

Determinism Through Path Diversity: Why Packet Replication Makes Sense

Jesica De Armas, Pere Tuset, Tengfei Chang, Ferran Adelantado, Thomas Watteyne, Xavier Vilajosana

► To cite this version:

Jesica De Armas, Pere Tuset, Tengfei Chang, Ferran Adelantado, Thomas Watteyne, et al.. Determinism Through Path Diversity: Why Packet Replication Makes Sense. International Conference on Intelligent Networking and Collaborative Systems (INCoS), Sep 2016, Ostrava, Czech Republic, France. <hal-01333316>

HAL Id: hal-01333316

<https://hal.inria.fr/hal-01333316>

Submitted on 20 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Determinism Through Path Diversity: Why Packet Replication Makes Sense

Jesica de Armas¹, Pere Tuset¹, Tengfei Chang²,
Ferran Adelantado¹, Thomas Watteyne², Xavier Vilajosana¹

Abstract—Industrial low-power wireless mesh standards, such as IEEE802.15.4-TSCH, WirelessHART and ISA100.10a, offer wire-like end-to-end reliability and a decade of battery lifetime. These technologies have become de-facto standards, used in the most demanding applications such as industrial process monitoring. In this paper, we explore what it takes to go from industrial process *monitoring* to industrial process *control*. The difference is that, in the latter case, the network needs to provide low and predictable latency, and deterministic operation. We explore the overall usefulness of packet replication, in which a source node sends multiple copies of a packet on disjoint multi-hop paths. We show, through extensive simulation, that packet replication allows for a reduction of end-to-end latency by 40%. In addition, packet replication significantly improves the network reliability through path diversity. This work is directly in line with standardization activities at the IETF 6TiSCH and DetNet working groups, to which it is being contributed.

I. INTRODUCTION

With an expectation of tens of billion of connected devices [1], the Internet of Things (IoT) is reshaping the industrial landscape with large economic opportunities [2], [3]. It is paving the way for improved control, and integration of industrial processes in quality and scope. This is materialized by the progressive extension of industrial automation from local communication buses, to factory and to factory ecosystem scopes. This extension must happen with minor impact and cost, while being backwards compatible with legacy infrastructure.

Wireless industrial communication emerged to facilitate this extension at a reduced cost. It has brought novel use cases due to its ability to be deployed almost everywhere. Several standard solutions have appeared in the last 7 years, including well-known wireless versions of industrial buses such as WirelessHART [4], proprietary wireless Profi-buses [5] and low power standards such as ISA100.11a [6] or IEEE802.15.4-2015 [7] managed by the IETF 6TiSCH architecture [8]. A key characteristic of these technologies, and in contrast with previous wireless technologies, is their scheduled operation. Those networks are based on the Time Division Multiple Access (TDMA), Frequency Division Multiple Access (FDMA), or a combination of both. The latest of these technologies are based on the concept of Time Synchronized Channel Hopping (TSCH) [9], dividing the access in timeslots and multiplexing them through different frequencies.

* This work is partially funded by the Spanish Ministry of Economy and FEDER Funds under SINERGIA (TEC2015-71303-R) project.

¹ IN3 - Computer Science and Multimedia Studies, Open University of Catalonia (UOC), Barcelona, Catalonia, Spain. <first.last>@uoc.edu

² Inria, EVA team, Paris, France. <first.last>@inria.fr




Fig. 1. Example of disjoint paths and packet replication in a multi-hop 6TiSCH network. Node a replicates packet P , which is forwarded through disjoint tracks along the multi-hop network to node f .

This enables a scheduler to allocate pairwise transmissions in dedicated cells, resulting in collision-free communication. Despite their scheduled nature, external interference and fading compromise the reliability of the network, specially in industrial scenarios, requiring extended functionality to overcome such unreliability. Channel hopping has proven to mitigate that effect [10] and successful industrial deployments demonstrated that 99.999% reliability can be achieved [11]. On top of wire-like reliability, certain critical industrial applications (e.g. automotive industry [2]), also require low and predictable latency.

In this paper, we propose to exploit end-to-end packet replication on disjoint network segments to lower end-to-end latency, while keeping wire-like end-to-end reliability. We want to quantify how much more energy is needed to make the latency lower and predictable, and discuss the suitability of the approach in current TSCH networks.

The remainder of this paper is organized as follows. Section II introduces the operation of an IEEE802.15.4-2015 TSCH network, with the scheduling solution defined by IETF 6TiSCH. Section III defines the packet replication and disjoint path selection approach we propose. Section IV presents the evaluation on the proposal on a set of large multi-hop network topologies. Finally, Section V concludes this paper.

II. IETF 6TiSCH SCHEDULING

An IETF Working Group called 6TiSCH [12] was created in 2013 to enable IPv6 over the TSCH mode of the IEEE802.15.4 standard [7]. In a 6TiSCH network, low-power wireless devices form a multi-hop Low-power Lossy Network (LLN). This LLN connects the constrained devices to the traditional Internet through one or more Low-power Border Routers (LBRs) [13].

In IEEE802.15.4-2015 TSCH, the nodes in the low-power mesh network communicate by following a schedule. A Time Division Multiple Access (TDMA) structure instructs each node what to do at each time slot of that schedule: transmit, receive, or sleep. The way this schedule is built determines the amount of traffic the network can produce, the latency of a packet, and the redundancy in the network. This schedule is a matrix with slotframe length and number of frequencies available as dimensions (see Fig. 2). It determines the amount of energy each node consumes, therefore the (battery) lifetime of the network. The goal of 6TiSCH is to develop a standard approach to manage this communication schedule, and match it against the traffic needs of the applications running on this network [14].

A cell is a single element in the TSCH schedule, identified by a slot offset and a channel offset. A cell is an atomic unit allocated by the scheduling algorithm. A cell can be scheduled or unscheduled. During an unscheduled cell, the node does not communicate, and keeps its radio off. A scheduled cell can be transmission-only, reception-only or transmission-reception. This enables a wide set of possibilities to build up a link-layer forwarding plane [15].

Scheduled cells can be grouped according to their destination and purpose. Groups of cells in the schedule of a node with the same destination are equivalent, and called “bundle” [16]. Since the slotframe repeats over time, each cell gives a “quantum” of bandwidth to a given neighbor. Modifying the number of cells in a bundle modifies the amount of resources allocated between two neighbors.

A sequence of cells scheduled along a multi-hop path is called “track” [16] (see Fig. 2). It is the result of a reservation, and belongs to the node that initializes the process of establishing the track. The track defines a multi-hop route that optimizes certain objective function defined by the scheduler, for example minimizing latency. The scheduler entirely defines the synchronization and communication tracks between nodes. By adding/removing cells between neighbors, the scheduler adapts the link-layer resources to the needs of the application.

The IEEE802.15.4 TSCH standard (and other TSCH technologies such as WirelessHART and ISA100.11a) define the mechanisms to execute a communication schedule. How the schedule is built, updated and maintained is out of the scope of the standard. Several scheduling approaches have been investigated by the research community.

A centralized Traffic Aware Scheduling Algorithm (TASA) has been proposed by Palattella *et al.* [17]. It uses map coloring to schedule cells and tracks in the network topology graph.

TASA sets up the TSCH schedule based on the network topology and the traffic load. It uses information from the routing protocol (e.g. RPL) and communication requirements of the applications to schedule cells and tracks. In that process, it provides some level of QoS (duty cycle, throughput, etc.) to each data flow traversing the network.

Decentralized TSCH scheduling is studied in the OpenWSN project [18] and 6TiSCH. A PID-based scheduling approach [19] proposes a control-loop like distributed scheduler that reacts to sudden traffic variations by adding/removing cells in the schedule. The Scheduling Function Zero (SF0) [20] uses a simple thresholds to react to variations in traffic load. SF0 relies on the 6top protocol (6P) [21] for neighbor nodes to negotiate adding/removing cells. Morell *et al.* introduce the concept of label switching in TSCH networks, and use reservation to establish and manage tracks between nodes in the network [15]. This approach computes the schedule of the network by collecting information along the track and installing it in a way similar to the RSVP protocol [22].

III. PROPOSAL: EXPLOITING DATA PLANE REDUNDANCY

This paper uses data plane redundancy to improve latency and reliability. The approach consists in replicating data plane information, and forwarding the different copies through disjoint tracks in order for it to reach the destination faster. This simple approach does increase the energy consumption of the network, but also increases the probability of the packet reaching the destination, and reduces the end-to-end latency.

The technique requires the network scheduler to form a connected graph of at least degree 2. This enables each node to have two candidate parents to communicate with, as per Fig. 2. Having two routing parents allows a node to exploit path diversity [23], [10].

The data forwarding plane, upon reception of a packet from an upper layer, replicates it N times, addressing each copy to a different routing parent. Those copies are then forwarded along the tracks until they reach the destination. A node relaying a copy does not replicate it further: the packet is forwarded through the preferred parent following the available track (hence the best route set up by the information from the routing protocol). In case a relay node has two copies of the same packet in its queue, it forwards each copy to a different parent to preserve diversity.

The best connected parents are selected according to a connectivity metric such as ETX [24], as per the RPL routing protocol [25]. The scheduler is in charge of setting up the tracks by using the routing and connectivity information obtained through the control plane links [26]. TASA computes the disjoint routes from every node to the sink, and allocates the corresponding cells of the nodes at each of the tracks. Track preference is updated dynamically by TASA according to the routing protocol information and the current connectivity status.

If a node has $P < N$ parents, only P copies are created, one forwarded to each parent to whom it has a link in the




Fig. 2. Example of TSCH schedule with disjoint tracks.

schedule. The nodes one hop from the DAGRoot only have one link to it; they hence send a single copy.

In the transmission process, packets are inserted in the link-layer queues and consumed by the MAC layer according to the network scheduler. In a scheduled network, time slots are dedicated by unicast communication between neighbors; only one of the copies is sent per slot.

It is normal for the DAGroot to receive several copies of the same data.

IV. PERFORMANCE EVALUATION

We evaluate the performance of the scheme by simulation.

We developed a TASA simulator¹. We implement both our proposed replication scheme and the TASA central scheduling approach [17].

The TASA scheduler computes disjoint tracks from any node to the DAGRoot. To each of the links installed by TASA, we associate a Packet Delivery Ratio (PDR) in $[0.85 \dots 1]$. This range has been selected to simulate well-connected dense networks, representative of a factory-floor deployment.

We simulate networks of 20, 40, 60 or 80 nodes. For each number of nodes, we generate 10 connected mesh network topologies, and on each topology execute 30 simulation runs. In each simulation run, each node sends a number of packets uniformly distribute in $[1 \dots 10]$ to the sink. The parameter $n = 3$ is used to configure the number of lin-layer retransmissions allowed at each hop. The TSCH slotframe contains 101 timeslots and 4 channel offsets. We plot the average results for each network size.

A. End-to-End Reliability

We measure the reduction in packet loss and therefore the improvement in reliability when more than one copy of a packet is sent.

Fig. 3 presents the percent improvement with respect to the case where a single copy of a packet is sent. We can see that an

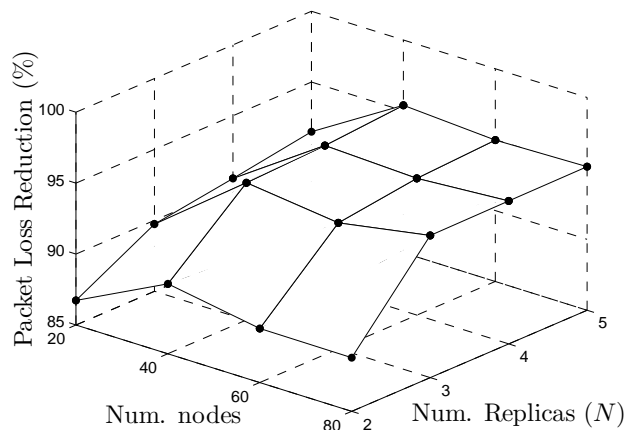


Fig. 3. Packet loss improvement according to the size of the network and number of copies of the packet.

almost 90% packet loss reduction² is obtained when using one extra copy. We also see that the gain obtained by replication increases with the number of copies. The gain provided by each extra copy decreases because each extra copy forces the use of a worse connected path, therefore its impact is smaller. Moreover, the larger the network (i.e. the more track diversity), the greater the benefits of replication.

B. Energy Consumption

Sending multiple copies costs more energy than sending a single copy. Table I summarizes the impact of replication on the energy consumption measured in number of extra active slots with respect to the single packet case.

¹ As an online addition to this paper, the source code of the simulator is available at <https://bitbucket.org/xvilajosana/wine>

²In our experiments from losing 5 packets every 100 to losing 1 packet every 350, which represents almost a 18 \times increase.

TABLE I
ENERGY CONSUMPTION IMPACT OF REPLICATION

Number of Copies	Normalize Energy Consumption
1 (reference)	1.000
2	1.863
3	2.434
4	2.756
5	2.914




Fig. 4. Packet reliability as a function of the maximum number of retransmissions for $N = \{1, 2, 3\}$.

C. End-to-End Latency

All paths from the nodes to the DAGRoot have been simulated with 4, 5 or 6 hops (i.e. the number of hops of each path is uniformly distributed in the set $[4,5,6]$) and the PDR of all links is 90%. The packet reliability as a function of the number of copies (N) and the maximum number of packet retransmissions allowed per link (n) is shown in Fig. 4. As can be seen, replication increases reliability.

A second benefit of replication is that end-to-end latency of a packet is reduced when several copies are sent. We define end-to-end latency as the time elapsed from the moment the data is generated at the source, until the moment the first copy reaches the destination. The multi-track diversity favors the delay reduction.

The number of slots allocated to a link in each frame also determines how fast copies reach the DAGRoot. Fig. 5 shows the latency when 1-5 time slots (out of the 101 slots in a slotframe) are scheduled to each link. It shows a 40% latency reduction when 2 copies are sent. Replication and scheduling are complementary, and should be jointly used to optimize the network performance.

V. CONCLUSION

This paper evaluates the improvement in reliability and latency when sending multiple copies of the same data on disjoint paths, in industrial wireless multi-hop networks. We show by simulation how adding a single extra copy of a data packet improves the reliability by $18\times$ while reducing the end-to-end latency by 40%, while only doubling the energy consumption. This paper contributes with a new mechanism to improve reliability and reducing latency by using redundant




Fig. 5. End-to-end latency when the number of copies (N) and the number of slots allocated in each slotframe changes. We set $n = 3$, the maximum number of link-layer re-transmissions.

data plane forwarding in a 6TiSCH network. It participates in the realization of the IETF DetNet vision of a deterministic, ultra-reliable end-to-end network architectures.

REFERENCES

- [1] J. Rivera and R. Van Der Meulen, "Forecast: The Internet of Things," *Gartner Report, Worldwide*, 2013.
- [2] P. C. Evans and M. Annunziata, "Industrial Internet: Pushing the Boundaries of Minds and Machines," *White Paper. General Electric*, 2012.
- [3] Bosch, "Creating connected manufacturing operations in the Internet of Things," *White Paper*, 2014.
- [4] H. C. Foundation, "WirelessHART Specification 75: TDMA Data-Link Layer (HCF SPEC-75)," 2008.
- [5] I. E. Commission, "IEC 61158: Digital data communications for measurement and control-Fieldbus for use in industrial control systems," 2003.
- [6] ISA, "ISA-100.11a-2009 Wireless Systems for Industrial Automation: Process Control and Related Applications," 2009.
- [7] IEEE, "02.15.4-2015: IEEE Approved Draft Standard for Low-Rate Wireless Personal Area Networks (WPANs)," 2015.
- [8] D. Dujovne, T. Watteyne, X. Vilajosana, and P. Thubert, "6TiSCH: Deterministic IP-enabled Industrial Internet (of Things)," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 36–41, 2014.
- [9] K. Pister and L. Doherty, "Tsmpt: Time synchronized mesh protocol," *IASTED Distributed Sensor Networks*, pp. 391–398, 2008.
- [10] T. Watteyne, A. Mehta, and K. Pister, "Reliability Through Frequency Diversity: Why Channel Hopping Makes Sense," in *Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, Tenerife, Canary Islands, Spain, 29-30 October 2009.
- [11] L. Doherty, W. Lindsay, and J. Simon, "Channel-specific wireless sensor network path data," in *International Conference on Computer Communications and Networks (ICCCN)*. IEEE, 2007, pp. 89–94.
- [12] M. R. Palattella, P. Thubert, X. Vilajosana, T. Watteyne, Q. Wang, and T. Engel, "6TiSCH Wireless Industrial Networks: Determinism Meets IPv6," in *Internet of Things*. Springer, 2014, pp. 111–141.
- [13] P. Thubert, T. Watteyne, M. R. Palattella, X. Vilajosana, and Q. Wang, "IETF 6TiSCH: Combining IPv6 Connectivity with Industrial Performance," in *International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS)*. IEEE, 2013, pp. 541–546.
- [14] T. Watteyne, M. R. Palattella, and L. A. Grieco, *Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement*, Internet Engineering Task Force Std. RFC7554, May 2015.
- [15] A. Morell, X. Vilajosana, J. L. Vicario, and T. Watteyne, "Label Switching over IEEE802.15.4e Networks," *Wiley Transactions on Emerging Telecommunications Technologies (ETT)*, vol. 24, no. 5, pp. 458–475, 3 May 2013.

- [16] M. R. Palattella, P. Thubert, T. Watteyne, and Q. Wang, "Terminology in IPv6 over the TSCH mode of IEEE 802.15.4e," Internet Engineering Task Force, Tech. Rep. draft-ietf-6tisch-terminology-07 [work in progress], 21 March 2016.
- [17] M. R. Palattella, N. Accettura, M. Dohler, L. A. Grieco, and G. Boggia, "Traffic Aware Scheduling Algorithm for Reliable Low-power Multi-hop IEEE 802.15.4e Networks," in *International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)*. Sydney, Australia: IEEE, 9-12 September 2012.
- [18] T. Watteyne, X. Vilajosana, B. Kerkez, F. Chraim, K. Weekly, Q. Wang, S. D. Glaser, and K. Pister, "OpenWSN: a Standards-based Low-power Wireless Development Environment," *Wiley Transactions on Emerging Telecommunications Technologies (ETT)*, vol. 23, no. 5, pp. 480–493, 2012.
- [19] M. Domingo-Prieto, T. Chang, X. Vilajosana, and T. Watteyne, "Distributed pid-based scheduling for 6tisch networks," *IEEE Communications Letters*, vol. 20, no. 5, pp. 1006–1009, May 2016.
- [20] D. Dujovne, L. A. Grieco, M. R. Palattella, and N. Accettura, "6TiSCH 6top Scheduling Function Zero (SF0)," Internet Engineering Task Force, Tech. Rep. draft-ietf-6tisch-6top-sf0-00 [work in progress], 13 May 2016.
- [21] Q. Wang and X. Vilajosana, "6top Protocol (6P)," Internet Engineering Task Force, Tech. Rep. draft-ietf-6tisch-6top-protocol-00 [work in progress], 27 April 2016.
- [22] L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSeRVation Protocol (RSVP) – Version 1 Functional Specification," Internet Engineering Task Force, Tech. Rep. RFC2205, 2 March 2013.
- [23] X. Vilajosana, J. Llosa, J. C. Pacho, I. Vilajosana, A. A. Juan, J. L. Vicario, and A. Morell, "Zero: Probabilistic routing for deploy and forget wireless sensor networks," *MDPI Sensors*, vol. 10, no. 10, p. 8920, 2010.
- [24] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A High-throughput Path Metric for Multi-hop Wireless Routing," in *International Conference on Mobile Computing and Networking (MobiCom)*. San Diego, CA, USA: ACM, 2003, pp. 134–146.
- [25] A. Brandt, J. Vasseur, J. Hui, K. Pister, P. Thubert, P. Levis, R. Struik, K. Richard, T. H. Clausen, and T. Winter, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks," 14 October 2015.
- [26] X. Vilajosana and K. Pister, "Minimal 6TiSCH Configuration," Internet Engineering Task Force, Tech. Rep. draft-ietf-6tisch-minimal [work in progress], June 2016.