

6TiSCH
Internet-Draft
Intended status: Experimental
Expires: September 6, 2018

D. Dujovne, Ed.
Universidad Diego Portales
LA. Grieco
Politecnico di Bari
MR. Palattella
Luxembourg Institute of Science and Technology (LIST)
N. Accettura
LAAS-CNRS
March 5, 2018

6TiSCH Experimental Scheduling Function (SFX)
[draft-ietf-6tisch-6top-sfx-01](#)

Abstract

This document defines a Scheduling Function called "Experimental Scheduling Function" (SFX). SFX dynamically adapts the number of scheduled cells between neighbor nodes, based on the amount of currently allocated cells and the neighbor nodes' cell requirements. Neighbor nodes negotiate in a distributed neighbor-to-neighbor basis the number of cell(s) to be added/deleted. SFX uses the 6P signaling messages to add/delete cells in the schedule. This function selects the candidate cells from the schedule, defines which cells will be added/deleted and triggers the allocation/deallocation process.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 6, 2018.

Copyright Notice

Copyright (c) 2018 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	3
2.	Scheduling Function Identifier	3
3.	Allocated and Used Cells	3
4.	Overprovisioning	3
5.	Scheduling Algorithm	4
5.1.	SFX Triggering Events	4
5.2.	SFX Cell Estimation Algorithm	4
5.3.	SFX Allocation Policy	5
6.	Rules for Celllist	7
7.	6P Timeout Value	7
8.	Meaning of Metadata Information	8
9.	Node Behavior at Boot	8
10.	Cell Type	8
11.	SFX Statistics	8
12.	Relocating Cells	8
13.	Forced Cell Deletion Policy	9
14.	6P Error Handling	9
15.	Experimental requirements	9
16.	Security Considerations	10
17.	IANA Considerations	10
17.1.	SFX Scheduling Function Identifiers	10
18.	Acknowledgments	10
19.	References	11
19.1.	Normative References	11
19.2.	Informative References	11
	Authors' Addresses	11

1. Introduction

This document defines a Scheduling Function using the 6P protocol [[I-D.ietf-6tisch-6top-protocol](#)], called "Experimental Scheduling Function" (SFX). SFX is designed to offer a number of functionalities to be usable in a wide range of applications. SFX defines two algorithms: the Scheduling Algorithm which defines the number of cells to allocate/delete between two neighbors, and the Relocation Algorithm defines when to relocate a cell.

To synthesize, a node running SFX determines when to add/delete cells in a three-step process:

1. It waits for a triggering event ([Section 5.1](#)).
2. It applies the Cell Estimation Algorithm (CEA) for a particular neighbor to determine how many cells are required to that neighbor ([Section 5.2](#)).
3. It applies the Allocation Policy to compare the number of required cells to the number of already scheduled cells, and determines the number of cells to add/delete ([Section 5.3](#)).

SFX addresses the requirements for a scheduling function listed in [Section 5.2](#) from [[I-D.ietf-6tisch-6top-protocol](#)], and follows the recommended outline listed in Section 5.3 of [[I-D.ietf-6tisch-6top-protocol](#)]. This document follows the terminology defined in [[I-D.ietf-6tisch-terminology](#)].

2. Scheduling Function Identifier

The Scheduling Function Identifier (SFID) of SFX is IANA_6TISCH_SFID_SFX.

3. Allocated and Used Cells

An allocated cell is assigned as a TX, RX or Shared cell on the schedule, as a reserved resource. This reservation does not imply that a packet will be transmitted during the scheduled cell time. A used cell is a cell where a packet has been transmitted during the scheduled cell time on the last slotframe.

4. Overprovisioning

Overprovisioning is the action and effect of increasing a value representing an amount of resources. In the case of SFX, overprovisioning is done as a provision to reduce traffic variability effects on packet loss, to the expense of artificially allocating a number of cells.

5. Scheduling Algorithm

A number of TX cells must be allocated between neighbor nodes in order to enable data transmission among them. A portion of these allocated cells will be used by neighbors, while the remaining cells can be over-provisioned to handle unanticipated increases in cell requirements. The Scheduling Algorithm collects the cell allocation/deallocation requests from the neighbors and the number of cells which are currently under usage. First, the Cell Estimation Algorithm calculates the number of required cells and second, the calculated number is transferred to the Allocation Policy. In order to reduce consumption, this algorithm is triggered only when there is a change on the number of used cells from a particular node.

5.1. SFX Triggering Events

We RECOMMEND SFX to be triggered by the following event: if there is a change on the number of used cells towards any of the neighbors. The exact mechanism of when SFX is triggered is implementation-specific.

5.2. SFX Cell Estimation Algorithm

The Cell Estimation Algorithm takes into account the number of currently used cells to the neighbor. This allows the algorithm to estimate a new number of cells to be scheduled to the neighbor. As a consequence, the Cell Estimation Algorithm for SFX follows the steps described below:

1. Collect the current number of used cells to the neighbor.
2. Calculate the new number of cells to be scheduled to the neighbor by adding the current number of used cells plus an OVERPROVISION number of cells.
3. Transfer the request to the allocation policy as REQUIREDCELLS.
4. Return to step 1 and wait for a triggering event.

The Cell Estimation Algorithm is depicted in Figure 1. The OVERPROVISION parameter is calculated as a percentage of the number of currently scheduled cells to the neighbor. OVERPROVISION is added to the amount of used cells to the neighbor to reduce the probability of packet loss given a sudden growth on the number of used cells to the neighbor. The OVERPROVISION value is implementation-specific. A value of OVERPROVISION equal to zero leads to queue growth and possible packet loss. In this case, there are no overprovisioned cells where a sudden growth on the number of cells can be absorbed and detected.

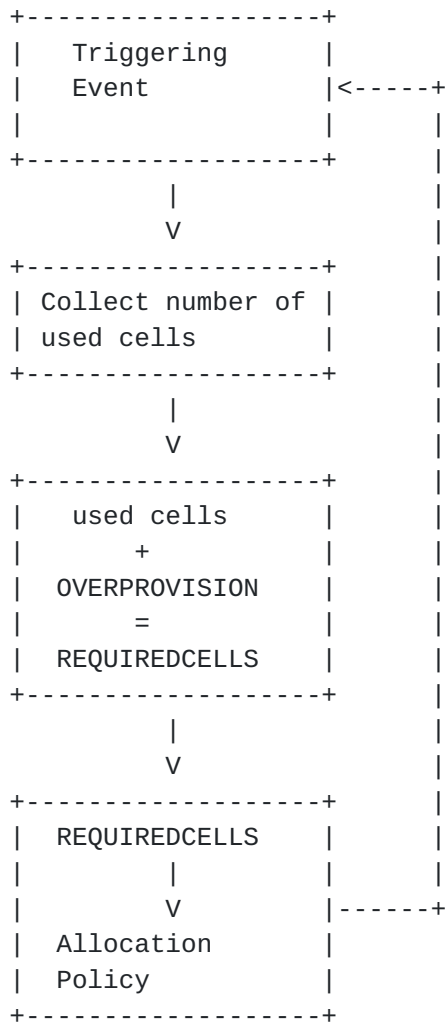


Figure 1: The SFX Estimation Algorithm

5.3. SFX Allocation Policy

The "Allocation Policy" is the set of rules used by SFX to decide when to add/delete cells to a particular neighbor to satisfy the cell requirements.

SFX uses the following parameters:

SCHEDULEDCELLS: The number of cells scheduled from the current node to a particular neighbor.

REQUIREDCELLS: The number of cells calculated by the Cell Estimation Algorithm from the current node to that neighbor.

SFXTHRESH: Threshold parameter introducing cell over-provisioning in the allocation policy. It is a non-negative value expressed as a number of cells. The definition of this value is implementation-

specific. A setting of $SFXTHRESH > 0$ causes the node to allocate at least $SFXTHRESH$ cells to each of its' neighbors.

The SFX allocation policy compares $REQUIREDCELLS$ with $SCHEDULEDCELLS$ and decides to add/delete cells taking into account $SFXTHRESH$. This is illustrated in Figure 2. The number of cells to be added/deleted is out of the scope of this document and it is implementation-dependent.




Figure 2: The SFX Allocation Policy

1. If $REQUIREDCELLS < (SCHEDULEDCELLS - SFXTHRESH)$, delete one or more cells.
2. If $(SCHEDULEDCELLS - SFXTHRESH) \leq REQUIREDCELLS \leq SCHEDULEDCELLS$, do nothing.
3. If $SCHEDULEDCELLS < REQUIREDCELLS$, add one or more cells.

When $SFXTHRESH$ equals 0, any discrepancy between $REQUIREDCELLS$ and $SCHEDULEDCELLS$ triggers an action to add/delete cells. Positive values of $SFXTHRESH$ reduce the number of 6P Transactions. The number of cells to add or delete is implementation-specific.

6. Rules for CellList

There are two methods to define the CellList. The Whitelist method fills the CellList with the number of proposed cells to the neighbor. The Blacklist method fills the CellList with the cells which cannot be used by the neighbor. The rule to select the method is implementation-specific. When issuing a 6top ADD Request, SFX executes the following sequence:

Whitelist case:

The Transaction Source node prepares the CellList field by selecting randomly the required cells, verifying that the slot offset is not occupied and choose channelOffset randomly for each cell.

The Transaction Destination node goes through the cells in the CellList in order, verifying whether there are no slotOffset conflicts.

Blacklist case:

The Transaction Source node prepares the CellList field by building a list of currently scheduled cells into the CellList. The Transaction Destination node selects randomly the required cells from the unallocated cells on the schedule, verifying that the slot offset is not occupied from the ones on the CellList.

SFX does not include any transaction retry process. If the transaction is not successful, SFX will be retriggered on the next slotframe if the number of used cells changes.

7. 6P Timeout Value

The timeout value is implementation-specific. The timeout value MAY be different for each transaction and each neighbor. The timeout range is from 0 to 128. The timeout MUST be added as an 7-bit on the Metadata header to the neighbor. There is no measurement unit associated to the timeout value. If the timeout expires, the node issues a RESET return code will be issued to the neighbor. SFX has no retry policy. Timeout examples are depicted on Figure 3 and Figure 4.

```
|Timeout Value-----|
|A|-----First Exchange----->|B|-----Second Exchange----->|A|
|Complete transaction-----|
```

Figure 3: Example Transaction where the timeout does not expire


```

|Timeout Value-----|
|A|-----First Exchange----->|B|-----Second Exchange----->|A|
|Non-Complete transaction-----|

```

Figure 4: Example Transaction where the timeout expires

8. Meaning of Metadata Information

The Metadata 16-bit field is used as follows:

BITS 0-7 [SLOTFRAME] are used to identify the slotframe number
 BITS 8-14 [TIMEOUT] represents the Timeout value
 BIT 15 [WBLIST] is used to indicate that the CellList provided is
 a Whitelist (value=0) or a Blacklist (value=1).

9. Node Behavior at Boot

In order to define a known state after the node is restarted, a CLEAR command is issued to each of the neighbor nodes to enable a new allocation process and at least a SFXTHRESH number of cells MUST be allocated to each of the neighbors.

10. Cell Type

SFX uses the TX (Transmission) cell type only, thus defining celloptions as TX=0, RX=1 and S=0 according to section 4.2.6 of [\[I-D.ietf-6tisch-6top-protocol\]](#).

11. SFX Statistics

Packet Delivery Rate (PDR) is calculated per cell, as the percentage of acknowledged packets, for the last 10 packet transmission attempts. There is no retransmission policy on SFX.

12. Relocating Cells

Allocated cells may experience packet loss from different sources, such as noise, interference or cell collision (after the same cell is allocated by other nodes in range on the network).

SFX uses Packet Delivery Rate (PDR) statistics to monitor the currently allocated cells for cell relocation (by changing their slotOffset and/or channelOffset). When the PDR of one or more softcells is below PDR_THRESHOLD, SFX relocates each of the cell(s) to a number of available cells selected randomly. PDR_THRESHOLD is out of the scope of this document and it is implementation-dependent.

13. Forced Cell Deletion Policy

When all the cells are scheduled, we need a policy to free cells, for example, under alarm conditions, or if a node disappears from the neighbor list. The action to follow this condition is out of scope of this document and it is implementation-dependent.

14. 6P Error Handling

A node implementing SFX handles a 6P Response depending on the Return Code it contains:

RC_SUCCESS:

If the number of elements in the CellList is the number of cells specified in the NumCells field of the 6P ADD Request, the operation is complete. The node does not take further action. If the number of elements in the CellList is smaller (possibly 0) than the number of cells specified in the NumCells field of the 6P ADD Request, the neighbor has received the request, but less than NumCells of the cells in the CellList were allocated. In that case, the node MAY retry immediately with a different CellList if the amount of storage space permits, or build a new (random) CellList.

RC_EOL: If an LIST command is issued and the RC_EOL is received, the node MUST understand what is specified on Section 3.3.5 of [\[I-D.ietf-6tisch-6top-protocol\]](#).

RC_ERR_VER: The node MUST NOT retry immediately. The node MAY add the neighbor node to a blacklist. The node MAY retry to contact this neighbor later.

RC_ERR_SFID: The node MUST NOT retry immediately. The node MAY add the neighbor node to a blacklist. The node MAY retry to contact this neighbor later.

RC_ERR_SEQNUM: The node MUST issue a CLEAR command to the neighbor.

RC_ERR_BUSY: Wait for a timeout and restart the scheduling process.

RC_ERR_CELLLIST: Wait for a timeout and restart the scheduling process.

RC_ERR_LOCKED: Wait for a timeout and restart the scheduling process.

RC_RESET: Abort 6P Transaction.

RC_ERR: Abort 6P Transaction. The node MAY retry to contact this neighbor later.

15. Experimental requirements

In order to evaluate the performance of this draft, we propose the following experimental work:

1. Define values for OVERPROVISION, SFXTHRESH and ranges to the number of cells to Add or Delete after the Allocation Policy is applied for typical use cases.
2. Analyze the scheduling stability (in terms of oscillation) and the hysteresis effect on scheduling using SFX. A tradeoff shall be found between the reactivity of the algorithm facing new scheduling requirements and the number of overprovisioned cells.
3. Define the PDR value below the Average which is most effective for blacklisting cells and a method to whitelist cells. Analyze the stability and long-term behavior of this algorithm.
4. Measure the distribution of cell scheduling delay (including the time taken by 6P) to estimate timeouts for different type of transactions.

16. Security Considerations

SFX is defined as an algorithm designed to efficiently fulfill bandwidth requirements between neighbour nodes and does not define a new protocol SFX uses the Minimal IPv6 over the TSCH Mode of IEEE 802.15.4e (6TiSCH) Configuration standardized on [\[RFC8180\]](#) and the 6top Protocol (6P): [\[I-D.ietf-6tisch-6top-protocol\]](#). SFX relies on the security framework described on [\[I-D.ietf-6tisch-minimal-security\]](#).

17. IANA Considerations

17.1. SFX Scheduling Function Identifiers

This document provide a new element to the "6P Scheduling Function Identifiers" sub-registry, which is part of the "IPv6 over the TSCH mode of IEEE 802.15.4e (6TiSCH) parameters" registry, as defined by [\[I-D.ietf-6tisch-6top-protocol\]](#). This Subtype is defined on figure Figure 5

SFID	Name	Reference
IANA_6TISCH_SFID_SFX	Experimental Scheduling Function (SFX)	RFCXXXX (NOTE:this)

Figure 5: IETF IE Subtype '6P'

18. Acknowledgments

Thanks to Kris Pister for his contribution in designing the default Bandwidth Estimation Algorithm. Thanks to Qin Wang and Thomas

Watteyne for their support in defining the interaction between SFX and the 6top sublayer.

This work is partially supported by the Fondecyt 1121475 Project, the Inria-Chile "Network Design" group, and the IoT6 European Project (STREP) of the 7th Framework Program (Grant 288445).

19. References

19.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8180] Vilajosana, X., Ed., Pister, K., and T. Watteyne, "Minimal IPv6 over the TSCH Mode of IEEE 802.15.4e (6TiSCH) Configuration", [BCP 210](#), [RFC 8180](#), DOI 10.17487/RFC8180, May 2017, <<https://www.rfc-editor.org/info/rfc8180>>.

19.2. Informative References

- [I-D.ietf-6tisch-6top-protocol] Wang, Q., Vilajosana, X., and T. Watteyne, "6top Protocol (6P)", [draft-ietf-6tisch-6top-protocol-09](#) (work in progress), October 2017.
- [I-D.ietf-6tisch-minimal-security] Vucinic, M., Simon, J., Pister, K., and M. Richardson, "Minimal Security Framework for 6TiSCH", [draft-ietf-6tisch-minimal-security-04](#) (work in progress), October 2017.
- [I-D.ietf-6tisch-terminology] Palattella, M., Thubert, P., Watteyne, T., and Q. Wang, "Terms Used in IPv6 over the TSCH mode of IEEE 802.15.4e", [draft-ietf-6tisch-terminology-10](#) (work in progress), March 2018.

Authors' Addresses

Diego Dujovne (editor)
Universidad Diego Portales
Escuela de Informatica y Telecomunicaciones
Av. Ejercito 441
Santiago, Region Metropolitana
Chile

Phone: +56 (2) 676-8121
Email: diego.dujovne@mail.udp.cl

Luigi Alfredo Grieco
Politecnico di Bari
Department of Electrical and Information Engineering
Via Orabona 4
Bari 70125
Italy

Phone: 00390805963911
Email: a.grieco@poliba.it

Maria Rita Palattella
Luxembourg Institute of Science and Technology (LIST)
Department 'Environmental Research and Innovation' (ERIN)
41, rue du Brill
Belvaux L-4422
Grand-duchy of Luxembourg

Phone: +352 275 888-5055
Email: mariarita.palattella@list.lu

Nicola Accettura
LAAS-CNRS
7, avenue du Colonel Roche
Toulouse 31400
France

Phone: +33 5 61 33 69 76
Email: nicola.accettura@laas.fr

