PhD in Information Technology and Electrical Engineering

Università degli Studi di Napoli Federico II

PhD Student: Flavio Cirillo

XXXIII Cycle

Training and Research Activities Report – Third Year

Tutor: Prof. Simon Pietro Romano

1. Information

Title: Flavio Cirillo, Master Degree in Computer Engineering, from University of Naples "Federico II" in 2014
Cycle: Ph.D. student of the XXXIII Cycle of the ITEE Course, at the University of Naples "Federico II"
Fellowship: No Fellowship
Tutor: Prof. Simon Pietro Romano

2. Study and Training Activities

External Courses

External Module	Lecturer	Date	CFU
Innovation management, entrepreneurship and intellectual property	Professor Pierlugi Rippa, University Federico II of Naples	2020-04-23 to 2020-06-11	5
IEEE Course. Cooperation in Autonomous Vehicles	Prof Dr. Alexander M. Wyglinski, Worcester Polytechnic Institute (WPI)	2020-05-04	0.3
IEEE Course. Developing and Validating Intelligent Vehicle Control Systems	Prof Dr. Alexander M. Wyglinski, Worcester Polytechnic Institute (WPI), and Associate Professor Dr. Nasser Lashgarian Azad, University of Waterloo.	2020-05-05	0.3
IEEE Course. Human Factors in Vehicle Automation	Prof Dr. Alexander M. Wyglinski, Worcester Polytechnic Institute (WPI), and Associate Professor Dr. Shan Bao, University of Michigan-Dearborn	2020-05-07	0.3
IEEE Course. Intelligent Control of Connected and Automated Vehicles	Prof Dr. Alexander M. Wyglinski, Worcester Polytechnic Institute (WPI), and Assistant Professor Dr. Yue Wang, Clemson University.	2020-05-08	0.3
IEEE Course. Object Visual Detection for Intelligent Vehicles.	Prof Dr. Alexander M. Wyglinski, Worcester Polytechnic Institute (WPI), and Full Professor Dr. Fabien Moutarde, MINES ParisTech (PSL Université Paris).	2020-05-08	0.3
IEEE Course. Object Visual Recognition for Intelligent Vehicles.	Prof Dr. Alexander M. Wyglinski, Worcester Polytechnic Institute (WPI), and Full Professor Dr. Fabien Moutarde, MINES ParisTech (PSL Université Paris).	2020-05-11	0.3
IEEE Course. Sensors for Autonomous Vehicles.	• Prof Dr. Alexander M. Wyglinski, Worcester Polytechnic Institute (WPI), and Dr. Steve Vozar, CTO and co-founder of May Mobility.	2020-05-12	0.3

IEEE Course. Enterprise Blockchain Overview	Dr. Steve Derezinski, Babson College	2020-05-13	0.3
IEEE Course. Enterprise Blockchain for Grid Modernization.	Dr. Claudio Lima, Blockchain Engineering Council.	2020-05-14	0.3
IEEE Course. Enterprise Blockchain for Healthcare	Dr. Steve Derezinski, Babson College, and Edward Bukstel, CEO of Clinical Blockchain LLC	2020-05-15	0.3
IEEE Course. Enterprise Blockchain for Supply Chain	Joseph Francis, Accenture Consulting's Communications, Media and Technology Practice.	2020-05-15	0.3
IEEE Course. Enterprise Blockchain for the Internet of Things	David Fragale, COO of Arwen, and Nancy Ranxing Li, Product Manager of Edge Computing and Blockchain at Verizon.	2020-05-18	0.3
IEEE Course. Application Scenarios of Edge Computing.	Assistant Professor Jie Cao, Eastern Michigan University, and Professor Weisong Shi, Wayne State University	2020-05-19	0.3
IEEE Course. Designing Security Solutions for Edge, Cloud, and IoT	Assistant Professor Kewei Sha, University of Houston - Clear Lake, and Professor Weisong Shi, Wayne State University	2020-05-20	0.3
IEEE Course. Overview of Edge Computing.	Professor Mahadev Satyanarayanan, Carnegie Mellon University, and Professor Weisong Shi, Wayne State University.	2020-05-21	0.3
IEEE Course. Research Challenges in Edge Computing.	Professor Songqing Chen, George Mason University, and Professor Weisong Shi, Wayne State University.	2020-05-21	0.3
IEEE Course. Tools and Software for Edge Computing Applications.	Dr. Quan Zhang, Salesforce, and Professor Weisong Shi, Wayne State University.	2020-05-21	0.3
IEEE Course. The Emerging Paradigm of the Social Internet of Things	Professor Antonio Iera, University of Calabria, Associate Professor Giacomo Morabito, University of Catania, and Associate Professor Luigi Atzori, University of Cagliari.	2020-05-22	0.3
IEEE Course. Social Internet of Things Reference Architecture and Use Cases.	Associate Professor Luigi Atzori, University of Cagliari	2020-05-25	0.3
IEEE Course. The Nature of Nudging	Professor John Sullins, Sonoma State University, and Prof. Laurence Y. Devillers, Paris-Sorbonne University	2020-05-26	0.3
3rd Advanced Course on Data Science & Machine Learning	Professor Giuseppe Nicosia, University of Catania and Professor Panos Paradalos, University of Florida	2020-07-13 to 2020-07-17	8
Digital Forensics' methods, practices and tools	Dr. Giovanni Cozzolino, University Federico II of Naples	2020-11-03 to 2020-11-10	3

Seminars

Seminar	Lecturer	Date	Place	H	CFU
Access the eLearning Library	Eszter Lukacs – IEEE Client Services Manager	2020/05/04 11:00-12:00	Online	1	0.2
Large Scale Training of Deep Neural Networks	Dr. Matthias Geihs, Giuseppe Fiameni, NVIDIA AI Technology Centre, Italy	2020/05/06 15:30-18:00	Online	2.5	0.5
IEEE Webinar, Lessons in Leadership: Preparing the Future Leaders of Your Engineering Workforce,	Braun C. Kiess, Sohaib Sheikh, and Jennie Fine	2020/05/06, 12:00-13:00	Online	1	0.2
Sensing	Maria Caterina Giordano, University of Genoa, Chiara Novara, Politecnico di Torino, Emiliano Descrovi, Politecnico di Torino, and Riccardo Sapienza, Imperial College of London	2020/05/20 14:30-18:40	Online	4	0.8

	Credits year 1								Credits year 2								Credits year 3									
		1	2	3	4	5	6			1	2	3	4	5	6			1	2	3	4	5	6			
	Estimated	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	Summary	Estimated	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	Summary	Estimated	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	Summary	Total	Check
Modules	15					6	6	12	15							0	14			11	8		3	22	34	30-70
Seminars	3	0,6	0,3	0,2	1,6			2,7	5	0,2	3,2	0,7	0,9	0,7	3,8	9,5	5			1,7				1,7	13,9	10-30
Research	42	4	4	9,7	7	12,6	8	45,3	40	10,5	6	9	6,5	11,5	7	50,5	41	5	5	5,3	6	6	9	36,3	132,1	80-140
	60	4,6	4,3	9,9	8,6	18,6	14	60	60	10,7	9,2	9,7	7,4	12,2	10,8	60	60	5	5	18	14	6	12	60	180	180

An amendment of credits has been applied: on Year 1 research credits have been shrunk from 46 to 45.3, on Year 2 research credits has been increased from 50 to 50.5.

3. Research Activities

During the third year of PhD I have worked towards advancing the IoT data analytics platforms on the path of the hyperconnected IoT. I have advanced the data access control and decentralized analytics with the data usage control that regulates how a consumer can use data respecting the data owner will. We implemented a data usage control system for streaming analytics. The system is decentralized among several domains that act as peers.

Finally, during my last months of the third year, I have written the thesis, organizing the contents I have written during my PhD to have a sturdy story and cross-linking my research subtopics together.

Data usage control

In the past decade, big data analytics has been often carried out by companies within their own individual centralized data infrastructure using existing open-source data processing frameworks such as Hadoop, Storm, Spark, and Flink. However, in the new business domains like smart cities, industry 4.0, and eHealth, data is often generated and managed by different organizations and there is a strong demand for data sharing across different organizations' domains in order to create new businesses or improve the efficiency

of existing ones. This demand triggers a fundamental change to the underlying data sharing infrastructure for supporting federated and trusted data analytics across domains. For example, Europe promotes a federated data infrastructure for joint businesses via two co-related initiatives, namely IDSA[1] and GAIA-X[2]. Due to the data protection and privacy regulation GDPR, one of the biggest open challenges to provide such a federated data sharing infrastructure is to enforce proper data usage control across different management domains regarding user-defined data usage control policies. Traditional data access control mechanisms (e.g., role-based or attribute-based) can limit data access rights by restricting who is allowed to access which data. Still, these mechanisms do not provide data providers control to regulate how data consumers can utilize their data after having access. To overcome such a limitation, data usage control[3] has been proposed as a complementary extension to traditional data access control by enforcing what must or must not happen to data even after the data is shared. However, to realize data usage control forfederated data analytics across domains, we face the following challenges:

1) Service orchestration, representing the mechanism to orchestrate the data processing pipelines of upper layer data analytics services[4] . For traditional data analytics in a cluster environment, data providers and consumers are usually the same users from the same organization. However, for federated data analytics across multiple domains, data providers and consumers are often from different organizations. In order to enable trusted data usage across domains transparently, the processing pipelines of data usage services must be made clear to both consumers and providers for federated data analytics. Therefore, it is highly desirable to have a common service model to define such data usage pipelines on a higher level. With such a common service model, we can outsource the complexity of programming data services from data providers/consumers to third-party service developers. Unfortunately, such a service model is still missing in state of the art.
2) Policy modelling, meaning how to express data usage policies from the perspective of data providers. Several policy models define data usage constraints flexibly and formally. Still, they generally represent the fine-grained low-level data usage constraints and usually require significant effort from data providers. For example, LUCON requires its users to specify data usage constraints of a data processing pipeline per flow, which is feasible only when the logic of the data processing pipeline is pre-defined and static. As pointed out by a survey[5], existing policy models are limited to express data usage control behaviors for dynamic data processing flows and also face the problem of balancing the trade-off between expressiveness and ease-of-use.

[1] IDSA, “IDS Reference Architecture Model - v2.0),” Tech. Rep., 2018.

[2] GAIA-X: A Federated Data Infrastructure for Europe,

[3] Usage Control In The Interation Data Spaces, https://www.internationaldataspaces.org/wp-content/uploads/2019/11/Usage-Control-in-IDS-V2.0 final.pdf

[4] 10. Cheng, Bin; Solmaz, Gürkan; Cirillo, Flavio; Kovacs, Ernö; Terasawa, Kazuyuki; Kitazawa, Atsushi; Fogflow: Easy programming of iot services over cloud and edges for smart cities, IEEE Internet of Things Journal, 5, 2, 696-707, 2017, IEEE

[5] A. Lazouski, F. Martinelli, and P. Mori, “Survey: Usage control in computer security: A survey,” Computer Science Review, 2010.

3) Policy enforcement, referring to the mechanism to continuously monitor and enforce pre-defined data usage constraints in the actual data usage process, which consists of three phases: before usage, ongoing usage, and after usage. Most existing data usage enforcement mechanisms can only support detective enforcement for static and pre-defined data processing flows, mainly due to their designs that fully decouple policy enforcement from service orchestration. They can monitor and check whether the orchestrated data processing flows are compliant with the defined usage policies after the data processing flows have been established; however, they cannot enforce data usage control preventively and proactively.

To overcome the limitations of existing data usage control approaches, we introduce an intent-oriented design principle to ease the modeling of both data usage control and service orchestration and, then, present a new data usage control approach called IntentKeeper (the name is inspired to ZooKeepr) to enable preventive data usage enforcement for federated data analytics in a proactive manner.

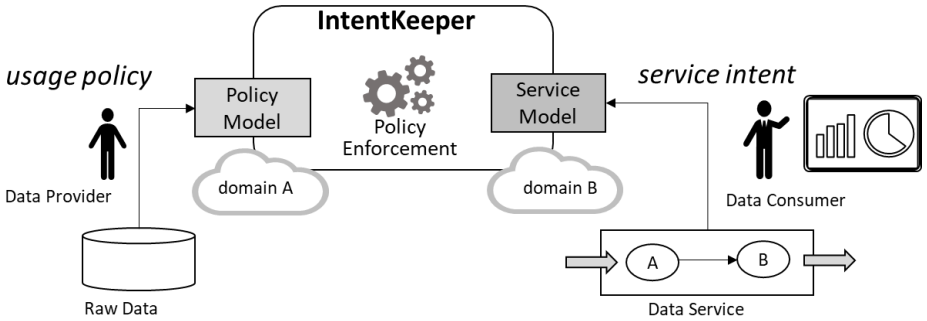

As illustrated in the figure above, the high level idea of IntentKeeper is to take two types of inputs, usage policy from the data provider and service intent from data consumer, and then translates both inputs directly into a set of orchestration actions that can set up the data processing flows compliant with the defined data usage policies.

Overall, we have brought the following contributions:

- We designed a new intent-oriented approach of modelling both usage policy and service intent to reduce the complexity for data providers and data consumers. Its effectiveness and efficiency are validated via a detailed use case study. The validation result show that our approach can reduce the specification complexity by 75% of moderately complex scenarios.
- We defined a new mechanism to enable preventive and proactive enforcement for data usage control by combining policy enforcement and service orchestration to make joint decisions. This new mechanism is realized in a decentralized way by leveraging permissioned blockchain to provide cross-domain synchronization and traceability.
- We implement the IntentKeeper system and test its performance. The experimental results show that the IntentKeeper system provides sufficiently short response times.
- We implement a demo based on an automotive use case in a smart city scenarios. We show the dynamic handling of IoT data service complying with the data usage policies of the data provider and the execution intent of the data consumer.

Collaborations:

- Prof. Simon Pietro Romano, Raffaele Porcellana, Marco Russo, University of Naples “Federico II”

- Dr. Bin Cheng, NEC Laboratories Europe, Germany,
- Dr Hisashi Sakamoto, NEC Solution Innovator, Japan

Contribution to the research community

- Serve as TPC member for IEEE Sensors Conference
- Serve as TPC member for IEEE WF-IoT Conference
- Serve as TPC member for IEEE WCNC Conference
- Serve as reviewer for IEEE Access journal
- Serve as reviewer for IEEE IoT Journal
- Serve as reviewer for Elsevier Evise Computer Networks Journal
- Serve as reviewer for Elsevier Smart Cities and Societies
- Serve as reviewer for IEEE IoT Magazine
- Serve as reviewer for IEEE Communication Standards Magazine
- Serve as reviewer for IEEE Transactions on Industrial Informatics
- Serve as reviewer for Wiley Software Practice and Experience
- Serve as reviewer for MDPI Smart Cities Journal
- Serve as reviewer for MDPI Sensors Journal

4. Products

Publications

- [Journal] **Flavio Cirillo**, David Gomez, Luis Diez, Ignacio Elicegui Maestro, Thomas Barrie Juel Gilbert, Reza Akhavan, Smart City IoT Services Creation through Large Scale Collaboration, IEEE IoT Journal.
- [Conference] **Flavio Cirillo**, Bin Cheng, Raffaele Porcellana, Marco Russo, Gurkan Solmaz, Hisashi Sakamoto and Simon Pietro Romano , IntentKeeper: Intent-oriented Data Usage Controlfor Federated Data Analytics, IEEE Local Computer Networks 2020

Patents

Filed

- US 16/384,989, **F. Cirillo**, Method and System for an Internet of Things Platform. Privacy-preserving IoT data availability announcement and regulating method for federating IoT platforms.
- US 16/892,322, **F. Cirillo**, B. Cheng, Method and System for a de-centralized Data Sharing Platform with Enforced Data Usage Control for Cross-Party Data Analytics

5. Conferences and Seminars

Attended Conferences, Seminars and Events

Event	Date	Place

ACM/IEEE IPSN 2020	2020/04/21-24	Virtually hosted at Sydney, Australia
IEEE LCN 2020	2020/11/16-18	Virtually hosted at Sydney, Australia

Presentations

Presentation	Event	Date	Place
IntentKeeper: Intent-oriented Data Usage Control for Federated Data Analytics	IEEE Local Computer Networks 2020	2020/11/16	Virtually hosted in Sydney, Australia
Edge Data Analytics	2nd IEEE Winter School on Fog/Edge Computing 2020	2020/12/29	Virtually hosted in Kanpur, India

6. Activities abroad

I have spent the 6 months, from 1st of July to 31st of December 2020, at the NEC Laboratories Europe, Heidelberg Germany.

7. Tutorship

I have tutored two Master degree students (Raffaele Porcellana and Marco Russo) to develop the project for their Master thesis.