMPLS flexi-grid label as described above, is also a parameter of the flexi-grid TSpec and Flowspec as described in [FLEXRSVP].

4.3. Composite Labels

The creation of a composite of multiple channels to support inverse multiplexing is already supported in GMPLS for TDM and OTN [RFC4606], [RFC6344], [RFC7139]. The mechanism used for flexi-grid is similar.

To signal an LSP that uses multiple flexi-grid slots a "compound label" is constructed. That is, the LABEL object is constructed from a concatenation of the 64-bit Flexi-Grid Labels shown in Figure 1. The number of elements in the label can be determined from the length of the LABEL object. The resulting LABEL object is shown in Figure 2 including the object header that is not normally shown in diagrammatic representations of RSVP-TE objects. Note that r is the count of component labels, and this is backward compatible with the label shown in Figure 1 where the value of r is 1.

The order of component labels MUST be presented in increasing order of the value n. Implementations MUST NOT infer anything about the encoding of a signal into the set of slots represented by a compound label from the label itself. Information about the encoding MAY be handled in other fields in signaling messages or through an out of band system, but such considerations are out of the scope of this document.

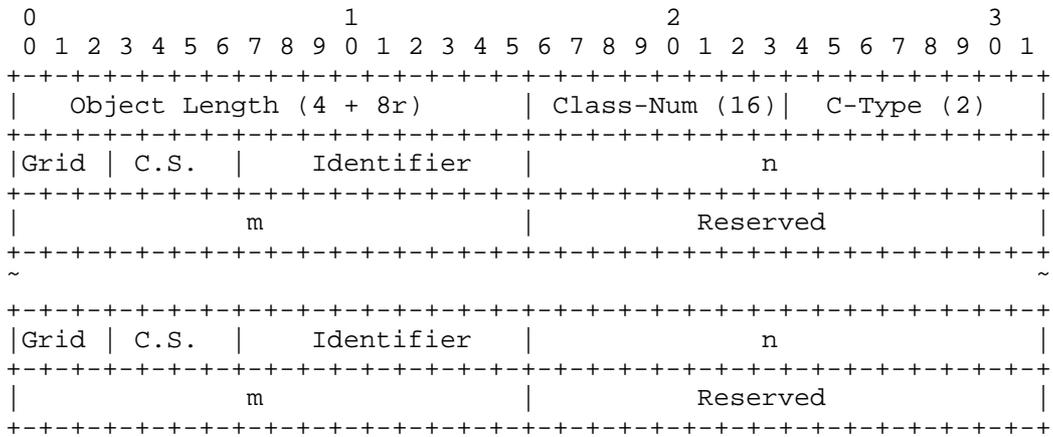


Figure 2 : A Compound Label for Virtual Concatenation

Note that specific rules must be applied as follows:

- Grid MUST show "ITU-T Flex" value 3 in each component label.
- C.S. MUST have the same value in each component label.
- Identifier in each component label may identify different physical equipment.
- Values of n and m in each component label define the slots that are concatenated.

At the time of writing [G.694.1] only supports only groupings of adjacent slots (i.e., without intervening unused slots that could be used for other purposes) of identical width (same value of m), and the component slots must be in increasing order of frequency (i.e., increasing order of the value n). The mechanism defined here MUST NOT be used for other forms of grouping unless and until those forms are defined and documented in Recommendations published by the ITU-T.

Note further that while the mechanism described here naturally means that all component channels are corouted, a composite channel can also be achieved by constructing individual LSPs from single flexi-grid slots and managing those LSPs as a group. A mechanism for achieving this for TDM is described in [RFC6344], but is out of scope for discussion in this document because the labels used are normal, single slot labels and require no additional definitions.

5. Manageability and Backward Compatibility Considerations

This section briefly considers issues of manageability and backward compatibility.

5.1. Control Plane Backward Compatibility

Labels are carried in two ways in GMPLS: for immediate use on the next hop and for use at remote hops.

It is an assumption of GMPLS that both ends of a link know what label types are supported and only use appropriate label types. If a label of an unknown type is received it will be processed as if it was of a known type since the Label Object and similar label-carrying objects do not contain a type identifier. Thus the introduction of a flexi-grid label in this document does not change the compatibility issues and a legacy node that does not support the new flexi-grid label should not expect to receive or handle such labels. If one is incorrectly used in communication with a legacy node it will attempt to process it as an expected label type with a potentially poor outcome.

It is possible that a GMPLS message transitting a legacy node will contain a flexi-grid label destined for or reported by a remote node. For example, an LSP that transits links of different technologies might record flexi-grid labels in a Record Route Object that is subsequently passed to a legacy node. Such labels will not have any impact on legacy implementations except as noted in the manageability considerations in the next section.

5.2. Manageability Considerations

This document introduces no new elements for management. That is, labels can continue to be used in the same way by the GMPLS protocols and where those labels were treated as opaque quantities with local or global significance, no change is needed to the management systems.

However, this document introduces some changes to the nature of a label that may require changes to management systems. Although Section 3.2 of [RFC3471] makes clear that a label is of variable length according to the type and that the type is supposed to be known a priori by both ends of a link, a management system is not guaranteed to be updated in step with upgrades or installations of new flexi-grid functionality in the network.

But an implementation expecting a 32 bit lambda label would not fail ungracefully because the first 32 bits follow the format of [RFC6205]. It would look at theses labels and read but not recognize the new grid type value. It would then give up trying to parse the label and (presumably) the whole of the rest of the message.

The management system can be upgraded in two steps:

- Firstly, systems that handle lambda labels as 32 bit quantities need to be updated to handle the increase length (64 bits) of labels as described in this document. These "unknown" 64 bit labels could be displayed as opaque 64 bit quantities and still add a lot of value for the operator (who might need to parse the label by hand). However, an implementation that already supports lambda labels as defined in [RFC6205] can safely continue to process the first 32 bits and display the fields defined in RFC 6205 as before leaving just the second 32 bits as opaque data.
- Second, a more sophisticated upgrade to a management system would fully parse the flex-gird labels and display them field-by-field as described in this document.

6. Implementation Status

[RFC Editor Note: Please remove this entire section prior to publication as an RFC.]

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in RFC 6982 [RFC6982]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to RFC 6982, "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit."

6.1. Centre Tecnologic de Telecomunicacions de Catalunya (CTTC)

Organization Responsible for the Implementation:

Centre Tecnologic de Telecomunicacions de Catalunya (CTTC)
Optical Networks and Systems Department

Implementation Name and Details:

ADRENALINE testbed
<http://networks.cttc.es/experimental-testbeds/>

Brief Description:

Experimental testbed implementation of GMPLS/PCE control plane.

Level of Maturity:

Implemented as extensions to a mature GMLPS/PCE control plane.
It is limited to research / prototyping stages but it has been used successfully for more than the last five years.

Coverage:

Support for the 64 bit label as described version 07 of this document.

This affects mainly the implementation of RSVP-TE and PCEP protocols:

- Generalized Label Support
 - Suggested Label Support
 - Upstream Label Support
 - ERO Label Subobjects and Explicit Label Control
- It is expected that this implementation will evolve to follow the evolution of this document.

Licensing:

Proprietary

Implementation Experience:

Implementation of this document reports no issues.
General implementation experience has been reported in a number of journal papers. Contact Ramon Casellas for more information or see http://networks.cttc.es/publications/?search=GMPLS&research_area=optical-networks-systems

Contact Information:

Ramon Casellas: ramon.casellas@cttc.es

Interoperability:

No report.

7. Security Considerations

[RFC6205] notes that the definition of a new label encoding does not introduce any new security considerations to [RFC3471] and [RFC3473]. That statement applies equally to this document.

For a general discussion on MPLS and GMPLS-related security issues,

see the MPLS/GMPLS security framework [RFC5920].

8. IANA Considerations

IANA maintains the "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters" registry that contains several subregistries.

8.1. Grid Subregistry

IANA is requested to allocate a new entry in this subregistry as follows:

Value	Grid	Reference
3	ITU-T Flex	[This.I-D]

8.2. DWDM Channel Spacing Subregistry

IANA is requested to allocate a new entry in this subregistry as follows:

Value	Channel Spacing (GHz)	Reference
5	6.25	[This.I-D]

9. Acknowledgments

This work was supported in part by the FP-7 IDEALIST project under grant agreement number 317999.

Very many thanks to Lou Berger for discussions of labels of more than 32 bits. Many thanks to Sergio Belotti and Pietro Vittorio Grandi for their support of this work. Thanks to Gabriele Galimberti for discussion of the size of the "m" field, and to Iftekhar Hussain for discussion of composite labels. Robert Sparks, Carlos Pignataro, and Paul Wouters provided review comments during IETF last call.

Special thanks to the Vancouver 2012 Pool Party for discussions and rough consensus: Dieter Beller, Ramon Casellas, Daniele Ceccarelli, Oscar Gonzalez de Dios, Iftekhar Hussain, Cyril Margaria, Lyndon Ong, Fatai Zhang, and Adrian Farrel.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC6205] Otani, T., and Li, D., "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, October 2011.
- [G.694.1] ITU-T Recommendation G.694.1 (revision 2), "Spectral grids for WDM applications: DWDM frequency grid", February 2012.

10.2. Informative References

- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
 - [RFC4606] Mannie, E., and Papadimitriou, D., "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 4606, August 2006.
 - [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
 - [RFC6344] Bernstein, G., Caviglia, D., Rabbat, R., and van Helvoort, H., "Operating Virtual Concatenation (VCAT) and the Link Capacity Adjustment Scheme (LCAS) with Generalized Multi-Protocol Label Switching (GMPLS)", RFC 6344, August 2011.
 - [RFC6982] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", RFC 6982, July 2013.
- [RFC Editor Note: This reference can be removed when Section 6 is removed]

- [RFC7139] Zhang, F., Zhang, G., Belotti, S., Ceccarelli, D., and Pithewan, K., "GMPLS Signaling Extensions for Control of Evolving G.709 Optical Transport Networks", RFC 7139, March 2014.

- [G.671] ITU-T Recommendation G.671, "Transmission characteristics of optical components and subsystems", 2009.

- [G.694.2] ITU-T Recommendation G.694.2, "Spectral grids for WDM applications: CWDM wavelength grid", December 2003.

- [FLEXFWRK] O. Gonzalez de Dios, et al., "Framework and Requirements for GMPLS based control of Flexi-grid DWDM networks", draft-ogrcetal-ccamp-flexi-grid-fwk, work in progress.

- [FLEXRSVP] Zhang, F., Gonzalez de Dios, O., and D. Ceccarelli, "RSVP-TE Signaling Extensions in support of Flexible Grid", draft-zhang-ccamp-flexible-grid-rsvp-te-ext, work in progress.

Appendix A. Flexi-Grid Example

Consider a fragment of an optical LSP between node A and node B using the flexible grid. Suppose that the LSP on this hop is formed:

- using the ITU-T Flexi-Grid
- the nominal central frequency of the slot 193.05 THz
- the nominal central frequency granularity is 6.25 GHz
- the slot width is 50 GHz.

In this case the label representing the switchable quantity that is the flexi-grid quantity is encoded as described in Section 4.1 with the following parameter settings. The label can be used in signaling or in management protocols to describe the LSP.

Grid = 3 : ITU-T Flexi-Grid

C.S. = 5 : 6.25 GHz nominal central frequency granularity

Identifier = local value indicating the laser in use

n = -8 :

Frequency (THz) = 193.1 THz + n * frequency granularity (THz)

193.05 (THz) = 193.1 (THz) + n * 0.00625 (THz)

n = (193.05-193.1)/0.00625 = -8

m = 4 :

Slot Width (GHz) = 12.5 GHz * m

50 (GHz) = 12.5 (GHz) * m

m = 50 / 12.5 = 4

Authors' Addresses

Adrian Farrel
Old Dog Consulting
EMail: adrian@olddog.co.uk

Daniel King
Old Dog Consulting
EMail: daniel@olddog.co.uk

Yao Li
Nanjing University
EMail: wsliguotou@hotmail.com

Fatai Zhang
Huawei Technologies
EMail: zhangfatai@huawei.com

Contributors' Addresses

Zhang Fei
Huawei Technologies
EMail: zhangfei7@huawei.com

Ramon Casellas
CTTC
EMail: ramon.casellas@cttc.es

