

none
Internet-Draft
Intended status: Informational
Expires: September 9, 2020

X. de Foy
A. Rahman
InterDigital Inc.
A. Galis
University College London
K. Makhijani
L. Qiang
Huawei Technologies
S. Homma
NTT
P. Martinez-Julia
NICT
March 8, 2020

Interconnecting (or Stitching) Network Slice Subnets
draft-defoy-coms-subnet-interconnection-04

Abstract

This document defines the network slice (NS) subnet as a general management plane concept that augments a baseline YANG network slice model with management attributes and operations enabling interconnections (or stitching) between network slices. The description of NS subnet interconnections is technology agnostic, and is not tied to a particular implementation of the interconnection in data plane.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 9, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://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	2
1.1.	Motivation and Roles of NS Subnet	3
1.2.	Usage of NS Subnets	3
1.3.	Terminology	5
2.	Information Model	5
2.1.	Base Information Model	5
2.2.	Interconnection Anchors	6
2.3.	Interconnection Instances	8
2.4.	Stitching Operation	9
2.4.1.	Operation Overview	9
2.4.2.	Stitching Scenarios	10
3.	Security Considerations	11
4.	IANA Considerations	11
5.	Informative References	11
	Authors' Addresses	12

[1.](#) Introduction

Network Slicing enables deployment and management of services with diverse requirements on end-to-end partitioned virtual networks over the same infrastructure, including networking, compute and storage resources. There were recent efforts in the IETF to define a transport slice ([\[I-D.nsdtd-teas-transport-slice-definition\]](#)) and to define a north-bound interface for such a transport slice ([\[I-D.contreras-teas-slice-nbi\]](#)). The mapping of transport slices in 5G mobile systems is also studied in [\[I-D.clt-dmm-tn-aware-mobility\]](#) and [\[I-D.geng-teas-network-slice-mapping\]](#).

Network slices may be managed through usage of YANG data models. For example, [\[I-D.liu-teas-transport-network-slice-yang\]](#) describes how existing YANG models can be augmented with network slice attributes.

Nevertheless, defining and managing a network slice (NS) end-to-end does not always have to be done directly. It may be convenient to define and manage separately subsets of an end-to-end slice. The concept of network slice subnet is defined originally in [\[NGMN Network Slicing\]](#), though we only need to retain its definition in the most universal form: network slice subnets are similar to network slices in most ways but cannot be operated in isolation as a complete network slice (e.g., a NS subnet can be seen as a network slice with unconnected links). NS subnets are interconnected with other NS subnets to form a complete, end-to-end network slice (i.e. interconnection and/or stitching of NS subnets). In the present draft, we describe a data model for describing interconnections between NS subnets, that enables assembling them in a hierarchical fashion.

[1.1.](#) Motivation and Roles of NS Subnet

NS subnet is a management plane concept that facilitates interconnections (also known as stitching) of network slices. It augments the base slice information model, that can be used to represent an end-to-end network slice. The extensions described in this document can be used to represent a slice subnet instead, and can also be used to represent an interconnection inside an end-to-end slice, i.e. they aim to represent interconnection points both "before" and "after" the interconnection takes place. Operations such as stitching subnets are also described.

The description of NS subnet interconnections is technology agnostic following the approach of the slice information model. Some interconnections may be implemented using the interplay between management plane and gateways in the data plane. [\[I-D.homma-rtgwg-slice-gateway\]](#) describes the requirements on such data plane network elements, and will provide input for the management plane mechanisms described in the present document.

[1.2.](#) Usage of NS Subnets

Using NS subnets can help:

- o Isolate management and maintenance of different portions of a network slice, over multiple infrastructure domains, or even within a single domain. For example, in Figure 1, NS orchestrator (NSO) 2 manages subnet A, in isolation from subnets B and C managed by NSO 3. NSO 1 can still manage the end-to-end slice as a whole, but it does not need to deal in detail with each subnet.
- o Isolate mapping towards different infrastructure technologies, even within the same domain. This can simplify NS orchestrator

implementation, since each NSO can specialize in managing a smaller set of technologies.

- o Enable advanced functions such as sharing a slice subnet between several slices, or substituting one slice subnet for another, e.g. for coping with load.

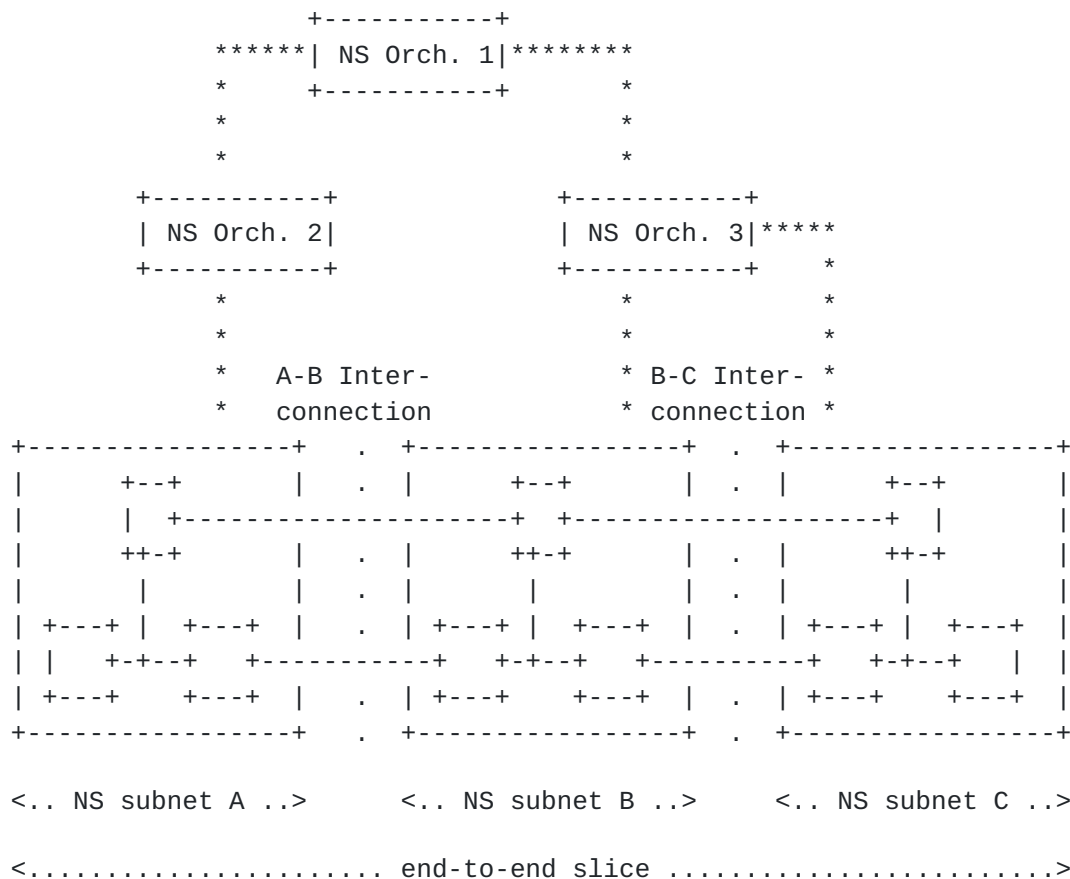


Figure 1: Overview of Network Slice Subnets Interconnection

Figure 1 illustrates how an end-to-end network slice may be composed of multiple slice subnets, each managed independently by a same or different NSO. In multi-administrative domain scenarios, using NS subnets can help limiting the information that needs to be shared between domains. At the infrastructure layer (i.e. in the data plane), the interconnection between NS subnets may involve:

- o a gateway, that performs protocol and/or identifier/label translation as needed,
- o two gateways, especially in cases where interconnected NS subnets are in different administrative domains,

- o nothing at all, in cases where the interconnection point can be abstracted away, e.g. when the NS subnets share a common infrastructure. In this case nodes from both NS subnets end up being directly interconnected between each other.

More detailed usage scenarios are described in [Section 2.4.2](#).

[1.3.](#) Terminology

Network slicing terminology, especially focusing on transport slices, is defined in [[I-D.nsd-t-teas-transport-slice-definition](#)].

Network Slice Subnet (NS subnet): a network slice designed to be interconnected with other network slices.

NS Stitching: a management operation consisting in creating an end-to-end NS or a larger NS subnet, by interconnecting a set of NS subnets together.

Interconnection Anchor: a management plane entity, part of a NS subnet model, representing an end point for use in future stitching operation.

Interconnection Instance (or Interconnect): a management plane entity, part of a NS subnet model, representing an interconnection realized by a stitching operation. It is distinct from a (data plane) gateway: an interconnect may be realized with or without using a gateway in the data plane.

[2.](#) Information Model

[2.1.](#) Base Information Model

The information model we use as base for network slicing is the network topology model ietf-network defined in [[RFC8345](#)], in which networks are composed of nodes and links, and in which termination points (TP), defined in nodes, are used to define source and destination of links.

A network slice data model instance, i.e. a YANG data model augmented using [[I-D.liu-teas-transport-network-slice-yang](#)]), represents a network slice. When such a data model instance includes at least an "interconnection anchor", as defined below, it represents a network slice subnet instance.

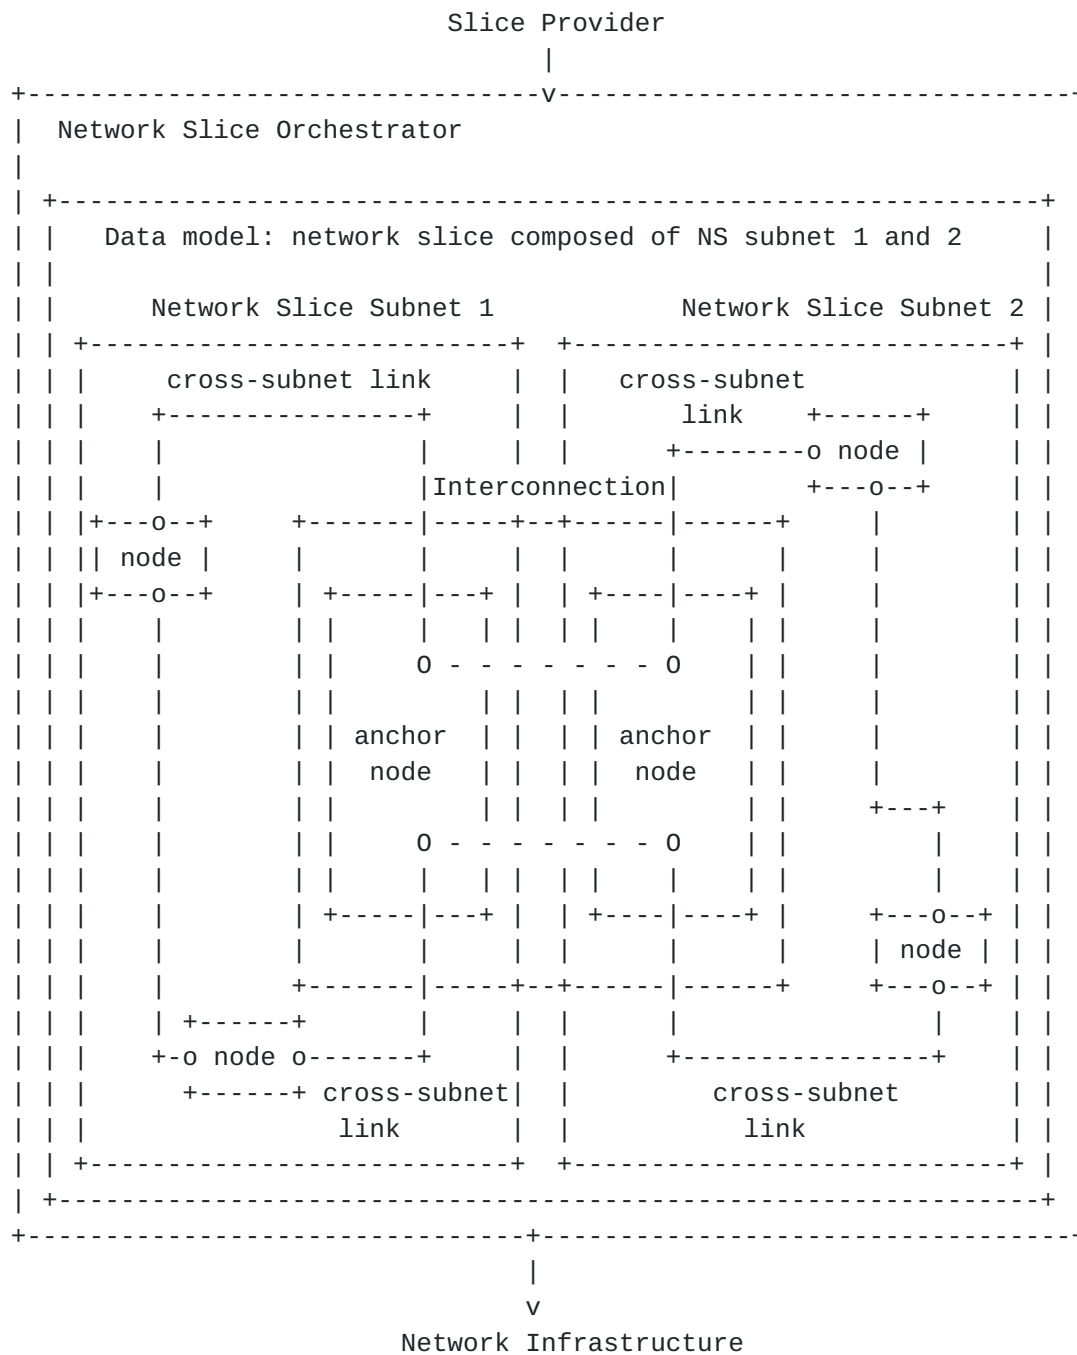
At high level, the extensions defined in this document will augment nodes and termination points:


```
module: ietf-network
+--rw networks
  +--rw network* [network-id]
    +--rw network-id
    +--rw network-types
    +--rw supporting-network* [network-ref]
      | +--rw network-ref
    +--rw node* [node-id]
      | +--... (augmented with attributes for
      | |      anchor/interconnection nodes)
      | +--rw nt:termination-point* [tp-id]
      | |   ... (augmented with attributes for
      | |      anchor/interconnection TP)
```

2.2. Interconnection Anchors

To represent an anchor point for future interconnections (i.e. an unconnected end of a link), a simple solution is to use an "interconnection anchor" termination point (or anchor TP). Within the data model describing a subnet, any link not entirely contained within the NS subnet must be terminated with such an anchor TP as source or destination. An anchor TP belongs to a "node" attribute, which we refer to as interconnection anchor node (or anchor node). Several anchor TPs can be grouped together in an anchor node, and such grouping may be used as a hint during a stitching operation (e.g. to place all interconnection points at a same location).

Figure 2 represents 2 interconnected network slice subnets.



Legend: o = termination point, O = anchor termination point

Figure 2: Network Slice Subnets Interconnection

Attributes of interconnection anchor nodes and termination points include:

- o Information enabling NS orchestrators to match anchor nodes and TPs from both NS during a stitching operation. A label may be a simple way to enable this.
- o Information to help locate the interconnection. For example, it could be a (sub-)domain name or geo-location information, that indicates where the interconnection point should be located. This can help for example in cases where the subnet is instantiated before stitching.
- o Information to help select the type of interconnection establishment: for example, this can indicate a preference for using interconnection over a gateway, or for abstracting away the interconnection point in the infrastructure plane.

```

+--rw node* [node-id]
  +-- (...)
  +-- anchor_node_config
    |   +-- label (and/or other auto stitching help)
    |   +-- hint for location (domain, geolocation, etc.)
    |   +-- hint for type (1 gateway, 2 gateways, ...)
  +--rw nt:termination-point* [tp-id]
    +-- (...)
    +-- anchor_tp_config
      +-- label (and/or other auto stitching help)
      +-- location (domain, geolocation, etc.)
      +-- type (1 gateway, 2 gateways, ...)

```

2.3. Interconnection Instances

There are two options for representing post-stitching network slices (or subnets). They are not mutually exclusive:

- o Option 1: subnet data models are updated with information describing the interconnection (e.g. anchor TPs and nodes are updated with new attributes representing the existing connection, if necessary).
- o Option 2: a new data model is generated to represent the resulting network slice (or subnet). In this composite data model, the interconnection may or may not be represented, this can be a choice made by the operator.

Option 1 and 2 can be used concurrently in a network. For example, a parent NS orchestrator may manage stitched NS subnets through underlying NS orchestrators, and at the same time expose to the NS operator a composite data model representing the resulting end-to-end slice.

To represent an existing interconnection in option 1, a simple solution is to add attributes to existing anchor nodes and anchor TPs. Those attributes will be described below. They aim to describe state and configuration associated with an active interconnection.

To represent an existing interconnection in option 2, a simple solution is to create new interconnection instance nodes and termination point. The same attributes as in option 1 may be associated with these nodes and TPs.

Attributes of interconnection instance nodes and termination points include:

- o State information (interconnection type, status, location...).
- o Service assurance related information: besides measurements (on throughput, loss rate, etc.), triggers depending on throughput, latency, etc. can be linked with a management action or event. A NS operator can use such events to take the decision to disable a NS subnet, replace a NS subnet with another, etc. to maintain overall service performance.

```

+--rw node* [node-id]
  +-- (...)
  +-- interconnection_instance_node_state
    | +-- status
    | +-- location (domain, geolocation, etc.)
    | +-- type (1 gateway, 2 gateways, ...)
  +-- interconnection_instance_node_service_assurance
    | +-- events (including triggers and event IDs)
    | +-- measurements
  +--rw nt:termination-point* [tp-id]
    +-- (...)
    +-- interconnection_instance_tp_state
      | +-- status
      | +-- location (domain, geolocation, etc.)
      | +-- type (1 gateway, 2 gateways, ...)
    +-- interconnection_instance_node_service_assurance
      +-- events (including triggers and event IDs)
      +-- measurements

```

2.4. Stitching Operation

2.4.1. Operation Overview

Stitching is an operation that takes two or more NS subnets as input, and produces a single composite NS subnet or end-to-end slice. It may occur when the slice subnets are being instantiated, or later.

The first step in this operation is to identify the anchors that will be used in the interconnection. This may be done by an automated algorithm that matches the possible interconnection points and decides which one will be used, according to the policies established by the NS operator. The operation in this case will require the presence of semantically-rich attributes in the candidate anchors to enable automatic matching without human intervention.

Other attributes of slices and anchors will also influence the operation and the resulting stitched (composite) object. For instance, network links that are interconnected must have compatible QoS attributes. Moreover, available networking protocols must also match among the underlying network elements that are being stitched. Otherwise, the operation will fail unless the NS operator (based on policy and/or NS subnet attributes) enables it to search for, and use, some "bridge" element in the underlying infrastructure.

2.4.2. Stitching Scenarios

This section briefly describes examples of usage for subnet stitching.

Traversal through a transport network.

Let's consider a network slice composed of (NS) subnet-A, and subnet-C (Figure 3). Subnet-A and subnet-C are deployed in independent domains and are mapped into a slice information model; in order to stitch these two together a transport segment is needed. N1 and N2 are anchor nodes within NS subnets A and C. Segment-B could be a simple link between the two NS subnets but it may also be a TE-link made available by a transport network provider. Segment-B may be involved in the stitching operation in one of several ways:

Segment-B may be set up as part of the stitching operation between NS subnets A and C, as a form of "bridge" mentioned in [Section 2.4](#). Segment-B will need to comply with service specific traffic constraints that are determined during the stitching operation, possibly using attributes from NS subnets A and C. In this case, the data plane implementation of N1 and N2 in the composite slice may be, for example, 2 distinct gateway functions terminating segment-B.

Segment-B may alternatively be represented as a distinct NS subnet, e.g. in cases where segment-B is complex and/or involves multiple network functions. In this case, the stitching operation may therefore involve 3 NS subnets A-B-C.




Figure 3: Example of NS subnets interconnection through transport network

Subnets in a single domain.

In this scenario multiple network slice subnets are defined as basic building blocks with specific service functions (or chains), topologies and traffic handling characteristics. These building blocks can be assembled through stitching to build end-to-end customized slices, but also to dynamically extend slices to adapt to traffic load. Additionally, stitching can also be used to share building blocks between multiple slices, e.g. to interconnect multiple slices with a shared function. In all these cases, interconnection instances may be entirely abstracted away, although they may also be implemented through one or multiple gateways, e.g. when stitched subnets belong to different sub-domains.

3. Security Considerations

Security aspects relative to network slices (e.g., for transport slices, in [[I-D.liu-teas-transport-network-slice-yang](#)]) are applicable to slice subnets, including transport security aspects, access control and protection of write operation on newly introduced nodes (e.g., termination-point).

4. IANA Considerations

This document has no actions for IANA.

5. Informative References

[I-D.clt-dmm-tn-aware-mobility]

Chunduri, U., Li, R., Bhaskaran, S., Kaippallimalil, J., Tantsura, J., Contreras, L., and P. Muley, "Transport Network aware Mobility for 5G", [draft-clt-dmm-tn-aware-mobility-05](#) (work in progress), November 2019.

[I-D.contreras-teas-slice-nbi]

Contreras, L., Homma, S., and J. Ordonez-Lucena,
"Considerations for defining a Transport Slice NBI",
[draft-contreras-teas-slice-nbi-00](#) (work in progress),
November 2019.

[I-D.geng-teas-network-slice-mapping]

Geng, X., Dong, J., Niwa, T., and J. Jin, "5G End-to-end
Network Slice Mapping from the view of Transport Network",
[draft-geng-teas-network-slice-mapping-00](#) (work in
progress), February 2020.

[I-D.homma-rtgwg-slice-gateway]

Homma, S., Foy, X., Galis, A., and L. Contreras, "Gateway
Function for Network Slicing", [draft-homma-rtgwg-slice-
gateway-01](#) (work in progress), November 2019.

[I-D.liu-teas-transport-network-slice-yang]

Liu, X., Tantsura, J., Bryskin, I., Contreras, L., and Q.
WU, "Transport Network Slice YANG Data Model", [draft-liu-
teas-transport-network-slice-yang-00](#) (work in progress),
November 2019.

[I-D.nsdt-teas-transport-slice-definition]

Rokui, R., Homma, S., and K. Makhijani, "IETF Definition
of Transport Slice", [draft-nsdt-teas-transport-slice-
definition-00](#) (work in progress), November 2019.

[NGMN_Network_Slicing]

NGMN, "Description of Network Slicing Concept", 10 2016,
<[https://www.ngmn.org/uploads/
media/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf](https://www.ngmn.org/uploads/media/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf)>.

[RFC8345]

Clemm, A., Medved, J., Varga, R., Bahadur, N.,
Ananthakrishnan, H., and X. Liu, "A YANG Data Model for
Network Topologies", [RFC 8345](#), DOI 10.17487/RFC8345, March
2018, <<https://www.rfc-editor.org/info/rfc8345>>.

Authors' Addresses

Xavier de Foy
InterDigital Inc.
1000 Sherbrooke West
Montreal
Canada

Email: Xavier.Defoy@InterDigital.com

Akbar Rahman
InterDigital Inc.
1000 Sherbrooke West
Montreal
Canada

Email: Akbar.Rahman@InterDigital.com

Alex Galis
University College London
Torrington Place
London WC1E 7JE
United Kingdom

Email: a.galis@ucl.ac.uk

Kiran Makhijani
Huawei Technologies
2890 Central Expressway
Santa Clara CA 95050
USA

Email: kiran.makhijani@huawei.com

Li Qiang
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: qiangli3@huawei.com

Shunsuke Homma
NTT, Corp.
3-9-11, Midori-cho
Musashino-shi, Tokyo 180-8585
Japan

Email: homma.shunsuke@lab.ntt.co.jp

Pedro Martinez-Julia

National Institute of Information and Communications Technology
Japan

Email: pedro@nict.go.jp