Using RAW as Control Plane for Wireless Deterministic Networks: Challenges Ahead

Carlos J. Bernardos[1], Alain Mourad[2], Milan Groshev[1], Luis M. Contreras[3], Marc Mollà Roselló[4],
Otilia Bularca[5], Valerio Frascolla[6], Peter Szilagyi[7], Sebastian Robitzsch[2]

[1]Universidad Carlos III de Madrid; [2]Interdigital Europe; [3]Telefonica I+D; [4]Ericsson; [5]SIMAVI; [6]Intel; [7]Nokia Bell Labs

Abstract—**This paper provides extensive analysis of RAW (Reliable and Available Wireless) enhancements and solutions needed to manage industrial environments more effectively. Starting from the description of the industrial use case, an analysis of gaps and potential new extensions is performed. Namely, the need to (*i*) support multi-domain operation, at both technology and administrative levels; (*ii*) integrate RAW with edge architectures; and, (*iii*) the support for mobility support in RAW networks, are analyzed. The identified gaps are indeed not yet tackled by the relevant standardization development organizations, mainly the Internet Engineering Task Force, and are thus object of our future work.**

Keywords—deteministic networking; TSN; edge computing

I. INTRODUCTION

In Industry 4.0 scenarios the latency requirements are of paramount importance for tasks involving real time operation, or accurate synchronization, as in the case of remote control of factory robots [1]. Having a network prone to suffer huge jitter and latency harnesses the correct behavior of industrial services. Although recent advances in access technologies [2] [3] have pushed the capabilities of wired and wireless solutions, the heterogeneity of involved SDOs (Standards Developing Organization) and technologies makes it challenging to have a system can that aggregate all that is needed to provide end-to-end (E2E) network guarantees across multiple sites. The European-funded research project PRogrammable AI-Enabled DeterminIstiC neTworking for 6G (PREDICT-6G[1]) addresses the above issues by: i) developing an E2E solution for deterministic (predictable, reliable and time sensitive) heterogeneous wireless networks; ii) impacting with its results several SDO, e.g, 3GPP, IETF, and IEEE 802.11; iii) using RAW [2] (Reliable and Available Wireless) as one of the key elements of the novel proposed control plane.

RAW is an effort to provide deterministic networking on a path that includes a wireless interface, with the scope of providing high reliability and availability of IP connectivity. The wireless medium presents challenges in achieving deterministic properties such as low packet error rate, limited consecutive losses, and limited latency. RAW, based on the concepts of the Deterministic Networking (DetNet[3]) Working Group (WG), addresses the above mentioned challenges in an IP network using scheduled wireless segments and other media such as frequency/time-sharing physical media resources with stochastic traffic, e.g., IEEE Std. 802.15.4 time slotted channel hopping (TSCH), 3GPP 5G ultra-reliable low latency communications (URLLC), IEEE 802.11ax/be, and L-band Digital Aeronautical Communications System (LDACS). Like DetNet, RAW technologies abstract the radio layers underneath, addressing the Layer 3 (L3) aspects in support of applications requiring high reliability and availability.

The RAW concept, discussed in [2], separates the path computation time scale, where complex paths are computed, from the path selection time scale (significantly lower than the path computation time), where the forwarding decision is taken for one or a few packets. The RAW approach operates at the path selection time scale to provide reliable and available wireless IP connectivity. The decision of which redundant forwarding path for each packet is selected, previously calculated by the Path Computation Element (PCE), is made by the Path Selection Engine (PSE). The PSE uses rapid local adjustments of forwarding tables within the diverse options defined by the PCE. This enables the use of Packet (hybrid) Automatic Repeat Request (ARQ), replication, elimination, and ordering (PAREO), and scheduled transmissions with the purpose of improving the utilization of resources.

In several use cases [4]-[8] reliability and availability are key requirements for wireless heterogeneous networks. For instance in residential gaming, endowed with eXtended Reality (XR), or in the manufacturing sector, where a plethora of devices are interconnected and generate data that need to be reliably delivered to the control and monitoring agents.

The structure of this paper is as follows. First, the use case of industrial manufacturing is presented, highlighting the need for additional RAW mechanisms. Then, the overall system design for an Industry 4.0 scenario is proposed. Next, three main gaps identified in current RAW/DetNet work are described, namely extensions for multidomain operation, integration of RAW, and edge and mobility in RAW domains. We conclude the paper with a summary and potential roadmap of contributions and adoptions at the IETF.

II. USE CASE : WIRELESS FOR INDUSTRIAL APPLICATIONS

A key use case in industrial environments is the one referring to the control of data networks, where periodic control loops operate on the collection of data from sensors that measure a physical property, e.g., a fluid temperature, a Programmable Logic Controller (PLC) that decides an action, e.g., warm up the mix, and several actuators that perform the required action, e.g., the power injection in a resistor.

A. Use case details

Control Loops

A Process Control designates continuous processing operations, like heating oil in a refinery or mixing drinking soda. Control loops in the Process Control industry operate at a very low rate,

---

typically four times per second. Factory Automation, on the other hand, deals with discrete goods such as individual automobile parts, and requires faster loops, on the order of milliseconds. Motion control that monitors dynamic activities may require even faster rates on the order of, or below, the millisecond. Finally, some industries exhibit hybrid behaviors, like canned soup by behaving as a process industry while mixing the food and then operating as a discrete manufacturing when putting the final product in cans and shipping them.

In all these cases, a packet must flow reliably between the sensor and the PLC, be processed by the PLC, being finally sent to the actuator within the control loop period. In some use-cases that inherit from analog operations, jitter might also alter the operation of the control loop. A rare packet loss is usually admissible, but typically four losses in a row cause an emergency halt of the production, a drawback that implies a high cost for the manufacturer.

Monitoring and diagnostics

Monitoring and diagnostics data is essential to improve the performance of a production line, e.g., by optimizing real-time processing or maintenance windows using Machine Learning predictions. Due to the lack of suitable and available wireless technologies over the previous decades, some specific industries such as Oil and Gas have been using serial cables, literally by the millions, to perform their process optimization. Nowadays wireless technologies are being deployed successfully in several vertical domains. Still, few industries would afford the high cost in the pursuit of the Holy Grail for the Industrial Internet of Things (IIoT), which is to provide the same benefits to all vertical industries, e.g., Smart Grid, Transportation, Manufacturing, and Medical. This requires a cheap, robust, secure and scalable IP-based access technology.

Inside the factory, wires may already be available to operate the Control Network, but monitoring and diagnostics data do not fit in that network for several reasons. On the one hand, traffic is abundant and asynchronous in factories, meaning that it may influence the deterministic nature of the control operations and impact the production. On the other hand, needed information must be reported to the carpeted floor over IP, which exposes to potential security breaches via the interconnection of the Operational Technology (OT) network with the Internet technology (IT) network, possibly also facilitating a potential disastrous and unwanted rogue access.

B. The Need for Wireless

Ethernet cables used on a robot arm are prone to breakage after a few thousands flexions, i.e., a lot faster than a power cable that is wider in diameter and more resilient. In general, wired networking and moving parts are not a good match, mostly in the case of fast and recurrent activities including rotation. When refurbishing older premises that were built before the Internet age, power is usually available everywhere, but data is not. It is often impractical, time consuming and expensive to deploy an Ethernet fabric across walls and between buildings.

Even when wiring is in place, e.g., in the case of an existing control network, asynchronous IP packets such as diagnostics may not be appropriate for operational and security reasons. For

those packets, the option to create a parallel wireless network offers a suitable solution that can scale with the many sensors and actuators that equip every robot, with valves and fans that are deployed on the factory floor. It may also help detect and prevent a failure that could impact the entire production, like the degradation (vibration) of a cooling fan on the ceiling. IEEE Std. 802.15.4 TSCH [9] is a promising technology for that purpose, if the scheduled operations enable to use the same network by asynchronous and deterministic flows in parallel.

C. Requirements for RAW

A Deterministic Network is backwards compatible with, and capable of transporting, statistically multiplexed traffic while preserving the properties of the accepted deterministic flows. The 6TiSCH Architecture [10] serves that requirement, however it focuses on best-effort IPv6 packet flows. RAW, instead, is able to lock so-called hard cells (i.e., scheduled cells [11] for use by a centralized scheduler) and leverage on time and spatial diversity over a graph of E2E paths, called a Track, that is based on those cells.

Over the course of the recent years, major Industrial Protocols (e.g., ODVA[4] with EtherNet/IP and Profinet[5]) have been migrating towards Ethernet and IP. To unleash the full power of the IP hourglass model [12], it should be possible to deploy any application over any network that has the physical capacity to transport across technologies the industrial flow, regardless of the MAC/PHY technology and of wired/wireless. RAW mechanisms should be able to setup a Track over a wireless access segment and a wired or wireless backbone to report both sensor data and critical monitoring within a bounded latency and maintaining the high reliability of the flows over time. It is also important to ensure that RAW solutions are interoperable with existing wireless capabilities, and that it can be extended using retrofitting. Maintainability, as a broader concept than reliability, is also important in industrial scenarios [13].

• *Non-latency critical considerations*

Monitoring and diagnostics applications do not require latency critical communications but demand reliable and scalable ones. On the other hand, process control applications involve control loops that require a bounded latency, thus being latency critical, but being E2E managed. Therefore, DetNet mechanisms can be applied in conjunction with RAW mechanisms.
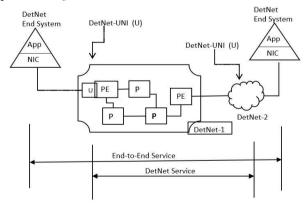

Fig. 1. DetNet Service Reference Model (Multidomain) (from RFC8655)

III. MULTI-DOMAIN EXTENSIONS

The DetNet WG focuses on deterministic data paths that operate over Layer 2 (bridged) and Layer 3 (routed) segments, where such paths can provide bounds in terms of latency, loss, and packet delay variation (jitter), as well as high reliability. The DetNet architecture document [14] includes the concept of multi-domain in the DetNet Service reference model (Fig. 5 of RFC 8655, reproduced here in Fig. 1 for convenience). However, the WG has not yet worked in detail on the necessary protocol operations to support multi-domain at control and data plane.

DetNet defines the Packet Replication, Elimination, and Ordering Functions (PREOF) to provide service protection. PREOF implies 4 capabilities:

- Sequencing information, by adding a sequence number or time stamp as part of DetNet.

- Replicating packets into multiple DetNet member flows, typically sending them along multiple different paths to the destination(s).

- Eliminating duplicate packets of a DetNet flow based on the sequencing information and track of received packets.

- Reordering DetNet flow's packets.

PAREO is a superset of DetNet's PREOF, defined in RAW, that includes radio-specific techniques such as short-range broadcast, Multi-Users Multiple Inputs Multiple Outputs (MUMIMO), constructive interference and overhearing, which can be leveraged separately or combined to increase reliability.

There are multiple scenarios and use cases that might involve multiple technology and/or administrative domains in DetNet and RAW. The next sections explore what are the main multi-domain aspects for the application, controller and network/data planes in DetNet and RAW, to identify some gaps that would require further work in SDO, e.g., IETF.

A. Application plane

The Application Plane incorporates the User Agent, which is a specialized application that interacts with end users and operators and performs requests for DetNet services via an abstract Flow Management Entity (FME), which may or may not be collocated with (one of) the end systems. At the Application Plane, a management interface enables the negotiation of flows between end systems.

In a multi-domain deployment, the User Agent might be aware of the existence of multiple domains. A multi-domain aware User Agent/application plane could take care of the negotiation of the flows in all involved domains, whereas a multi-domain unaware one will have to rely on the network to take care of it transparently.

B. Controller plane

We refer to the Controller Plane as the aggregation of the Control and Management planes. The term Controller Plane Function (CPF) refers to any device operating in that plane, whether a PCE [15], a Network Management Entity (NME), or a distributed control protocol. The CPF is a core element of a controller, in charge of computing deterministic paths to be applied in the Network Plane. A (Northbound) Service Interface enables applications in the Application Plane to communicate with the entities in the Controller Plane.

In DetNet, one or more CPFs collaborate to implement the requests from the FME creating per-flow, per-hop behaviors installed in the DetNet nodes for each individual flow. Adding multi-domain support might require some support at the CPF. For example, CPFs sitting at different domains need to discover and authenticate each other, and then negotiate per-hop behaviors. Depending on the multi-domain support provided by the Application Plane, the Controller Plane might be relieved from some responsibilities (e.g., if the Application Plane takes care of splitting what needs to be provided by each domain).

Solutions like the ones above are not sufficient alone to solve the multi-domain RAW problem, as the PSEs need to have some additional information from the other involved domains to be sensitive/reactive to transient changes, and to ensure a certain level of reliability and availability in a multi-domain wireless heterogeneous mesh network. Section III.D explores in more detail the RAW-specific multi-domain problem and proposes some initial solutions.

C. Network/Data plane

The Network Plane represents the network devices and protocols, regardless of the layer at which the network devices operate. It includes the Data Plane and Operational Plane (e.g., Operation, Administration and Maintenance (OAM)) [16] aspects. A Southbound (Network) Interface enables the entities in the Controller Plane to communicate with devices in the Network Plane.

At the Network Plane, DetNet nodes may exchange information regarding the state of the paths, between adjacent DetNet nodes and possibly with the end systems. In a multi-domain environment, nodes belonging to different domains might need to exchange information. This might require protocol translations and/or abstractions, as the different domains might not offer the same capabilities nor use the same network protocols. Additionally, OAM protocols [16] might also need to be extended to support multi-domain operation.
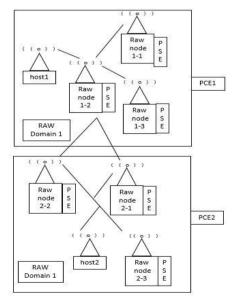

Fig. 2. Scenario showing multiple RAW domains

3

Note as well, that performing PREOF or PAREO across multiple domains poses additional challenges, as knowledge of all the involved domains might not be available and/or the data planes at each domain could also be different.

D. RAW specific analysis

We can refer to domains managed by a single PCE as single-domain RAW, where nodes are typically operated and managed by a single administrative entity. In this scenario, the PSE can make use of "tracks" and paths involving only the nodes belonging to the RAW domain.

There are scenarios where hosts are connected to different RAW domains and they need to communicate to each other with certain reliability and/or availability guarantees, for example in large factories where networks might be organized in domains (e.g., per production lines or building/sites), in residential environments where there are different networks (e.g., one at home and one in the garden), or even vehicular scenarios (e.g., hosts connected to different vehicles).

Fig. 2 shows an example of communication involving two RAW domains. As opposed to a single-domain scenario, where a single PCE may compute all possible "tracks" at longer time scale, and the PSE functionality may perform using all information available at the domain, multi-domain scenarios pose additional challenges that are not solved yet.

Each RAW domain operates independently from the other domains. While there exist inter-PCE solutions today, allowing one domain's PCE to learn some inter-domain paths, such solutions are still not much effective, as the PSE of one domain would not have full visibility nor capability to act on the other domains (e.g., there are no multi-domain OAM solutions in place yet), limiting its capability to guarantee any given Service Level Agreement (SLA). Therefore, there is a need to define inter-PSE coordination mechanisms across domains.

Solutions like Hierarchical PCE (H-PCE) are not sufficient alone to solve the multi-domain RAW problem, as the PSEs need to have some additional information from the other involved domains to be sensitive/reactive to transient changes, to ensure a certain level of reliability and availability in a multi-domain wireless heterogeneous mesh network.

Within a single domain, the RAW framework architecture works by having the PCE in charge of computing the paths (tracks) and the PSE(s) taking the short time decisions of which sub-tracks to use. Note that the PSE is assumed to be either a distributed functionality (performed by every RAW router of the path, which takes forwarding decisions based on the local and OAM information that they have), or a centralized functionality executed by the entry (ingress) router in the domain (note that if there are multiple ingress nodes, then there might be multiple PSEs), which then performs source routing.

In scenarios with multiple connected RAW domains, running uncoordinatedly RAW solutions in each domain is not an effective solution. PSEs would need to have global E2E information as well as be capable of running OAM mechanisms [16] [17] to monitor the quality of the selected paths.

IV. RAW IN MULTI-ACCESS EDGE DEPLOYMENTS

Multi-access Edge Computing (MEC) capabilities deployed at the edge of the mobile network can facilitate the efficient and dynamic provision of services to mobile users. The ETSI MEC WG intends to specify an open environment for integrating MEC capabilities with service providers' networks, also including applications from third parties. These distributed computing capabilities will make the IT infrastructure available for the deployment of functions in mobile access networks.

One relevant exemplary scenario showing the need for an integration of RAW and MEC is introduced next. One of the main (and differential) use cases of 5G is URLLC. Among the many so-called "verticals" that require URLLC features, Industry 4.0 (also referred to as Wireless for Industrial Applications) is probably the one with more short-term potential. As identified in [8], this scenario also calls for RAW solutions, as cables are not that suitable for the robots and autonomous guided vehicles typically used in factories. This is also a very natural scenario for MEC deployments, as bounded and very low latency are needed between the robots and physical actuators including the control logic managing them. This scenario assumes a wireless domain, involving multiple MEC platforms to ensure low latency to applications, by being able to host MEC applications in several locations, and dynamically migrating them when the terminals (e.g.., end-user devices) move around and the serving MEC platform might no longer be capable of meeting the latency requirements.

A. Problem Statement

According to current standards, MEC platform(s) would have to interact with a PCE for data plane requests and updates. This tremendously limits the capabilities to guarantee real-time forwarding decisions, as it makes not possible/very challenging to manage forwarding decisions in real/near-real time.

RAW solutions being explored today consider the role of the PSE, which computes at a short time scale which of the available paths (called tracks) -- computed by a PCE -- should be used per flow/packet and also which PAREO functions can be used, in order to provide the flow with the required availability and reliability features. The PSE interacts with the PCE and with the RAW nodes so they can setup the required per-flow state, to recognize the flow and determine the forwarding policy to be applied. These RAW forwarding decisions can be distributed among the necessary nodes using in-band signaling (e.g., extending Segment Routing, BIER-TE or DETNET tagging) or can be taken autonomously by each forwarding node locally (based on its knowledge of the status of the network, gained via OAM RAW-specific mechanisms). Industrial environments are a good example of applications where reliability and availability are critical. Ensuring this in wireless heterogeneous and multi-hop networks requires using multiple paths, using PAREO functions and even using dual or multiple connectivity. Terminal mobility makes it even more challenging to guarantee certain reliability and availability levels, due to the dynamic and fast changes that are needed at the data plane level. The short-time scale forwarding decisions that are required to ensure reliability and availability with terminal mobility cannot be managed if MEC platforms can only interact with the data plane through the PCE.

The PCE is responsible for routing computation and maintenance in a network, and it is typically a centralized entity that can even reside outside the network. It is meant to compute and establish redundant paths, but not to be sensitive/reactive to transient changes, and therefore it is not capable of ensuring a certain level of reliability and availability in a wireless heterogeneous mesh network, especially if some of the nodes

(e.g., the end terminals) might be mobile. With currently standardized solutions, a MEC platform could only interact with the RAW network through the PCE. This reference point is defined between the MEC platform and the data plane of the virtualization infrastructure to instruct the data plane on how to route traffic among applications, networks, services, etc. This reference point is not further specified by ETSI MEC, but it would be the one that could be used by current solutions to allow for MEC to request, from the data plane on the RAW network, a certain behavior (in terms of availability and reliability) for MEC applications traffic flows. Note that as the PCE might reside outside the RAW network, the path between the RAW network and the PCE might be expensive and slow (e.g., it might need to traverse the whole RAW network) and reaching to the PCE can in this case be slow regarding the speed of events that affect the forwarding operation at the radio layer.
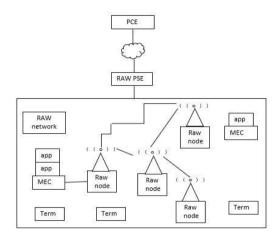

Fig. 3. RAW and MEC integration: Block diagram

Additionally, the MEC architecture as currently defined by ETSI does not have any component designed to deal with the specifics of a heterogeneous wireless multi-hop networks (e.g., RAW), and therefore, it is very limited in terms of what a MEC application (through the MEC platform) can request to the data plane of a heterogeneous wireless multi-hop network. Besides this, the lack of RAW-aware component at the ETSI MEC architecture prevents any enhancement at either the MEC side (e.g., MEC application migrations triggered by RAW status updates) or the RAW side (e.g., PAREO function updates triggered by MEC app/terminal mobility).

Because of all these reasons, there is a growing need to define a new RAW-enabled component at the ETSI MEC architecture, aimed at enabling a more direct interaction between the MEC platform and the RAW network, allowing the MEC platform to notify events and/or request actions to the RAW network quick enough. This involves some challenges, as the RAW PSE must understand the needs from the running MEC applications, so it can properly configure the RAW nodes to provide the required reliability and availability.

B. RAW and MEC integration

This paper proposes a new entity inside the MEC platform: the RAW control function (RAW ctrl) as shown in Fig. 4. This entity is responsible for computing what to instruct the RAW PSE, based on the requirements of the MEC applications, as

well as to take decisions at the MEC side (e.g., migration of applications) based on information about the RAW network status.

As a result of the introduction of the RAW ctrl and the actions it is responsible of, new interactions and interface semantics are added. These interactions and semantics can be terminated at either the PCE or the RAW PSE, depending on the requirements of the MEC applications. For near real-time coordination and control between MEC and RAW mechanisms, the interactions are between the RAW ctrl and the RAW PSE. We mostly refer to this deployment model from now on, as it is the one that allows for near real-time updates on the forwarding plane. However, it is worth noting that an alternative deployment model in which the RAW ctrl interacts with the PCE is also possible, though only supporting non-real-time interactions.

The MEC-RAW new interface semantics/extensions are depicted in Fig. 3 and allow the MEC platform to issue requests to the RAW network, through the RAW PSE, so it can behave as required by the MEC applications. The new semantics of the interface between the MEC platform and the RAW PSE do not only serve to convey the requests, but also to synchronize the status and topology of the RAW (relevant portion of the) network, enabling real-time or near-real time forwarding decisions. Next, we show an exemplary signaling diagram for the most relevant procedures.

C. Terminal-based joint selection and configuration of MEC host and RAW network

We need extensions to enable: i) a terminal to discover any RAW-enabled network on the path between the terminal and the MEC application host, and the RAW network associated capabilities; ii) the terminal to request desired reliability and availability requirements to be met simultaneously by the MEC+RAW system; and, iii) direct notifications from the RAW network to the terminal, to help with E2E application-level optimization. Most of the required extensions are related to ETSI MEC components and interfaces, and therefore are out of the scope of the IETF. However, we still briefly describe them for completeness, focusing on the IETF RAW components and interactions.
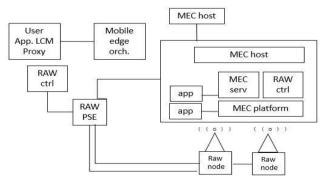

Fig. 4. Terminal-based joint selection and configuration of MEC host and RAW network: Block diagram

Fig. 4 shows the components and interfaces impacted by the extensions mentioned in this paper. The MEC Use Application Life Cycle Manager Proxy block (UALCMP) is logically extended with a RAW controller (RAW ctrl) functionality, to

enable a terminal to learn about the RAW and MEC capabilities of the 5G network it is attached to and specify its requirements in terms of availability and reliability for joint MEC application instantiation and RAW network configuration.

V. MOBILITY IN SINGLE RAW DOMAIN

As opposed to static scenarios, where possible "tracks" do not change due to mobility, mobility scenarios pose additional complexity that has not been tackled yet. Control plane solutions need to cope with mobility, by proactively preparing the network for the change of point of attachment of the mobile node, and the impact that this has in terms of new tracks used for the traffic. This requires inter-PSE coordination for the preparation of the handover. L2-specific extensions can be used to assist the mobile node in determining where to roam if stringent conditions that require RAW support need to be maintained.

The IETF DETNET and RAW WGs are responsible for the definition of data and control plane mechanisms to support deterministic networking in wired and wireless multi-hop networks. Current solutions are limited to static scenarios, where neither the mobile nodes nor the internal/local network nodes move. Therefore, solutions are needed to solve the mobile node mobility problem in single domain RAW networks. For example, it is needed to enable mechanisms allowing a terminal to signal an imminent handover and convey its QoS requirements. The signaling messages among RAW nodes (PSEs) to prepare and coordinate an imminent handover –so application's QoS can be maintained– need to be specified.

VI. CONCLUSION

This paper has stressed the need for new RAW extensions in order to being able to deploy deterministic networks in several vertical domains, specifically in industrial scenarios. Among the identified gaps, we have focused on support for multi-domain, integration with edge, and mobility support in RAW domains. We have touched on the involved SDOs, focusing mainly on IETF and ETSI MEC aspects. Finally, we have described in detail the motivations for these extensions and outlined the required needed extensions/solutions. Part of those will be addressed by the ongoing work of the PREDICT-6G research project.

ACKNOWLEDGMENT

REFERENCES

[1] D. R. Mafioletti, et al, "*Programmable Data Planes as the Next Frontier for Networked Robotics Security: A ROS Use Case*," 2021 17th International Conference on Network and Service Management (CNSM), Izmir, Turkey, 2021.

[2] Thubert, P. (2022, December). *Reliable and Available Wireless Architecture.* Internet-Draft, IETF.

[3] Frascolla, V. ., Cavalcanti, D. ., & Shah, R. . (2023). *Wi-Fi Evolution: The Path Towards Wi-Fi 7 and Its Impact on IIoT.* Journal of Mobile Multimedia, 19(01), 263–276.

[4] Bernardos, C. J., & Mourad, A. (2022, September). *Extensions to enable wireless reliability and availability in multi-access edge deployments.* Internet-Draft, IETF.

[5] Bernardos, C. J., & Mourad, A. (2022, September). *RAW multidomain extensions.* Internet-Draft, IETF.

[6] Bernardos, C. J., & Mourad, A. (2022, July). *DETNET multidomain extensions.* Internet-Draft, IETF.

[7] Bernardos, C. J., & Mourad, A. (2022, September). *Terminal-based joint selection and configuration of MEC host and RAW network.* Internet-Draft, IETF.

[8] Bernardos, C. J., Papadopoulos, G. Z., Thubert, P., & Theoleyre, F. (2022, October). *RAW Use-Cases*. Internet-Draft, IETF.

[9] Watteyne, T., Palattella, M., & Grieco, L. (2015, May). *Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement.* RFC 7554.

[10] Thubert, P. (2021, May). *An Architecture for IPv6 over the Time-Slotted Channel Hopping Mode of IEEE 802.15.4 (6TiSCH).* RFC 9030.

[11] Palattella, M. R., Thubert, P., Watteyne, T., & Wang, Q. (2018, March). *Terms Used in IPv6 over the TSCH mode of IEEE 802.15.4e.* Internet-Draft, IETF.

[12] Micah Beck. 2019. On the hourglass model. Commun. ACM 62, 7 (July 2019), 48–57. https://doi.org/10.1145/3274770.

[13] Martinez, B., Cano, C., & Vilajosana, X. (2019). *A Square Peg in a Round Hole: The Complex Path for Wireless in the Manufacturing Industry.* IEEE Communications Magazine, 57, 109-115. doi:10.1109/MCOM.2019.1800570

[14] Finn, N., Thubert, P., Varga, B., & Farkas, J. (2019, October). *Deterministic Networking Architecture.* RFC 8655.

[15] Farrel, A., Vasseur, J.-P., & Ash, J. (2006, August). *A Path Computation Element (PCE)-Based Architecture.* RFC 4655.

[16] Mirsky, G., Theoleyre, F., Papadopoulos, G. Z., Bernardos, C. J., Varga, B., & Farkas, J. (2022, October). *Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet).* Internet-Draft, IETF.

[17] Theoleyre, F., Papadopoulos, G. Z., Mirsky, G., & Bernardos, C. J. (2022, September). *Operations, Administration and Maintenance (OAM) features for RAW.* Internet-Draft, IETF.