

COINRG
Internet-Draft
Intended status: Informational
Expires: 6 May 2021

I. Kunze
K. Wehrle
RWTH Aachen University
D. Trossen
Huawei Technologies Duesseldorf GmbH
2 November 2020

Use Cases for In-Network Computing
draft-kunze-coin-industrial-use-cases-04

Abstract

Computing in the Network (COIN) comes with the prospect of deploying functionality on networking devices, such as switches and network interface cards. While such functionality can be beneficial in several contexts, it has to be carefully placed into the context of the general Internet communication. This document discusses some use cases to demonstrate how real applications can benefit from COIN and to showcase essential requirements that have to be fulfilled by COIN applications.

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 6 May 2021.

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](#) 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 |
| 2. | Terminology | 3 |
| 3. | Industrial Use Cases | 3 |
| 3.1. | IIoT Network Scenario | 4 |
| 3.2. | In-Network Control / Time-sensitive applications | 5 |
| 3.2.1. | Characterization and Requirements | 5 |
| 3.2.2. | Approaches | 6 |
| 3.3. | Large Volume Applications/ Traffic Filtering | 7 |
| 3.3.1. | Characterization and Requirements | 7 |
| 3.3.2. | Approaches | 8 |
| 3.4. | Industrial Safety (Dead Man's Switch) | 9 |
| 3.4.1. | Characterization and Requirements | 10 |
| 3.4.2. | Approaches | 10 |
| 4. | Security Considerations | 11 |
| 5. | Immersive Devices | 11 |
| 5.1. | Mobile Application Offloading | 11 |
| 5.2. | Edge AR/VR | 11 |
| 6. | Infrastructure Services | 11 |
| 6.1. | Distributed AI | 11 |
| 6.2. | Content Delivery Networks | 11 |
| 6.3. | CFaaS | 12 |
| 7. | Taxonomy | 12 |
| 8. | IANA Considerations | 12 |
| 9. | Conclusion | 12 |
| 10. | Informative References | 12 |
| | Authors' Addresses | 14 |

[1.](#) Introduction

The Internet bases on a best-effort network with limited guarantees regarding the timely and successful transmission of packets. Functionality is generally provided by the end-hosts while the network is kept simple and only intended to forward the packets. This design-choice is suitable for general Internet-based applications and has helped in the rapid growth of the Internet. However, there are several domains which, e.g., demand a number of strict performance guarantees that cannot be provided over regular best-effort networks. In this context, flexibly distributing the computation tasks across the network can help to achieve the guarantees and increase the overall performance.

The different domains, however, have different requirements and it is unclear whether there can be a common solution to all COIN scenarios or if solutions have to be tailored to each scenario.

This document first presents applications and requirements of some domains to illustrate the importance of COIN for realizing advanced applications. Based on these discussion, the draft then creates a taxonomy of COIN scenarios with the goal of guiding future work.

2. Terminology

Programmable network devices (PNDs): Network devices, such as network interface cards and switches, which are programmable, e.g., using P4

3. Industrial Use Cases

The industrial domain is characterized by diverse sets of requirements which often cannot be provided over regular best-effort networks. Consequently, there is a large number of specialized applications and protocols designed to give the required strict performance guarantees, e.g., regarding real-time capabilities. Time-Sensitive-Networking [TSN] as an enhancement to the standard Ethernet, e.g., tries to achieve these requirements on the link layer by statically reserving shares of the bandwidth. In the Industrial Internet of Things (IIoT), however, more and more parts of the industrial production domain are interconnected. This increases the complexity of the industrial networks, makes them more dynamic, and creates more diverse sets of requirements. In these scenarios, solutions on the link layer alone are not sufficient.

The challenge is to develop concepts that can satisfy the dynamic performance requirements of modern industrial networks. COIN presents a promising starting point because it allows to flexibly distribute computation tasks across the network which can help to manage dynamic changes. As specifying general requirements for the industrial production domain is difficult due to the mentioned diversity, this document next characterizes and analyzes three distinct scenarios to showcase potential requirements for the industrial production domain, thereby illustrating how COIN can be helpful.

3.1. IIoT Network Scenario

Common components of the IIoT can be divided into three categories as illustrated in Figure 1. Following [\[I-D.mcbride-edge-data-discovery-overview\]](#), EDGE DEVICES, such as sensors and actuators, constitute the boundary between the physical and digital world. They communicate the current state of the physical world to the digital world by transmitting sensor data or let the digital world interact with the physical world by executing actions after receiving (simple) control information. The processing of the sensor data and the creation of the control information is done on COMPUTING DEVICES. They range from small-powered controllers close to the EDGE DEVICES, to more powerful edge or remote clouds in larger distances. The connection between the EDGE and COMPUTING DEVICES is established by NETWORKING DEVICES. In the industrial domain, they range from standard devices, e.g., typical Ethernet switches, which can interconnect all Ethernet-capable hosts, to proprietary equipment with proprietary protocols only supporting hosts of specific vendors.

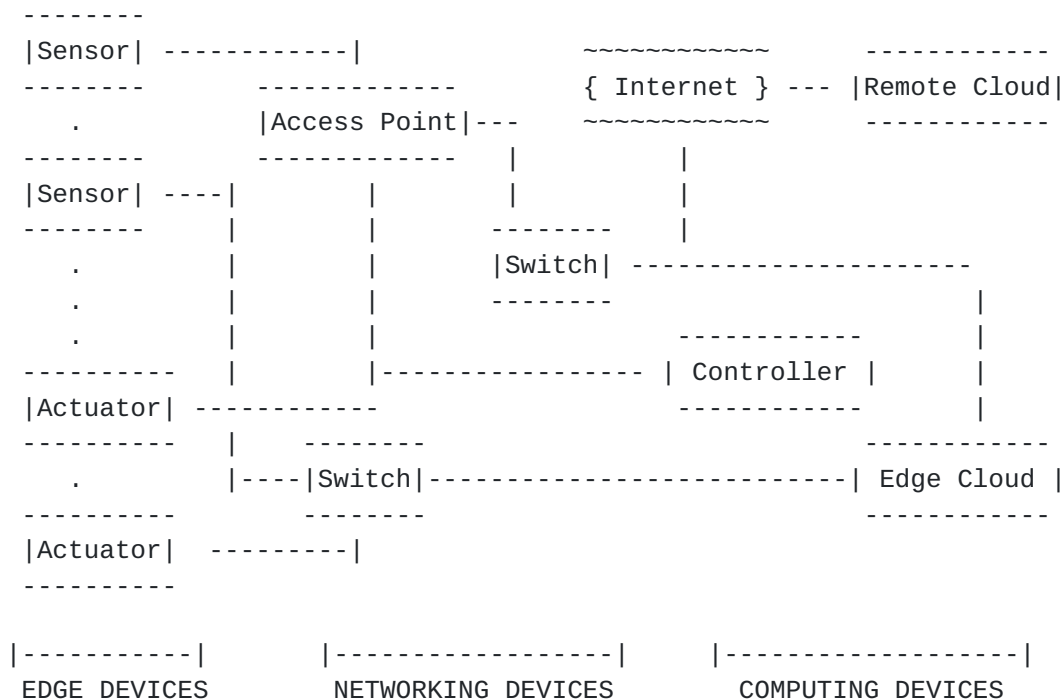


Figure 1: Industrial networks show a high level of heterogeneity.

3.2. In-Network Control / Time-sensitive applications

The control of physical processes and components of a production line is essential for the growing automation of production and ideally allows for a consistent quality level. Traditionally, the control has been exercised by control software running on programmable logic controllers (PLCs) located directly next to the controlled process or component. This approach is best-suited for settings with a simple model that is focused on a single or few controlled components.

Modern production lines and shop floors are characterized by an increasing amount of involved devices and sensors, a growing level of dependency between the different components, and more complex control models. A centralized control is desirable to manage the large amount of available information which often has to be pre-processed or aggregated with other information before it can be used. PLCs are not designed for this array of tasks and computations could theoretically be moved to more powerful devices. These devices are no longer close to the controlled objects and induce additional latency.

It is worthwhile to investigate whether the outsourcing of control functionality to distant computation platforms is viable because these platforms have a high level of flexibility and scalability. In the following, we describe the requirements and characteristics of the control setting in more detail.

3.2.1. Characterization and Requirements

A control process consists of two main components as illustrated in Figure 2: a system under control and a controller. In feedback control, the current state of the system is monitored, e.g., using sensors and the controller influences the system based on the difference between the current and the reference state to keep it close to this reference state.

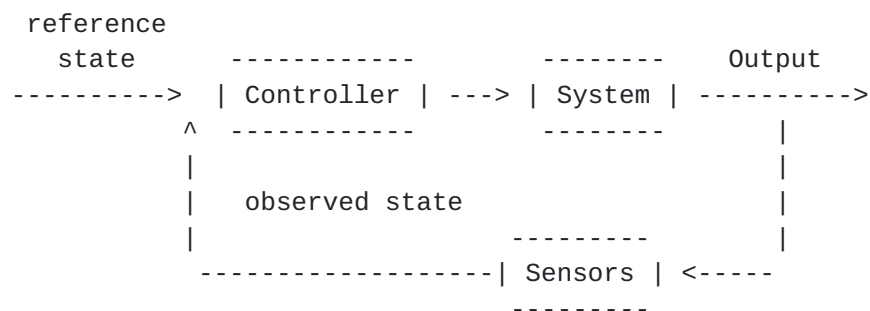


Figure 2: Simple feedback control model.

Apart from the control model, the quality of the control primarily depends on the timely reception of the sensor feedback, because the controller can only react if it is notified of changes in the system state. Depending on the dynamics of the controlled system, the control can be subject to tight latency constraints, often in the single-digit millisecond range. While low latencies are essential, there is an even greater need for stable and deterministic levels of latency, because controllers can generally cope with different levels of latency, if they are designed for them, but they are significantly challenged by dynamically changing or unstable latencies. The unpredictable latency of the Internet exemplifies this problem if off-premise cloud platforms are included.

The main requirements for the industrial control scenario are low and stable latencies to ensure that processes can work continuously and that no machines are damaged.

3.2.2. Approaches

Control models, in general, can become involved but there is a variety of control algorithms that are composed of simple computations such as matrix multiplication. These are supported by some PNDs and it is thus possible to compose simplified approximations of the more complex algorithms and deploy them in the network. While the simplified versions induce a more inaccurate control, they allow for a quicker response and might be sufficient to operate a basic tight control loop while the overall control can still be exercised from the cloud. The problem, however, is that networking devices typically only allow for integer precision computation while floating-point precision is needed by most control algorithms. Additionally, computational capabilities vary for different available PNDs. Yet, early approaches like [\[RUETH\]](#) and [\[VESTIN\]](#) have already shown the general applicability of such ideas, but there are still a lot of open research questions not limited to the following:

- * How can one derive the simplified versions of the overall controller?
 - How complex can they become?
 - How can one take the limited computational precision of networking devices into account when making them?
- * How does one distribute the simplified versions in the network?
- * How does the overall controller interact with the simplified versions?

3.3. Large Volume Applications/ Traffic Filtering

In the IIoT, processes and machines can be monitored more effectively resulting in more available information. This data can be used to deploy machine learning (ML) techniques and consequently help to find previously unknown correlations between different components of the production which in turn helps to improve the overall production system. Newly gained knowledge can be shared between different sites of the same company or even between different companies [[PENNEKAMP](#)].

Traditional company infrastructure is neither equipped for the management and storage of such large amounts of data nor for the computationally expensive training of ML approaches. Similar to the considerations in [Section 3.2](#), off-premise cloud platforms offer cost-effective solutions with a high degree of flexibility and scalability. While the unpredictable latency of the Internet is only a subordinate problem for this use case, moving all data to off-premise locations primarily poses infrastructural challenges which are presented in more detail in the following.

3.3.1. Characterization and Requirements

Processes in the industrial domain are monitored by distributed sensors which range from simple binary (e.g., light barriers) to sophisticated sensors measuring the system with varying degrees of resolution. Sensors can further serve different purposes, as some might be used for time-critical process control while others are only used as redundant fallback platforms. Overall, there is a high level of heterogeneity which makes managing the sensor output a challenging task.

Depending on the deployed sensors and the complexity of the observed system, the resulting overall data volume can easily be in the range of several Gbit/s [[GLEBKE](#)]. Using off-premise clouds for managing the data requires uploading or streaming the growing volume of sensor data using the companies' Internet access which is typically limited to a few hundred of Mbit/s. While large networking companies can simply upgrade their infrastructure, most industrial companies rely on traditional ISPs for their Internet access. Higher access speeds are hence tied to higher costs and, above all, subject to the supply of the ISPs and consequently not always available. A major challenge is thus to devise a methodology that is able to handle such amounts of data over limited access links.

Another aspect is that business data leaving the premise and control of the company further comes with security concerns, as sensitive information or valuable business secrets might be contained in it. Typical security measures such as encrypting the data make in-network

computing techniques hardly applicable as they typically work on unencrypted data. Adding security to in-network computing approaches, either by adding functionality for handling encrypted data or devising general security measures, is thus an auspicious field for research which we describe in more detail in [Section 4](#).

3.3.2. Approaches

There are at least two concepts which might be suitable for reducing the amount of transmitted data in a meaningful way:

1. filtering out redundant or unnecessary data
2. aggregating data by applying pre-processing steps within the network

Both concepts require detailed knowledge about the monitoring infrastructure at the factories and the purpose of the transmitted data.

3.3.2.1. Traffic Filters

Sensors are often set up redundantly, i.e., part of the collected data might also be redundant. Moreover, they are often hard to configure or not configurable at all which is why their resolution or sampling frequency is often larger than required. Consequently, it is likely that more data is transmitted than is needed or desired. A trivial idea for reducing the amount of data is to filter out redundant or undesired data before it leaves the premise using simple traffic filters that are deployed in the on-premise network. There are different approaches to how this topic can be tackled. A first step would be to scale down the available sensor data to the data rate that is needed. For example, if a sensor transmits with a frequency of 5 kHz, but the control entity only needs 1 kHz, only every fifth packet containing sensor data is let through. Alternatively, sensor data could be filtered down to a lower frequency while the sensor value is in an uninteresting range, but let through with higher resolution once the sensor value range becomes interesting. It is important that end-hosts are informed about the filtering so that they can distinguish between data loss and data filtered out on purpose.

In this context, the following research questions can be of interest:

- * How can traffic filters be designed?
- * How can traffic filters be coordinated and deployed?

- * How can traffic filters be changed dynamically?
- * How can traffic filtering be signaled to the end-hosts?

3.3.2.2. In-Network (Pre-)Processing

There are manifold computations that can be performed on the sensor data in the cloud. Some of them are very complex or need the complete sensor data during the computation, but there are also simpler operations which can be done on subsets of the overall dataset or earlier on the communication path as soon as all data is available. One example is finding the maximum of all sensor values which can either be done iteratively at each intermediate hop or at the first hop, where all data is available.

Using expert knowledge about the exact computation steps and the concrete transmission path of the sensor data, simple computation steps can be deployed in the on-premise network to reduce the overall data volume and potentially speed up the processing time in the cloud.

Related work has already shown that in-network aggregation can help to improve the performance of distributed ML applications [[SAPIO](#)]. Investigating the applicability of stream data processing techniques to programmable networking devices is also interesting, because sensor data is usually streamed. In this context, the following research questions can be of interest:

- * Which (pre-)processing steps can be deployed in the network?
 - How complex can they become?
- * How can applications incorporate the (pre-)processing steps?
- * How can the programming of the techniques be streamlined?

3.4. Industrial Safety (Dead Man's Switch)

Despite increasing automation in production processes, human workers are still often necessary. Consequently, safety measures have a high priority to ensure that no human life is endangered. In traditional factories, the regions of contact between humans and machines are well-defined and interactions are simple. Simple safety measures like emergency switches at the working positions are enough to provide a decent level of safety.

Modern factories are characterized by increasingly dynamic and complex environments with new interaction scenarios between humans and robots. Robots can either directly assist humans or perform tasks autonomously. The intersect between the human working area and the robots grows and it is harder for human workers to fully observe the complete environment.

Additional safety measures are essential to prevent accidents and support humans in observing the environment. The increased availability of sensor data and the detailed monitoring of the factories can help to build additional safety measures if the corresponding data is collected early at the correct position.

3.4.1. Characterization and Requirements

Industrial safety measures are typically hardware solutions because they have to pass rigorous testing before they are certified and deployment-ready. Standard measures include safety switches and light barriers. Additionally, the working area can be explicitly divided into 'contact' and 'safe' areas, indicating when workers have to watch out for interactions with machinery.

These measures are static solutions, potentially relying on specialized hardware, and are challenged by the increased dynamics of modern factories where the factory configuration can be changed on demand. Software solutions offer higher flexibility as they can dynamically respect new information gathered by the sensor systems, but in most cases they cannot give guaranteed safety. Yet, it is worthwhile to investigate whether such solutions can introduce additional safety measures.

3.4.2. Approaches

Software-based solutions can take advantage of the large amount of available sensor data. Different safety indicators within the production hall can be combined within the network so that programmable networking devices can give early responses if a potential safety breach is detected. A rather simple possibility could be to track the positions of human workers and robots. Whenever a robot gets too close to a human in a non-working area or if a human enters a defined safety zone, robots are stopped to prevent injuries. More advanced concepts could also include image data or combine arbitrary sensor data.

In this context, the following research questions can be of interest:

- * Which additional safety measures can be provided?

- Do these measures actually improve safety?
- * Which sensor information can be combined and how?

4. Security Considerations

Note: This section will need consolidation once new use cases are added to the draft. Current in-network computing approaches typically work on unencrypted plain text data because today's networking devices usually do not have crypto capabilities. As is already mentioned in [Section 3.3.1](#), this above all poses problems when business data, potentially containing business secrets, is streamed into remote computing facilities and consequently leaves the control of the company. Insecure on-premise communication within the company and on the shop-floor is also a problem as machines could be intruded from the outside. It is thus crucial to deploy security and authentication functionality on on-premise and outgoing communication although this might interfere with in-network computing approaches. Ways to implement and combine security measures with in-network computing are described in more detail in [[I-D.fink-coin-sec-priv](#)].

5. Immersive Devices

5.1. Mobile Application Offloading

NOTE: Will be moved here from
[I-D.[draft-sarathchandra-coin-appcentres-03](#)]

5.2. Edge AR/VR

NOTE: Could be transfered from (expired)
[I-D.[draft-montpetit-coin-xr-03](#)]

Additional potential sources: [I-D.[draft-geng-rtgwg-cfn-req-00](#)]

6. Infrastructure Services

6.1. Distributed AI

NOTE: Will be moved here from
[I-D.[draft-sarathchandra-coin-appcentres-03](#)]

6.2. Content Delivery Networks

NOTE: Will be moved here from
[I-D.[draft-sarathchandra-coin-appcentres-03](#)]

6.3. CFaaS

NOTE: Will be moved here from
[I-D.[draft-sarathchandra-coin-appcentres-03](#)]

7. Taxonomy

NOTE: The taxonomy is intended to generalize characteristics of the different presented use cases and work on it will start once more use cases are added to the draft.

8. IANA Considerations

N/A

9. Conclusion

There are several domains that can profit from COIN.

Industrial scenarios have unique sets of requirements mostly focusing around tight latency constraints with high required bandwidths.

NOTE: Further aspects will be added once more use cases are added to the draft.

10. Informative References

[GLEBKE] Glebke, R., Henze, M., Wehrle, K., Niemietz, P., Trauth, D., Mattfeld MBA, P., and T. Bergs, "A Case for Integrated Data Processing in Large-Scale Cyber-Physical Systems", Proceedings of the 52nd Hawaii International Conference on System Sciences, DOI 10.24251/hicss.2019.871, 2019, <<https://doi.org/10.24251/hicss.2019.871>>.

[I-D.[draft-geng-rtgwg-cfn-req-00](#)]
Geng, L. and P. Willis, "Compute First Networking (CFN) Scenarios and Requirements", Work in Progress, Internet-Draft, [draft-geng-rtgwg-cfn-req-00](#), 4 November 2019, <<http://www.ietf.org/internet-drafts/draft-geng-rtgwg-cfn-req-00.txt>>.

[I-D.[draft-montpetit-coin-xr-03](#)]
Montpetit, M., "In Network Computing Enablers for Extended Reality", Work in Progress, Internet-Draft, [draft-montpetit-coin-xr-03](#), 8 July 2019, <<http://www.ietf.org/internet-drafts/draft-montpetit-coin-xr-03.txt>>.

[I-D.[draft-sarathchandra-coin-appcentres-03](#)]

Trossen, D., Sarathchandra, C., and M. Boniface, "In-Network Computing for App-Centric Micro-Services", Work in Progress, Internet-Draft, [draft-sarathchandra-coin-appcentres-03](#), 23 October 2020, <<http://www.ietf.org/internet-drafts/draft-sarathchandra-coin-appcentres-03.txt>>.

[I-D.fink-coin-sec-priv]

Fink, I. and K. Wehrle, "Enhancing Security and Privacy with In-Network Computing", Work in Progress, Internet-Draft, [draft-fink-coin-sec-priv-01](#), 8 September 2020, <<http://www.ietf.org/internet-drafts/draft-fink-coin-sec-priv-01.txt>>.

[I-D.mcbride-edge-data-discovery-overview]

McBride, M., Kutscher, D., Schooler, E., Bernardos, C., Lopez, D., and X. Foy, "Edge Data Discovery for COIN", Work in Progress, Internet-Draft, [draft-mcbride-edge-data-discovery-overview-05](#), 1 November 2020, <<http://www.ietf.org/internet-drafts/draft-mcbride-edge-data-discovery-overview-05.txt>>.

[PENNEKAMP]

Pennekamp, J., Henze, M., Schmidt, S., Niemietz, P., Fey, M., Trauth, D., Bergs, T., Brecher, C., and K. Wehrle, "Dataflow Challenges in an Internet of Production: A Security & Privacy Perspective", Proceedings of the ACM Workshop on Cyber-Physical Systems Security & Privacy - CPS-SPC'19, DOI 10.1145/3338499.3357357, 2019, <<https://doi.org/10.1145/3338499.3357357>>.

[RUETH]

Rueth, J., Glebke, R., Wehrle, K., Causevic, V., and S. Hirche, "Towards In-Network Industrial Feedback Control", Proceedings of the 2018 Morning Workshop on In-Network Computing - NetCompute '18, DOI 10.1145/3229591.3229592, 2018, <<https://doi.org/10.1145/3229591.3229592>>.

[SAPIO]

Sapio, A., "Scaling Distributed Machine Learning with In-Network Aggregation", 2019, <<https://arxiv.org/abs/1903.06701>>.

[TSN]

IEEE, ., "Time-Sensitive Networking (TSN) Task Group", 2019, <<https://1.ieee802.org/tsn/>>.

[VESTIN]

Vestin, J., Kassler, A., and J. Akerberg, "FastReact: In-Network Control and Caching for Industrial Control Networks using Programmable Data Planes", 2018 IEEE 23rd

International Conference on Emerging Technologies and
Factory Automation (ETFA), DOI 10.1109/etfa.2018.8502456,
September 2018,
<<https://doi.org/10.1109/etfa.2018.8502456>>.

Authors' Addresses

Ike Kunze
RWTH Aachen University
Ahornstr. 55
D-52074 Aachen
Germany

Email: kunze@comsys.rwth-aachen.de

Klaus Wehrle
RWTH Aachen University
Ahornstr. 55
D-52074 Aachen
Germany

Email: wehrle@comsys.rwth-aachen.de

Dirk Trossen
Huawei Technologies Duesseldorf GmbH
Riesstr. 25C
D-80992 Munich
Germany

Email: Dirk.Trossen@Huawei.com

