

COINRG
Internet-Draft
Intended status: Informational
Expires: 25 April 2022

I. Fink
K. Wehrle
RWTH Aachen University
22 October 2021

Enhancing Security and Privacy with In-Network Computing
draft-fink-coin-sec-priv-03

Abstract

With the growing interconnection of devices, cyber security and data protection are of increasing importance. This is especially the case regarding cyber-physical systems due to their close entanglement with the physical world. Misbehavior and information leakage can lead to financial and physical damage and endanger human lives and well-being. Thus, hard security and privacy requirements are necessary to be met. Furthermore, a thorough investigation of incidents is essential for ultimate protection. Computing in the Network (COIN) allows the processing of traffic and data directly in the network and at line-rate. Thus, COIN presents a promising solution for efficiently providing security and privacy mechanisms as well as network monitoring. This document discusses select mechanisms to demonstrate how COIN concepts can be applied to counter existing shortcomings of cyber security and data privacy.

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 25 April 2022.

Copyright Notice

Copyright (c) 2021 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. Protection Mechanisms	4
2.1. Encryption and Integrity Checks	4
2.2. Authentication and Authorization	5
2.3. Policies	6
2.4. In-Network Vulnerability Patches	7
2.5. Anonymization	7
3. Intrusion and Anomaly Detection	8
3.1. Intrusion Detection	8
3.2. Dead Man's Switch	9
4. Network Monitoring	9
5. Security Considerations	9
6. IANA Considerations	10
7. Conclusion	10
8. Informative References	10
Authors' Addresses	12

1. Introduction

With the ongoing digitalization, previously isolated devices and systems are increasingly connected to the Internet, concerning all aspects of life. In particular, in the context of Cyber-Physical Systems (CPS) and the (Industrial) Internet of Things, machines and infrastructure are equipped with additional sensors and CPUs to allow for automatization and higher processing efficiency. The entanglement of the sensors with the physical world leads to high sensitivity of the transmitted and collected data.

Consequently, digitalization expands the attack surface and the possible impacts of cyber attacks, increasing the importance of proper protection mechanisms.

Devices in CPS are often resource-constrained and do not offer the possibility to implement elaborate security mechanisms. Furthermore, legacy devices and communication protocols are often still used in industrial networks but were not designed to face the security and privacy challenges the new interconnection brings. Thus, communication and access are often unprotected.

Upgrading legacy devices with protection mechanisms is an effortful and expensive procedure. A promising approach for retrofitting security is the deployment of suitable mechanisms within the network. To date, this is mainly realized using middle-boxes, leading to overhead and the need for additional hardware.

One general and widespread security component is Intrusion Detection Systems (IDS) to detect and, ideally, prevent undesired events in a network. However, IDS are usually implemented in software, again running on middle-boxes or edge devices in the same network. Thus, their reaction time is limited as well as their information gain, which is usually addressed by deploying additional IDS components for traffic monitoring.

Last, the after-treatment of incidents in networks is critical to detect exploited vulnerabilities and prevent future attacks. Network forensics serves to retrace and comprehend the origin and course of malicious events. However, to provide high performance, the underlying monitoring of network traffic requires dedicated networking devices and/or expensive subscriptions to respective features, leading to high costs.

One common problem is that security is usually provided by software solutions. These often require additional hardware and lead to performance overhead, which is especially unfavorable in the context of time-sensitive applications, e.g., in industry. Existing high-performance solutions, e.g., running on traditional networking devices, require dedicated, costly, and unsustainable hardware.

Computing in the Network (COIN) covers these shortfalls by using programmable networking devices to conduct dynamic and custom processing of network packets at line-rate. Thus, security-related functions and packet inspection can be implemented and applied centrally in the network, e.g., at a programmable switch.

This draft explores the opportunities of COIN for improving security and privacy as follows: We first describe feasible mechanisms for preventing attacks and intrusion in the first place. Then, we present which mechanisms we can implement with COIN for detecting intrusion and undesired behavior when it has already taken place. Last, we explore how COIN can improve network monitoring for detecting, analyzing and following up incidents, thus fixing vulnerabilities.

2. Protection Mechanisms

The common ground for providing security and data privacy is to protect against unauthorized access. That protection is primarily provided by basic security mechanisms such as encryption, integrity checks, authentication, and authorization. These are often missing in resource-constrained environments or regarding legacy industrial devices. [RFC7744] thoroughly discusses the need for authentication and authorization in resource-restrained environments. [RFC8576] presents security and privacy risks and challenges specific to the IoT. In the following, we describe how COIN can help to retrofit suitable mechanisms.

2.1. Encryption and Integrity Checks

Encryption is critical to preserve confidentiality when transmitting data. Integrity checks prevent undetected manipulation, which can remain unnoticed even despite encryption, e.g., in case of flipped bits. Due to resource-constraints, many devices in CPS do not provide encryption or calculation of check-sums.

By default, secure cryptographic functions are not supported by current programmable switches either and hard to realize due to their computational constraints. However, there are efforts by researchers to implement AES encryption with scrambled lookup tables [CHEN] and cryptographically secure keyed hash functions [YOO] on existing programmable hardware switches. Furthermore, future generations of programmable hardware switches might provide secure cryptographic functions by design.

With secure cryptography at hand, COIN would allow to encrypt and hash packet contents efficiently while passing a respective programmable networking device. Concretely, data could be encrypted and supplemented with a check-sum directly at the first networking device passed by a packet before forwarding it through the Internet to its designated destination. Subsequent decryption and integrity checks could be executed at the last networking device before the destination. Alternatively, this could be implemented at the destination if supported by the respective device. This approach

does not require deployment of or forwarding to additional middle-boxes. Thus, no additional attack surface or processing overhead is introduced, presenting a promising foundation for retrofitting security in time-sensitive scenarios such as industrial processes.

Another use-case is the implementation of whole standards for secure communication on programmable networking devices, offering new flexibility. For example, researchers examined the benefits of deploying IPsec and MACsec on programmable switches [HAUSER-IPsec],[HAUSER-MACsec] but their implementations only target software switches due to the missing cryptographic capabilities of existing programmable hardware switches.

Future research is needed to clarify if and at which costs hardware for enabling cryptographic calculations could and should be embedded in future generations of programmable networking devices.

2.2. Authentication and Authorization

Authentication and authorization mechanisms are needed to avoid unauthorized access to devices and their manipulation in the first place. With COIN, networking devices can flexibly decide whether to forward packets, thus offer efficient and fine-granular access control.

One possibility is to conduct a handshake between the sender and networking device before starting the communication with industrial devices. Cryptographic calculations (e.g. required for certificate-based authentication) can be offloaded to the control plane if not feasible in the data plane of the networking device due to computational constraints. Existing research also proposed and implemented authentication in the data plane, e.g., using port-knocking [ALMAINI]. Authorization information can be stored in tables in the data plane or requested from the control plane. Since authentication and authorization is only needed when starting or refreshing a connection, the necessity and overhead for consulting the control plane are limited. Subsequent to the authentication and authorization process, the respective decisions can be flexibly enforced by the networking device. For example, different fine-granular policies can be linked to different authorization levels and different devices. In the case of unsuccessful authorization or authentication, networking devices can inform the network administrator about possible intrusion of the system.

Overall, COIN can realize efficient and flexible access control, reducing overhead and attack surface.

2.3. Policies

Control processes can include communication between various parties. Even despite authorization and authentication mechanisms, undesired behavior can occur. For instance, malicious third-party software might be installed on the approved device and thus implicitly gain access. Depending on the involved devices and their capabilities, proper authorization and authentication might not be possible at all. An effective way to exclude malicious behavior nevertheless is policy-based access control.

[RFC8520] proposes the Manufacturer Usage Description (MUD), a standard for defining the communication behavior of IoT devices, which use specific communication patterns. The definition is primarily based on domain names, ports, and protocols (e.g., TCP and UDP). Further characteristics as the TLS usage [I-D.draft-ietf-opsawg-mud-tls-05] or the required bandwidth of a device [I-D.draft-lear-opsawg-mud-bw-profile-01] can help to define connections more narrowly. By defining the typical behavior, we can exclude deviating communication, including undesired behavior. Likewise to IoT devices, industrial devices usually serve a specific purpose. Thus, applying MUD or similar policies is viable in industrial scenarios as well.

The problem that remains is the efficient enforcement of such policies through fine-granular and flexible traffic filtering. While middle-boxes increase costs and processing overhead, primary SDN approaches as OpenFlow allow only filtering based on match-action rules regarding fixed protocol header fields. Evaluation of traffic statistics for, e.g., limiting the bandwidth, requires consultation of the remote controller. This leads to latency overheads, which are not acceptable in time-sensitive scenarios.

In contrast, the COIN paradigm allows flexible filtering even concerning the content of packets and connection metadata. Furthermore, traffic filtering can be executed by programmable networking devices at line-rate.

Last, not only network communication behavior of devices can be defined in policies. For example, COIN can be also used to consider additional (contextual) parameters, e.g., the time of day or activity of other devices in the network [KANG].

2.4. In-Network Vulnerability Patches

Resource-constrained devices are typically hard to update. Thus, device vulnerabilities often cannot be fixed after deployment. As a remedy and special case of policies, rules could be defined to describe known attack signatures. By enforcing these rules at programmable networking devices, e.g., by dropping matching traffic, COIN would offer an efficient way to avoid exploitation of device vulnerabilities. Another potential advantage is the easy and extensive roll-out of such "in-network patches" in the form of (automatic) software updates of the networking device.

Future research is needed to evaluate the potential and benefits of in-network patches compared to traditional security measures, e.g., firewalls, and provide proof of concepts using existing devices and vulnerabilities.

2.5. Anonymization

Due to their interconnection with the physical world, the generation of sensitive data is inherent to CPS. Smart infrastructure leads to the collection of sensitive (user) data. In industrial networks, information about confidential processes is gathered. Such data is increasingly shared with other entities to increase production efficiency or enable automatic processing.

Despite the benefits of data exchange, manufacturers and individuals might not want to share sensitive information. While deployment of privacy mechanisms is usually not possible at resource-constrained or legacy devices, COIN has the potential to apply privacy mechanisms in a flexible and comprehensive manner.

Data could be pseudonymized at networking devices by, e.g., extracting and replacing specific values. Furthermore, elaborate anonymization techniques could be implemented in the network by sensibly decreasing the data accuracy. For example, concepts like k-Anonymity could be applied by aggregating the values of multiple packets before forwarding the result. Noise addition could be implemented by adding a random number to values. Similarly, the state-of-the-art technique differential privacy could be implemented by adding noise to responses to statistical requests.

Indeed, recent research exploits programmable hardware switches to implement performant and light-weight anonymization of communication by rewriting source addresses and information and hiding path information, e.g., using randomization [MOGHADDAM]. Similarly, [WANG] realizes traffic obfuscation by encrypting IPv4 addresses in the data plane.

Future research is needed to implement and evaluate further privacy mechanisms and COIN's potential for this field.

3. Intrusion and Anomaly Detection

Ideally, attacks are prevented from the outset. However, in the case of incidents, fast detection is critical for limiting damage. Deployment of sensors, e.g., in industrial control systems, can help to monitor the system state and detect anomalies. This can be used in combination with COIN to detect intrusion and to provide advanced safety measures, as described in the following.

3.1. Intrusion Detection

Data of sensors or monitored communication behavior can be compared against expected patterns to detect intrusion. Even if intrusion prevention is deployed and connections are allowed when taken individually, subtle attacks might still be possible. For example, a series of values might be out of line if put into context even though the individual values are unobtrusive. Anomaly detection can be used to detect such abnormalities and notify the network administrator for further assessment.

While intrusion detection systems (IDS) are usually deployed at middle-boxes or external servers, COIN provides the possibility to detect anomalies at-line rate, e.g., by maintaining statistics about traffic flows. This decreases costs and latency, which is valuable for a prompt reaction. Another advantage is that one central networking device can monitor traffic from multiple devices. In contrast, multiple distributed middle boxes are usually needed to achieve the same information gain. Last, programmable networking devices can be used to preprocess traffic and reduce load on subsequent in IDS components as examined by [LEWIS].

Besides intrusion, anomalies can also imply safety risks. In the following, we pick up the potential of COIN to support safety.

3.2. Dead Man's Switch

[I-D.draft-irtf-coinrg-use-cases-00] addresses the potential of COIN for improving industrial safety. Detection of an anomaly in the sensor data or operational flow can be used to automatically trigger an emergency shutdown of a system or single system components if the data indicates an actual hazard. Apart from that, other safety measures like warning systems or isolation of areas can be implemented. While we do not aim at replacing traditional dead man's switches, we see the potential of COIN to accelerate the detection of failures. Thus, COIN can valuably complement existing safety measures.

4. Network Monitoring

After detecting an incident, it is essential to investigate its origin and scope. The results of this analysis can be used to allow for consistent recovery, to adapt protection mechanisms, and prevent similar events in the future. For enabling potential investigation, traffic is constantly captured (e.g., in the form of flow records), which requires dedicated hardware in large networks. Furthermore, it might be preferable to exclude traffic, e.g., from specific subnets, from the analysis. Dynamic and fine-granular traffic filtering is not possible with traditional networking devices, leading to storage and processing overhead.

With COIN, networking devices can be programmed to export traffic characteristics without significant overhead when forwarding a packet. Furthermore, monitoring can be done more flexibly, e.g., by applying fine-granular traffic filtering. Also, header fields of particular interest can be efficiently extracted. Therefore, COIN can considerably decrease the load and increase the efficiency of security-related network monitoring.

The presented prospects are underlined by recent work. For example, in [SONCHACK], flow records are created in the control plane of programmable hardware switches while expensive packet preprocessing is offloaded to the data plane.

5. Security Considerations

When implementing security and privacy measures in networking devices, their security and failure resistance is critical. Related research questions to clarify in the future are stated in [I-D.draft-kutscher-coinrg-dir-02].

6. IANA Considerations

N/A

7. Conclusion

COIN has the potential to improve and retrofit security and privacy, especially with regard to resource-restrained and legacy devices.

First, COIN can provide intrusion prevention mechanisms like authentication and efficient enforcement of (context-based) policies. Easily deployable in-network patches of device vulnerabilities could further improve security. Encryption and integrity checks are limited by the current hardware but already targeted by recent research.

Second, COIN allows examining packet contents at networking devices, which can help implement fast and comprehensive anomaly and intrusion detection.

Last, COIN can contribute to an efficient and targeted traffic monitoring for incident analysis.

Further investigation of the feasibility and potential of the presented mechanisms is subject to future research.

8. Informative References

- [ALMAINI] Almaini, A., Al-Dubai, A., Romdhani, I., Schramm, M., and A. Alsarhan, "Lightweight edge authentication for software defined networks", Computing 103, 291-311 (2021), August 2020, <<https://link.springer.com/article/10.1007/s00607-020-00835-4>>.
- [CHEN] Chen, X., "Implementing AES Encryption on Programmable Switches via Scrambled Lookup Tables", In Proceedings of the Workshop on Secure Programmable Network Infrastructure, August 2020, <<https://dl.acm.org/doi/abs/10.1145/3405669.3405819>>.
- [HAUSER-IPsec] Hauser, F., Häberle, M., Schmidt, M., and M. Menth, "P4-IPsec: Site-to-Site and Host-to-Site VPN With IPsec in P4-Based SDN", In IEEE Access, vol. 8, July 2020, <<https://ieeexplore.ieee.org/document/9151942>>.

[HAUSER-MACsec]

Hauser, F., Häberle, M., Schmidt, M., and M. Menth,
"P4-MACsec: Dynamic Topology Monitoring and Data Layer
Protection With MACsec in P4-Based SDN", In IEEE Access,
vol. 8, March 2020,
<<https://ieeexplore.ieee.org/document/9044731>>.

[I-D.draft-ietf-opsawg-mud-tls-05]

Reddy, T., Wing, D., and B. Anderson, "Manufacturer Usage
Description (MUD) (D)TLS Profiles for IoT Devices", Work
in Progress, Internet-Draft, draft-ietf-opsawg-mud-tls-05,
27 July 2021, <[https://datatracker.ietf.org/doc/html/
draft-ietf-opsawg-mud-tls-05](https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-mud-tls-05)>.

[I-D.draft-irtf-coinrg-use-cases-00]

Kunze, I., Wehrle, K., Trossen, D., and M.J. Montpetit,
"Use Cases for In-Network Computing", Work in Progress,
Internet-Draft, draft-irtf-coinrg-use-cases-00, 17
February 2021, <[https://tools.ietf.org/html/draft-irtf-
coinrg-use-cases-00](https://tools.ietf.org/html/draft-irtf-coinrg-use-cases-00)>.

[I-D.draft-kutscher-coinrg-dir-02]

Kutscher, D., Karkkainen, T., and J. Ott, "Directions for
Computing in the Network", Work in Progress, Internet-
Draft, draft-kutscher-coinrg-dir-02, 31 July 2020,
<[https://datatracker.ietf.org/doc/html/draft-kutscher-
coinrg-dir-02](https://datatracker.ietf.org/doc/html/draft-kutscher-coinrg-dir-02)>.

[I-D.draft-lear-opsawg-mud-bw-profile-01]

Lear, E. and O. Friel, "Bandwidth Profiling Extensions for
MUD", Work in Progress, Internet-Draft, draft-lear-opsawg-
mud-bw-profile-01, 8 July 2019,
<[https://datatracker.ietf.org/doc/html/draft-lear-opsawg-
mud-bw-profile-01](https://datatracker.ietf.org/doc/html/draft-lear-opsawg-mud-bw-profile-01)>.

[KANG]

Kang, Q., Morrison, A., Tang, Y., Chen, A., and X. Luo,
"Programmable In-Network Security for Context-aware BYOD
Policies", In Proceedings of the 29th USENIX Security
Symposium (USENIX Security 20), August 2020,
<[https://www.usenix.org/conference/usenixsecurity20/
presentation/kang](https://www.usenix.org/conference/usenixsecurity20/presentation/kang)>.

[LEWIS]

Lewis, B., Broadbent, A., and N. Race, "P4ID: P4 Enhanced
Intrusion Detection", 2019 IEEE Conference on Network
Function Virtualization and Software Defined Networks
(NFV-SDN), November 2019,
<<https://ieeexplore.ieee.org/document/9040044>>.

[MOGHADDAM]

Moghaddam, H. and A. Mosenia, "Programmable In-Network Obfuscation of Traffic", arXiv:1911.09642 [cs.CR], November 2019, <<https://arxiv.org/abs/1911.09642>>.

[RFC7744] Seitz, L., Ed., Gerdes, S., Ed., Selander, G., Mani, M., and S. Kumar, "Use Cases for Authentication and Authorization in Constrained Environments", RFC 7744, DOI 10.17487/RFC7744, January 2016, <<https://www.rfc-editor.org/info/rfc7744>>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.

[RFC8576] Garcia-Morchon, O., Kumar, S., and M. Sethi, "Internet of Things (IoT) Security: State of the Art and Challenges", RFC 8576, DOI 10.17487/RFC8576, April 2019, <<https://www.rfc-editor.org/info/rfc8576>>.

[SONCHACK] Sonchack, J., Aviv, A., Keller, E., and J. Smith, "Turboflow: Information Rich Flow Record Generation on Commodity Switches", In Proceedings of the Thirteenth EuroSys Conference, April 2018, <<https://dl.acm.org/doi/abs/10.1145/3190508.3190558>>.

[WANG] Wang, L., Kim, H., Mittal, P., and J. Rexford, "Programmable In-Network Obfuscation of Traffic", arXiv:2006.00097 [cs.NI], 2020, <<https://arxiv.org/abs/2006.00097>>.

[YOO] Yoo, S. and X. Chen, "Secure Keyed Hashing on Programmable Switches", In Proceedings of the ACM SIGCOMM 2021 Workshop on Secure Programmable Network Infrastructure, August 2021, <<https://arxiv.org/abs/1911.09642>>.

Authors' Addresses

Ina Berenice Fink
RWTH Aachen University
Ahornstr. 55
D-52062 Aachen
Germany

Phone: +49-241-80-21419
Email: fink@comsys.rwth-aachen.de

Klaus Wehrle
RWTH Aachen University
Ahornstr. 55
D-52062 Aachen
Germany

Phone: +49-241-80-21401
Email: wehrle@comsys.rwth-aachen.de

COINRG
Internet-Draft
Intended status: Informational
Expires: 28 April 2022

I. Kunze
K. Wehrle
RWTH Aachen
D. Trossen
Huawei
M.J. Montpetit
Concordia
X. de Foy
InterDigital Communications, LLC
D. Griffin
M. Rio
UCL
25 October 2021

Use Cases for In-Network Computing
draft-irtf-coinrg-use-cases-01

Abstract

Computing in the Network (COIN) comes with the prospect of deploying processing 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 28 April 2022.

Copyright Notice

Copyright (c) 2021 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	4
2. Terminology	5
3. Providing New COIN Experiences	6
3.1. Mobile Application Offloading	6
3.1.1. Description	6
3.1.2. Characterization	7
3.1.3. Existing Solutions	9
3.1.4. Opportunities and Research Questions for COIN	9
3.1.5. Requirements	10
3.2. Extended Reality (XR)	11
3.2.1. Description	11
3.2.2. Characterization	11
3.2.3. Existing Solutions	11
3.2.4. Opportunities and Research Questions for COIN	12
3.2.5. Requirements	14
3.3. Personalised and interactive performing arts	15
3.3.1. Description	15
3.3.2. Characterization	15
3.3.3. Existing solutions	17
3.3.4. Opportunities	17
3.3.5. Research Questions:	17
3.3.6. Requirements	18
4. Supporting new COIN Systems	19
4.1. Industrial Network Scenario	19
4.2. In-Network Control / Time-sensitive applications	20
4.2.1. Description	20
4.2.2. Characterization	21
4.2.3. Existing Solutions	21
4.2.4. Opportunities for COIN	22
4.2.5. Research Questions for COIN	22
4.2.6. Requirements	23
4.3. Large Volume Applications - Filtering	23

4.3.1.	Description	23
4.3.2.	Characterization	24
4.3.3.	Existing Solutions	25
4.3.4.	Opportunities for COIN	25
4.3.5.	Research Questions for COIN	25
4.3.6.	Requirements	26
4.4.	Large Volume Applications - (Pre-)Preprocessing	26
4.4.1.	Description	26
4.4.2.	Characterization	26
4.4.3.	Existing Solutions	26
4.4.4.	Opportunities for COIN	27
4.4.5.	Research Questions for COIN	27
4.4.6.	Requirements	27
4.5.	Industrial Safety	27
4.5.1.	Description	28
4.5.2.	Characterization	28
4.5.3.	Existing Solutions	28
4.5.4.	Opportunities for COIN	29
4.5.5.	Research Questions for COIN	29
4.5.6.	Requirements	29
5.	Improving existing COIN capabilities	29
5.1.	Content Delivery Networks	29
5.1.1.	Description	29
5.1.2.	Characterization	30
5.1.3.	Existing Solutions	30
5.1.4.	Opportunities and Research Questions for COIN	30
5.1.5.	Requirements	31
5.2.	Compute-Fabric-as-a-Service (CFaaS)	31
5.2.1.	Description	31
5.2.2.	Characterization	31
5.2.3.	Existing Solutions	32
5.2.4.	Opportunities and Research Questions for COIN	32
5.2.5.	Requirements	33
5.3.	Virtual Networks Programming	33
5.3.1.	Description	33
5.3.2.	Characterization	34
5.3.3.	Existing Solutions	36
5.3.4.	Opportunities	36
5.3.5.	Research Questions for COIN	37
5.3.6.	Requirements	38
6.	Enabling new COIN capabilities	38
6.1.	Distributed AI	38
6.1.1.	Description	38
6.1.2.	Characterization	39
6.1.3.	Existing Solutions	39
6.1.4.	Opportunities and Research Questions for COIN	39
6.1.5.	Requirements	40
7.	Analysis	40

8. Security Considerations	41
9. IANA Considerations	41
10. Conclusion	41
11. List of Use Case Contributors	41
12. References	42
12.1. Normative References	42
12.2. Informative References	42
Authors' Addresses	45

1. Introduction

The Internet is a best-effort network that offers limited guarantees regarding the timely and successful transmission of packets. Data manipulation and protocol functionality is generally provided by the end-hosts while the network is kept simple and only intended as a "store and forward" packet facility. This design choice is suitable for a wide variety of 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 or require more closed loop integration to manage data flows. In this context, allowing for a tighter integration of compute and network resources, enabling the more flexible distribution of computation tasks across the network, i.e., beyond 'just' endpoints, may help to achieve the desired guarantees and behaviours as well as increase the overall performance. The vision of 'in-network computing' and the provisioning of capabilities that capitalize on such joint computation and communication resource usage throughout the network is core to the efforts in the COIN RG; we term those capabilities as 'COIN capabilities' in the remainder of the document.

We believe that such vision of 'in-network computing' can be best outlined along four dimensions of use cases, namely those that (i) provide new user experiences through the utilization of COIN capabilities (termed 'COIN experiences'), (ii) enable new COIN systems, e.g., through new interactions between communication and compute providers, (iii) improve on already existing COIN capabilities as well as (iv) enable new such COIN capabilities. Sections 3 through 6 capture those categories of use cases as the main structure of this document.

Through delving into individual examples within each of the above categories, we aim to outline opportunities and possible research questions that may need consideration by the wider community when pushing forward the 'in-network computing' vision. Furthermore, insights into possible requirements for an evolving solution space of

collected COIN capabilities is another objective of the individual use case descriptions. This results in the following taxonomy used to describe each of the use cases:

1. Description: Purpose of the use case and explanation of the use case behavior
2. Characterization: Explanation of the services that are being utilized and realized as well as the semantics of interactions in the use case.
3. Existing solutions: Describe, if existing, current methods that may realize the use case.
4. Opportunities: Outline how COIN capabilities may support or improve on the use case in terms of performance and other metrics.
5. Research questions: State essential questions that are suitable for guiding research to achieve the outlined opportunities
6. Requirements: Describe the requirements for any solutions for COIN capabilities that may need development along the opportunities outlined in item 4; here, we limit requirements to those COIN capabilities, recognizing that any use case will realistically hold many additional requirements for its realization.

In order to provide a useful input into future roadmapping on what COIN capabilities may emerge and how solutions of such capabilities may look like as well as what questions remain to realize such solutions, we will analyze the use cases in Section 7 by providing an overview of key research questions across all use cases, while similarly gathering key requirements identified across all use cases. Through this, we intent to converge to key aspects in the form of possible open questions as well as requirements that may steer future (COIN) research work.

2. Terminology

The following terminology has been partly aligned with [I-D.draft-kutscher-coinrg-dir]:

(COIN) Program: a set of computations requested by a user

(COIN) Program Instance: one currently executing instance of a program

(COIN) Function: a specific computation that can be invoked as part of a program

COIN Capability: a feature enabled through the joint processing of computation and communication resources in the network

COIN Experience: a new user experience brought about through the utilization of COIN capabilities

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

(COIN) Execution Environment: a class of target environments for function execution, for example, a JVM-based execution environment that can run functions represented in JVM byte code

COIN System: the PNDs (and end systems) and their execution environments, together with the communication resources interconnecting them, operated by a single provider or through interactions between multiple providers that jointly offer COIN capabilities

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

3. Providing New COIN Experiences

3.1. Mobile Application Offloading

3.1.1. Description

The scenario can be exemplified in an immersive gaming application, where a single user plays a game using a VR headset. The headset hosts functions that "display" frames to the user, as well as the functions for VR content processing and frame rendering combining with input data received from sensors in the VR headset.

Once this application is partitioned into constituent (COIN) programs and deployed throughout a COIN system, utilizing the COIN execution environment, only the "display" (COIN) programs may be left in the headset, while the compute intensive real-time VR content processing (COIN) programs can be offloaded to a nearby resource rich home PC or a PND in the operator's access network, for a better execution (faster and possibly higher resolution generation).

3.1.2. Characterization

Partitioning a mobile application into several constituent (COIN) programs allows for denoting the application as a collection of (COIN) functions for a flexible composition and a distributed execution. In our example above, most functions of a mobile application can be categorized into any of three, "receiving", "processing" and "displaying" function groups.

Any device may realize one or more of the (COIN) programs of a mobile application and expose them to the (COIN) system and its constituent (COIN) execution environments. When the (COIN) program sequence is executed on a single device, the outcome is what you see today as applications running on mobile devices.

However, the execution of (COIN) functions may be moved to other (e.g., more suitable) devices, including PNDs, which have exposed the corresponding (COIN) programs as individual (COIN) program instances to the (COIN) system by means of a 'service identifier'. The result of the latter is the equivalent to 'mobile function offloading', for possible reduction of power consumption (e.g., offloading CPU intensive process functions to a remote server) or for improved end user experience (e.g., moving display functions to a nearby smart TV) by selecting more suitable placed (COIN) program instances in the overall (COIN) system.

Figure 1 shows one realization of the above scenario, where a 'DPR app' is running on a mobile device (containing the partitioned Display(D), Process(P) and Receive(R) COIN programs) over an SDN network. The packaged applications are made available through a localized 'playstore server'. The mobile application installation is realized as a 'service deployment' process, combining the local app installation with a distributed (COIN) program deployment (and orchestration) on most suitable end systems or PNDs ('processing server').

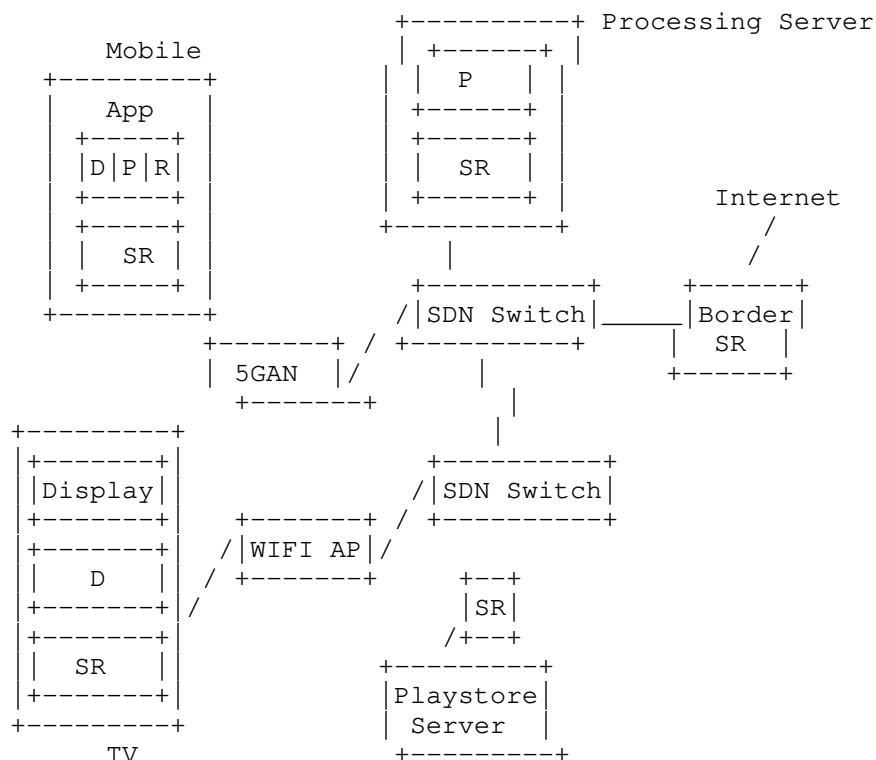


Figure 1: Application Function Offloading Example.

Such localized deployment could, for instance, be provided by a visiting site, such as a hotel or a theme park. Once the 'processing' (COIN) program is terminated on the mobile device, the 'service routing' (SR) elements in the network route (service) requests instead to the (previously deployed) 'processing' (COIN) program running on the processing server over an existing SDN network. Here, capabilities and other constraints for selecting the appropriate (COIN) program, in case of having deployed more than one, may be provided both in the advertisement of the (COIN) program and the service request itself.

As an extension to the above scenarios, we can also envision that content from one processing (COIN) program may be distributed to more than one display (COIN) program, e.g., for multi/many-viewing scenarios, thereby realizing a service-level multicast capability towards more than one (COIN) program.

3.1.3. Existing Solutions

NOTE: material on solutions like ETSI MEC will be added here later

3.1.4. Opportunities and Research Questions for COIN

Opportunities:

- * The packaging of (COIN) programs into existing mobile application packaging may enable the migration from current (mobile) device-centric execution of those mobile application towards a possible distributed execution of the constituent (COIN) programs that are part of the overall mobile application.
- * The orchestration for deploying (COIN) program instances in specific end systems and PNDs alike may open up the possibility for localized infrastructure owners, such as hotels or venue owners, to offer their compute capabilities to their visitors for improved or even site-specific experiences.
- * The execution of (current mobile) app-level (COIN) programs may speed up the execution of said (COIN) program by relocating the execution to more suitable devices, including PNDs.
- * The support for service-level routing of requests (service routing in [APPCENTRES]) may support higher flexibility when switching from one (COIN) program instance to another, e.g., due to changing constraints for selecting the new (COIN) program instance.
- * The ability to identifying service-level in-network computing elements will allow for routing service requests to those COIN elements, including PNDs, therefore possibly allowing for new in-network functionality to be included in the mobile application.
- * The support for constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) may allow for a more flexible and app-specific selection of (COIN) program instances, thereby allowing for better meeting the app-specific and end user requirements.

Research Questions:

- * RQ 3.1.1: How to combine service-level orchestration frameworks with app-level packaging methods?
- * RQ 3.1.2: How to reduce latencies involved in (COIN) program interactions where (COIN) program instance locations may change quickly?

- * RQ 3.1.3: How to signal constraints used for routing requests towards (COIN) program instances in a scalable manner?
- * RQ 3.1.4: How to identify (COIN) programs and program instances?
- * RQ 3.1.5: How to identify specific choice of (COIN) program instances over others?
- * RQ 3.1.6: How to provide affinity of service requests towards (COIN) program instances, i.e., longer-term transactions with ephemeral state established at a specific (COIN) program instance?
- * RQ 3.1.7: How to provide constraint-based routing decisions at packet forwarding speed?
- * RQ 3.1.8: What in-network capabilities may support the execution of (COIN) programs and their instances?

3.1.5. Requirements

Req 3.1.1: Any COIN system MUST provide means for routing of service requests between resources in the distributed environment.

Req 3.1.2: Any COIN system MUST provide means for identifying services exposed by (COIN) programs for directing service requests

Req 3.1.3: Any COIN system MUST provide means for identifying (COIN) program instances for directing (affinity) requests to a specific (COIN) program instance

Req 3.1.4: Any COIN system MUST provide means for dynamically choosing the best possible service sequence of one or more (COIN) programs for a given application experience, i.e., support for chaining (COIN) program executions.

Req 3.1.5: Means for discovering suitable (COIN) programs SHOULD be provided.

Req 3.1.6: Any COIN system MUST provide means for pinning the execution of a service of a specific (COIN) program to a specific resource, i.e., (COIN) program instance in the distributed environment.

Req 3.1.7: Any COIN system SHOULD provide means for packaging micro-services for deployments in distributed networked computing environments.

Req 3.1.8: The packaging MAY include any constraints regarding the deployment of (COIN) program instances in specific network locations or compute resources, including PNDs.

Req 3.1.9: Such packaging SHOULD conform to existing application deployment models, such as mobile application packaging, TOSCA orchestration templates or tar balls or combinations thereof.

Req 3.1.10: Any COIN system MUST provide means for real-time synchronization and consistency of distributed application states.

3.2. Extended Reality (XR)

3.2.1. Description

Virtual Reality (VR) and Augmented Reality (AR) taken together as Extended Reality (XR) are at the center of a number of advances in interactive technologies. While initially associated with gaming and entertainment, XR applications now include remote diagnosis, maintenance, telemedicine, manufacturing and assembly, autonomous systems, smart cities, and immersive classrooms.

3.2.2. Characterization

XR is one example of the Multisource-Multidestination Problem that combines video, haptics, and tactile experiences in interactive or networked multi-party and social interactions. Thus, XR is difficult to deliver with a client-server cloud-based solution as it requires a combination of: stream synchronization, low delays and delay variations, means to recover from losses and optimized caching and rendering as close as possible to the user at the network edge. Many XR services that involve video holography and haptics, require very low delay or generate large amounts of data, both requiring a careful look at data filtering and reduction, functional distribution and partitioning. Hence, XR uses recent advances in in-network programming, distributed networks, orchestration and resource discovery to support the XR advanced immersive requirements. It is important to note that the use of in-network computing for XR does not imply a specific protocol but targets an architecture enabling the deployment of the services. This includes computing in the nodes from content source to destination.

3.2.3. Existing Solutions

Related XR or XR-enabling solutions using in-network computation or related technologies include:

- * Enabling Scalable Edge Video Analytics with Computing-In-Network (Jun Chen Jiang of the University of Chicago): this work brings a periodical re-profiling to adapt the video pipeline to the dynamic video content that is a characteristic of XR. The implication is that "need tight network-app coupling" for real time video analytics.
- * VR journalism, interactive VR movies and meetings in cyberspace (many projects PBS, MIT interactive documentary lab, Huawei research - references to be provided): typical VR is not made for multiparty and these applications require a tight coupling of the local and remote rendering and data capture and combinations of cloud (for more static information) and edge (for dynamic content).
- * Local rendering of holographic content using near field computation (heritage from advances cockpit interactions - looking for non military papers): a lot has been said recently of the large amounts of data necessary to transmit and use holographic imagery in communications. Transmitting the near field information and rendering the image locally allows to reduce the data rates by 1 or 2.
- * ICE-AR [ICE] project at UCLA (Jeff Burke): while this project is a showcase of the NDN network architecture it also uses a lot of edge-cloud capabilities for example for inter-server games and advanced video applications.

3.2.4. Opportunities and Research Questions for COIN

Opportunities:

In-network computing for XR profits from the heritage of extensive research in the past years on Information Centric Networking, Machine Learning, network telemetry, imaging and IoT as well as distributed security and in-network coding. The opportunities include:

- * Reduced latency: the physical distance between the content cloud and the users must be short enough to limit the propagation delay to the 20 ms usually cited for XR applications; the use of local CPU and IoT devices for range of interest (RoI) detection and dynamic rendering may enable this.
- * Video transmission: better transcoding and use of advanced context-based compression algorithms, pre-fetching and pre-caching and movement prediction not only in the cloud.

- * **Monitoring:** telemetry is a major research topic for COIN and it enables to monitor and distribute the XR services.
- * **Network access:** push some networking functions in the kernel space into the user space to enable the deployment of stream specific algorithms for congestion control and application-based load balancing based on machine learning and user data patterns.
- * **Functional decomposition:** functional decomposition, localization and discovery of computing and storage resources in the network. But it is not only finding the best resources but qualifying those resources in terms of reliability especially for mission critical services in XR (medicine for example). This could include intelligence services.

Research Questions:

There is a need for more research resource allocation problems at the edge to enable interactive operation and quality of experience in VR. These include multi-variate and heterogeneous goal optimization problems requiring advanced analysis. Image rendering and video processing in XR leverages different HW capabilities combinations of CPU and GPU. Research questions include:

- * RQ 3.2.1: Can current programmable network entities be sufficient to provide the speed required to provide and execute complex filtering operations that includes metadata analysis for complex and dynamic scene rendering?
- * RQ 3.2.2: How can the interoperability of CPU/GPU be optimized to combine low level packet filtering with the higher layer processors needed for image processing and haptics?
- * RQ 3.2.3: Can the use of joint learning algorithms across both data center and edge computers be used to create optimal functionality allocation and the creation of semi-permanent datasets and analytics for usage trending resulting in better localization of XR functions?
- * RQ 3.2.4: Can COIN improve the dynamic distribution of control, forwarding and storage resources and related usage models in XR?

3.2.5. Requirements

XR requirements include the need to provide real-time interactivity for immersive and increasingly mobile immersive applications with tactile and time-sensitive data and high bandwidth for high resolution images and local rendering for 3D images and holograms. Since XR deals with personal information and potentially protected content XR must also provide a secure environment and ensure user privacy. Additionally, the sheer amount of data needed for and generated by the XR applications can use recent trend analysis and mechanisms, including machine learning to find these trends and reduce the size of the data sets. The requirements can be summarized as:

Req 3.2.1: Allow joint collaboration.

Req 3.2.2: Provide multi-views.

Req 3.2.3: Include extra streams dynamically for data intensive services, manufacturing and industrial processes.

Req 3.2.4: Enable multistream, multidevice, multideestination applications.

Req 3.2.5: Use new Internet Architectures at the edge for improved performance and performance management.

Req 3.2.6: Integrate with holography, 3D displays and image rendering processors.

Req 3.2.7: All the use of multicast distribution and processing as well as peer to peer distribution in bandwidth and capacity constrained environments.

Req 3.2.8: Evaluate the integration local and fog caching with cloud-based pre-rendering.

Req 3.2.9: Evaluate ML-based congestion control to manage XR sessions quality of service and to determine how to prioritize data.

Req 3.2.10: Consider higher layer protocols optimization to reduce latency especially in data intensive applications at the edge.

Req 3.2.11: Provide trust, including blockchains and smart-contracts to enable secure community building across domains.

Req 3.2.12: Support nomadicity and mobility (link to mobile edge).

Req 3.2.13: Use 5G slicing to create independent session-driven processing/rendering.

Req 3.2.14: Provide performance optimization by data reduction, tunneling, session virtualization and loss protection.

Req 3.2.15: Use AI/ML for trend analysis and data reduction when appropriate.

3.3. Personalised and interactive performing arts

3.3.1. Description

This use case covers live productions of the performing arts where the performers and audience are in different physical locations. The performance is conveyed to the audience through multiple networked streams which may be tailored to the requirements of individual audience members; and the performers receive live feedback from the audience.

There are two main aspects: i) to emulate as closely as possible the experience of live performances where the performers and audience are co-located in the same physical space, such as a theatre; and ii) to enhance traditional physical performances with features such as personalisation of the experience according to the preferences or needs of the audience members.

Examples of personalisation include:

- * viewpoint selection such as choosing a specific seat in the theatre or for more advanced positioning of the audience member's viewpoint outside of the traditional seating - amongst, above or behind the performers (but within some limits which may be imposed by the performers or the director for artistic reasons);
- * augmentation of the performance with subtitles, audio-description, actor-tagging, language translation, advertisements/product-placement, other enhancements/filters to make the performance accessible to disabled audience members (removal of flashing images for epileptics, alternative colour schemes for colour-blind audience members, etc.).

3.3.2. Characterization

There are several chained functional entities which are candidates for being deployed as (COIN) Programs.

- * Performer aggregation and editing functions

- * Distribution and encoding functions
- * Personalisation functions
 - to select which of the existing streams should be forwarded to the audience member
 - to augment streams with additional metadata such as subtitles
 - to create new streams after processing existing ones: to interpolate between camera angles to create a new viewpoint or to render point clouds from the audience member's chosen perspective
 - to undertake remote rendering according to viewer position, e.g. creation of VR headset display streams according to audience head position – when this processing has been offloaded from the viewer's end-system to the in-network function due to limited processing power in the end-system, or to limited network bandwidth to receive all of the individual streams to be processed.
- * Audience feedback sensor processing functions
- * Audience feedback aggregation functions

These are candidates for deployment as (COIN) Programs in PNDs rather than being located in end-systems (at the performers' site, the audience members' premises or in a central cloud location) for several reasons:

- * personalisation of the performance according to audience preferences and requirements makes it unfeasible to be done in a centralised manner at the performer premises: the computational resources and network bandwidth would need to scale with the number of audience members' personalised streams.
- * rendering of VR headset content to follow viewer head movements has an upper bound on lag to maintain viewer QoE, which requires the processing to be undertaken sufficiently close to the viewer to avoid large network latencies.
- * viewer devices may not have the processing-power to undertake the personalisation or the viewers' network may not have the capacity to receive all of the constituent streams to undertake the personalisation functions.

- * there are strict latency requirements for live and interactive aspects that require the deviation from the direct network path from performers to audience to be minimised, which reduces the opportunity to route streams via large-scale processing capabilities at centralised data-centres.

3.3.3. Existing solutions

Note: Existing solutions for some aspects of this use case are covered in the Mobile Application Offloading, Extended Reality, and Content Delivery Networks use cases.

3.3.4. Opportunities

- * Executing media processing and personalisation functions on-path as (COIN) Programs in PNDs will avoid detour/stretch to central servers which increases latency as well as the consumption of bandwidth on more network resources (links and routers). For example, in this use case the chain of (COIN) Programs and propagation over the interconnecting network segments for performance capture, aggregation, distribution, personalisation, consumption, capture of audience response, feedback processing, aggregation, rendering should be achieved within an upper bound of latency (the tolerable amount is to be defined, but in the order of 100s of ms to mimic performers perceiving audience feedback, such as laughter or other emotional responses in a theatre setting).
- * Processing of media streams allows (COIN) Programs, PNDs and the wider (COIN) System/Environment to be contextual aware of flows and their requirements which can be used for determining network treatment of the flows, e.g. path selection, prioritisation, multi-flow coordination, synchronisation & resilience.

3.3.5. Research Questions:

- * RQ 3.3.1: In which PNDs should (Coin) Programs for aggregation, encoding and personalisation functions be located? Close to the performers or close to the audience members?
- * RQ 3.3.2: How far from the direct network path from performer to audience should (COIN) programs be located, considering the latency implications of path-stretch and the availability of processing capacity at PNDs? How should tolerances be defined by users?

- * RQ 3.3.3: Should users decide which PNDs should be used for executing (COIN) Programs for their flows or should they express requirements and constraints that will direct decisions by the orchestrator/manager of the COIN System?
- * RQ 3.3.4: How to achieve network synchronisation across multiple streams to allow for merging, audio-video interpolation and other cross-stream processing functions that require time synchronisation for the integrity of the output? How can this be achieved considering that synchronisation may be required between flows that are: i) on the same data pathway through a PND/router, ii) arriving/leaving through different ingress/egress interfaces of the same PND/router, iii) routed through disjoint paths through different PNDs/routers?
- * RQ 3.3.5: Where will COIN Programs will be executed? In the data-plane of PNDs, in other on-router computational capabilities within PNDs, or in adjacent computational nodes?
- * RQ 3.3.6: Are computationally-intensive tasks - such as video stitching or media recognition and annotation - considered as suitable candidate (COIN) Programs or should they be implemented in end-systems?
- * RQ 3.3.7: If the execution of COIN Programs is offloaded to computational nodes outside of PNDs, e.g. for processing by GPUs, should this still be considered as in-network processing? Where is the boundary between in-network processing capabilities and explicit routing of flows to endsystems?

3.3.6. Requirements

- * Req 3.3.1: Users should be able to specify requirements on network and processing metrics (such as latency and throughput bounds) and the COIN System should be able to respect those requirements and constraints when routing flows and selecting PNDs for executing (COIN) Programs.
- * Req 3.3.2: A COIN System should be able to synchronise flow treatment and processing across multiple related flows which may be on disjoint paths.

4. Supporting new COIN Systems

While the best-effort nature of the Internet enables a wide variety of applications, there are several domains whose requirements are hard to satisfy over regular best-effort networks. Consequently, there is a large number of specialized appliances and protocols designed to provide 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. However, solutions on the link layer alone are not always sufficient.

The industrial domain, e.g., currently evolves towards increasingly interconnected systems in turn increasing the complexity of the underlying networks, making them more dynamic, and creating more diverse sets of requirements. Concepts satisfying the dynamic performance requirements of modern industrial applications thus become harder to develop. In this context, COIN offers new possibilities as it allows to flexibly distribute computation tasks across the network and enables novel forms of interaction between communication and computation providers.

This document illustrates the potential for new COIN systems using the example of the industrial domain by characterizing and analyzing specific scenarios to showcase potential requirements, as specifying general requirements is difficult due to the domain's mentioned diversity.

4.1. Industrial Network Scenario

Common components of industrial networks can be divided into three categories as illustrated in Figure 2. 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.

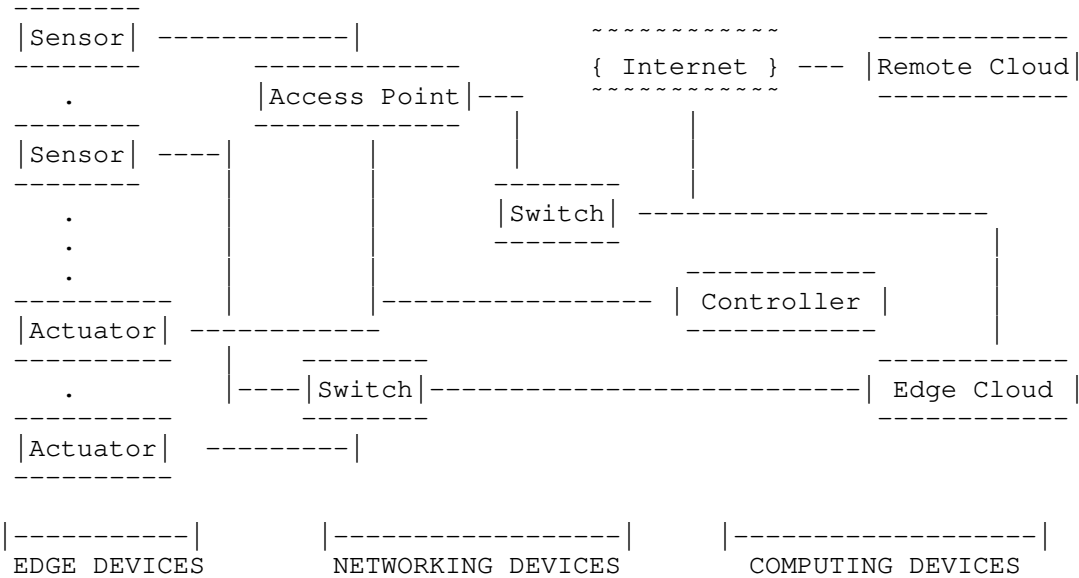


Figure 2: Industrial networks show a high level of heterogeneity.

4.2. In-Network Control / Time-sensitive applications

4.2.1. Description

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. Moving compute functionality onto COIN execution environments inside the network offers a new solution space to these challenges.

4.2.2. Characterization

A control process consists of two main components as illustrated in Figure 3: 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.

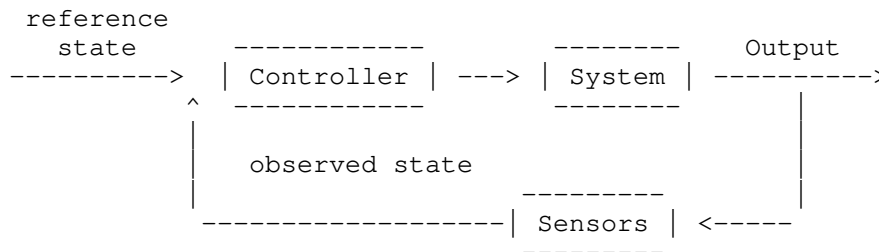


Figure 3: Simple feedback control model.

Apart from the control model, the quality of the control primarily depends on the timely reception of the sensor feedback which 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, e.g., off-premise cloud platforms are included.

4.2.3. Existing Solutions

Control functionality is traditionally executed on PLCs close to the machinery. These PLCs typically require vendor-specific implementations and are often hard to upgrade and update which makes such control processes inflexible and difficult to manage. Moving computations to more freely programmable devices thus has the potential of significantly improving the flexibility. In this context, directly moving control functionality to (central) cloud environments is generally possible, yet only feasible if latency constraints are lenient.

4.2.4. Opportunities for COIN

COIN offers the possibility of bringing the system and the controller closer together, thus possibly satisfying the latency requirements, by performing simple control logic on PNDs and/or in COIN execution environments. While control models, in general, can become involved, 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.

Opportunities:

- * Execute simple (end-host) COIN functions on PNDs to satisfy tight latency constraints of control processes

4.2.5. Research Questions for COIN

Bringing the required computations to PNDs is challenging as these 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 [KUNZE]. 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:

Research Questions:

- * RQ 4.2.1: How to derive simplified versions of the global (control) function?
 - How to account for the limited computational precision of PNDs?
 - How to find suitable tradeoffs regarding simplicity of the control function ("accuracy of the control") and implementation complexity ("implementability")?
- * RQ 4.2.2: How to distribute the simplified versions in the network?

- Can there be different control levels, e.g., "quite inaccurate & very low latency" (PNDs, deep in the network), "more accurate & higher latency" (more powerful COIN execution environments, farer away), "very accurate & very high latency" (cloud environments, far away)?
- Who decides which control instance is executed and how?
- How do the different control instances interact?

4.2.6. Requirements

Req 4.2.1: The interaction between the COIN execution environments and the global controller SHOULD be explicit.

Req 4.2.2: The interaction between the COIN execution environments and the global controller MUST NOT negatively impact the control quality.

Req 4.2.3: Actions of the COIN execution environments MUST be overridable by the global controller.

Req 4.2.4: Functions in COIN execution environments SHOULD be executed with predictable delay.

Req 4.2.5: Functions in COIN execution environments MUST be executed with predictable accuracy.

4.3. Large Volume Applications - Filtering

4.3.1. Description

In modern industrial networks, processes and machines can be monitored closely resulting in large volumes of available information. This data can be used to find previously unknown correlations between different parts of the value chain, e.g., by deploying machine learning (ML) techniques, 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. Off-premise cloud platforms offer cost-effective solutions with a high degree of flexibility and scalability, however, moving all data to off-premise locations poses infrastructural challenges. Pre-processing or filtering the data already in COIN execution environments can be a new solution to this challenge.

4.3.2. Characterization

4.3.2.1. General Characterization of Large Volume Applications

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 COIN techniques hardly applicable as they typically work on unencrypted data. Adding security to COIN 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 8.

4.3.2.2. Specific Characterization for Filtering Solutions

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.

4.3.3. Existing Solutions

Current approaches for handling such large amounts of information typically build upon stream processing frameworks such as Apache Flink. While they allow for handling large volume applications, they are tied to performant server machines and upscaling the information density also requires a corresponding upscaling of the compute infrastructure.

4.3.4. Opportunities for COIN

PNDs and COIN execution environments are in a unique position to reduce the data rates due to their line-rate packet processing capabilities. Using these capabilities, it is possible 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 could 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. While the former variant is oblivious to the semantics of the sensor data, the latter variant requires an understanding of the current sensor levels. In any case, 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.

Opportunities:

- * (Semantic) packet filtering based on packet header and payload, as well as multi-packet information

4.3.5. Research Questions for COIN

- * RQ 4.3.1: How to design COIN programs for (semantic) packet filtering?

- Which criteria for filtering make sense?
- * RQ 4.3.2: How to distribute and coordinate COIN programs?
- * RQ 4.3.3: How to dynamically change COIN programs?
- * RQ 4.3.4: How to signal traffic filtering by COIN programs to end-hosts?

4.3.6. Requirements

Req 4.3.1: Filters MUST conform to application-level syntax and semantics.

Req 4.3.2: Filters MAY leverage packet header and payload information.

Req 4.3.3: Filters SHOULD be reconfigurable at run-time.

4.4. Large Volume Applications - (Pre-)Preprocessing

4.4.1. Description

See Section 4.3.1.

4.4.2. Characterization

4.4.2.1. General Characterization of Large Volume Applications

See Section 4.3.2.1.

4.4.2.2. Specific Characterization for Preprocessing Solutions

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.

4.4.3. Existing Solutions

See Section 4.3.3.

4.4.4. Opportunities for COIN

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 PNDs is also interesting, because sensor data is usually streamed.

Opportunities:

- * (Semantic) data (pre-)processing, e.g., in the form of computations across multiple packets and potentially leveraging packet payload

4.4.5. Research Questions for COIN

- * RQ 4.4.1: Which kinds of COIN programs can be leveraged for (pre-)processing steps?
 - How complex can they become?
- * RQ 4.4.2: How to distribute and coordinate COIN programs?
- * RQ 4.4.3: How to dynamically change COIN programs?
- * RQ 4.4.4: How to incorporate the (pre-)processing steps into the overall system?

4.4.6. Requirements

Req 4.4.1: Preprocessors **MUST** conform to application-level syntax and semantics.

Req 4.4.2: Preprocessors **MAY** leverage packet header and payload information.

Req 4.4.3: Preprocessors **SHOULD** be reconfigurable at run-time.

4.5. Industrial Safety

4.5.1. Description

Despite an 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.

4.5.2. Characterization

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.

4.5.3. Existing Solutions

Due to the importance of safety, there is a wide range of software-based approaches aiming at enhancing security. One example are tag-based systems, e.g., using RFID, where drivers of forklifts can be warned if pedestrian workers carrying tags are nearby. Such solutions, however, require setting up an additional system and do not leverage existing sensor data.

4.5.4. Opportunities for COIN

COIN systems could leverage the increased availability of sensor data and the detailed monitoring of the factories to enable additional safety measures. Different safety indicators within the production hall can be combined within the network so that PNDs can give early responses if a potential safety breach is detected.

One 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.

Opportunities:

- * Execute simple (end-host) COIN functions on PNDs to create early emergency reactions based on diverse sensor feedback

4.5.5. Research Questions for COIN

- * RQ 4.5.1: Which additional safety measures can be provided?
 - Do these measures actually improve safety?
- * RQ 4.5.2: Which sensor information can be combined and how?

4.5.6. Requirements

Req 4.5.1: COIN-based safety measures **MUST NOT** degrade existing safety measures.

Req 4.5.2: COIN-based safety measures **MAY** enhance existing safety measures.

5. Improving existing COIN capabilities

5.1. Content Delivery Networks

5.1.1. Description

Delivery of content to end users often relies on Content Delivery Networks (CDNs) storing said content closer to end users for latency reduced delivery with DNS-based indirection being utilized to serve the request on behalf of the origin server.

5.1.2. Characterization

From the perspective of this draft, a CDN can be interpreted as a (network service level) set of (COIN) programs, implementing a distributed logic for distributing content from the origin server to the CDN ingress and further to the CDN replication points which ultimately serve the user-facing content requests.

5.1.3. Existing Solutions

NOTE: material on solutions will be added here later

Studies such as those in [FCDN] have shown that content distribution at the level of named content, utilizing efficient (e.g., Layer 2) multicast for replication towards edge CDN nodes, can significantly increase the overall network and server efficiency. It also reduces indirection latency for content retrieval as well as reduces required edge storage capacity by benefiting from the increased network efficiency to renew edge content more quickly against changing demand.

5.1.4. Opportunities and Research Questions for COIN

Opportunities:

- * The support for service-level routing of requests (service routing in [APPCENTRES]) to specific (COIN) program instances may improve on end user experience in faster retrieving (possibly also more, e.g., better quality) content.
- * Supporting the constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) may improve the overall end user experience by selecting a 'more suitable' (COIN) program instance over another, e.g., avoiding/reducing overload situation in specific (COIN) program instances.
- * Supporting Layer 2 capabilities for multicast (compute interconnection and collective communication in [APPCENTRES]) may increase the network utilization and therefore increase the overall system utilization.

Research Questions: in addition to those request question for Section 3.1,

- * RQ 5.1.1: How to utilize L2 multicast to improve on CDN designs? How to utilize in-network capabilities in those designs?

- * RQ 5.1.2: What forwarding methods may support the required multicast capabilities (see [FCDN])
- * RQ 5.1.3: What are the right routing constraints that reflect both compute and network capabilities?
- * RQ 5.1.4: Could traffic steering be performed at the data path and per service request? If so, what would be performance improvements?
- * RQ 5.1.5: How could storage be traded off against frequent, multicast-based, replication (see [FCDN])?
- * RQ 5.1.6: What scalability limits exist for L2 multicast capabilities? How to overcome them?

5.1.5. Requirements

Requirements 3.1.1 through 3.1.6 also apply for CDN service access. In addition:

Req 5.1.1: Any solution SHOULD utilize Layer 2 multicast transmission capabilities for responses to concurrent service requests.

5.2. Compute-Fabric-as-a-Service (CFaaS)

5.2.1. Description

Layer 2 connected compute resources, e.g., in regional or edge data centres, base stations and even end-user devices, provide the opportunity for infrastructure providers to offer CFaaS type of offerings to application providers. App and service providers may utilize the compute fabric exposed by this CFaaS offering for the purposes defined through their applications and services. In other words, the compute resources can be utilized to execute the desired (COIN) programs of which the application is composed, while utilizing the inter-connection between those compute resources to do so in a distributed manner.

5.2.2. Characterization

We foresee those CFaaS offerings to be tenant-specific, a tenant here defined as the provider of at least one application. For this, we foresee an interaction between CFaaS provider and tenant to dynamically select the appropriate resources to define the demand side of the fabric. Conversely, we also foresee the supply side of the fabric to be highly dynamic with resources being offered to the fabric through, e.g., user-provided resources (whose supply might

depend on highly context-specific supply policies) or infrastructure resources of intermittent availability such as those provided through road-side infrastructure in vehicular scenarios.

The resulting dynamic demand-supply matching establishes a dynamic nature of the compute fabric that in turn requires trust relationships to be built dynamically between the resource provider(s) and the CFaaS provider. This also requires the communication resources to be dynamically adjusted to interconnect all resources suitably into the (tenant-specific) fabric exposed as CFaaS.

5.2.3. Existing Solutions

NOTE: material on solutions will be added here later

5.2.4. Opportunities and Research Questions for COIN

Opportunities:

- * Supporting service-level routing of compute resource requests (service routing in [APPCENTRES]) may allow for utilizing the wealth of compute resources in the overall CFaaS fabric for execution of distributed applications, where the distributed constituents of those applications are realized as (COIN) programs and executed within a COIN system as (COIN) program instances.
- * Supporting the constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) will allow for optimizing both the CFaaS provider constraints as well as tenant-specific constraints.
- * Supporting Layer 2 capabilities for multicast (compute interconnection and collective communication in [APPCENTRES]) will allow for increasing both network utilization but also possible compute utilization (due to avoiding unicast replication at those compute endpoints), thereby decreasing total cost of ownership for the CFaaS offering.

Research Questions: similar to those for Section 3.1, in addition

- * RQ 5.2.1: How to convey tenant-specific requirements for the creation of the L2 fabric?
- * RQ 5.2.2: How to dynamically integrate resources, particularly when driven by tenant-level requirements and changing service-specific constraints?

- * RQ 5.2.3: How to utilize in-network capabilities to aid the availability and accountability of resources, i.e., what may be (COIN) programs for a CFaaS environment that in turn would utilize the distributed execution capability of a COIN system?

5.2.5. Requirements

For the provisioning of services atop the CFaaS, requirements 3.1.1 through 3.1.6 should be addressed, too. In addition:

Req 5.2.1: Any solution SHOULD expose means to specify the requirements for the tenant-specific compute fabric being utilized for the service execution.

Req 5.2.2: Any solution SHOULD allow for dynamic integration of compute resources into the compute fabric being utilized for the app execution; those resources include, but are not limited to, end user provided resources. From a COIN system perspective, new resources must be possible to be exposed as possible (COIN) execution environments.

Req 5.2.3: Any solution MUST provide means to optimize the inter-connection of compute resources, including those dynamically added and removed during the provisioning of the tenant-specific compute fabric.

Req 5.2.4: Any solution MUST provide means for ensuring availability and usage of resources is accounted for.

5.3. Virtual Networks Programming

5.3.1. Description

The term "virtual network programming" is proposed to describe mechanisms by which tenants deploy and operate COIN programs in their virtual network. Such COIN programs can for example be P4 programs, OpenFlow rules, or higher layer programs. This feature can enable other use cases described in this draft to be deployed using virtual networks services, over underlying networks such as datacenters, mobile networks, or other fixed or wireless networks.

For example COIN programs could perform the following on a tenant's virtual network:

- * Allow or block flows, and request rules from an SDN controller for each new flow, or for flows to or from specific hosts that needs enhanced security

- * Forward a copy of some flows towards a node for storage and analysis
- * Update counters based on specific sources/destinations or protocols, for detailed analytics
- * Associate traffic between specific endpoints, using specific protocols, or originated from a given application, to a given slice, while other traffic use a default slice
- * Experiment with a new routing protocol (e.g., ICN), using a P4 implementation of a router for this protocol

5.3.2. Characterization

To provide a concrete example of virtual COIN programming, we consider a use case using a 5G underlying network, the 5GLAN virtualization technology, and the P4 programming language and environment. Section 5.1 of [I-D.ravi-icnrg-5gc-icn] provides a description of the 5G network functions and interfaces relevant to 5GLAN, which are otherwise specified in [TS23.501] and [TS23.502]. From the 5GLAN service customer/tenant standpoint, the 5G network operates as a switch.

In the use case depicted in Figure 4, the tenant operates a network including a 5GLAN network segment (seen as a single logical switch), as well as fixed segments. This can be in a plant or enterprise network, using for an example a 5G Non-Public Network (NPN). The tenant uses P4 programs to determine the operation of the fixed and 5GLAN switches. The tenant provisions a 5GLAN P4 program into the mobile network, and can also operate a controller. The mobile devices (or User Equipment nodes) UE1, UE2, UE3 and UE4 are in the same 5GLAN, as well as Device1 and Device2 (through UE4).

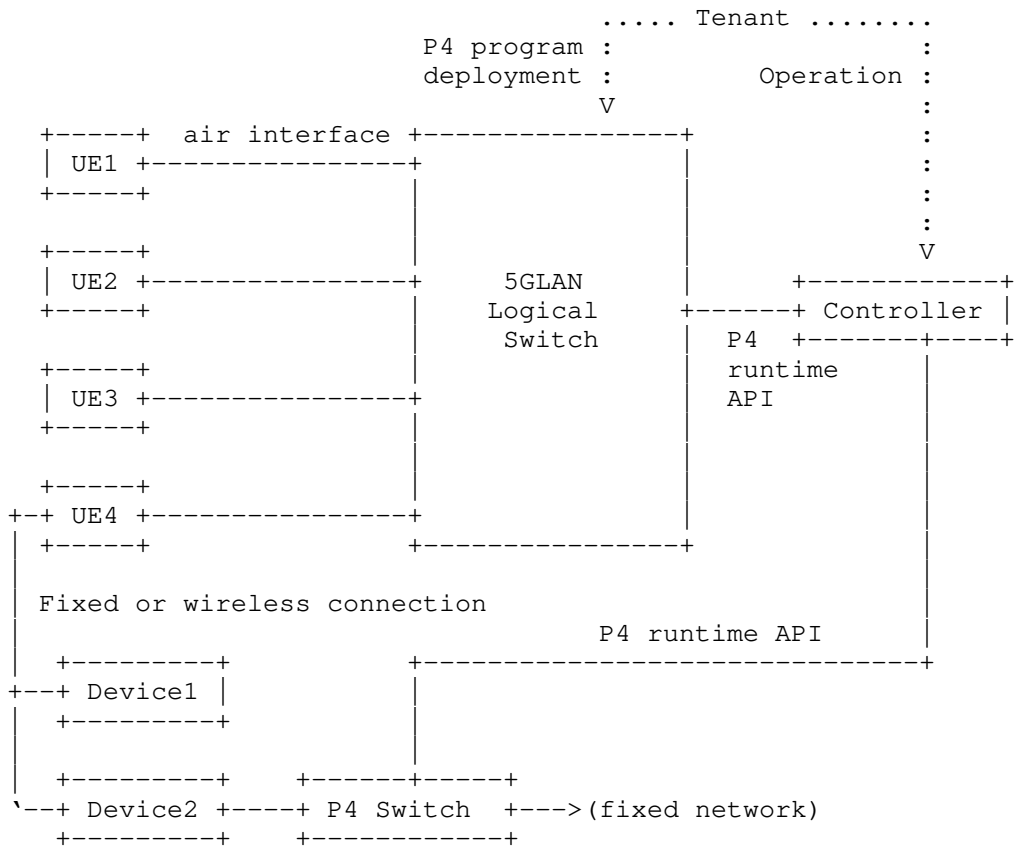


Figure 4: 5G Virtual Network Programming Overview

Looking in more details in Figure 5, the 5GLAN P4 program can be split between multiple data plane nodes (PDU Session Anchor (PSA) User Plane Functions (UPF), other UPFs, or even mobile devices), although in some cases the P4 program may be hosted on a single node. In the most general case, a distributed deployment is useful to keep traffic on optimal paths, because, except in simple cases, within a 5GLAN all traffic will not pass through a single node. In this example, P4 programs could be deployed in UPF1, UPF2, UPF3, UE3 and UE4. UE1-UE2 traffic is using a local switch on PSA UPF1, UE1-UE3 traffic is tunneled between PSA UPF1 and PSA UPF2 through the N19 interface, and UE1-UE4 traffic is forwarded through an external Data Network (DN). Traffic between Device1 and Device2 is forwarded through UE4.

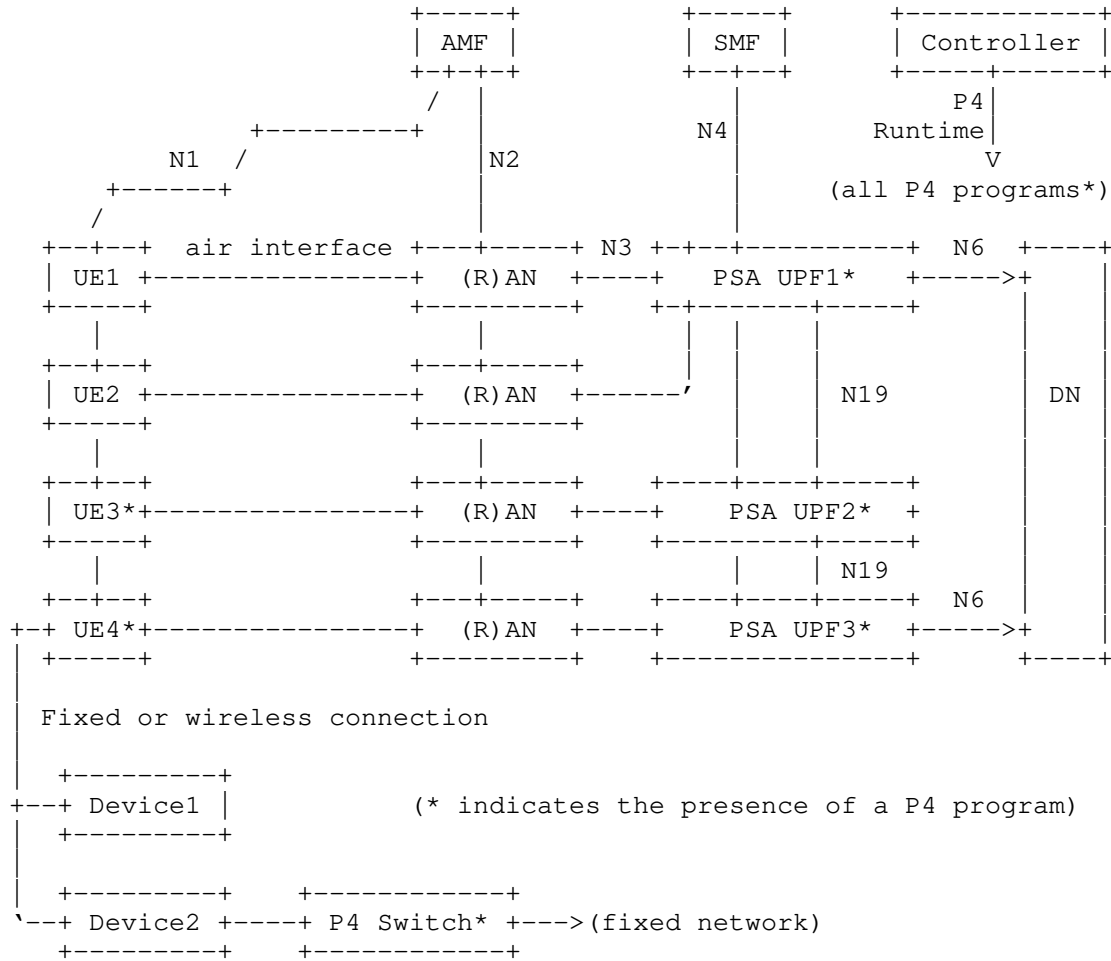


Figure 5: 5G Virtual Network Programming Details

5.3.3. Existing Solutions

Research has been conducted, for example by [Stoyanov], to enable P4 network programming of individual virtual switches. To our knowledge, no complete solution has been developed for deploying virtual COIN programs over mobile or datacenter networks.

5.3.4. Opportunities

Virtual network programming by tenants could bring benefits such as:

- * A unified programming model, which can facilitate porting in-network computing between data centers, 5G networks, and other fixed and wireless networks, as well as sharing controller, code and expertise.
- * Increasing the level of customization available to customers/tenants of mobile networks or datacenters, when compared with typical configuration capabilities. For example, 5G network evolution points to an ever increasing specialization and customization of private mobile networks, which could be handled by tenants using a programming model similar to P4.
- * Using network programs to influence underlying network service (e.g., request specific QoS for some flows in 5G or datacenters), to increases the level of in-depth customization available to tenants.

5.3.5. Research Questions for COIN

- * RQ 5.3.1: Underlying Network Awareness: a virtual COIN program can be able to influence, and be influenced by, the underling network (e.g., the 5G network or data center). For example, a virtual COIN program may be aware of the slice used by a flow, and possibly influence slice selection. Since some information and actions may be available on some nodes and not others, underlying network awareness may impose additional constraints on distributed network programs location.
- * RQ 5.3.2: Splitting/Distribution: a virtual COIN program may need to be deployed across multiple computing nodes, leading to research questions around instance placement and distribution. As a primary reason for this, program logic should be applied exactly once or at least once per packet, while allowing optimal forwarding path by the underlying network. For example, a 5GLAN P4 program may need to run on multiple UPFs. Research challenges include defining manual (by the programmer) or automatic methods to distribute COIN programs that use a low or minimal amount of resources. Distributed P4 programs are studied in [I-D.hsingh-coinrg-reqs-p4comp] and [Sultana].
- * RQ 5.3.3: Multi-Tenancy Support: multiple virtual COIN program instances can run on the same compute node. While mechanism were proposed for P4 multi-tenancy in a switch [Stoyanov], research questions remains, about isolation between tenants, fair repartition of resources.

- * RQ 5.3.4: Security: how can tenants and underlying networks be protected against security risks, including overuse or misuse of network resources, injection of traffic, access to unauthorized traffic?
- * RQ 5.3.5: Higher layer processing: can a virtual network model facilitate the deployment of COIN programs acting on application layer data? This is an open question since the present section focused on packet/flow processing.

5.3.6. Requirements

- * Req 5.3.1: A COIN system supporting virtualization should enable tenants to deploy COIN programs onto their virtual networks.
- * Req 5.3.2: A virtual COIN program should process flows/packets once and only once (or at least once for idempotent operations), even if the program is distributed over multiple PNDs.
- * Req 5.3.3: Multi-tenancy should be supported for virtual COIN programs, i.e., instances of virtual COIN programs from different tenants can share underlying PNDs. This includes requirements for secure isolation between tenants, and fair (or policy-based) sharing of computing resources.
- * Req 5.3.4: Virtual COIN programs should support mobility of endpoints.

6. Enabling new COIN capabilities

6.1. Distributed AI

6.1.1. Description

There is a growing range of use cases demanding for the realization of AI capabilities among distributed endpoints. Such demand may be driven by the need to increase overall computational power for large-scale problems. From a COIN perspective, those capabilities may be realized as (COIN) programs and executed throughout the COIN system, including in PNDs.

Some solutions may desire the localization of reasoning logic, e.g., for deriving attributes that better preserve privacy of the utilized raw input data. Quickly establishing (COIN) program instances in nearby compute resources, including PNDs, may even satisfy such localization demands on-the-fly (e.g., when a particular use is being realized, then terminated after a given time).

6.1.2. Characterization

Examples for large-scale AI problems include biotechnology and astronomy related reasoning over massive amounts of observational input data. Examples for localizing input data for privacy reasons include radar-like application for the development of topological mapping data based on (distributed) radio measurements at base stations (and possibly end devices), while the processing within radio access networks (RAN) already constitute a distributed AI problem to a certain extent albeit with little flexibility in distributing the execution of the AI logic.

6.1.3. Existing Solutions

Reasoning frameworks, such as TensorFlow, may be utilized for the realization of the (distributed) AI logic, building on remote service invocation through protocols such as gRPC [GRPC] or MPI [MPI] with the intention of providing an on-chip NPU (neural processor unit) like abstraction to the AI framework.

NOTE: material on solutions like ETSI MEC and 3GPP work will be added here later

6.1.4. Opportunities and Research Questions for COIN

Opportunities:

- * Supporting service-level routing of requests (service routing in [APPCENTRES]), with AI services being exposed to the network and executed as part of (COIN) programs in selected (COIN) program instances, may provide a highly distributed execution of the overall AI logic, thereby addressing, e.g., localization but also computational concerns (scale-in/out).
- * The support for constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) may allow for utilizing the most suitable HW capabilities (e.g., support for specific AI HW assistance in the COIN element, including a PND), while also allowing to select resources, e.g., based on available compute ability such as number of cores to be used.
- * Supporting collective communication between multiple instances of AI services, i.e., (COIN) program instances, may positively impact network but also compute utilization by moving from unicast replication to network-assisted multicast operation.

Research Questions:

- * RQ 6.1.1: similar to use case in Section 3.1
- * RQ 6.1.2: What are the communication patterns that may be supported by collective communication solutions?
- * RQ 6.1.3: How to achieve scalable multicast delivery with rapidly changing receiver sets?
- * RQ 6.1.4: What in-network capabilities may support the collective communication patterns found in distributed AI problems?
- * RQ 6.1.5: How to provide a service routing capability that supports any invocation protocol (beyond HTTP)?

6.1.5. Requirements

Requirements 3.1.1 through 3.1.6 also apply for general distributed AI capabilities. In addition:

Req 6.1.1: Any COIN system MUST provide means to specify the constraints for placing (AI) execution logic in the form of (COIN) programs in certain logical execution points (and their associated physical locations), including PNDs.

Req 6.1.2: Any COIN system MUST provide support for app/micro-service specific invocation protocols for requesting (COIN) program services exposed to the COIN system.

7. Analysis

The goal of this analysis is to identify aspects that are relevant across all use cases to help in shaping the research agenda of COINRG. For this purpose, this section will condense the opportunities, research questions, as well as requirements of the different presented use cases and analyze these for similarities across the use cases.

Through this, we intend to identify cross-cutting opportunities, research questions as well as requirements (for COIN system solutions) that may aid the future work of COINRG as well as the larger research community.

8. 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 4.3.2, 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].

9. IANA Considerations

N/A

10. Conclusion

There are several domains that can profit from capabilities that are provided by in-network and generally distributed compute capabilities. In this draft, we differentiated use cases in which COIN capabilities may enable new experiences, while others may expose new or improve on existing system capabilities, while yet other use cases may see COIN capabilities enable new environments, such as in industrial networking.

Beyond the mere description and characterization of those use cases, we identified opportunities arising from utilizing COIN capabilities as well as research questions that may need to be addressed to reap those opportunities. Lastly, we also outlined possible requirements for realizing a COIN system at some point.

Our analysis across all use cases in those dimensions of opportunities, research questions and requirements targets the support for future work in this space and is therefore directly positioned as input into the initial milestones of the COIN RG.

11. List of Use Case Contributors

- * Dirk Trossen has contributed the following use cases: Section 3.1, Section 5.1, Section 5.2, Section 6.1.

- * Marie-Jose Montpetit has contributed the XR use case (Section 3.2).
- * David Griffin and Miguel Rio have contributed the use case on performing arts (Section 3.3).
- * Ike Kunze and Klaus Wehrle have contributed the industrial use cases (Section 4).
- * Xavier De Foy has contributed the use case on virtual networks programming (Section 5.3)

12. References

12.1. 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>>.

12.2. Informative References

- [APPCENTRES] Trossen, D., Sarathchandra, C., and M. Boniface, "In-Network Computing for App-Centric Micro-Services", Work in Progress, Internet-Draft, draft-sarathchandra-coin-appcentres-04, 26 January 2021, <<https://www.ietf.org/internet-drafts/draft-sarathchandra-coin-appcentres-04.txt>>.
- [FCDN] Al-Naday, M., Reed, M.J., Riihijarvi, J., Trossen, D., Thomos, N., and M. Al-Khalidi, "A Flexible and Efficient CDN Infrastructure without DNS Redirection of Content Reflection", <<https://arxiv.org/pdf/1803.00876.pdf>>.
- [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>>.
- [GRPC] "High performance open source universal RPC framework", <<https://grpc.io/>>.

- [I-D.draft-kutscher-coinrg-dir]
Kutscher, D., Kaerkkainen, T., and J. Ott, "Directions for Computing in the Network", Work in Progress, Internet-Draft, draft-kutscher-coinrg-dir-02, 31 July 2020, <<https://www.ietf.org/archive/id/draft-kutscher-coinrg-dir-02.txt>>.
- [I-D.fink-coin-sec-priv]
Fink, I. B. and K. Wehrle, "Enhancing Security and Privacy with In-Network Computing", Work in Progress, Internet-Draft, draft-fink-coin-sec-priv-03, 22 October 2021, <<https://www.ietf.org/archive/id/draft-fink-coin-sec-priv-03.txt>>.
- [I-D.hsingh-coinrg-reqs-p4comp]
Singh, H. and M. Montpetit, "Requirements for P4 Program Splitting for Heterogeneous Network Nodes", Work in Progress, Internet-Draft, draft-hsingh-coinrg-reqs-p4comp-03, 18 February 2021, <<https://www.ietf.org/archive/id/draft-hsingh-coinrg-reqs-p4comp-03.txt>>.
- [I-D.mcbride-edge-data-discovery-overview]
McBride, M., Kutscher, D., Schooler, E., Bernardos, C. J., Lopez, D. R., and X. D. Foy, "Edge Data Discovery for COIN", Work in Progress, Internet-Draft, draft-mcbride-edge-data-discovery-overview-05, 1 November 2020, <<https://www.ietf.org/archive/id/draft-mcbride-edge-data-discovery-overview-05.txt>>.
- [I-D.ravi-icnrg-5gc-icn]
Ravindran, R., Suthar, P., Trossen, D., Wang, C., and G. White, "Enabling ICN in 3GPP's 5G NextGen Core Architecture", Work in Progress, Internet-Draft, draft-ravi-icnrg-5gc-icn-04, 31 May 2019, <<https://www.ietf.org/archive/id/draft-ravi-icnrg-5gc-icn-04.txt>>.
- [ICE]
Burke, J., "ICN-Enabled Secure Edge Networking with Augmented Reality: ICE-AR.", ICE-AR Presentation at NDNCOM. , 2018, <<https://www.nist.gov/news-events/events/2018/09/named-data-networking-community-meeting-2018>>.

- [KUNZE] Kunze, I., Glebke, R., Scheiper, J., Bodenbenner, M., Schmitt, R., and K. Wehrle, "Investigating the Applicability of In-Network Computing to Industrial Scenarios", 2021 4th IEEE International Conference on Industrial Cyber-Physical Systems (ICPS), DOI 10.1109/icps49255.2021.9468247, May 2021, <<https://doi.org/10.1109/icps49255.2021.9468247>>.
- [MPI] Vishnu, A., Siegel, C., and J. Daily, "Scaling Distributed Machine Learning with In-Network Aggregation", <<https://arxiv.org/pdf/1603.02339.pdf>>.
- [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, DOI 10.1145/3229591.3229592, August 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>>.
- [Stoyanov] Stoyanov, R. and N. Zilberman, "MTPSA: Multi-Tenant Programmable Switches", ACM P4 Workshop in Europe (EuroP4'20) , 2020, <<https://eng.ox.ac.uk/media/6354/stoyanov2020mtpsa.pdf>>.
- [Sultana] Sultana, N., Sonchack, J., Giesen, H., Pedisich, I., Han, Z., Shyamkumar, N., Burad, S., DeHon, A., and B.T. Loo, "Flightplan: Dataplane Disaggregation and Placement for P4 Programs", 2020, <<https://flightplan.cis.upenn.edu/flightplan.pdf>>.
- [TS23.501] 501, 3gpp-23., "Technical Specification Group Services and System Aspects; System Architecture for the 5G System; Stage 2 (Rel.17)", 3GPP , 2021, <<https://www.3gpp.org/DynaReport/23501.htm>>.

- [TS23.502] 502, 3gpp-23., "Technical Specification Group Services and System Aspects; Procedures for the 5G System; Stage 2 (Rel.17)", 3GPP , 2021,
<<https://www.3gpp.org/DynaReport/23502.htm>>.
- [TSN] "IEEE Time-Sensitive Networking (TSN) Task Group",
<<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

Marie-Jose Montpetit
Concordia University
Montreal
Canada

Email: marie@mjmontpetit.com

Xavier de Foy
InterDigital Communications, LLC
1000 Sherbrooke West
Montreal H3A 3G4
Canada

Email: xavier.defoy@interdigital.com

David Griffin
University College London
Gower St
London
WC1E 6BT
United Kingdom

Email: d.griffin@ucl.ac.uk

Miguel Rio
University College London
Gower St
London
WC1E 6BT
United Kingdom

Email: miguel.rio@ucl.ac.uk

COINRG
Internet-Draft
Intended status: Informational
Expires: 8 September 2022

I. Kunze
K. Wehrle
RWTH Aachen
D. Trossen
Huawei
M.J. Montpetit
Concordia
X. de Foy
InterDigital Communications, LLC
D. Griffin
M. Rio
UCL
7 March 2022

Use Cases for In-Network Computing
draft-irtf-coinrg-use-cases-02

Abstract

Computing in the Network (COIN) comes with the prospect of deploying processing 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 8 September 2022.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	4
2. Terminology	6
3. Providing New COIN Experiences	6
3.1. Mobile Application Offloading	6
3.1.1. Description	6
3.1.2. Characterization	7
3.1.3. Existing Solutions	9
3.1.4. Opportunities	9
3.1.5. Research Questions	9
3.1.6. Requirements	10
3.2. Extended Reality and Immersive Media	11
3.2.1. Description	11
3.2.2. Characterization	11
3.2.3. Existing Solutions	12
3.2.4. Opportunities	13
3.2.5. Research Questions	13
3.2.6. Requirements	14
3.3. Personalised and interactive performing arts	14
3.3.1. Description	15
3.3.2. Characterization	15
3.3.3. Existing solutions	17
3.3.4. Opportunities	17
3.3.5. Research Questions:	17
3.3.6. Requirements	18
4. Supporting new COIN Systems	18
4.1. Industrial Network Scenario	19
4.2. In-Network Control / Time-sensitive applications	20
4.2.1. Description	20
4.2.2. Characterization	21
4.2.3. Existing Solutions	21
4.2.4. Opportunities	22
4.2.5. Research Questions	22

4.2.6.	Requirements	23
4.3.	Large Volume Applications - Filtering	23
4.3.1.	Description	23
4.3.2.	Characterization	24
4.3.3.	Existing Solutions	25
4.3.4.	Opportunities	25
4.3.5.	Research Questions	26
4.3.6.	Requirements	26
4.4.	Large Volume Applications - (Pre-)Preprocessing	26
4.4.1.	Description	26
4.4.2.	Characterization	26
4.4.3.	Existing Solutions	27
4.4.4.	Opportunities	27
4.4.5.	Research Questions	27
4.4.6.	Requirements	27
4.5.	Industrial Safety	28
4.5.1.	Description	28
4.5.2.	Characterization	28
4.5.3.	Existing Solutions	28
4.5.4.	Opportunities	29
4.5.5.	Research Questions	29
4.5.6.	Requirements	29
5.	Improving existing COIN capabilities	29
5.1.	Content Delivery Networks	29
5.1.1.	Description	29
5.1.2.	Characterization	30
5.1.3.	Existing Solutions	30
5.1.4.	Opportunities	30
5.1.5.	Research Questions	30
5.1.6.	Requirements	31
5.2.	Compute-Fabric-as-a-Service (CFaaS)	31
5.2.1.	Description	31
5.2.2.	Characterization	31
5.2.3.	Existing Solutions	32
5.2.4.	Opportunities	32
5.2.5.	Research Questions	32
5.2.6.	Requirements	33
5.3.	Virtual Networks Programming	33
5.3.1.	Description	33
5.3.2.	Characterization	34
5.3.3.	Existing Solutions	36
5.3.4.	Opportunities	36
5.3.5.	Research Questions	37
5.3.6.	Requirements	38
6.	Enabling new COIN capabilities	38
6.1.	Distributed AI	38
6.1.1.	Description	38
6.1.2.	Characterization	39

6.1.3. Existing Solutions	39
6.1.4. Opportunities	39
6.1.5. Research Questions	40
6.1.6. Requirements	40
7. Analysis	40
7.1. Opportunities	40
7.2. Research Questions	41
7.2.1. Categorization	41
7.2.2. Analysis	42
7.3. Requirements	49
8. Security Considerations	49
9. IANA Considerations	49
10. Conclusion	49
11. List of Use Case Contributors	50
12. References	50
12.1. Normative References	50
12.2. Informative References	50
Authors' Addresses	53

1. Introduction

The Internet was designed as a best-effort packet network that offers limited guarantees regarding the timely and successful transmission of packets. Data manipulation, computation, and more complex protocol functionality is generally provided by the end-hosts while network nodes are kept simple and only offer a "store and forward" packet facility. This design choice has shown suitable for a wide variety of applications and has helped in the rapid growth of the Internet.

However, with the expansion of the Internet, there are more and more fields that require more than best-effort forwarding including strict performance guarantees or closed-loop integration to manage data flows. In this context, allowing for a tighter integration of computing and networking resources, enabling a more flexible distribution of computation tasks across the network, e.g., beyond 'just' endpoints, may help to achieve the desired guarantees and behaviors as well as increase overall performance. The vision of 'in-network computing' and the provisioning of such capabilities that capitalize on joint computation and communication resource usage throughout the network is core to the efforts in the COIN RG; we refer to those capabilities as 'COIN capabilities' in the remainder of the document.

We believe that such vision of 'in-network computing' can be best outlined along four dimensions of use cases, namely those that (i) provide new user experiences through the utilization of COIN capabilities (referred to as 'COIN experiences'), (ii) enable new

COIN systems, e.g., through new interactions between communication and compute providers, (iii) improve on already existing COIN capabilities and (iv) enable new COIN capabilities. Sections 3 through 6 capture those categories of use cases and provide the main structure of this document. The goal is to present how the presence of computing resources inside the network impacts existing services and applications or allows for innovation in emerging fields.

Through delving into some individual examples within each of the above categories, we aim to outline opportunities and propose possible research questions for consideration by the wider community when pushing forward the 'in-network computing' vision. Furthermore, insights into possible requirements for an evolving solution space of collected COIN capabilities is another objective of the individual use case descriptions. This results in the following taxonomy used to describe each of the use cases:

1. Description: Purpose of the use case and explanation of the use case behavior
2. Characterization: Explanation of the services that are being utilized and realized as well as the semantics of interactions in the use case.
3. Existing solutions: Describe, if existing, current methods that may realize the use case.
4. Opportunities: Outline how COIN capabilities may support or improve on the use case in terms of performance and other metrics.
5. Research questions: State essential questions that are suitable for guiding research to achieve the outlined opportunities
6. Requirements: Describe the requirements for any solutions for COIN capabilities that may need development along the opportunities outlined in item 4; here, we limit requirements to those COIN capabilities, recognizing that any use case will realistically hold many additional requirements for its realization.

In Section 7, we will summarize the key research questions across all use cases and identify key requirements across all use cases. This will provide a useful input into future roadmapping on what COIN capabilities may emerge and how solutions of such capabilities may look like. It will also identify what open questions remain for these use cases to materialize as well as define requirements to steer future (COIN) research work.

2. Terminology

The following terminology has been partly aligned with [I-D.draft-kutscher-coinrg-dir]:

(COIN) Program: a set of computations requested by a user

(COIN) Program Instance: one currently executing instance of a program

(COIN) Function: a specific computation that can be invoked as part of a program

COIN Capability: a feature enabled through the joint processing of computation and communication resources in the network

COIN Experience: a new user experience brought about through the utilization of COIN capabilities

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

(COIN) Execution Environment: a class of target environments for function execution, for example, a JVM-based execution environment that can run functions represented in JVM byte code

COIN System: the PNDs (and end systems) and their execution environments, together with the communication resources interconnecting them, operated by a single provider or through interactions between multiple providers that jointly offer COIN capabilities

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

3. Providing New COIN Experiences

3.1. Mobile Application Offloading

3.1.1. Description

The scenario can be exemplified in an immersive gaming application, where a single user plays a game using a VR headset. The headset hosts functions that "display" frames to the user, as well as the functions for VR content processing and frame rendering combining with input data received from sensors in the VR headset.

Once this application is partitioned into constituent (COIN) programs and deployed throughout a COIN system, utilizing the COIN execution environment, only the "display" (COIN) programs may be left in the headset, while the compute intensive real-time VR content processing (COIN) programs can be offloaded to a nearby resource rich home PC or a PND in the operator's access network, for a better execution (faster and possibly higher resolution generation).

3.1.2. Characterization

Partitioning a mobile application into several constituent (COIN) programs allows for denoting the application as a collection of (COIN) functions for a flexible composition and a distributed execution. In our example above, most functions of a mobile application can be categorized into any of three, "receiving", "processing" and "displaying" function groups.

Any device may realize one or more of the (COIN) programs of a mobile application and expose them to the (COIN) system and its constituent (COIN) execution environments. When the (COIN) program sequence is executed on a single device, the outcome is what you see today as applications running on mobile devices.

However, the execution of (COIN) functions may be moved to other (e.g., more suitable) devices, including PNDs, which have exposed the corresponding (COIN) programs as individual (COIN) program instances to the (COIN) system by means of a 'service identifier'. The result of the latter is the equivalent to 'mobile function offloading', for possible reduction of power consumption (e.g., offloading CPU intensive process functions to a remote server) or for improved end user experience (e.g., moving display functions to a nearby smart TV) by selecting more suitable placed (COIN) program instances in the overall (COIN) system.

Figure 1 shows one realization of the above scenario, where a 'DPR app' is running on a mobile device (containing the partitioned Display(D), Process(P) and Receive(R) COIN programs) over an SDN network. The packaged applications are made available through a localized 'playstore server'. The mobile application installation is realized as a 'service deployment' process, combining the local app installation with a distributed (COIN) program deployment (and orchestration) on most suitable end systems or PNDs ('processing server').

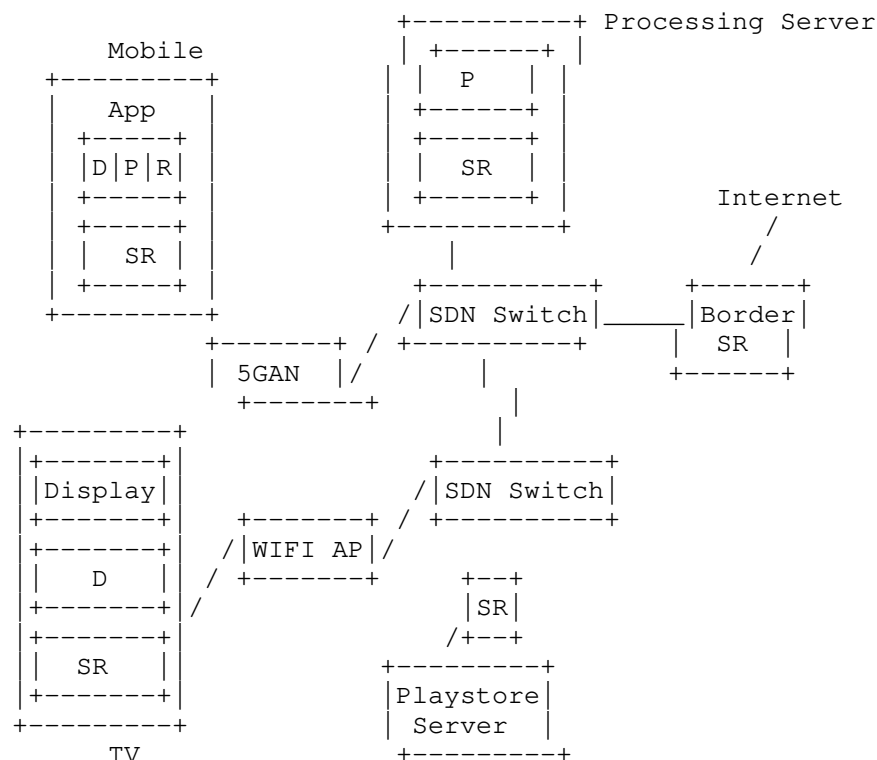


Figure 1: Application Function Offloading Example.

Such localized deployment could, for instance, be provided by a visiting site, such as a hotel or a theme park. Once the 'processing' (COIN) program is terminated on the mobile device, the 'service routing' (SR) elements in the network route (service) requests instead to the (previously deployed) 'processing' (COIN) program running on the processing server over an existing SDN network. Here, capabilities and other constraints for selecting the appropriate (COIN) program, in case of having deployed more than one, may be provided both in the advertisement of the (COIN) program and the service request itself.

As an extension to the above scenarios, we can also envision that content from one processing (COIN) program may be distributed to more than one display (COIN) program, e.g., for multi/many-viewing scenarios, thereby realizing a service-level multicast capability towards more than one (COIN) program.

3.1.3. Existing Solutions

NOTE: material on solutions like ETSI MEC will be added here later

3.1.4. Opportunities

- * The packaging of (COIN) programs into existing mobile application packaging may enable the migration from current (mobile) device-centric execution of those mobile application towards a possible distributed execution of the constituent (COIN) programs that are part of the overall mobile application.
- * The orchestration for deploying (COIN) program instances in specific end systems and PNDs alike may open up the possibility for localized infrastructure owners, such as hotels or venue owners, to offer their compute capabilities to their visitors for improved or even site-specific experiences.
- * The execution of (current mobile) app-level (COIN) programs may speed up the execution of said (COIN) program by relocating the execution to more suitable devices, including PNDs.
- * The support for service-level routing of requests (service routing in [APPCENTRES]) may support higher flexibility when switching from one (COIN) program instance to another, e.g., due to changing constraints for selecting the new (COIN) program instance.
- * The ability to identifying service-level in-network computing elements will allow for routing service requests to those COIN elements, including PNDs, therefore possibly allowing for new in-network functionality to be included in the mobile application.
- * The support for constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) may allow for a more flexible and app-specific selection of (COIN) program instances, thereby allowing for better meeting the app-specific and end user requirements.

3.1.5. Research Questions

- * RQ 3.1.1: How to combine service-level orchestration frameworks with app-level packaging methods?
- * RQ 3.1.2: How to reduce latencies involved in (COIN) program interactions where (COIN) program instance locations may change quickly?

- * RQ 3.1.3: How to signal constraints used for routing requests towards (COIN) program instances in a scalable manner?
- * RQ 3.1.4: How to identify (COIN) programs and program instances?
- * RQ 3.1.5: How to identify specific choice of (COIN) program instances over others?
- * RQ 3.1.6: How to provide affinity of service requests towards (COIN) program instances, i.e., longer-term transactions with ephemeral state established at a specific (COIN) program instance?
- * RQ 3.1.7: How to provide constraint-based routing decisions at packet forwarding speed?
- * RQ 3.1.8: What in-network capabilities may support the execution of (COIN) programs and their instances?

3.1.6. Requirements

- * Req 3.1.1: Any COIN system MUST provide means for routing of service requests between resources in the distributed environment.
- * Req 3.1.2: Any COIN system MUST provide means for identifying services exposed by (COIN) programs for directing service requests
- * (Req 3.1.3: Any COIN system MUST provide means for identifying (COIN) program instances for directing (affinity) requests to a specific (COIN) program instance
- * Req 3.1.4: Any COIN system MUST provide means for dynamically choosing the best possible service sequence of one or more (COIN) programs for a given application experience, i.e., support for chaining (COIN) program executions.
- * Req 3.1.5: Means for discovering suitable (COIN) programs SHOULD be provided.
- * Req 3.1.6: Any COIN system MUST provide means for pinning the execution of a service of a specific (COIN) program to a specific resource, i.e., (COIN) program instance in the distributed environment.
- * Req 3.1.7: Any COIN system SHOULD provide means for packaging micro-services for deployments in distributed networked computing environments.

- * Req 3.1.8: The packaging MAY include any constraints regarding the deployment of (COIN) program instances in specific network locations or compute resources, including PNDs.
- * Req 3.1.9: Such packaging SHOULD conform to existing application deployment models, such as mobile application packaging, TOSCA orchestration templates or tar balls or combinations thereof.
- * Req 3.1.10: Any COIN system MUST provide means for real-time synchronization and consistency of distributed application states.

3.2. Extended Reality and Immersive Media

3.2.1. Description

Virtual Reality (VR), Augmented Reality (AR) and immersive media (the metaverse) taken together as Extended Reality (XR) are the drivers of a number of advances in interactive technologies. XR is one example of the Multisource-Multidestination Problem that combines video, haptics, and tactile experiences in interactive or networked multi-party and social interactions. While initially associated with gaming and entertainment, XR applications now include remote diagnosis, maintenance, telemedicine, manufacturing and assembly, autonomous systems, smart cities, and immersive classrooms.

Because XR requirements include the need to provide real-time interactivity for immersive and increasingly mobile immersive applications with tactile and time-sensitive data and high bandwidth for high resolution images and local rendering for 3D images and holograms, they are difficult to run over traditional networks; in consequence innovation is needed to deploy the full potential of the applications.

3.2.2. Characterization

Collaborative XR experiences are difficult to deliver with a client-server cloud-based solution as they require a combination of: stream synchronization, low delays and delay variations, means to recover from losses and optimized caching and rendering as close as possible to the user at the network edge. XR deals with personal information and potentially protected content this an XR application must also provide a secure environment and ensure user privacy. Additionally, the sheer amount of data needed for and generated by the XR applications can use recent trend analysis and mechanisms, including machine learning to find these trends and reduce the size of the data sets. Video holography and haptics require very low delay or generate large amounts of data, both requiring a careful look at data filtering and reduction, functional distribution and partitioning.

The operation of XR over networks requires some computing in the nodes from content source to destination. But a lot of these remain in the realm of research to resolve the resource allocation problem and provide adequate quality of experience. These include multi-variate and heterogeneous goal optimization problems at merging nodes requiring advanced analysis. Image rendering and video processing in XR leverages different HW capabilities combinations of CPU and GPU at the edge (even at the mobile edge) and in the fog network where the content is consumed. It is important to note that the use of in-network computing for XR does not imply a specific protocol but targets an architecture enabling the deployment of the services.

3.2.3. Existing Solutions

In-network computing for XR profits from the heritage of extensive research in the past years on Information Centric Networking, Machine Learning, network telemetry, imaging and IoT as well as distributed security and in-network coding.

- * Enabling Scalable Edge Video Analytics with Computing-In-Network (Jun Chen Jiang of the University of Chicago): this work brings a periodical re-profiling to adapt the video pipeline to the dynamic video content that is a characteristic of XR. The implication is that we "need tight network-app coupling" for real time video analytics.
- * VR journalism, interactive VR movies and meetings in cyberspace (many projects PBS, MIT interactive documentary lab, Huawei research - references to be provided): typical VR is not made for multiparty and these applications require a tight coupling of the local and remote rendering and data capture and combinations of cloud (for more static information) and edge (for dynamic content).
- * Local rendering of holographic content using near field computation (heritage from advances cockpit interactions - looking for non military papers): a lot has been said recently of the large amounts of data necessary to transmit and use holographic imagery in communications. Transmitting the near field information and rendering the image locally allows to reduce the data rates by 1 or 2.
- * ICE-AR [ICE] project at UCLA (Jeff Burke): while this project is a showcase of the NDN network architecture it also uses a lot of edge-cloud capabilities for example for inter-server games and advanced video applications.

3.2.4. Opportunities

- * **Reduced latency:** the physical distance between the content cloud and the users must be short enough to limit the propagation delay to the 20 ms usually cited for XR applications; the use of local CPU and IoT devices for range of interest (RoI) detection and dynamic rendering may enable this.
- * **Video transmission:** better transcoding and use of advanced context-based compression algorithms, pre-fetching and pre-caching and movement prediction not only in the cloud.
- * **Monitoring:** telemetry is a major research topic for COIN and it enables to monitor and distribute the XR services.
- * **Network access:** push some networking functions in the kernel space into the user space to enable the deployment of stream specific algorithms for congestion control and application-based load balancing based on machine learning and user data patterns.
- * **Functional decomposition:** functional decomposition, localization and discovery of computing and storage resources in the network. But it is not only finding the best resources but qualifying those resources in terms of reliability especially for mission critical services in XR (medicine for example). This could include intelligence services.

3.2.5. Research Questions

- * **RQ 3.2.1:** Can current programmable network entities be sufficient to provide the speed required to provide and execute complex filtering operations that includes metadata analysis for complex and dynamic scene rendering?
- * **RQ 3.2.2:** How can the interoperability of CPU/GPU be optimized to combine low level packet filtering with the higher layer processors needed for image processing and haptics?
- * **RQ 3.2.3:** Can the use of joint learning algorithms across both data center and edge computers be used to create optimal functionality allocation and the creation of semi-permanent datasets and analytics for usage trending resulting in better localization of XR functions?
- * **RQ 3.2.4:** Can COIN improve the dynamic distribution of control, forwarding and storage resources and related usage models in XR?

3.2.6. Requirements

- * Req 3.2.1: Allow joint collaboration.
- * Req 3.2.2: Provide multi-views.
- * Req 3.2.3: Include extra streams dynamically for data intensive services, manufacturing and industrial processes.
- * Req 3.2.4: Enable multistream, multidevice, multideestination applications.
- * Req 3.2.5: Use new Internet Architectures at the edge for improved performance and performance management.
- * Req 3.2.6: Integrate with holography, 3D displays and image rendering processors.
- * Req 3.2.7: All the use of multicast distribution and processing as well as peer to peer distribution in bandwidth and capacity constrained environments.
- * Req 3.2.8: Evaluate the integration local and fog caching with cloud-based pre-rendering.
- * Req 3.2.9: Evaluate ML-based congestion control to manage XR sessions quality of service and to determine how to prioritize data.
- * Req 3.2.10: Consider higher layer protocols optimization to reduce latency especially in data intensive applications at the edge.
- * Req 3.2.11: Provide trust, including blockchains and smart-contracts to enable secure community building across domains.
- * Req 3.2.12: Support nomadicity and mobility (link to mobile edge).
- * Req 3.2.13: Use 5G slicing to create independent session-driven processing/rendering.
- * Req 3.2.14: Provide performance optimization by data reduction, tunneling, session virtualization and loss protection.
- * Req 3.2.15: Use AI/ML for trend analysis and data reduction when appropriate.

3.3. Personalised and interactive performing arts

3.3.1. Description

This use case covers live productions of the performing arts where the performers and audience are in different physical locations. The performance is conveyed to the audience through multiple networked streams which may be tailored to the requirements of individual audience members; and the performers receive live feedback from the audience.

There are two main aspects: i) to emulate as closely as possible the experience of live performances where the performers and audience are co-located in the same physical space, such as a theatre; and ii) to enhance traditional physical performances with features such as personalisation of the experience according to the preferences or needs of the audience members.

Examples of personalisation include:

- * Viewpoint selection such as choosing a specific seat in the theatre or for more advanced positioning of the audience member's viewpoint outside of the traditional seating - amongst, above or behind the performers (but within some limits which may be imposed by the performers or the director for artistic reasons);
- * Augmentation of the performance with subtitles, audio-description, actor-tagging, language translation, advertisements/product-placement, other enhancements/filters to make the performance accessible to disabled audience members (removal of flashing images for epileptics, alternative colour schemes for colour-blind audience members, etc.).

3.3.2. Characterization

There are several chained functional entities which are candidates for being deployed as (COIN) Programs.

- * Performer aggregation and editing functions
- * Distribution and encoding functions
- * Personalisation functions
 - to select which of the existing streams should be forwarded to the audience member
 - to augment streams with additional metadata such as subtitles

- to create new streams after processing existing ones: to interpolate between camera angles to create a new viewpoint or to render point clouds from the audience member's chosen perspective
 - to undertake remote rendering according to viewer position, e.g. creation of VR headset display streams according to audience head position – when this processing has been offloaded from the viewer's end-system to the in-network function due to limited processing power in the end-system, or to limited network bandwidth to receive all of the individual streams to be processed.
- * Audience feedback sensor processing functions
 - * Audience feedback aggregation functions

These are candidates for deployment as (COIN) Programs in PNDs rather than being located in end-systems (at the performers' site, the audience members' premises or in a central cloud location) for several reasons:

- * Personalisation of the performance according to audience preferences and requirements makes it unfeasible to be done in a centralised manner at the performer premises: the computational resources and network bandwidth would need to scale with the number of audience members' personalised streams.
- * Rendering of VR headset content to follow viewer head movements has an upper bound on lag to maintain viewer QoE, which requires the processing to be undertaken sufficiently close to the viewer to avoid large network latencies.
- * Viewer devices may not have the processing-power to undertake the personalisation or the viewers' network may not have the capacity to receive all of the constituent streams to undertake the personalisation functions.
- * There are strict latency requirements for live and interactive aspects that require the deviation from the direct network path from performers to audience to be minimised, which reduces the opportunity to route streams via large-scale processing capabilities at centralised data-centres.

3.3.3. Existing solutions

Note: Existing solutions for some aspects of this use case are covered in the Mobile Application Offloading, Extended Reality, and Content Delivery Networks use cases.

3.3.4. Opportunities

- * Executing media processing and personalisation functions on-path as (COIN) Programs in PNDs will avoid detour/stretch to central servers which increases latency as well as the consumption of bandwidth on more network resources (links and routers). For example, in this use case the chain of (COIN) Programs and propagation over the interconnecting network segments for performance capture, aggregation, distribution, personalisation, consumption, capture of audience response, feedback processing, aggregation, rendering should be achieved within an upper bound of latency (the tolerable amount is to be defined, but in the order of 100s of ms to mimic performers perceiving audience feedback, such as laughter or other emotional responses in a theatre setting).
- * Processing of media streams allows (COIN) Programs, PNDs and the wider (COIN) System/Environment to be contextual aware of flows and their requirements which can be used for determining network treatment of the flows, e.g. path selection, prioritisation, multi-flow coordination, synchronisation & resilience.

3.3.5. Research Questions:

- * RQ 3.3.1: In which PNDs should (COIN) Programs for aggregation, encoding and personalisation functions be located? Close to the performers or close to the audience members?
- * RQ 3.3.2: How far from the direct network path from performer to audience should (COIN) programs be located, considering the latency implications of path-stretch and the availability of processing capacity at PNDs? How should tolerances be defined by users?
- * RQ 3.3.3: Should users decide which PNDs should be used for executing (COIN) Programs for their flows or should they express requirements and constraints that will direct decisions by the orchestrator/manager of the COIN System?
- * RQ 3.3.4: How to achieve network synchronisation across multiple streams to allow for merging, audio-video interpolation and other cross-stream processing functions that require time

synchronisation for the integrity of the output? How can this be achieved considering that synchronisation may be required between flows that are: i) on the same data pathway through a PND/router, ii) arriving/leaving through different ingress/egress interfaces of the same PND/router, iii) routed through disjoint paths through different PNDs/routers?

- * RQ 3.3.5: Where will COIN Programs will be executed? In the data-plane of PNDs, in other on-router computational capabilities within PNDs, or in adjacent computational nodes?
- * RQ 3.3.6: Are computationally-intensive tasks - such as video stitching or media recognition and annotation - considered as suitable candidate (COIN) Programs or should they be implemented in end-systems?
- * RQ 3.3.7: If the execution of COIN Programs is offloaded to computational nodes outside of PNDs, e.g. for processing by GPUs, should this still be considered as in-network processing? Where is the boundary between in-network processing capabilities and explicit routing of flows to endsystems?

3.3.6. Requirements

- * Req 3.3.1: Users should be able to specify requirements on network and processing metrics (such as latency and throughput bounds) and the COIN System should be able to respect those requirements and constraints when routing flows and selecting PNDs for executing (COIN) Programs.
- * Req 3.3.2: A COIN System should be able to synchronise flow treatment and processing across multiple related flows which may be on disjoint paths.

4. Supporting new COIN Systems

While the best-effort nature of the Internet enables a wide variety of applications, there are several domains whose requirements are hard to satisfy over regular best-effort networks.

Consequently, there is a large number of specialized appliances and protocols designed to provide 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. However, solutions on the link layer alone are not always sufficient.

The industrial domain, e.g., currently evolves towards increasingly interconnected systems in turn increasing the complexity of the underlying networks, making them more dynamic, and creating more diverse sets of requirements. Concepts satisfying the dynamic performance requirements of modern industrial applications thus become harder to develop. In this context, COIN offers new possibilities as it allows to flexibly distribute computation tasks across the network and enables novel forms of interaction between communication and computation providers.

This document illustrates the potential for new COIN systems using the example of the industrial domain by characterizing and analyzing specific scenarios to showcase potential requirements, as specifying general requirements is difficult due to the domain's mentioned diversity.

4.1. Industrial Network Scenario

Common components of industrial networks can be divided into three categories as illustrated in Figure 2. 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.

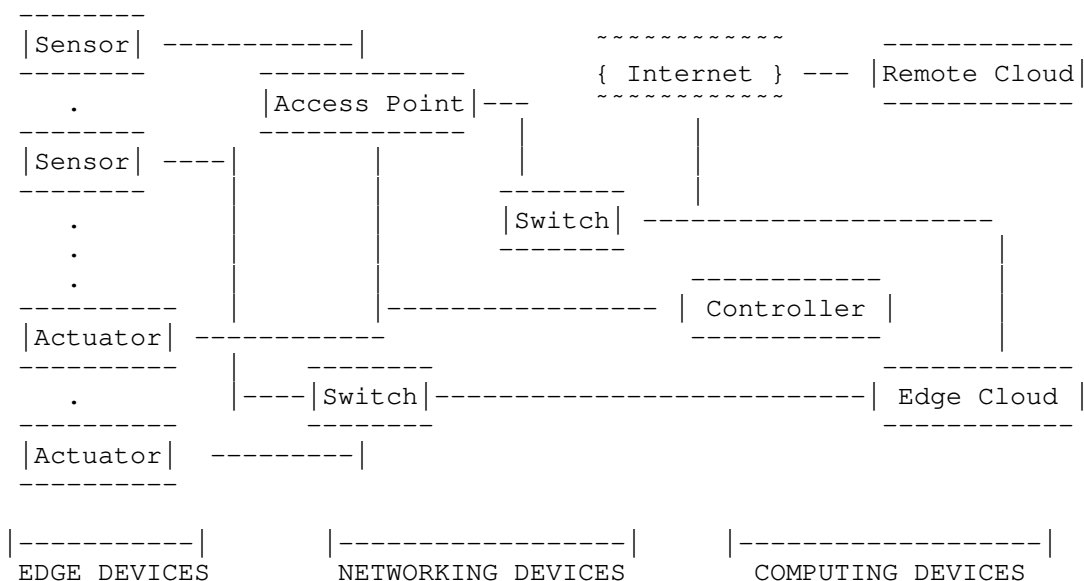


Figure 2: Industrial networks show a high level of heterogeneity.

4.2. In-Network Control / Time-sensitive applications

4.2.1. Description

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. Moving compute functionality onto COIN execution environments inside the network offers a new solution space to these challenges.

4.2.2. Characterization

A control process consists of two main components as illustrated in Figure 3: 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.

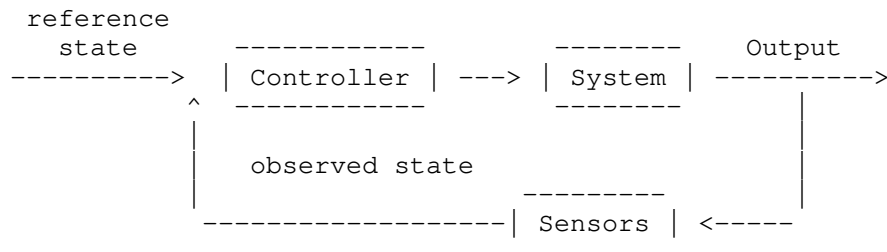


Figure 3: Simple feedback control model.

Apart from the control model, the quality of the control primarily depends on the timely reception of the sensor feedback which 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, e.g., off-premise cloud platforms are included.

4.2.3. Existing Solutions

Control functionality is traditionally executed on PLCs close to the machinery. These PLCs typically require vendor-specific implementations and are often hard to upgrade and update which makes such control processes inflexible and difficult to manage. Moving computations to more freely programmable devices thus has the potential of significantly improving the flexibility. In this context, directly moving control functionality to (central) cloud environments is generally possible, yet only feasible if latency constraints are lenient.

4.2.4. Opportunities

COIN offers the possibility of bringing the system and the controller closer together, thus possibly satisfying the latency requirements, by performing simple control logic on PNDs and/or in COIN execution environments.

While control models, in general, can become involved, 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.

Opportunities:

- * Execute simple (end-host) COIN functions on PNDs to satisfy tight latency constraints of control processes

4.2.5. Research Questions

Bringing the required computations to PNDs is challenging as these 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 [KUNZE]. 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:

Research Questions:

- * RQ 4.2.1: How to derive simplified versions of the global (control) function?
 - How to account for the limited computational precision of PNDs?
 - How to find suitable tradeoffs regarding simplicity of the control function ("accuracy of the control") and implementation complexity ("implementability")?
- * RQ 4.2.2: How to distribute the simplified versions in the network?

- Can there be different control levels, e.g., "quite inaccurate & very low latency" (PNDs, deep in the network), "more accurate & higher latency" (more powerful COIN execution environments, farer away), "very accurate & very high latency" (cloud environments, far away)?
- Who decides which control instance is executed and how?
- How do the different control instances interact?

4.2.6. Requirements

- * Req 4.2.1: The interaction between the COIN execution environments and the global controller SHOULD be explicit.
- * Req 4.2.2: The interaction between the COIN execution environments and the global controller MUST NOT negatively impact the control quality.
- * Req 4.2.3: Actions of the COIN execution environments MUST be overridable by the global controller.
- * Req 4.2.4: Functions in COIN execution environments SHOULD be executed with predictable delay.
- * Req 4.2.5: Functions in COIN execution environments MUST be executed with predictable accuracy.

4.3. Large Volume Applications - Filtering

4.3.1. Description

In modern industrial networks, processes and machines can be monitored closely resulting in large volumes of available information. This data can be used to find previously unknown correlations between different parts of the value chain, e.g., by deploying machine learning (ML) techniques, 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. Off-premise cloud platforms offer cost-effective solutions with a high degree of flexibility and scalability, however, moving all data to off-premise locations poses infrastructural challenges. Pre-processing or filtering the data already in COIN execution environments can be a new solution to this challenge.

4.3.2. Characterization

4.3.2.1. General Characterization of Large Volume Applications

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 COIN techniques hardly applicable as they typically work on unencrypted data. Adding security to COIN 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 8.

4.3.2.2. Specific Characterization for Filtering Solutions

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.

4.3.3. Existing Solutions

Current approaches for handling such large amounts of information typically build upon stream processing frameworks such as Apache Flink. While they allow for handling large volume applications, they are tied to performant server machines and upscaling the information density also requires a corresponding upscaling of the compute infrastructure.

4.3.4. Opportunities

PNDs and COIN execution environments are in a unique position to reduce the data rates due to their line-rate packet processing capabilities. Using these capabilities, it is possible 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 could 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.

While the former variant is oblivious to the semantics of the sensor data, the latter variant requires an understanding of the current sensor levels. In any case, 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.

Opportunities:

- * (Semantic) packet filtering based on packet header and payload, as well as multi-packet information

4.3.5. Research Questions

- * RQ 4.3.1: How to design COIN programs for (semantic) packet filtering?
 - Which criteria for filtering make sense?
- * RQ 4.3.2: How to distribute and coordinate COIN programs?
- * RQ 4.3.3: How to dynamically change COIN programs?
- * RQ 4.3.4: How to signal traffic filtering by COIN programs to end-hosts?

4.3.6. Requirements

- * Req 4.3.1: Filters MUST conform to application-level syntax and semantics.
- * Req 4.3.2: Filters MAY leverage packet header and payload information.
- * Req 4.3.3: Filters SHOULD be reconfigurable at run-time.

4.4. Large Volume Applications - (Pre-)Preprocessing

4.4.1. Description

See Section 4.3.1.

4.4.2. Characterization

4.4.2.1. General Characterization of Large Volume Applications

See Section 4.3.2.1.

4.4.2.2. Specific Characterization for Preprocessing Solutions

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.

4.4.3. Existing Solutions

See Section 4.3.3.

4.4.4. Opportunities

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 PNDs is also interesting, because sensor data is usually streamed.

Opportunities:

- * (Semantic) data (pre-)processing, e.g., in the form of computations across multiple packets and potentially leveraging packet payload

4.4.5. Research Questions

- * RQ 4.4.1: Which kinds of COIN programs can be leveraged for (pre-)processing steps?
 - How complex can they become?
- * RQ 4.4.2: How to distribute and coordinate COIN programs?
- * RQ 4.4.3: How to dynamically change COIN programs?
- * RQ 4.4.4: How to incorporate the (pre-)processing steps into the overall system?

4.4.6. Requirements

- * Req 4.4.1: Preprocessors MUST conform to application-level syntax and semantics.
- * Req 4.4.2: Preprocessors MAY leverage packet header and payload information.
- * Req 4.4.3: Preprocessors SHOULD be reconfigurable at run-time.

4.5. Industrial Safety

4.5.1. Description

Despite an 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.

4.5.2. Characterization

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.

4.5.3. Existing Solutions

Due to the importance of safety, there is a wide range of software-based approaches aiming at enhancing security. One example are tag-based systems, e.g., using RFID, where drivers of forklifts can be warned if pedestrian workers carrying tags are nearby. Such solutions, however, require setting up an additional system and do not leverage existing sensor data.

4.5.4. Opportunities

COIN systems could leverage the increased availability of sensor data and the detailed monitoring of the factories to enable additional safety measures. Different safety indicators within the production hall can be combined within the network so that PNDs can give early responses if a potential safety breach is detected.

One 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.

Opportunities:

- * Execute simple (end-host) COIN functions on PNDs to create early emergency reactions based on diverse sensor feedback

4.5.5. Research Questions

- * RQ 4.5.1: Which additional safety measures can be provided?
 - Do these measures actually improve safety?
- * RQ 4.5.2: Which sensor information can be combined and how?

4.5.6. Requirements

- * Req 4.5.1: COIN-based safety measures MUST NOT degrade existing safety measures.
- * Req 4.5.2: COIN-based safety measures MAY enhance existing safety measures.

5. Improving existing COIN capabilities

5.1. Content Delivery Networks

5.1.1. Description

Delivery of content to end users often relies on Content Delivery Networks (CDNs) storing said content closer to end users for latency reduced delivery with DNS-based indirection being utilized to serve the request on behalf of the origin server.

5.1.2. Characterization

From the perspective of this draft, a CDN can be interpreted as a (network service level) set of (COIN) programs, implementing a distributed logic for distributing content from the origin server to the CDN ingress and further to the CDN replication points which ultimately serve the user-facing content requests.

5.1.3. Existing Solutions

NOTE: material on solutions will be added here later

Studies such as those in [FCDN] have shown that content distribution at the level of named content, utilizing efficient (e.g., Layer 2) multicast for replication towards edge CDN nodes, can significantly increase the overall network and server efficiency. It also reduces indirection latency for content retrieval as well as reduces required edge storage capacity by benefiting from the increased network efficiency to renew edge content more quickly against changing demand.

5.1.4. Opportunities

- * Supporting service-level routing of requests (service routing in [APPCENTRES]) to specific (COIN) program instances may improve on end user experience in faster retrieving (possibly also more, e.g., better quality) content.
- * Supporting the constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) may improve the overall end user experience by selecting a 'more suitable' (COIN) program instance over another, e.g., avoiding/reducing overload situation in specific (COIN) program instances.
- * Supporting Layer 2 capabilities for multicast (compute interconnection and collective communication in [APPCENTRES]) may increase the network utilization and therefore increase the overall system utilization.

5.1.5. Research Questions

In addition to those request question for Section 3.1:

- * RQ 5.1.1: How to utilize L2 multicast to improve on CDN designs? How to utilize in-network capabilities in those designs?

- * RQ 5.1.2: What forwarding methods may support the required multicast capabilities (see [FCDN])
- * RQ 5.1.3: What are the right routing constraints that reflect both compute and network capabilities?
- * RQ 5.1.4: Could traffic steering be performed at the data path and per service request? If so, what would be performance improvements?
- * RQ 5.1.5: How could storage be traded off against frequent, multicast-based, replication (see [FCDN])?
- * RQ 5.1.6: What scalability limits exist for L2 multicast capabilities? How to overcome them?

5.1.6. Requirements

Requirements 3.1.1 through 3.1.6 also apply for CDN service access. In addition:

- * Req 5.1.1: Any solution SHOULD utilize Layer 2 multicast transmission capabilities for responses to concurrent service requests.

5.2. Compute-Fabric-as-a-Service (CFaaS)

5.2.1. Description

Layer 2 connected compute resources, e.g., in regional or edge data centres, base stations and even end-user devices, provide the opportunity for infrastructure providers to offer CFaaS type of offerings to application providers. App and service providers may utilize the compute fabric exposed by this CFaaS offering for the purposes defined through their applications and services. In other words, the compute resources can be utilized to execute the desired (COIN) programs of which the application is composed, while utilizing the inter-connection between those compute resources to do so in a distributed manner.

5.2.2. Characterization

We foresee those CFaaS offerings to be tenant-specific, a tenant here defined as the provider of at least one application. For this, we foresee an interaction between CFaaS provider and tenant to dynamically select the appropriate resources to define the demand side of the fabric. Conversely, we also foresee the supply side of the fabric to be highly dynamic with resources being offered to the

fabric through, e.g., user-provided resources (whose supply might depend on highly context-specific supply policies) or infrastructure resources of intermittent availability such as those provided through road-side infrastructure in vehicular scenarios.

The resulting dynamic demand-supply matching establishes a dynamic nature of the compute fabric that in turn requires trust relationships to be built dynamically between the resource provider(s) and the CFaaS provider. This also requires the communication resources to be dynamically adjusted to interconnect all resources suitably into the (tenant-specific) fabric exposed as CFaaS.

5.2.3. Existing Solutions

NOTE: material on solutions will be added here later

5.2.4. Opportunities

- * Supporting service-level routing of compute resource requests (service routing in [APPCENTRES]) may allow for utilizing the wealth of compute resources in the overall CFaaS fabric for execution of distributed applications, where the distributed constituents of those applications are realized as (COIN) programs and executed within a COIN system as (COIN) program instances.
- * Supporting the constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) will allow for optimizing both the CFaaS provider constraints as well as tenant-specific constraints.
- * Supporting Layer 2 capabilities for multicast (compute interconnection and collective communication in [APPCENTRES]) will allow for increasing both network utilization but also possible compute utilization (due to avoiding unicast replication at those compute endpoints), thereby decreasing total cost of ownership for the CFaaS offering.

5.2.5. Research Questions

Similar to those for Section 3.1. In addition:

- * RQ 5.2.1: How to convey tenant-specific requirements for the creation of the L2 fabric?
- * RQ 5.2.2: How to dynamically integrate resources, particularly when driven by tenant-level requirements and changing service-specific constraints?

- * RQ 5.2.3: How to utilize in-network capabilities to aid the availability and accountability of resources, i.e., what may be (COIN) programs for a CFaaS environment that in turn would utilize the distributed execution capability of a COIN system?

5.2.6. Requirements

For the provisioning of services atop the CFaaS, requirements 3.1.1 through 3.1.6 should be addressed, too. In addition:

- * Req 5.2.1: Any solution SHOULD expose means to specify the requirements for the tenant-specific compute fabric being utilized for the service execution.
- * Req 5.2.2: Any solution SHOULD allow for dynamic integration of compute resources into the compute fabric being utilized for the app execution; those resources include, but are not limited to, end user provided resources. From a COIN system perspective, new resources must be possible to be exposed as possible (COIN) execution environments.
- * Req 5.2.3: Any solution MUST provide means to optimize the inter-connection of compute resources, including those dynamically added and removed during the provisioning of the tenant-specific compute fabric.
- * Req 5.2.4: Any solution MUST provide means for ensuring availability and usage of resources is accounted for.

5.3. Virtual Networks Programming

5.3.1. Description

The term "virtual network programming" is proposed to describe mechanisms by which tenants deploy and operate COIN programs in their virtual network. Such COIN programs can for example be P4 programs, OpenFlow rules, or higher layer programs. This feature can enable other use cases described in this draft to be deployed using virtual networks services, over underlying networks such as datacenters, mobile networks, or other fixed or wireless networks.

For example COIN programs could perform the following on a tenant's virtual network:

- * Allow or block flows, and request rules from an SDN controller for each new flow, or for flows to or from specific hosts that needs enhanced security

- * Forward a copy of some flows towards a node for storage and analysis
- * Update counters based on specific sources/destinations or protocols, for detailed analytics
- * Associate traffic between specific endpoints, using specific protocols, or originated from a given application, to a given slice, while other traffic use a default slice
- * Experiment with a new routing protocol (e.g., ICN), using a P4 implementation of a router for this protocol

5.3.2. Characterization

To provide a concrete example of virtual COIN programming, we consider a use case using a 5G underlying network, the 5GLAN virtualization technology, and the P4 programming language and environment. Section 5.1 of [I-D.ravi-icnrg-5gc-icn] provides a description of the 5G network functions and interfaces relevant to 5GLAN, which are otherwise specified in [TS23.501] and [TS23.502]. From the 5GLAN service customer/tenant standpoint, the 5G network operates as a switch.

In the use case depicted in Figure 4, the tenant operates a network including a 5GLAN network segment (seen as a single logical switch), as well as fixed segments. This can be in a plant or enterprise network, using for an example a 5G Non-Public Network (NPN). The tenant uses P4 programs to determine the operation of the fixed and 5GLAN switches. The tenant provisions a 5GLAN P4 program into the mobile network, and can also operate a controller. The mobile devices (or User Equipment nodes) UE1, UE2, UE3 and UE4 are in the same 5GLAN, as well as Device1 and Device2 (through UE4).

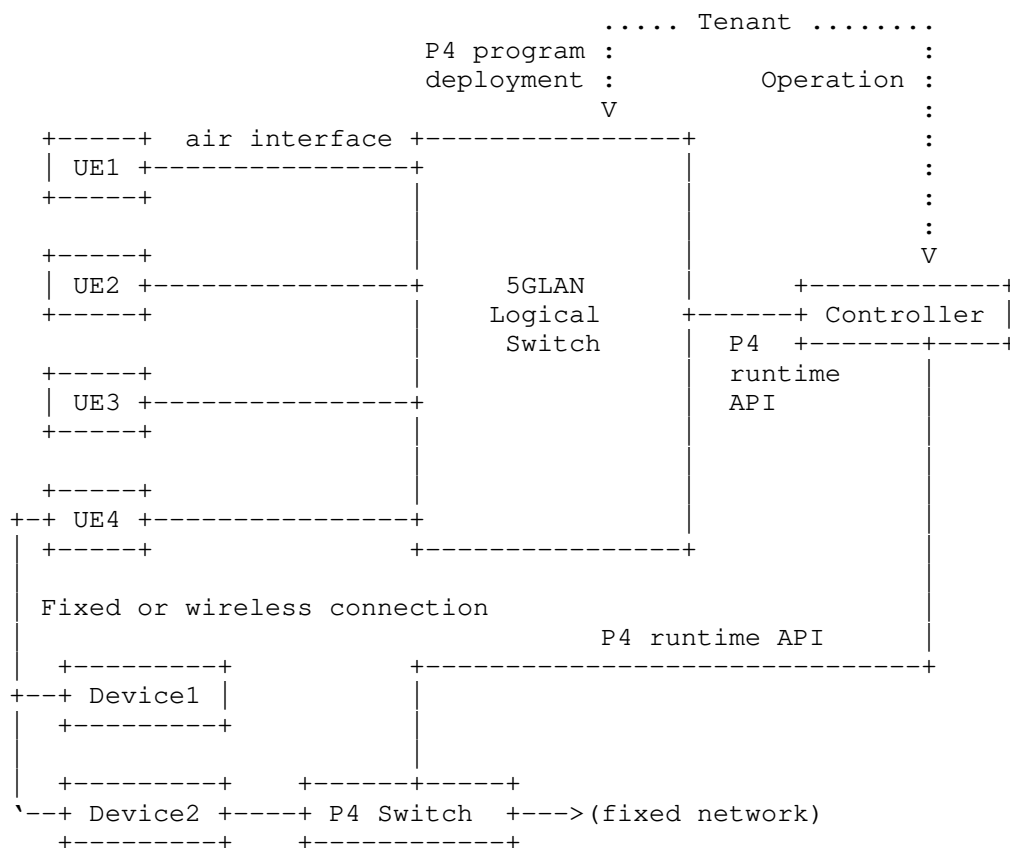


Figure 4: 5G Virtual Network Programming Overview

Looking in more details in Figure 5, the 5GLAN P4 program can be split between multiple data plane nodes (PDU Session Anchor (PSA) User Plane Functions (UPF), other UPFs, or even mobile devices), although in some cases the P4 program may be hosted on a single node. In the most general case, a distributed deployment is useful to keep traffic on optimal paths, because, except in simple cases, within a 5GLAN all traffic will not pass through a single node. In this example, P4 programs could be deployed in UPF1, UPF2, UPF3, UE3 and UE4. UE1-UE2 traffic is using a local switch on PSA UPF1, UE1-UE3 traffic is tunneled between PSA UPF1 and PSA UPF2 through the N19 interface, and UE1-UE4 traffic is forwarded through an external Data Network (DN). Traffic between Device1 and Device2 is forwarded through UE4.

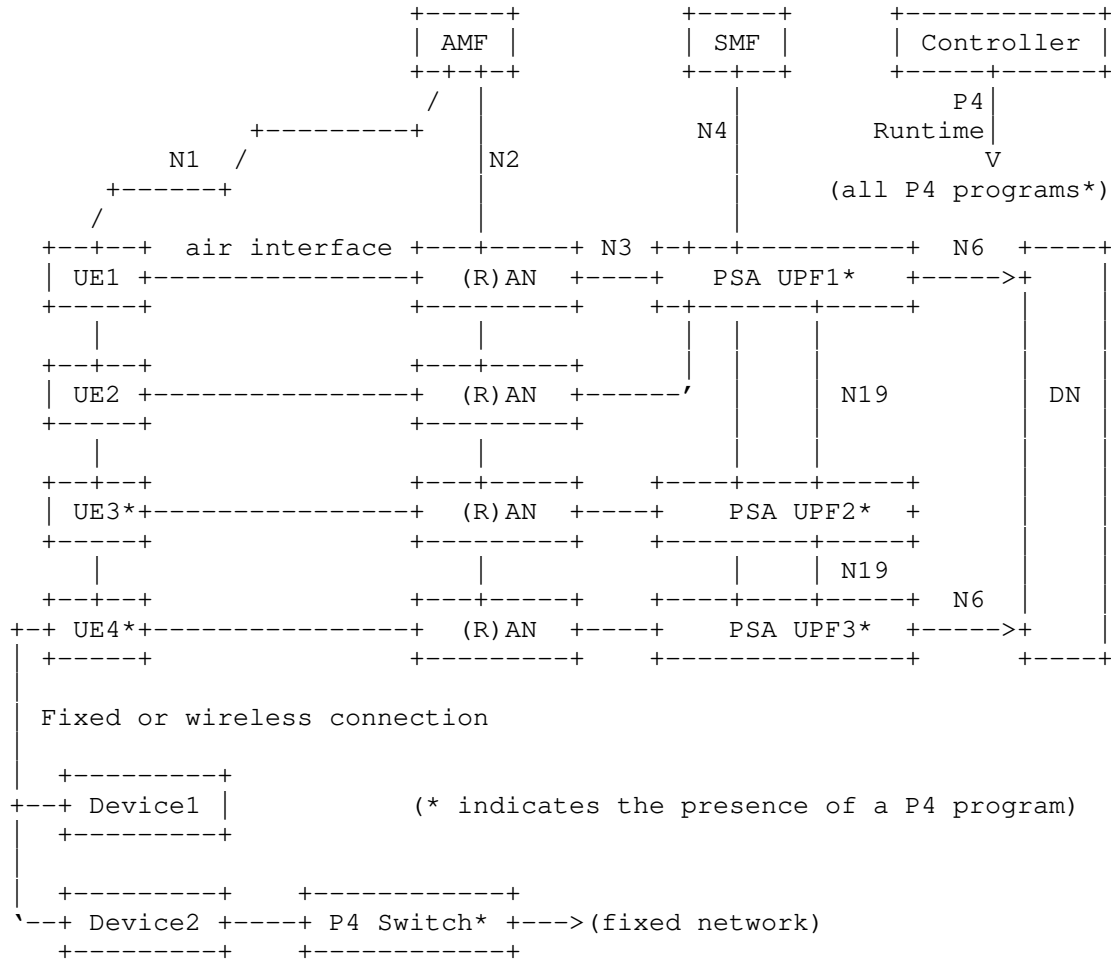


Figure 5: 5G Virtual Network Programming Details

5.3.3. Existing Solutions

Research has been conducted, for example by [Stoyanov], to enable P4 network programming of individual virtual switches. To our knowledge, no complete solution has been developed for deploying virtual COIN programs over mobile or datacenter networks.

5.3.4. Opportunities

Virtual network programming by tenants could bring benefits such as:

- * A unified programming model, which can facilitate porting in-network computing between data centers, 5G networks, and other fixed and wireless networks, as well as sharing controller, code and expertise.
- * Increasing the level of customization available to customers/tenants of mobile networks or datacenters, when compared with typical configuration capabilities. For example, 5G network evolution points to an ever increasing specialization and customization of private mobile networks, which could be handled by tenants using a programming model similar to P4.
- * Using network programs to influence underlying network service (e.g., request specific QoS for some flows in 5G or datacenters), to increases the level of in-depth customization available to tenants.

5.3.5. Research Questions

- * RQ 5.3.1: Underlying Network Awareness: a virtual COIN program can be able to influence, and be influenced by, the underling network (e.g., the 5G network or data center). For example, a virtual COIN program may be aware of the slice used by a flow, and possibly influence slice selection. Since some information and actions may be available on some nodes and not others, underlying network awareness may impose additional constraints on distributed network programs location.
- * RQ 5.3.2: Splitting/Distribution: a virtual COIN program may need to be deployed across multiple computing nodes, leading to research questions around instance placement and distribution. As a primary reason for this, program logic should be applied exactly once or at least once per packet, while allowing optimal forwarding path by the underlying network. For example, a 5GLAN P4 program may need to run on multiple UPFs. Research challenges include defining manual (by the programmer) or automatic methods to distribute COIN programs that use a low or minimal amount of resources. Distributed P4 programs are studied in [I-D.hsingh-coinrg-reqs-p4comp] and [Sultana].
- * RQ 5.3.3: Multi-Tenancy Support: multiple virtual COIN program instances can run on the same compute node. While mechanism were proposed for P4 multi-tenancy in a switch [Stoyanov], research questions remains, about isolation between tenants, fair repartition of resources.

- * RQ 5.3.4: Security: how can tenants and underlying networks be protected against security risks, including overuse or misuse of network resources, injection of traffic, access to unauthorized traffic?
- * RQ 5.3.5: Higher layer processing: can a virtual network model facilitate the deployment of COIN programs acting on application layer data? This is an open question since the present section focused on packet/flow processing.

5.3.6. Requirements

- * Req 5.3.1: A COIN system supporting virtualization should enable tenants to deploy COIN programs onto their virtual networks.
- * Req 5.3.2: A virtual COIN program should process flows/packets once and only once (or at least once for idempotent operations), even if the program is distributed over multiple PNDs.
- * Req 5.3.3: Multi-tenancy should be supported for virtual COIN programs, i.e., instances of virtual COIN programs from different tenants can share underlying PNDs. This includes requirements for secure isolation between tenants, and fair (or policy-based) sharing of computing resources.
- * Req 5.3.4: Virtual COIN programs should support mobility of endpoints.

6. Enabling new COIN capabilities

6.1. Distributed AI

6.1.1. Description

There is a growing range of use cases demanding for the realization of AI capabilities among distributed endpoints. Such demand may be driven by the need to increase overall computational power for large-scale problems. From a COIN perspective, those capabilities may be realized as (COIN) programs and executed throughout the COIN system, including in PNDs.

Some solutions may desire the localization of reasoning logic, e.g., for deriving attributes that better preserve privacy of the utilized raw input data. Quickly establishing (COIN) program instances in nearby compute resources, including PNDs, may even satisfy such localization demands on-the-fly (e.g., when a particular use is being realized, then terminated after a given time).

6.1.2. Characterization

Examples for large-scale AI problems include biotechnology and astronomy related reasoning over massive amounts of observational input data. Examples for localizing input data for privacy reasons include radar-like application for the development of topological mapping data based on (distributed) radio measurements at base stations (and possibly end devices), while the processing within radio access networks (RAN) already constitute a distributed AI problem to a certain extent albeit with little flexibility in distributing the execution of the AI logic.

6.1.3. Existing Solutions

Reasoning frameworks, such as TensorFlow, may be utilized for the realization of the (distributed) AI logic, building on remote service invocation through protocols such as gRPC [GRPC] or MPI [MPI] with the intention of providing an on-chip NPU (neural processor unit) like abstraction to the AI framework.

NOTE: material on solutions like ETSI MEC and 3GPP work will be added here later

6.1.4. Opportunities

- * Supporting service-level routing of requests (service routing in [APPCENTRES]), with AI services being exposed to the network and executed as part of (COIN) programs in selected (COIN) program instances, may provide a highly distributed execution of the overall AI logic, thereby addressing, e.g., localization but also computational concerns (scale-in/out).
- * The support for constraint-based selection of a specific (COIN) program instance over others (constraint-based routing in [APPCENTRES]) may allow for utilizing the most suitable HW capabilities (e.g., support for specific AI HW assistance in the COIN element, including a PND), while also allowing to select resources, e.g., based on available compute ability such as number of cores to be used.
- * Supporting collective communication between multiple instances of AI services, i.e., (COIN) program instances, may positively impact network but also compute utilization by moving from unicast replication to network-assisted multicast operation.

6.1.5. Research Questions

- * RQ 6.1.1: similar to use case in Section 3.1
- * RQ 6.1.2: What are the communication patterns that may be supported by collective communication solutions?
- * RQ 6.1.3: How to achieve scalable multicast delivery with rapidly changing receiver sets?
- * RQ 6.1.4: What in-network capabilities may support the collective communication patterns found in distributed AI problems?
- * RQ 6.1.5: How to provide a service routing capability that supports any invocation protocol (beyond HTTP)?

6.1.6. Requirements

Requirements 3.1.1 through 3.1.6 also apply for general distributed AI capabilities. In addition:

- * Req 6.1.1: Any COIN system MUST provide means to specify the constraints for placing (AI) execution logic in the form of (COIN) programs in certain logical execution points (and their associated physical locations), including PNDs.
- * Req 6.1.2: Any COIN system MUST provide support for app/micro-service specific invocation protocols for requesting (COIN) program services exposed to the COIN system.

7. Analysis

The goal of this analysis is to identify aspects that are relevant across all use cases to help in shaping the research agenda of COINRG. For this purpose, this section will condense the opportunities, research questions, as well as requirements of the different presented use cases and analyze these for similarities across the use cases.

Through this, we intend to identify cross-cutting opportunities, research questions as well as requirements (for COIN system solutions) that may aid the future work of COINRG as well as the larger research community.

7.1. Opportunities

To be added later.

7.2. Research Questions

After carefully considering the different use cases along with their research questions, we propose the following layered categorization to structure the content of the research questions which we illustrate in Figure 6.

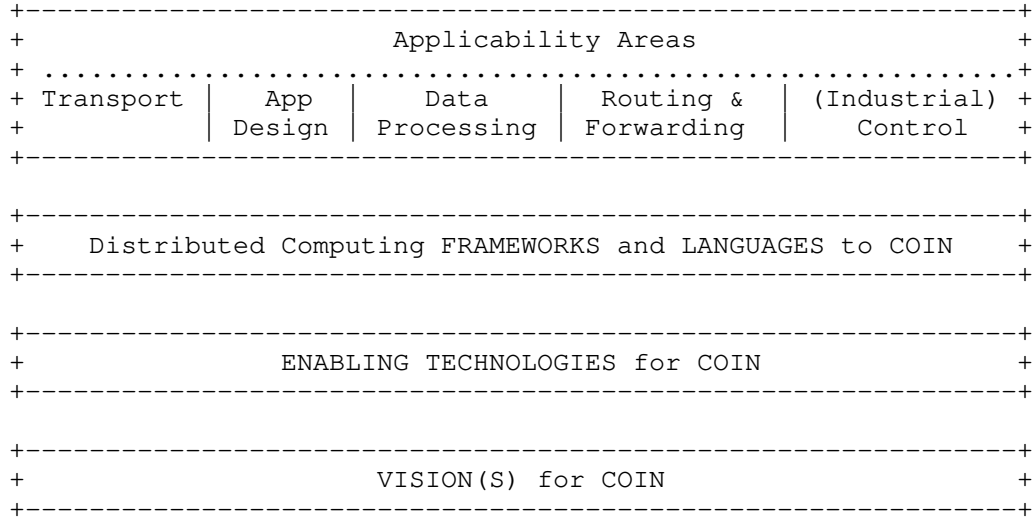


Figure 6: Research Questions Categorization

7.2.1. Categorization

Three categories deal with concretizing fundamental building blocks of COIN and COIN itself.

- * VISION(S) for COIN: Questions that aim at defining and shaping the exact scope of COIN.
- * ENABLING TECHNOLOGIES for COIN: Questions that target the capabilities of the technologies and devices intended to be used in COIN.
- * Distributed Computing FRAMEWORKS and LANGUAGES to COIN: Questions that aim at concretizing how a framework or languages for deploying and operating COIN systems might look like.

Additionally, there are use-case near research questions that are heavily influenced by the specific constraints and goals of the use cases. We call this category "applicability areas" and refine it into the following subgroups:

- * Transport:
- * App Design:
- * Data Processing:
- * Routing & Forwarding:
- * (Industrial) Control

7.2.2. Analysis

7.2.2.1. VISION(S) for COIN

The following research questions presented in the use cases belong to this category:

3.1.8, 3.2.1, 3.3.5, 3.3.6, 3.3.7, 5.3.3, 6.1.2, 6.1.4

The research questions centering around the COIN VISION dig into what is considered COIN and what scope COIN functionality should have. In contrast to the ENABLING TECHNOLOGIES, this section looks at the problem from a more philosophical perspective.

7.2.2.1.1. Where to perform computations

The first aspect of this is where/on which devices COIN programs will/should be executed (3.3.5). In particular, it is debatable whether COIN programs will/should only be executed in PNDs or whether other "adjacent" computational nodes are also in scope. In case of the latter, an arising question is whether such computations are still to be considered as "in-network processing" and where the exact line is between "in-network processing" and "routing to end systems" (3.3.7). In this context, it is also interesting to reason about the desired feature sets of PNDs (and other COIN execution environments) as these will shift the line between "in-network processing" and "routing to end systems" (3.1.8).

7.2.2.1.2. Are tasks suitable for COIN

Digging deeper into the desired feature sets, some research questions address the question of which domains are to be considered of interest/relevant to COIN. For example, whether computationally-intensive tasks are suitable candidates for (COIN) Programs (3.3.6).

7.2.2.1.3. (Is COIN)/(What parts of COIN are) suitable for the tasks

Turning the previous aspect around, some questions try to reason whether COIN can be sensibly used for specific tasks. For example, it is a question of whether current PNDs are fast and expressive enough for complex filtering operations (3.2.1).

There are also more general notions of this question, e.g., what "in-network capabilities" might be used to address certain problem patterns (6.1.4) and what new patterns might be supported (6.1.2). What is interesting about these different questions is that the former raises the question of whether COIN can be used for specific tasks while the latter asks which tasks in a larger domain COIN might be suitable for.

7.2.2.1.4. What are desired forms for deploying COIN functionality

The final topic addressed in this part deals with the deployment vision for COIN programs (5.3.3).

In general, multiple programs can be deployed on a single PND/COIN element. However, to date, multi-tenancy concepts are, above all, available for "end-host-based" platforms, and, as such, there are manifold questions centering around (1) whether multi-tenancy is desirable for PNDs/COIN elements and (2) how exactly such functionality should be shaped out, e.g., which (new forms of) hardware support needs to be provided by PNDs/COIN elements.

7.2.2.2. ENABLING TECHNOLOGIES for COIN

The following research questions presented in the use cases belong to this category:

3.1.7, 3.1.8, 3.2.2, 4.3.4, 4.4.4, 5.1.1, 5.1.2, 5.1.6, 5.3.1, 6.1.3, 6.1.4,

The research questions centering around the ENABLING TECHNOLOGIES for COIN dig into what technologies are needed to enable COIN, which of the existing technologies can be reused for COIN and what might be needed to make the VISION(S) for COIN a reality. In contrast to the VISION(S), this section looks at the problem from a practical perspective.

7.2.2.2.1. COIN compute technologies

Picking up on the topics discussed in Section 7.2.2.1.1 and Section 7.2.2.1.2, this category deals with how such technologies might be realized in PNDs and with which functionality should even be realized (3.1.8).

7.2.2.2.2. Forwarding technology

Another group of research questions focuses on "traditional" networking tasks, i.e., L2/L3 switching and routing decisions.

For example, how COIN-powered routing decisions can be provided at line-rate (3.1.7). Similarly, how (L2) multicast can be used for COIN (vice versa) (5.1.1), which (new) forwarding capabilities might be required within PNDs to support the concepts (5.1.2), and how scalability limits of existing multicast capabilities might be overcome using COIN (5.1.6).

In this context, it is also interesting how these technologies can be used to address quickly changing receiver sets (6.1.3), especially in the context of collective communication (6.1.4).

7.2.2.2.3. Incorporating COIN in existing systems

Some research questions deal with questions around how COIN (functionality) can be included in existing systems.

For example, if COIN is used to perform traffic filtering, how end-hosts can be made aware that data/information/traffic is deliberately withheld (4.3.4). Similarly, if data is pre-processed by COIN, how can end-hosts be signaled the new semantics of the received data (4.4.4).

In particular, these are not only questions concerning the functionality scope of PNDs or protocols but might also depend on how programming frameworks for COIN are designed. Overall, this category deals with how to handle knowledge and action imbalances between different nodes within COIN networks (5.3.1).

7.2.2.2.4. Enhancing device interoperability

Finally, the increasing diversity of devices within COIN raises interesting questions of how the capabilities of the different devices can be combined and optimized (3.2.2).

7.2.2.3. Distributed Computing FRAMEWORKS and LANGUAGES to COIN

The following research questions presented in the use cases belong to this category:

3.1.1, 3.2.3, 3.3.1, 3.3.2, 3.3.3, 3.3.5, 4.2.1, 4.2.2, 4.3.2/4.4.2, 4.3.3/4.4.3, 4.3.4, 4.4.4, 5.2.1, 5.2.2, 5.2.3, 5.3.1, 5.3.2, 5.3.3, 5.3.4, 5.3.5,

This category mostly deals with how COIN programs can be deployed and orchestrated.

7.2.2.3.1. COIN program composition

One aspect of this topic is how the exact functional scope of COIN programs can/should be defined. For example, it might be an idea to define an "overall" program that then needs to be deployed to several devices (5.3.2). In that case, how should this composition be done: manually or automatically? Further aspects to consider here are how the different computational capabilities of the available devices can be taken into account and how these can be leveraged to obtain suitable distributed versions of the overall program (4.2.1).

In particular, it is an open question of how "service-level" frameworks can be combined with "app-level" packaging methods (3.1.1) or whether virtual network models can help facilitate the composition of COIN programs (5.3.5). This topic also again includes the considerations regarding multi-tenancy support (5.3.3, cf. Section 7.2.2.1.4) as such function distribution might necessitate deploying functions of several entities on a single device.

7.2.2.3.2. COIN function placement

In this context, another interesting aspect is where exactly functions should be placed and who should influence these decisions. Such function placement could, e.g., be guided by the available devices (3.3.5, c.f. Section 7.2.2.1.1) and their position with regards to the communicating entities (3.3.1), and it could also be specified in terms of the "distance" from the "direct" network path (3.3.2).

However, it might also be an option to leave the decision to users or at least provide means to express requirements/constraints (3.3.3). Here, the main question is how tenant-specific requirements can actually be conveyed (5.2.1).

7.2.2.3.3. COIN function deployment

Once the position for deployment is fixed, a next problem that arises is how the functions can actually be deployed (4.3.2,4.4.2). Here, first relevant questions are how COIN programs/program instances can be identified (3.1.4) and how preferences for specific COIN program instances can be noted (3.1.5). It is then interesting to define how different COIN program can be coordinated (4.3.2,4.4.2), especially if there are program dependencies (4.2.2, cf. Section 7.2.2.3.1).

7.2.2.3.4. COIN dynamic system operation

In addition to static solutions to the described problems, the increasing dynamics of today's networks will also require dynamic solutions. For example, it might be necessary to dynamically change COIN programs at run-time (4.3.3, 4.4.3) or to include new resources, especially if service-specific constraints or tenant requirements change (5.2.2). It will be interesting to see if COIN frameworks can actually support the sometimes required dynamic changes (3.2.4). In this context, providing availability and accountability of resources can also be an important aspect.

7.2.2.3.5. COIN system integration

COIN systems will potentially not only exist in isolation, but will have to interact with existing systems. Thus, there are also several questions addressing the integration of COIN systems into existing ones. As already described in Section 7.2.2.2.3, the semantics of changes made by COIN programs, e.g., filtering packets or changing payload, will have to be communicated to end-hosts (4.3.4,4.4.4). Overall, there has to be a common middleground so that COIN systems can provide new functionality while not breaking "legacy" systems. How to bridge different levels of "network awareness" (5.3.1) in an explicit and general manner might be a crucial aspect to investigate.

7.2.2.3.6. COIN system properties - optimality, security and more

A final category deals with meta objectives that should be tackled while thinking about how to realize the new concepts. In particular, devising strategies for achieving an optimal function allocation/placement are important to effectively the high heterogeneity of the involved devices (3.2.3).

On another note, security in all its facets needs to be considered as well, e.g., how to protect against misuse of the systems, unauthorized traffic and more (5.3.4). We acknowledge that these issues are not yet discussed in detail in this document.

7.2.2.4. Applicability Area - Transport

The following research questions presented in the use cases belong to this category:

3.1.2

Further research questions concerning transport solutions are discussed in more detail in [TRANSPORT].

Today's transport protocols are generally intended for end-to-end communications. Thus, one important question is how COIN program interactions should be handled, especially if the deployment locations of the program instances change (quickly) (3.1.2).

7.2.2.5. Applicability Area - App Design

The following research questions presented in the use cases belong to this category:

4.3.1, 5.1.1, 5.1.3, 5.1.5

The possibility of incorporating COIN resources into application programs increases the scope for how applications can be designed and implemented. In this context, the general question of how the applications can be designed and which (low-level) triggers could be included in the program logic comes up (4.3.1). Similarly, providing sensible constraints to route between compute and network capabilities (when both kinds of capabilities are included) is also important (5.1.3). Many of these considerations boil down to a question of trade-off, e.g, between storage and frequent updates (5.1.5), and how (new) COIN capabilities can be sensibly used for novel application design (5.1.1).

7.2.2.6. Applicability Area - Data Processing

The following research questions presented in the use cases belong to this category:

3.2.3, 4.4.1, 4.5.2

Many of the use cases deal with novel ways of processing data using COIN. Interesting questions in this context are which types of COIN programs can be used to (pre-)process data (4.4.1) and which parts of packet information can be used for these processing steps, e.g., payload vs. header information (4.5.2). Additionally, data processing within COIN might even be used to support a better localization of the COIN functionality (3.2.3).

7.2.2.7. Applicability Area - Routing & Forwarding

The following research questions presented in the use cases belong to this category:

3.1.2, 3.1.3, 3.1.4, 3.1.5, 3.1.6, 5.1.2, 5.1.3, 5.1.4, 6.1.5,

Being a central functionality of traditional networking devices, routing and forwarding are also prime candidates to profit from enhanced COIN capabilities. In this context, a central question, also raised as part of the framework in Section 7.2.2.3.3, is how different COIN entities can be identified (3.1.4) and how the choice for a specific instance can be signalled (3.1.5). Building upon this, next questions are which constraints could be used to make the forwarding/routing decisions (5.1.3), how these constraints can be signalled in a scalable manner (3.1.3), and how quickly changing COIN program locations can be included in these concepts, too (3.1.2).

Once specific instances are chosen, higher-level questions revolve around "affinity". In particular, how affinity on service-level can be provided (3.1.6), whether traffic steering should actually be performed on this level of granularity or rather on a lower level (5.1.4) and how invocation for arbitrary application-level protocols, e.g., beyond HTTP, can be supported (6.1.5). Overall, a question is what specific forwarding methods should or can be supported using COIN (5.1.2).

7.2.2.8. Applicability Area - (Industrial) Control

The following research questions presented in the use cases belong to this category:

3.2.4, 3.3.1, 3.3.4, 4.2.1, 4.4.1, 4.5.1

The final applicability area deals with use cases exercising some kind of control functionality. These processes, above all, require low latencies and might thus especially profit from COIN functionality. Consequently, the aforementioned question of function placement (cf. Section 7.2.2.3.2, e.g., close to one of the end-points or deep in the network, is also a very relevant question for this category of applications (3.3.1).

Focusing more explicitly on control processes, one idea is to deploy different controllers with different control granularities within a COIN system. On the one hand, it is an interesting question how these controllers with different granularities can be derived based on one original controller (4.2.1). On the other hand, how to achieve synchronisation between these controllers or, more generally,

between different entities or flows/streams within the COIN system is also a relevant problem (3.3.4). Finally, it is still to be found out whether using COIN for such control processes indeed improves the existing systems, e.g., in terms of safety (4.5.1) or in terms of performance (3.2.4).

7.3. Requirements

To be added later.

8. 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 4.3.2, 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].

9. IANA Considerations

N/A

10. Conclusion

This draft presented use cases gathered from several fields that can and could profit from capabilities that are provided by in-network and, more generally, distributed compute capabilities. We distinguished between use cases in which COIN may (i) enable new experiences, (ii) expose new features or (iii) improve on existing system capabilities, and (iv) other use cases where COIN capabilities enable totally new applications, for example, in industrial networking.

Beyond the mere description and characterization of those use cases, we identified opportunities arising from utilizing COIN capabilities as well as research questions that may need to be addressed to reap those opportunities. We also outlined possible requirements for realizing a COIN system addressing these use cases.

But of course this is only a snapshot of the potential COIN use cases. In fact, the decomposition of many current client server applications into node by node transit could identify other opportunities for adding computing to forwarding notably in supply-chain, health care, intelligent cities and transportation and even financial services (amongst others). As these become better defined they will be added to the list presented here. We are, however, confident that our analysis across all use cases in those dimensions of opportunities, research questions, and requirements has identified commonalities that will support future work in this space. Hence, the use cases presented are directly positioned as input into the milestones of the COIN RG in terms of required functionalities.

11. List of Use Case Contributors

- * Dirk Trossen has contributed the following use cases: Section 3.1, Section 5.1, Section 5.2, Section 6.1.
- * Marie-Jose Montpetit has contributed the XR use case (Section 3.2).
- * David Griffin and Miguel Rio have contributed the use case on performing arts (Section 3.3).
- * Ike Kunze and Klaus Wehrle have contributed the industrial use cases (Section 4).
- * Xavier De Foy has contributed the use case on virtual networks programming (Section 5.3)

12. References

12.1. 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>>.

12.2. Informative References

[APPCENTRES]

Trossen, D., Sarathchandra, C., and M. Boniface, "In-Network Computing for App-Centric Micro-Services", Work in Progress, Internet-Draft, draft-sarathchandra-coin-appcentres-04, 26 January 2021, <<https://www.ietf.org/internet-drafts/draft-sarathchandra-coin-appcentres-04.txt>>.

[FCDN]

Al-Naday, M., Reed, M.J., Riihijarvi, J., Trossen, D., Thomos, N., and M. Al-Khalidi, "A Flexible and Efficient CDN Infrastructure without DNS Redirection of Content Reflection", <<https://arxiv.org/pdf/1803.00876.pdf>>.

[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 Annual Hawaii International Conference on System Sciences, DOI 10.24251/hicss.2019.871, 2019, <<https://doi.org/10.24251/hicss.2019.871>>.

[GRPC]

"High performance open source universal RPC framework", <<https://grpc.io/>>.

[I-D.draft-kutscher-coinrg-dir]

Kutscher, D., Kaerkkainen, T., and J. Ott, "Directions for Computing in the Network", Work in Progress, Internet-Draft, draft-kutscher-coinrg-dir-02, 31 July 2020, <<https://www.ietf.org/archive/id/draft-kutscher-coinrg-dir-02.txt>>.

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

Fink, I. B. and K. Wehrle, "Enhancing Security and Privacy with In-Network Computing", Work in Progress, Internet-Draft, draft-fink-coin-sec-priv-03, 22 October 2021, <<https://www.ietf.org/archive/id/draft-fink-coin-sec-priv-03.txt>>.

[I-D.hsingh-coinrg-reqs-p4comp]

Singh, H. and M. Montpetit, "Requirements for P4 Program Splitting for Heterogeneous Network Nodes", Work in Progress, Internet-Draft, draft-hsingh-coinrg-reqs-p4comp-03, 18 February 2021, <<https://www.ietf.org/archive/id/draft-hsingh-coinrg-reqs-p4comp-03.txt>>.

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

McBride, M., Kutscher, D., Schooler, E., Bernardos, C. J., Lopez, D. R., and X. D. Foy, "Edge Data Discovery for COIN", Work in Progress, Internet-Draft, draft-mcbride-

edge-data-discovery-overview-05, 1 November 2020,
<<https://www.ietf.org/archive/id/draft-mcbride-edge-data-discovery-overview-05.txt>>.

- [I-D.ravi-icnrg-5gc-icn] Ravindran, R., Suthar, P., Trossen, D., Wang, C., and G. White, "Enabling ICN in 3GPP's 5G NextGen Core Architecture", Work in Progress, Internet-Draft, draft-ravi-icnrg-5gc-icn-04, 31 May 2019, <<https://www.ietf.org/archive/id/draft-ravi-icnrg-5gc-icn-04.txt>>.
- [ICE] Burke, J., "ICN-Enabled Secure Edge Networking with Augmented Reality: ICE-AR.", ICE-AR Presentation at NDNCOM. , 2018, <<https://www.nist.gov/news-events/events/2018/09/named-data-networking-community-meeting-2018>>.
- [KUNZE] Kunze, I., Glebke, R., Scheiper, J., Bodenbenner, M., Schmitt, R., and K. Wehrle, "Investigating the Applicability of In-Network Computing to Industrial Scenarios", 2021 4th IEEE International Conference on Industrial Cyber-Physical Systems (ICPS), DOI 10.1109/icps49255.2021.9468247, May 2021, <<https://doi.org/10.1109/icps49255.2021.9468247>>.
- [MPI] Vishnu, A., Siegel, C., and J. Daily, "Scaling Distributed Machine Learning with In-Network Aggregation", <<https://arxiv.org/pdf/1603.02339.pdf>>.
- [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, DOI 10.1145/3229591.3229592, August 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>>.

- [Stoyanov] Stoyanov, R. and N. Zilberman, "MTPSA: Multi-Tenant Programmable Switches", ACM P4 Workshop in Europe (EuroP4'20) , 2020, <<https://eng.ox.ac.uk/media/6354/stoyanov2020mtpsa.pdf>>.
- [Sultana] Sultana, N., Sonchack, J., Giesen, H., Pedisich, I., Han, Z., Shyamkumar, N., Burad, S., DeHon, A., and B.T. Loo, "Flightplan: Dataplane Disaggregation and Placement for P4 Programs", 2020, <<https://flightplan.cis.upenn.edu/flightplan.pdf>>.
- [TRANSPORT] Kunze, I., Wehrle, K., and D. Trossen, "Transport Protocol Issues of In-Network Computing Systems", Work in Progress, Internet-Draft, draft-kunze-coinrg-transport-issues-05, 25 October 2021, <<https://www.ietf.org/archive/id/draft-kunze-coinrg-transport-issues-05.txt>>.
- [TS23.501] 501, 3gpp-23., "Technical Specification Group Services and System Aspects; System Architecture for the 5G System; Stage 2 (Rel.17)", 3GPP , 2021, <<https://www.3gpp.org/DynaReport/23501.htm>>.
- [TS23.502] 502, 3gpp-23., "Technical Specification Group Services and System Aspects; Procedures for the 5G System; Stage 2 (Rel.17)", 3GPP , 2021, <<https://www.3gpp.org/DynaReport/23502.htm>>.
- [TSN] "IEEE Time-Sensitive Networking (TSN) Task Group", <<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

Marie-Jose Montpetit
Concordia University
Montreal
Canada
Email: marie@mjmontpetit.com

Xavier de Foy
InterDigital Communications, LLC
1000 Sherbrooke West
Montreal H3A 3G4
Canada
Email: xavier.defoy@interdigital.com

David Griffin
University College London
Gower St
London
WC1E 6BT
United Kingdom
Email: d.griffin@ucl.ac.uk

Miguel Rio
University College London
Gower St
London
WC1E 6BT
United Kingdom
Email: miguel.rio@ucl.ac.uk

COINRG
Internet-Draft
Intended status: Informational
Expires: 28 April 2022

I. Kunze
K. Wehrle
RWTH Aachen University
D. Trossen
Huawei
25 October 2021

Transport Protocol Issues of In-Network Computing Systems
draft-kunze-coinrg-transport-issues-05

Abstract

Today's transport protocols offer a variety of functionality based on the notion that the network is to be treated as an unreliable communication medium. Some, like TCP, establish a reliable connection on top of the unreliable network while others, like UDP, simply transmit datagrams without a connection and without guarantees into the network. These fundamental differences in functionality have a significant impact on how COIN approaches can be designed and implemented. Furthermore, traditional transport protocols are not designed for the multi-party communication principles that underlie many COIN approaches. This document raises several questions regarding the use of transport protocols in connection with COIN.

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 28 April 2022.

Copyright Notice

Copyright (c) 2021 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	3
2. Terminology	3
3. Technology Areas	3
3.1. Addressing	4
3.1.1. Research questions and challenges	4
3.1.2. Related concepts and efforts	5
3.2. Flow granularity	6
3.2.1. Research questions and challenges	7
3.2.2. Related concepts and efforts	8
3.3. Collective Communication	8
3.3.1. Research questions and challenges	9
3.3.2. Related concepts and efforts	9
3.4. Authentication	9
3.4.1. Research questions and challenges	10
3.4.2. Related concepts and efforts	10
3.5. Security	10
3.5.1. Research questions and challenges	10
3.5.2. Related concepts and efforts	10
3.6. Transport Features	11
3.6.1. Reliability	11
3.6.2. Flow/Congestion Control	14
4. Summary of related research and standardization efforts	15
5. Gap Analysis	16
5.1. Addressing	17
5.2. Flow granularity	17
5.3. Collective Communication	17
5.4. Authentication	17
5.5. Security	17
5.6. Transport Features	17
5.6.1. Reliability	17
5.6.2. Flow/Congestion Control	17
6. Summary of Issues Identified	17
7. Security Considerations	17
8. IANA Considerations	18
9. Conclusion	18
10. Informative References	18
Authors' Addresses	22

1. Introduction

A fundamental consideration for the Internet's design is that functions can be implemented correctly and completely only with the knowledge of the applications, as formulated in [E2E]. This choice is reflected in the end-to-end (E2E) principle [RFC1958],[RFC2775] in that end-hosts perform most, if not all, relevant computations. The network only performs transparent, reasonable operations such as delivering the packets without modifying them with transport protocols designed to facilitate the direct communication between those end-hosts.

[E2E], however, does consider that "sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement". We link this consideration to the field of computing in the network (COIN), which encourages explicit computations in the network, introducing an intertwined complexity as the computations on the end-hosts depend on the functionality deployed in the network.

Such thinking, to some extent, challenges traditional ``end-to-end'' transport protocols as they are not designed to address in-network computation entities or to include more than two devices into a communication, even for inherent functionalities provided by the transport protocol. Some of the resulting problems when considering in-network computation in the context of an overall E2E problem are already presented in [I-D.draft-kutscher-coinrg-dir-02].

This draft focusses on the potential opportunities and research questions for the design of transport protocols that may assume the availability of in-network computing capabilities, including the possible collaboration with other IRTF and IETF groups, such as PAN RG, the IETF transport area in general, or the LOOPS BOF, for finding suitable solutions. In particular, the draft first describes how different aspects of transport protocols might be affected by in-network computing functionality before it is analyzed how existing transport protocols map to the identified questions and challenges.

2. Terminology

COIN element: Device on which COIN functionality can be deployed

3. Technology Areas

3.1. Addressing

The traditional addressing concept of the Internet is that end-hosts directly address each other with all computational intelligence residing at the network edges. With COIN, computations move into the network and need to be integrated into the established infrastructure. In systems where the whole network is under the control of the network operator this integration can be implemented by explicitly adjusting the communication schemes based on the COIN functionality. Considering larger scales, this approach of manually adjusting traffic patterns and applications to correctly incorporate changes made by the network is not feasible.

What is needed are ways to specify which kind of functionality should be applied to the transmitted data on the path inside the network and maybe even where or by whom the execution should take place. From an identification perspective, addressing may not only need to specify the set of functionality that is being desired but also enable to provide affinity to a member of the set of computational nodes that provide said functionality.

For instance, orchestration functionality may be implemented using an indirection mechanism which routes a packet along a pre-defined or dynamically chosen path which then realizes the desired functionality. One possibility is to directly route on service or functionality identifiers instead of sending individual packets between locator-addressed network elements

[I-D.draft-irtf-coinrg-use-cases]. While this aligns the routing more clearly with the communication between computational elements, selecting the 'right' computational endpoint (out of possibly several ones) becomes critical to the proper functioning of the overall service.

3.1.1. Research questions and challenges

1. How should end-hosts address the COIN functionality?
2. How can the treatment of the transmitted data, i.e., which COIN functionality to execute, be represented in the addressing of the request?
3. How can end-hosts direct computational requests to different computational endpoints of the same service in different network locations, i.e., decide where the COIN functionality is executed?
4. How to decide (and encode the decision) which computational endpoint to choose (from possibly several ones existing in the network)?

5. How can devices which do not implement COIN functionality be integrated into the systems without breaking the COIN or legacy functionality?

3.1.2. Related concepts and efforts

- * Segment and Source Routing (see [SPRING-WG])

Source Routing allows a sender to (partially) define the route of a packet through the network. This mechanism can be leveraged to steer the traffic along COIN nodes and thus trigger desired COIN functionality. The SPRING WG is scoped to define procedures around Segment Routing [SR], a modern variant of Source Routing for IPv6 and MPLS.

- * (Service/Network) Function Chaining/Composition (see [SFC-WG])

Service Function Chaining (SFC) describes a process to first define an ordered list of service functions (e.g., firewalls) and then steer traffic through these functions [SFC-PS]. The SFC WG is tasked with defining suitable orchestration techniques for SFC. The existing SFC architecture [SFC-Arch] and the Network Service Header [SFC-NSH] already provide fundamental mechanisms. Interpreting COIN functionality as service functions could make SFC applicable to COIN at Layer 2 and Layer 3, but also at name-based, e.g., HTTP level [RFC8677]

- * Internet services over ICN (see [ICNIP])

Work in the ICN RG [ICNRG] has generally studied the addressing of information rather than endpoints, opening up the possibility for providing information from different sources, including in-network elements, such as for caching purposes. The work in [ICNIP] utilizes the ICN capabilities to address services directly as a named entity, including IP endpoints, in order to support concepts like virtualization of service endpoints and provisioning within edge and in-network locations. The solution in [ICNIP] proposes the use of a Layer 2 path-based forwarding with service identifiers used to address the specific service endpoint.

- * Flexible Addressing (see [I-D.draft-jia-intarea-scenarios-problems-addressing])

Although the work in the INT Area of the IETF does not postulate a specific solution, it outlines a number of communication scenarios and challenges, some of which aligns with those outlined above. The companion draft [I-D.draft-jia-intarea-internet-addressing-gap-analysis] provides a

deeper gap analysis of existing solutions (the above mentioned solutions are presented here, too), identifying a number of issues that arise from the specific point solutions realized by those solutions. The authors argue for both flexibility and extensibility of addressing; key aspects that any solution addressing the research questions outlined above would benefit from.

* Semantic Routing (see
[I-D.draft-king-irtf-semantic-routing-survey])

The survey at [I-D.draft-king-irtf-semantic-routing-survey] provides an overview of efforts on addressing and routing that incorporate a semantic beyond the one defined by IPv6, covering both existing IETF solutions as well as ongoing research, defined as 'semantic routing' in the draft. The companion draft at [I-D.draft-king-irtf-challenges-in-routing] outlines a number of challenges that exist for such extension of the addressing semantic, some of which align with the issues identified in this document. More importantly, the draft discusses the possible deployment of semantic routing solutions, e.g., as an overlay or limited to a single domain (following the Limited Domain concept of [RFC8799]). Some of the challenges identified in [I-D.draft-king-irtf-challenges-in-routing] apply to a COIN environment, while not being limited to it. For instance, the intended scope of any enhanced addressing (e.g., identifying COIN elements on-path in a scenario) or the description of path characteristics that COIN traffic would need to adhere to.

3.2. Flow granularity

Core networking hardware pipelines such as backbone switches are built to process incoming packets on a per-packet basis, keeping little to no state between them. This is appropriate for the general task of forwarding packets, but might not be sufficient for COIN as information that is needed for the computations can be spread across several packets. In a TCP stream, for example, data is dynamically distributed across different segments which means that the data needed for application-level computations might also be split up. In contrast to that the content of UDP datagrams is defined by the application itself which is why the datagrams could either be self-contained or information can be cleverly distributed onto different datagrams. Summarizing, different transport protocols induce different meanings to the packets that they send out which needs to be accounted for in COIN elements as they have to know how the received data is to be interpreted. There are at least three options for this.

1. Every packet is treated individually. This maps to the capabilities of existing networking equipment.
2. Every packet is treated as part of a message. In this setting, the packet alone does not have enough information for the computations. Instead, it is important to know the content of the surrounding packets which together form the overall message.
3. Every packet is treated as part of a byte stream. Here, all previous packets and potentially even all subsequent packets need to be taken into consideration for the computations as the current packet could, e.g., be the first of a group of packets, a packet in the middle, or the final packet.

Along those options above, the question arises how shorter-term 'messages' (or transactions) of the computation should be handled compared to the often longer-term management of the network resources needed to transmit the packets across one or more such messages. For instance, error control may be best applied to the individual messages between computational endpoints, while congestion control may be applied across several messages at the level of the relation between the network elements hosting the computational endpoints. In this view, the notion of a 'flow' may separate message or transaction handling from the resource management aspect, where a flow may be divided into sub-flows (said messages or transactions) with error control being applied to those sub-flows but resource management being applied to the overall flow. Such choice of flow granularity would consequently have a significant impact on how and where computations can be performed as well as ensuring that end-hosts know who has altered the data and how.

3.2.1. Research questions and challenges

1. Which flow granularities are sensible for which scenarios and upper layer protocols?
2. How do the different flow granularities map to error and congestion control?
3. How is flow granularity used for creating affinity in, e.g., routing choices?
4. How may flow granularity information be used in COIN elements, e.g., to support routing and transport protocol realizations?

3.2.2. Related concepts and efforts

As mentioned above, flow granularities are defined in transport protocols through their semantic for the unit of transfer, which can be a 'datagram' or a 'flow'. Upper layer protocols, such as HTTP, map their application data into this semantic, resulting, for instance, in a flow of HTTP requests. Note that the flow identified by the 5-tuple for the transport connection usually also carries the reverse direction of communication, e.g., in the form of HTTP responses. The introduction of 'TCP re-use' in HTTP/1.1 introduced the capability of sending many HTTP request/response interactions in a single TCP flow. The notion of flow granularity is being used in [DYNCAST] to link the relation of one or more application level interactions to a specific service instance in deployment scenarios where more than one service instance may serve requests for a given service; [DYNCAST] refers to the problem of 'instance affinity', i.e., the need to send one or more such interactions to the same instance before being able to choose another instance (e.g., based on computing or network metrics, as suggested in [DYNCAST]). At this point of the work, the potential realization of such 'instance affinity' and the relation to transport (as well as application) protocols has not been discussed yet.

Within the concept of Service Function Chaining (SFC) [SFC-Arch], a chain of services is formed and expressed through the next service header (NSH) [SFC-NSH], which provides entries into a next hop table maintained at each Service Function Forwarder (SFF) [SFC-Arch]. Packet classification takes place at the entry point of the chain, therefore providing a notion of flow granularity where the chain is treated as the 'unit of transfer'. Chaining can take place at Layer 2 or Layer 3, but also at a name-based layer (such as HTTP), as proposed in [RFC8677].

3.3. Collective Communication

COIN scenarios may exhibit a collective communication semantic, i.e., a communication between one and more computational endpoints, as is for example illustrated by use cases in [I-D.draft-sarathchandra-coin-appcentres-04]. With this, unicast and multicast transmissions become almost equal forms of communication, as also observed in work on information-centric networking (ICN) [ICNRG].

Yet, these many-point relations may be ephemeral down to the granularity of individual service requests between computational endpoints which questions the viability of stateful routing and transport approaches used for long-lived multicast scenarios such as liveTV transmissions.

This is particularly pertinent for the transport layer where reliability and flow control among a quickly changing set of receivers is a challenging problem. The ability to divide receiver groups with the support of in-network COIN elements may provide solutions that will cater to the possible dynamics of collective communication among computational endpoints.

3.3.1. Research questions and challenges

1. How to handle ephemeral transport relations at the request level across more than one endpoint?
2. How to separate longer-term resource management from shorter-term transaction handling for, e.g., error and flow control?
3. What role could COIN elements play in improving on solutions for questions 1 and 2?

3.3.2. Related concepts and efforts

As stated above, work in the [ICNRG] has long considered multicast and unicast delivery as two communication models, realized by the same communication method that utilizes the interest-data model of ICN. The work in [ICNIP] utilizes a different approach by relying on path-based forwarding of packets identified through service-level identifiers (such as URLs but also IP addresses), where return path multicast is achieved through binary operations over the path information of incoming service requests. The utilized transport network technology is that of 5GLAN or SDN, where the latter uses an OpenFlow-compatible approach to path-based forwarding with constant state requirements for the in-network forwarders. A similar approach is used in [BIER-MC] albeit at the level of a BIER overlay network. [ICNIP] also discusses, albeit briefly only, the separation of longer-lived resource management from shorter-lived transaction handling to increase efficiency of the ephemeral return path communication at the transport level.

3.4. Authentication

The realisation of COIN legitimizes and actively promotes that data transmitted from one host to another can be altered on the way inside the network. This opens the door for foul play as all intermediate network elements – no matter if they are malicious or misbehaving by accident, COIN elements, or 'traditional' middleboxes – could simply start altering parts of the original data and potentially cause harm to the end-hosts. What is needed are mechanisms with which the receiving host can verify (a) how and (b) by whom the data has been altered on the way. In fact, these might very well be two distinct

mechanisms as one (a) only focusses on the changes that are made to the data while (b) requires a scheme with which COIN elements can be uniquely identified (could very well relate to Section 3.1) and subsequently authenticated.

3.4.1. Research questions and challenges

1. How are changes to the data within the network communicated to the end-hosts?
2. How are the COIN elements that are responsible for the changes communicated to the end-hosts?
3. How are changes made by the COIN elements authenticated?

3.4.2. Related concepts and efforts

* Proof of Transit [SFC-PoT]

The Proof of Transit concept of the SFC WG allows for proving that packets have passed a defined path. Using this concept, it could at least be possible to make sure that a packet has indeed passed the desired COIN elements. However, it does not provide means to validate which changes were made by the known nodes.

3.5. Security

Many early COIN concepts require an unencrypted transmission of data. At the same time, there is a general tendency towards more and more security features in communication protocols. QUIC, e.g., encrypts all payload data and almost all header content already inside the transport layer. This makes current COIN concepts infeasible in settings where QUIC connections are used as the COIN elements do not have access to any packet content. Using COIN thus also depends on how well security mechanisms like encryption can be integrated into COIN frameworks.

3.5.1. Research questions and challenges

To be added.

3.5.2. Related concepts and efforts

To be added.

3.6. Transport Features

Depending on application needs, different transport protocols provide different features. These features shape the behavior of the protocol and have to be taken into account when developing COIN functionality. In this section, we focus on the impact of reliability as well as flow and congestion control to create awareness for the multifaceted interaction between the transport protocols and COIN elements.

3.6.1. Reliability

Applications require a reliable transport whenever it is important that all data is transmitted successfully. TCP[TCP] provides such a reliable communication as it first sets up a dedicated connection and then ensures the successful reception of all data. In contrast, UDP[UDP] is a connectionless protocol without guarantees and COIN elements working on UDP transmissions must be robust to lost information. This is not the case for applications on top of TCP, but the retransmissions and the TCP state, which TCP uses to achieve the reliability, make packet processing for COIN more complex due to at least three reasons.

The concept of retransmissions bases on the end-to-end principle as retransmissions are performed by the sender if it has determined that the receiver did not receive the corresponding original message. Both participants can then act knowing that parts of the overall data are still missing. For simple COIN elements, which are not aware of the involved TCP states and which do not track sequence numbers, it is difficult to identify (a) that a packet in the sequence is missing and (b) that a packet is a retransmission. One question is whether COIN elements should incorporate an understanding for retransmissions on the basis of existing transport mechanisms or if a COIN-capable transport should include dedicated signals for the COIN elements.

Apart from challenges in identifying retransmissions, there is also the fact that they are sent out of order with the original packet sequence. Depending on the chosen flow granularity (see Section 3.2), COIN elements might have to hold contextual information for a prolonged time once they identify an impending retransmission. Moreover, they might have to postpone or cancel computations if data is missing and instead schedule later computations. The main question arising from this is: to what extent should COIN elements be capable of incorporating retransmissions into their computation schemes and how much additional storage capabilities are required for this?

When incorporating COIN elements into the retransmission mechanisms, it is also an interesting question whether it should be possible to request or perform retransmissions from COIN elements. Considering a setting with COIN elements that are capable of detecting missing packets and retransmission requests, it might improve the overall performance if the COIN element directly requests or performs the retransmission instead of forwarding the packet/request through the complete sequence of elements. This is especially interesting in the context of collective communication where reliability mechanisms could make use of the multi-source nature of the communication and leverage the presence of many COIN elements in the network, for instance by using network coding techniques, which in turn may benefit from COIN elements participating in the reliability mechanism. In all cases, the aforementioned storage capabilities are relevant so that the COIN elements can store enough information. The general question, i.e., which nodes in the sequence should do the retransmission, has already been worked on in the context of multicast transport protocols.

Depending on the extent of realization of the presented retransmission features, COIN elements might almost have to implement some of TCP's state to fulfil their tasks. Considering that different COIN elements have different computational and storage capacities, it is very likely that not every form of transport integration into COIN can be supported by every available COIN platform. The choice of devices included into the communication will hence certainly affect the types of transport protocols that can be operated on the COIN networks.

Another aspect to consider is the 'unit' that needs to be reliably transferred. In stream-based transport protocols, such as TCP, packets represent the smallest unit of transfer. However, different choices in the flow granularity and a possible move to larger-than-a-packet messages or transactions, as suggested in Section 3.2, might make other approaches to reliability viable that operate on the basis of such messages.

3.6.1.1. Research questions and challenges

1. What is the unit of reliable transfer?
2. How to utilize more than one computational endpoint in the reliability mechanism?
3. Should COIN elements be aware of retransmissions?
4. How can COIN elements identify missing packets or retransmissions?

5. Should COIN elements be explicitly notified about retransmissions?
6. To what extent should COIN elements be capable of incorporating retransmissions into their computation schemes?
7. How much storage capabilities are required for incorporating retransmissions?
8. How can COIN elements incorporate missing packets into their computations?
9. How to deal with state changes in COIN elements caused by data lost later in the communication chain and then retransmitted?
10. Should COIN elements be capable of requesting retransmissions/ answering retransmission requests?
11. Which devices should perform retransmissions?
12. Do COIN elements have to keep transport state?
13. How much transport state do COIN elements have to keep?

3.6.1.2. Related concepts and efforts

* Transmission Control Protocol [TCP]

TCP provides reliable, ordered, and error-checked delivery of a byte stream. As such, TCP does not allow for payload changes. This means that COIN elements could only make changes to lower header information.

* Stream Control Transmission Protocol [SCTP]

SCTP provides ensures a reliable exchange of messages. In contrast to TCP, it decouples reliability from in-order delivery and thus allows for sending messages without ordering. Additionally, it has also been extended to provide partial reliability, i.e., controlling the desired reliability on a per-message basis [SCTP-PR].

* Constrained Application Protocol [CoAP]

CoAP is a specialized protocol targeting nodes that are constrained, e.g., in terms of compute power or available bandwidth. It is message-based and distinguishes between confirmable and non-confirmable messages, i.e., similar to SCTP, allows for controlling the reliability on a per-message basis.

* User Datagram Protocol [UDP]

UDP is a message-based protocol that does not provide any guarantees regarding reliability to the application layer.

3.6.2. Flow/Congestion Control

TCP incorporates mechanisms to avoid overloading the receiving host (flow control) and the network (congestion control) and determines its sending rate as the minimum value of what both mechanisms determine as feasible for the system. This approach is based on the notion that computing and forwarding hosts are separated and is challenged by the inclusion of COIN elements, i.e., computing nodes in the network.

Flow control bases on explicit end-host information as the participating end-hosts notify each other about how much data they are capable of processing and consequently do not transmit more data as the other host can handle. This only changes if one of the end-hosts updates its flow control information.

Congestion control, on the other hand, interprets volatile feedback from the network to guess a sending rate that is possible given the current network conditions. Most congestion control algorithms hereby follow a cyclical procedure where the sending end-hosts constantly increase their sending rate until they detect network congestion. They then decrease their sending rate once and start to increase it again.

In this traditional two-fold approach, loss, delay, or any other congestion signal (depending on the congestion control algorithm) induced by COIN elements (only in case that they are the bottleneck) is interpreted as network congestion and thus accounted for in the congestion control mechanism. This means that the sending end-host may repeatedly overload the computational capabilities of the COIN elements when probing for the current network conditions instead of respecting general device capabilities as is done by flow control.

In the context of COIN, the granularity of flows may see a division into sub-flows or messages to better represent the used computational semantic as discussed in Section 3.2. This raises the question whether flow and congestion control should be applied to longer term flows (of many sub-flows or messages) or directly to sub-flows. Eventually, this could possibly lead to a separation of error control (for sub-flows) and flow control (for longer-term flows). A subsequent challenge is then how to reconcile the possible volatile nature of sub-flow relations (between computational endpoints) with the longer-term relationship between network endpoints that will see

a flow of messages between them. This is particularly pertinent in collective communication scenarios, where many forward unicast sub-flows may lead to a single multicast sub-flow response albeit only for that one response message. Reconciling the various unicast resource regimes into a single (ephemeral) multicast one poses a significant challenge.

Consequently, the question arises whether COIN elements should be able to participate in end-to-end flow control.

3.6.2.1. Research questions and challenges

1. Should COIN elements be covered by congestion control?
2. Should COIN elements be able to participate in end-to-end flow control?
3. How could a resource constraint scheme similar to flow control be realized for COIN elements?
4. How to reconcile message-level flexibility in transport relations between computational endpoints with longer-term resource stability between network elements participating in the computational scenario?

3.6.2.2. Related concepts and efforts

* Transmission Control Protocol [TCP]

TCP implements flow and congestion control. The traffic is controlled using TCP's receiver and congestion windows.

* Separation of Data Path and Data Flow

[I-D.draft-asai-tsvwg-transport-review-02] proposes to explicitly divide transport protocols into two parts: a data path and a data flow layer. Essentially, the data path layer is responsible for handling path-related tasks, such as congestion control, and as such spans multiple flows. The data flow layer on top then handles flow-related tasks such as retransmissions and flow control. As indicated by the early stage of the document, the concrete structure is still up for debate. Yet, explicitly dividing congestion and flow control could give the opportunity to devise more sophisticated approaches to incorporate COIN elements.

4. Summary of related research and standardization efforts

Issue	Efforts
Addressing	Segment and Source Routing: <ul style="list-style-type: none"> - [SPRING-WG] - Segment Routing [SR] (Service/Network) Function Chaining/Composition: <ul style="list-style-type: none"> - [SFC-WG] - SFC Problem Statement [SFC-PS] - SFC Architecture [SFC-Arch] - SFC Network Service Header [SFC-NSH] - Internet Services over IP [ICNIP]
Flow Granularity	Service Function Chaining [SFC-Arch], [SFC-NSH] Use cases and problem statement for dynamic anycast [DYNCAST]
Collective Communication	Information-centric networking [ICNRG] Internet Services over IP [ICNIP] HTTP multicast over BIER [BIER-MC]
Authentication	SFC Proof of Transit [SFC-PoT]
Reliability	Transmission Control Protocol [TCP] Stream Control Transmission Protocol [SCTP] Constrained Application Protocol [CoAP] User Datagram Protocol [UDP]
Flow/Congestion Control	[I-D.draft-asai-tsvwg-transport-review-02]

Figure 1: Related research and standardization efforts.

5. Gap Analysis

This section provides a gap analysis within the identified technology areas with respect to existing IETF solutions and ongoing efforts in transport protocols that were summarized in the previous section.

The goal of such analysis is to identify issues with those existing solutions through and within COIN environments. From the viewpoint of structuring the gap analysis, approaches such as those taken in [I-D.draft-jia-intarea-internet-addressing-gap-analysis] may be used as examples as well as direct (if suitable) input into the gap analysis performed here.

5.1. Addressing

TBD

5.2. Flow granularity

TBD

5.3. Collective Communication

TBD

5.4. Authentication

TBD

5.5. Security

TBD

5.6. Transport Features

5.6.1. Reliability

TBD

5.6.2. Flow/Congestion Control

TBD

6. Summary of Issues Identified

This section will summarize the main issues across the investigated transport technology areas of the previous section.

7. Security Considerations

COIN changes the traditional paradigm of a simple network and the corresponding end-to-end principle as it encourages computations in/by the network. Approaches designed to protect transmitted data, such as Transport Layer Security (TLS), which is even embedded into newer transport protocols like QUIC, rely on the end-to-end principle and are thus conceptually not compatible with COIN without a consistent view on how in-network compute elements would fit into the traditional end-to-end model.

Additionally, COIN elements often do not support required cryptographic functionality.

Thus, there may be a need for new security concepts specific to COIN environment that may have to be developed to allow for a secure use within COIN environments.

8. IANA Considerations

N/A

9. Conclusion

The advent of COIN may bring many new use cases, as documented in [I-D.draft-irtf-coinrg-use-cases], with promises of improved solutions for various problems. The concept of in-network computing capabilities, however, is not directly compatible with the end-to-end nature of transport protocols, thereby posing a number of key questions regarding COIN and transport protocols.

Those key questions, positioned across key technology areas for transport protocols, lead us to look at possible gaps that may be found in existing solutions when it comes to suitably supporting COIN environments and scenarios. The gap analysis performed in this document and the issues identified as a result of this analysis are positioned as possible input into shaping a research agenda for new transport protocols that have in-network computing capabilities in mind and support them securely as well as the best performance possible.

10. Informative References

- [BIER-MC] Trossen, D., Rahman, A., Wang, C., and T. T. Eckert, "Applicability of BIER Multicast Overlay for Adaptive Streaming Services", Work in Progress, Internet-Draft, draft-ietf-bier-multicast-http-response-06, 10 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-bier-multicast-http-response-06.txt>>.
- [CoAP] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [DYNCAST] Liu, P., Willis, P., and D. Trossen, "Dynamic-Anycast (Dyncast) Use Cases and Problem Statement", Work in Progress, Internet-Draft, February 2021, <<https://datatracker.ietf.org/doc/draft-liu-dyncast-ps-usecases/>>.

- [E2E] Saltzer, J., Reed, D., and D. Clark, "End-to-end arguments in system design", ACM Transactions on Computer Systems Vol. 2, pp. 277-288, DOI 10.1145/357401.357402, November 1984, <<https://doi.org/10.1145/357401.357402>>.
- [I-D.draft-asai-tsvwg-transport-review-02]
Asai, H., "Separation of Data Path and Data Flow Sublayers in the Transport Layer", Work in Progress, Internet-Draft, draft-asai-tsvwg-transport-review-02, 15 September 2021, <<https://www.ietf.org/archive/id/draft-asai-tsvwg-transport-review-02.txt>>.
- [I-D.draft-irtf-coinrg-use-cases]
Kunze, I., Wehrle, K., Trossen, D., and M. Montpetit, "Use Cases for In-Network Computing", Work in Progress, Internet-Draft, draft-irtf-coinrg-use-cases-00, 17 February 2021, <<https://www.ietf.org/archive/id/draft-irtf-coinrg-use-cases-00.txt>>.
- [I-D.draft-jia-intarea-internet-addressing-gap-analysis]
Jia, Y., Trossen, D., Iannone, L., Shenoy, N., and P. Mendes, "Gap Analysis in Internet Addressing", Work in Progress, Internet-Draft, draft-jia-intarea-internet-addressing-gap-analysis-00, 12 July 2021, <<https://www.ietf.org/archive/id/draft-jia-intarea-internet-addressing-gap-analysis-00.txt>>.
- [I-D.draft-jia-intarea-scenarios-problems-addressing]
Jia, Y., Trossen, D., Iannone, L., Shenoy, N., Mendes, P., 3rd, D. E. E., and P. Liu, "Challenging Scenarios and Problems in Internet Addressing", Work in Progress, Internet-Draft, draft-jia-intarea-scenarios-problems-addressing-01, 12 July 2021, <<https://www.ietf.org/archive/id/draft-jia-intarea-scenarios-problems-addressing-01.txt>>.
- [I-D.draft-king-irtf-challenges-in-routing]
King, D. and A. Farrel, "Challenges for the Internet Routing Infrastructure Introduced by Changes in Address Semantics", Work in Progress, Internet-Draft, draft-king-irtf-challenges-in-routing-03, 14 June 2021, <<https://www.ietf.org/archive/id/draft-king-irtf-challenges-in-routing-03.txt>>.

- [I-D.draft-king-irtf-semantic-routing-survey]
King, D. and A. Farrel, "A Survey of Semantic Internet Routing Techniques", Work in Progress, Internet-Draft, draft-king-irtf-semantic-routing-survey-02, 28 June 2021, <<https://www.ietf.org/archive/id/draft-king-irtf-semantic-routing-survey-02.txt>>.
- [I-D.draft-kutscher-coinrg-dir-02]
Kutscher, D., Karkkainen, T., and J. Ott, "Directions for Computing in the Network", Work in Progress, Internet-Draft, draft-kutscher-coinrg-dir-02, 31 July 2020, <<http://www.ietf.org/internet-drafts/draft-kutscher-coinrg-dir-02.txt>>.
- [I-D.draft-sarathchandra-coin-appcentres-04]
Trossen, D., Sarathchandra, C., and M. Boniface, "In-Network Computing for App-Centric Micro-Services", Work in Progress, Internet-Draft, draft-sarathchandra-coin-appcentres-04, 26 January 2021, <<https://www.ietf.org/internet-drafts/draft-sarathchandra-coin-appcentres-04.txt>>.
- [ICNIP] Trossen, D., Robitzsch, S., Essex, U., AL-Naday, M., and J. Riihijarvi, "Internet Services over ICN in 5G LAN Environments", Work in Progress, Internet-Draft, draft-trossen-icnrg-internet-icn-5glan-04, 1 October 2020, <<http://www.ietf.org/internet-drafts/draft-trossen-icnrg-internet-icn-5glan-04.txt>>.
- [ICNRG] "Information-Centric Networking, IRTF Research Group", <<https://datatracker.ietf.org/rg/icnrg/about/>>.
- [RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <<https://www.rfc-editor.org/info/rfc1958>>.
- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, DOI 10.17487/RFC2775, February 2000, <<https://www.rfc-editor.org/info/rfc2775>>.
- [RFC8677] Trossen, D., Purkayastha, D., and A. Rahman, "Name-Based Service Function Forwarder (nSFF) Component within a Service Function Chaining (SFC) Framework", RFC 8677, DOI 10.17487/RFC8677, November 2019, <<https://www.rfc-editor.org/info/rfc8677>>.

- [RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.
- [SCTP] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<https://www.rfc-editor.org/info/rfc4960>>.
- [SCTP-PR] Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP) Partial Reliability Extension", RFC 3758, DOI 10.17487/RFC3758, May 2004, <<https://www.rfc-editor.org/info/rfc3758>>.
- [SFC-Arch] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.
- [SFC-NSH] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <<https://www.rfc-editor.org/info/rfc8300>>.
- [SFC-PoT] Brockners, F., Bhandari, S., Mizrahi, T., Dara, S., and S. Youell, "Proof of Transit", Work in Progress, Internet-Draft, draft-ietf-sfc-proof-of-transit-08, 1 November 2020, <<https://www.ietf.org/internet-drafts/draft-ietf-sfc-proof-of-transit-08.txt>>.
- [SFC-PS] Quinn, P., Ed. and T. Nadeau, Ed., "Problem Statement for Service Function Chaining", RFC 7498, DOI 10.17487/RFC7498, April 2015, <<https://www.rfc-editor.org/info/rfc7498>>.
- [SFC-WG] "Service Function Chaining, IETF Working Group", <<https://datatracker.ietf.org/wg/sfc/about/>>.
- [SPRING-WG] "Source Packet Routing in Networking, IETF Working Group", <<https://datatracker.ietf.org/wg/spring/about/>>.
- [SR] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

- [TCP] Postel, J., "Transmission Control Protocol", STD 7,
RFC 793, DOI 10.17487/RFC0793, September 1981,
<<https://www.rfc-editor.org/info/rfc793>>.
- [UDP] Postel, J., "User Datagram Protocol", STD 6, RFC 768,
DOI 10.17487/RFC0768, August 1980,
<<https://www.rfc-editor.org/info/rfc768>>.

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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 26 April 2022

Z. Li
K. Yao
Y. Li
China Mobile
23 October 2021

A Compute Resources Oriented Scheduling Mechanism based on Dataplane
Programmability
draft-li-coinrg-compute-resource-scheduling-00

Abstract

With massive data growing in the internet, how to effectively use the compute resources has become a quite hot topic. In order to cool down the pressure in today's large data centers, some compute resources have been moved towards the edge, gradually forming a distributed Compute Force Network. Force is a physical cause which can change the state of a motion or an object. We refer the definition from physics and extend its philosophy to network that in future, the network can be a compute force which can facilitate the integration of different kinds of compute resources, no matter hardware or software, making the computation fast and effective. In this draft, we present a compute resources oriented scheduling mechanism based on dataplane programmability, which can effectively schedule and manage compute resources in the network.

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 26 April 2022.

Copyright Notice

Copyright (c) 2021 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. Conventions Used in This Document	3
2.1. Terminology	3
2.2. Requirements Language	3
3. Design	3
3.1. Network Topology	3
3.2. Mechanism Statement	5
4. Typical Way of Realization	7
5. Security Considerations	8
6. IANA Considerations	8
7. Normative References	8
Authors' Addresses	8

1. Introduction

As Moore's law has been gradually reaching its limitation, the computation of massive data and diverse computational requirements can not be satisfied by simply upgrading the computation resources on a single chip. There become an emerging trend that domain specific computation resources like GPU, DPU and programmable switches are becoming more and more popular, generating diverse use cases in the network. For example, in network computing and in memory computing. In network computing means using programmable switches or DPUs to offload network functions so as to accelerate network speed. And in memory computing means that the computer memory does not only serve as the storage, but also provide the computation. With the development of these domain specific architectures, network should serve as a force which could facilitate the integration of all these different types of computation resources, in turn forming a Compute Force Network. In CFN, how to effectively schedule these computation resources is a topic that's worthy of studying.

Current ways to do compute resources allocation include extending protocols like DNS so as to realize the awareness and scheduling of compute resources, but the management of these compute resources must be done in the centralized controller. a DNS client wants to do some computing tasks, e.g. Machine learning models training, and the

client will send a request to DNS server. Then, DNS server will inform the client which compute node is available at the moment. However, activating and deactivate this compute node to work, e.g. creating a virtual machine, is done by centralized controller, which we think is not very efficient and timely, considering massive data waits to be computed in the network. The weakness above has provoked an idea to realize the scheduling and management of compute resources by extending current routing protocols like SRv6 with the help of programmable network elements. The detailed design is presented in this draft.

2. Conventions Used in This Document

2.1. Terminology

CFN Compute Force Network

DNS Domain Name Service

SRv6 Segment Routing over IPv6

GPU Graphics Processing Unit

DPU Data Processing Unit

2.2. 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 BCP 14[RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Design

The detailed design of the mechanism is presented in this section. A typical topology will be shown below and the definition of each part of the network topology will be given, and then the whole procedure will be explained clearly the second subsection.

3.1. Network Topology

The network topology is shown in figure below where there are several major parts inside, namely consumer, computation manament node, compute node with programmable DPU, and programmable network element.

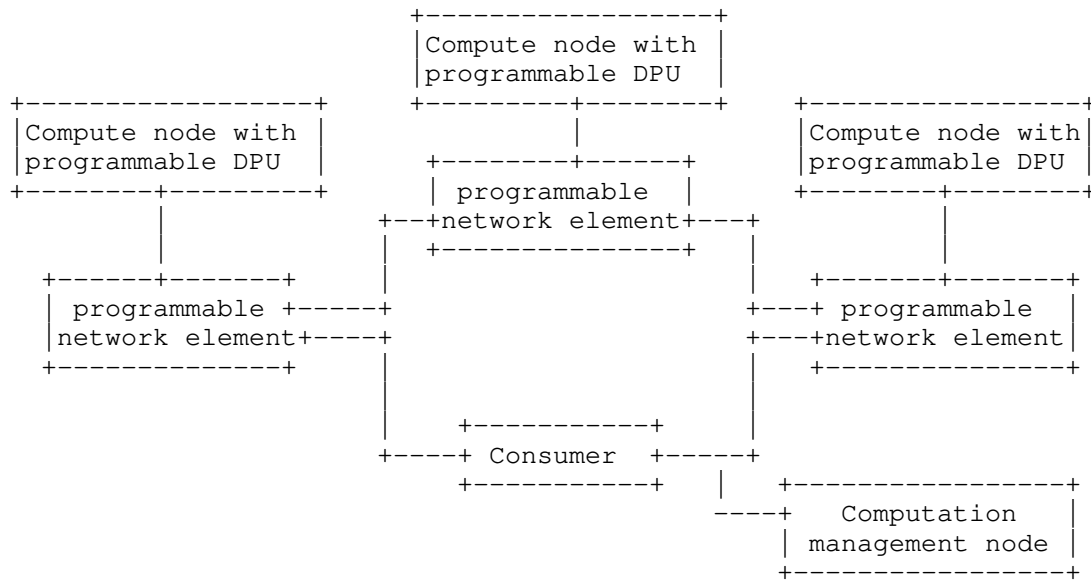


Figure 1: Figure 1: Network Topology

- Consumer: End node generating computing tasks which need to be done by compute resources
- Compute node: A network node that has the resources to finish computing tasks generated by consumers, e.g. a server or a cluster of servers.
- Programmable DPU: An unit that is connected to a compute node and a programmable element, responsible for the lifetime management of compute node and the communication with programmable element.
- Programmable network element: A network device which communicates with customers and programmable DPU, forwarding messages bidirectionally including requests for computing resources, activating or deactivating specific compute resource, and other routing messages.
- Computation management node: A network node that has the full view of the computation resources in the network, dynamically managing these resources and generate consuming receipt.

3.2. Mechanism Statement

In this section, the detailed procedure of the communication between the consumer and the compute management node which passes through programmable DPU, programmable network element, and compute node will be declared step by step .

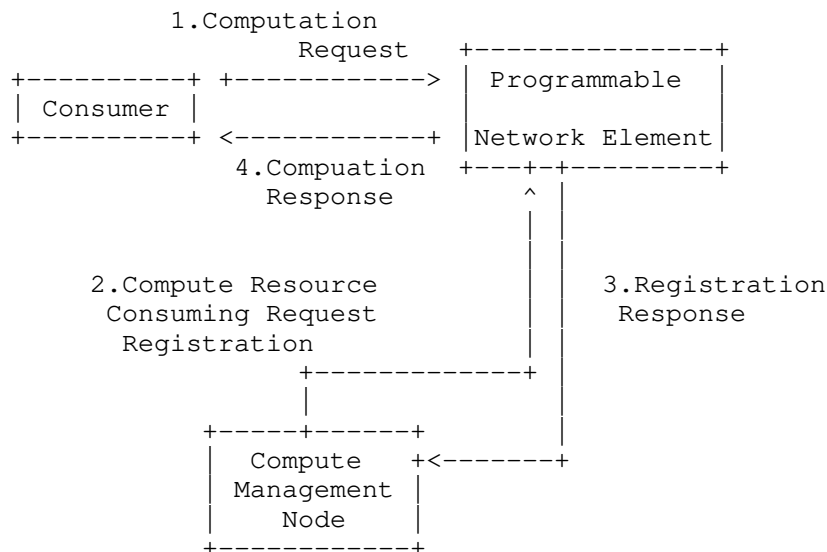


Figure 2: Figure 2: Computation Request Procedure

* Step1: computation request registration. When a consumer wants to do some computing tasks, e.g. machine learning model training, it first needs to send a request message to the compute management node for computation resource pre-allocation. The message is passed through programmable network element where some modification on the packet header can be done on the dataplane. Information like computation category, configuration template can be added into packet header, which could notify the compute management node that what kind of computation resource it needs to shedule, e.g. how many GPUs are needed in the task. Afterwards, The management node will send back a message in which the specific computation node IP address is inserted. If no such computation node is available at the moment, the managment node will send back a refusal. And at last, the programmable network element will forward the message to the consumer.

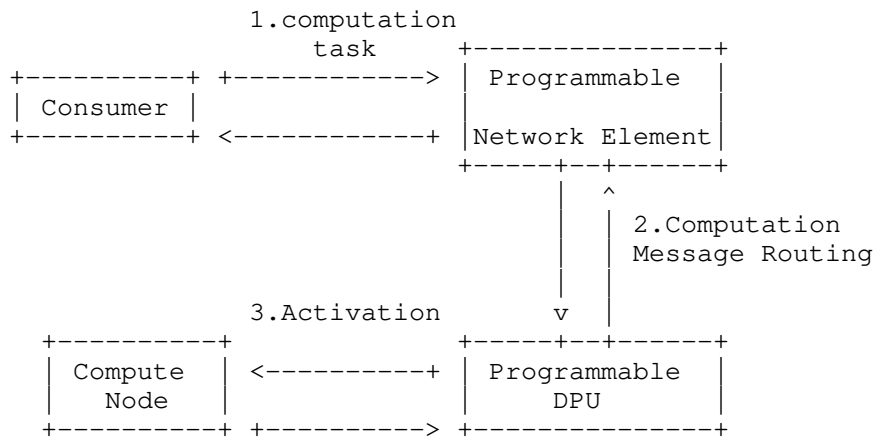


Figure 3: Figure 3: Computation Activation

* Step 2:Computation activation. Consumer will send the actual computation task to programmable network element which will do some modification on the packet. The activation message of the compute node will be encapsulated into the packet which could enable the lifetime management of the computation and the working progress of the compute node. And then, the message will be forwarded to the programmable DPU directly connected to the compute node where the decapsulation of the packet will be done. The DPU will tell the compute node to work and dynamically monitor the state of the compute node until the task is finished.

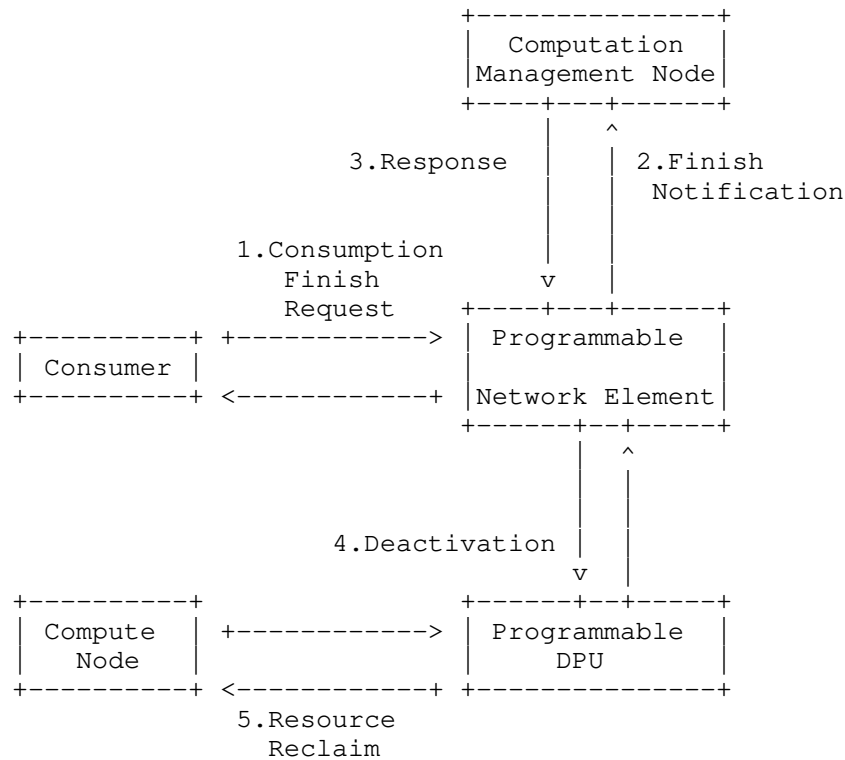


Figure 4: Figure 4: Consumption Finish

* Step 3: When the compute node notify the consumer that the task has been finished, the consumer will decide whether there is any waiting task, if not, the consumer will send a consumption finish request to the computation management node. Like computation request registration, the programmable network element will then insert information of the compute node and forward the notification message to the computation management node. when the programmable network element receives a response message, it will start deactivation procedure and tell the compute node to collect back the resource used for previous computation. This is the end the lifetime of computation of a single task.

4. Typical Way of Realization

The mechanism stated in above section can be realized by extending protocols like SRv6. The lifetime management message can be inserted dynamically in dataplane with the help of those programmable hardware. Such modification can be done flexibly and in line rate.

5. Security Considerations

TBD.

6. IANA Considerations

TBD.

7. 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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

Authors' Addresses

Zhiqiang Li
China Mobile
Beijing
100053
China

Email: lizhiqiangyjy@chinamobile.com

Kehan Yao
China Mobile
Beijing
100053
China

Email: yaokehan@chinamobile.com

Yang Li
China Mobile
Beijing
100053
China

Email: liyangzn@chinamobile.com