POLITECNICO MILANO 1863
DIPARTIMENTO DI ELETTRONICA INFORMAZIONE E BIOINGEGNERIA

IT PhD

PHD IN INFORMATION TECHNOLOGY

Synthetic Information for Advisory Board, November 2019

COMPUTER SCIENCE AND ENGINEERING
ELECTRONICS
SYSTEMS AND CONTROL
TELECOMMUNICATIONS

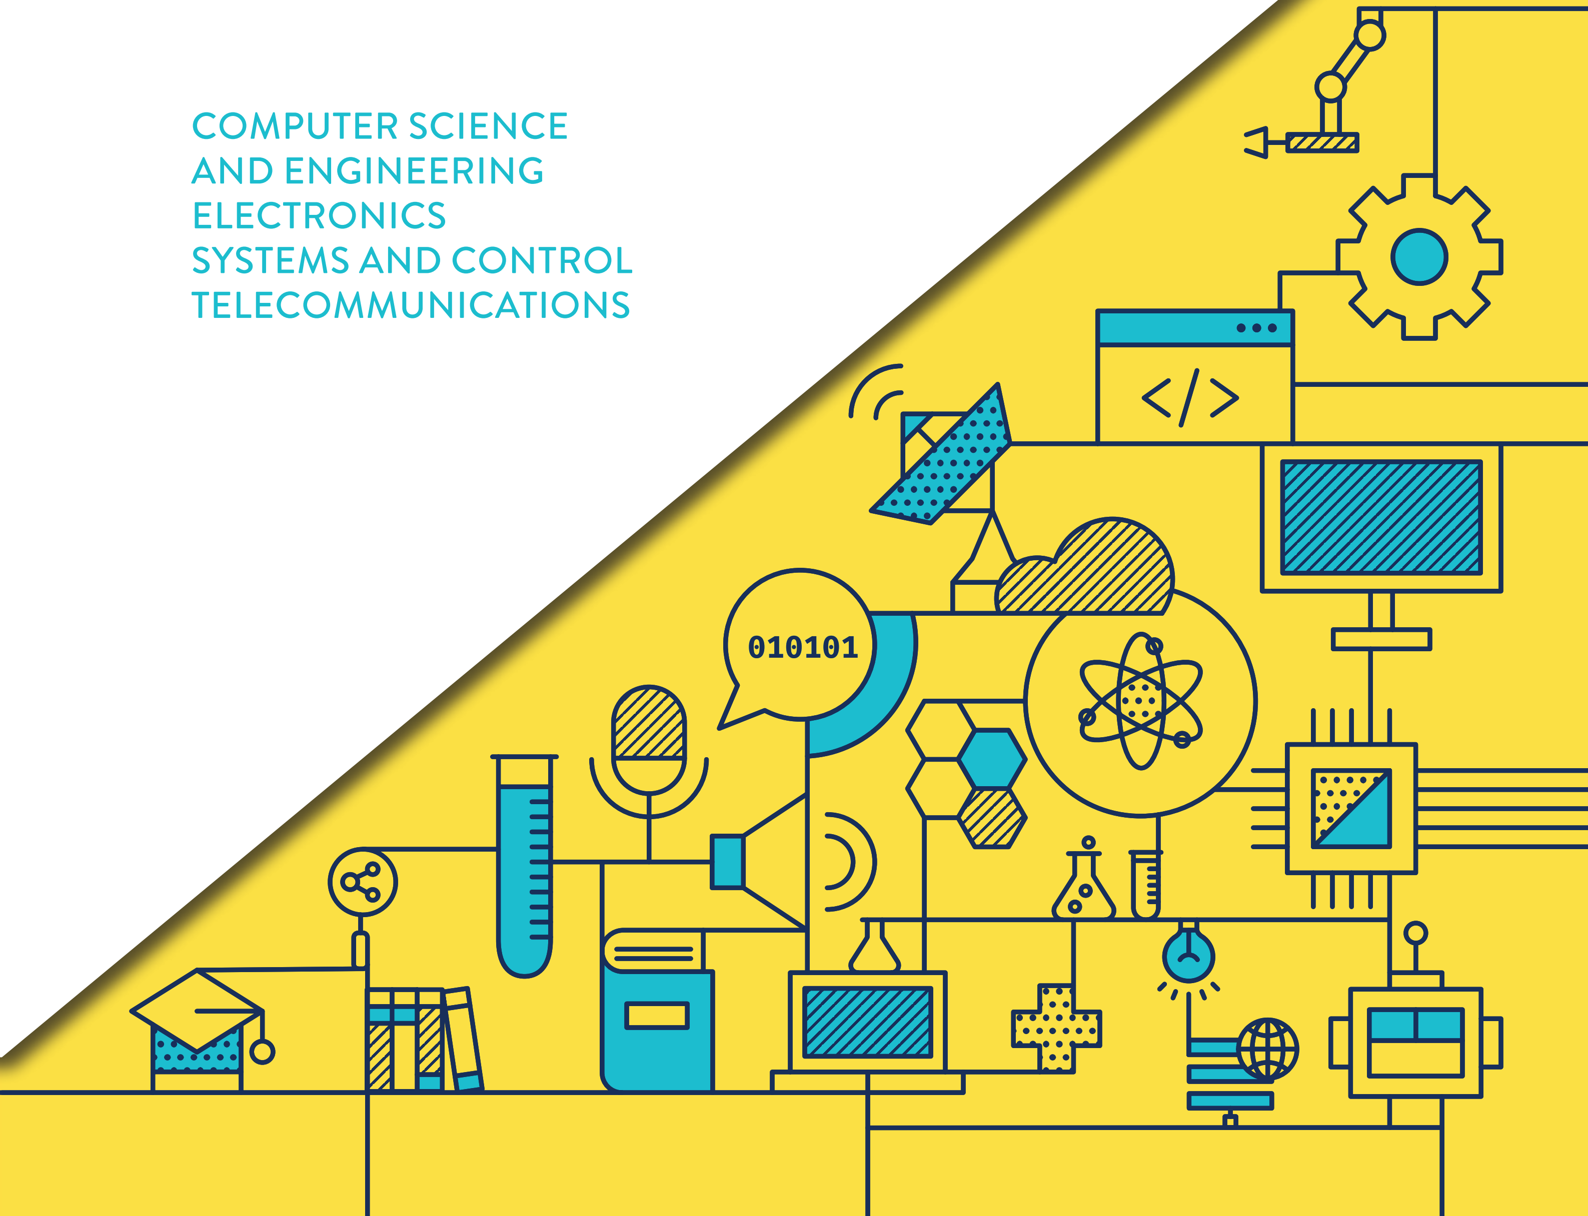

PHD IN INFORMATION TECHNOLOGY

Synthetic Information for Advisory Board, November 2019

INDEX

PREFACE

The PhD program in Information Technology (IT) is organized within the Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB) and offers the opportunity to develop a PhD thesis in the research areas of Computer Science and Engineering, Electronics, Systems and Control, Telecommunications.

Nowadays, these fields of research are of enormous scientific and technical interest to companies, governmental organizations, and the society in general. Scientific collaboration of DEIB with prestigious research institutes abroad facilitates the entrance into the world of international research through meetings with scientists, common research, and visits to laboratories worldwide.

Intense industrial collaboration of DEIB in application research allows the doctoral student to become acquainted with the activities of technologically advanced companies, thus acquiring the necessary elements for an informed choice between a career in industrial or academic organizations.

The goal of this report is to present some summary information about the status of the programme and its teaching and research activites.

PHD IN IT

THE NUMBERS

Academic year 2019 - XXXII, XXXIII, XXXIV cycles.
PhD students enrolled in the Program: 192 including 31 foreign students.

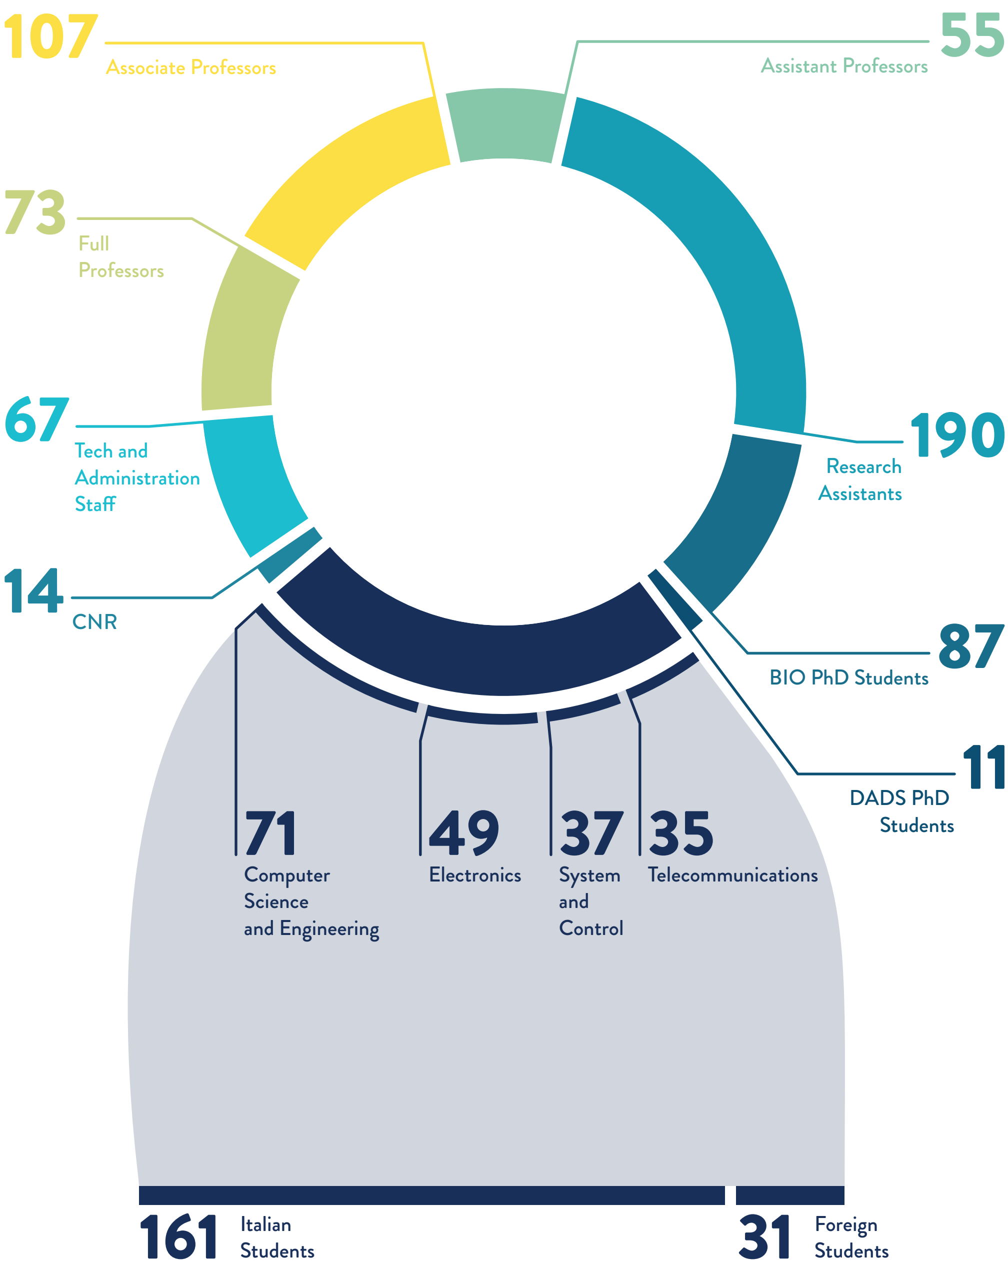

ENROLLMENT TREND

PhD student enrollment trend in the last three years

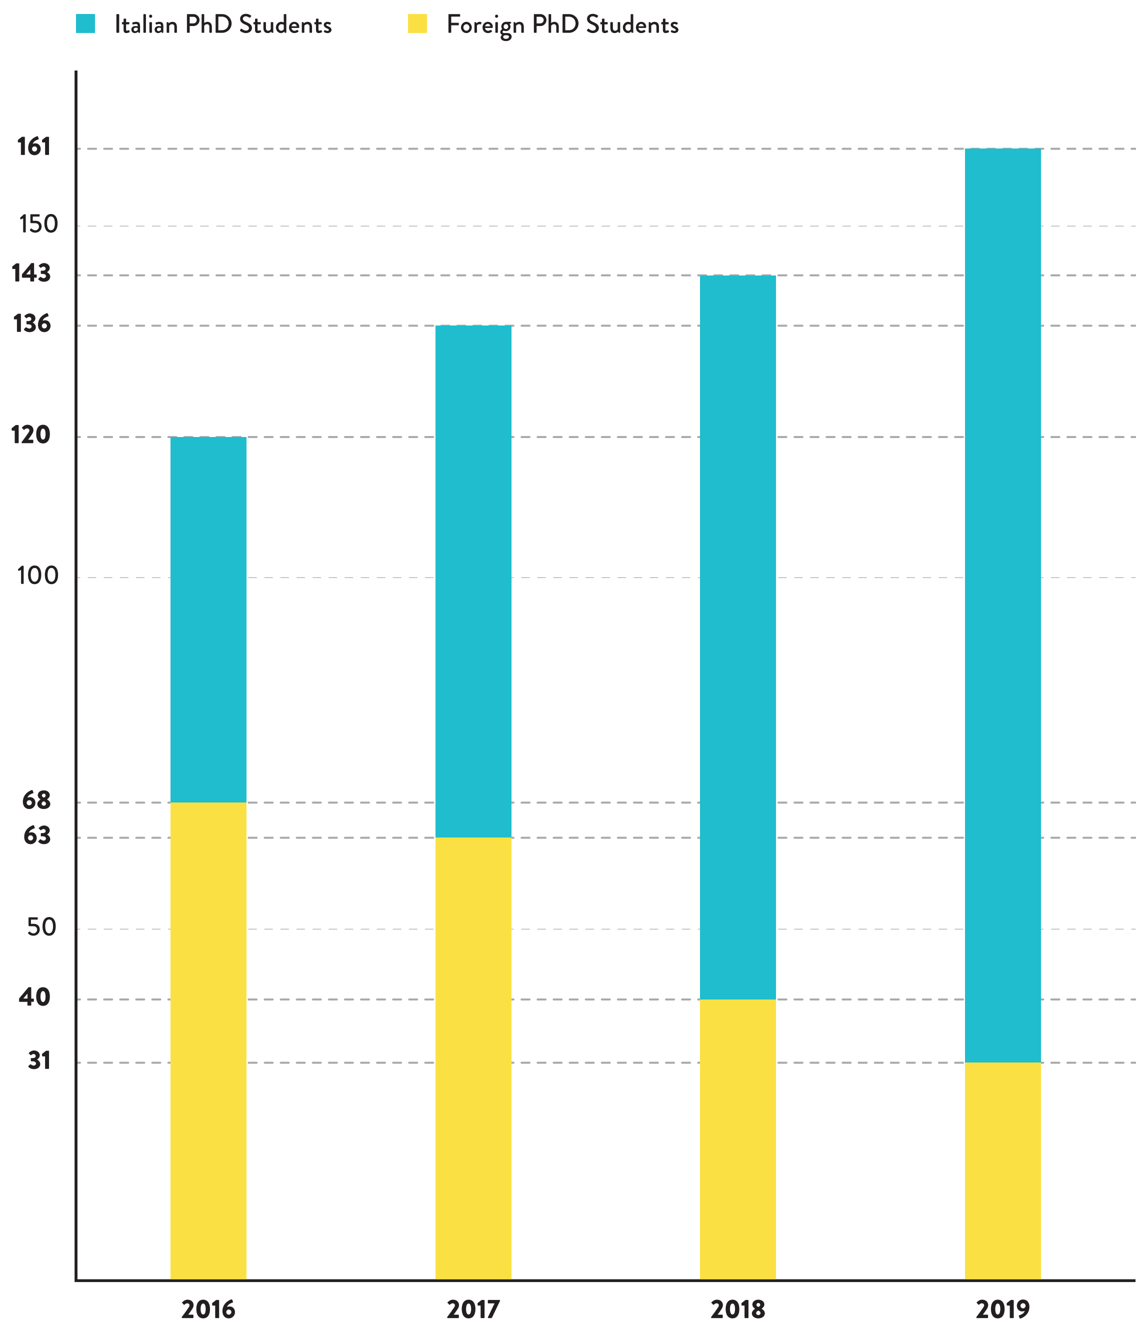

PHD IN IT

COURSES

PLANNED FOUNDATIONS COURSES FOR EACH DEIB DEPARTMENT AREA

SECTION OF COMPUTER SCIENCE AND ENGINEERING

2017/2018

- Advanced Topics on Heterogeneous System Architectures
- Automated Verification of Timed Systems
- Business Process Management
- Data and Results Visualization
- Deep Learning: Theory, Techniques and Applications
- Designing Interaction
- Engineering Complex Systems with Big Data And Information Technology- Idealeague Doctoral School
- Image Classification: Modern Approaches
- Integration and Computational Analysis of Genomic Information
- Parallel Computing Using MPI and OpenMP

2018/2019

- Advanced Topic on Reconfigurable FPGA-Based Systems Design
- Advanced Topics in Computer Security
- Advances in Deep Learning With Applications in Text and Image Processing
- Artificial Intelligence on the Web: Web Search, Recommender Systems and Crowdsourcing
- Complex Networks
- Concurrent Object-Oriented Programming
- Data and Information Quality
- Formal Languages and Automata to Model Complex Structures: Two-Dimensional and Operator Precedence Languages
- Genomic Computing
- Intelligent Multiagent Systems
- Internet Economics
- Learning Sparse Representations for Image and Signal Modeling
- Stream and Complex Event Processing in the Big Data Era
- Virtual and Mixed Reality

2019/2020

- Advances in System-on-Chip Design
- Automated Verification and Monitoring of Timed Systems
- Business Process Management
- Computer Architectures for Deep Neural Networks
- Data and Results Visualization
- Machine Learning for Non-matrix Data
- Online Learning and Monitoring
- Parallel Computing on Traditional (Core-Based) and Emerging (GPU-Based) Architectures Through Openmp and Openacc / Opencl
- Selected Topics in Cryptography

SECTION OF ELECTRONICS

2017/2018

- Advanced MEMS Gyroscopes
- Microcontrollers for Embedded Systems
- Organic Electronics: Principles, Devices and Applications
- Signal Integrity in Very-High Speed Digital Circuits
- Single-Photon Detectors for Advanced Scientific and Consumer Applications

2018/2019

- Beyond CMOS Computing
- Digital Circuits and Systems for DSP and FPGA-Based Processing
- Embedded Sensor Systems
- High Resolution Electronic Measurements in Nano-Bio Science
- Nuclear Microelectronics

2019/2020

- Advanced MEMS Gyroscopes
- Microcontrollers for Embedded Systems
- Organic Electronics: Principles, Devices and Applications
- Signal Integrity in Very-High Speed Digital Circuits

SECTION OF SYSTEMS AND CONTROL

2017/2018

- Constrained Numerical Optimization with Control Applications - Theory and Algorithms
- Cooperative and Noncooperative Optimization and Control
- Model Predictive Control
- Object-Oriented Modelling and Simulation
- Sliding Mode Control: Theory and Applications

2018/2019

- Data-Driven Approaches to Uncertain Optimization: Theory and Applications
- Hybrid Systems
- Nonlinear System Identification
- Stochastic Programming

2019/2020

- Analysis of Complex Networks: Theory and Applications
- Distributed Algorithms for Optimization and Control Over Networks
- Model Predictive Control
- Sliding Mode Control: Theory and Applications

SECTION OF TELECOMMUNICATIONS

2017/2018

- Machine Learning Methods for Communication Networks and Systems
- Molecular and Nanoscale Communication
- Network Traffic Measurement and Analysis
- Statistical Signal Processing

2018/2019

- Advanced Topics in Music Informatics
- Fiber Optic Sensing
- Location, Navigation and Cooperative Systems for Smart Mobility
- Numerical Methods for Electromagnetics

2019/2020

- Biological Communication System Design and Modelling
- Introduction to Quantum Mechanics for ICT
- Machine Learning Methods for Communication Networks and Systems
- Statistical Signal Processing in Engineering

PHD IN IT

THE ACTIVITIES

The attainment of the PhD title in Information Technology requires a study and research activity of at least three years equivalent to full time study, research and development of PhD thesis.

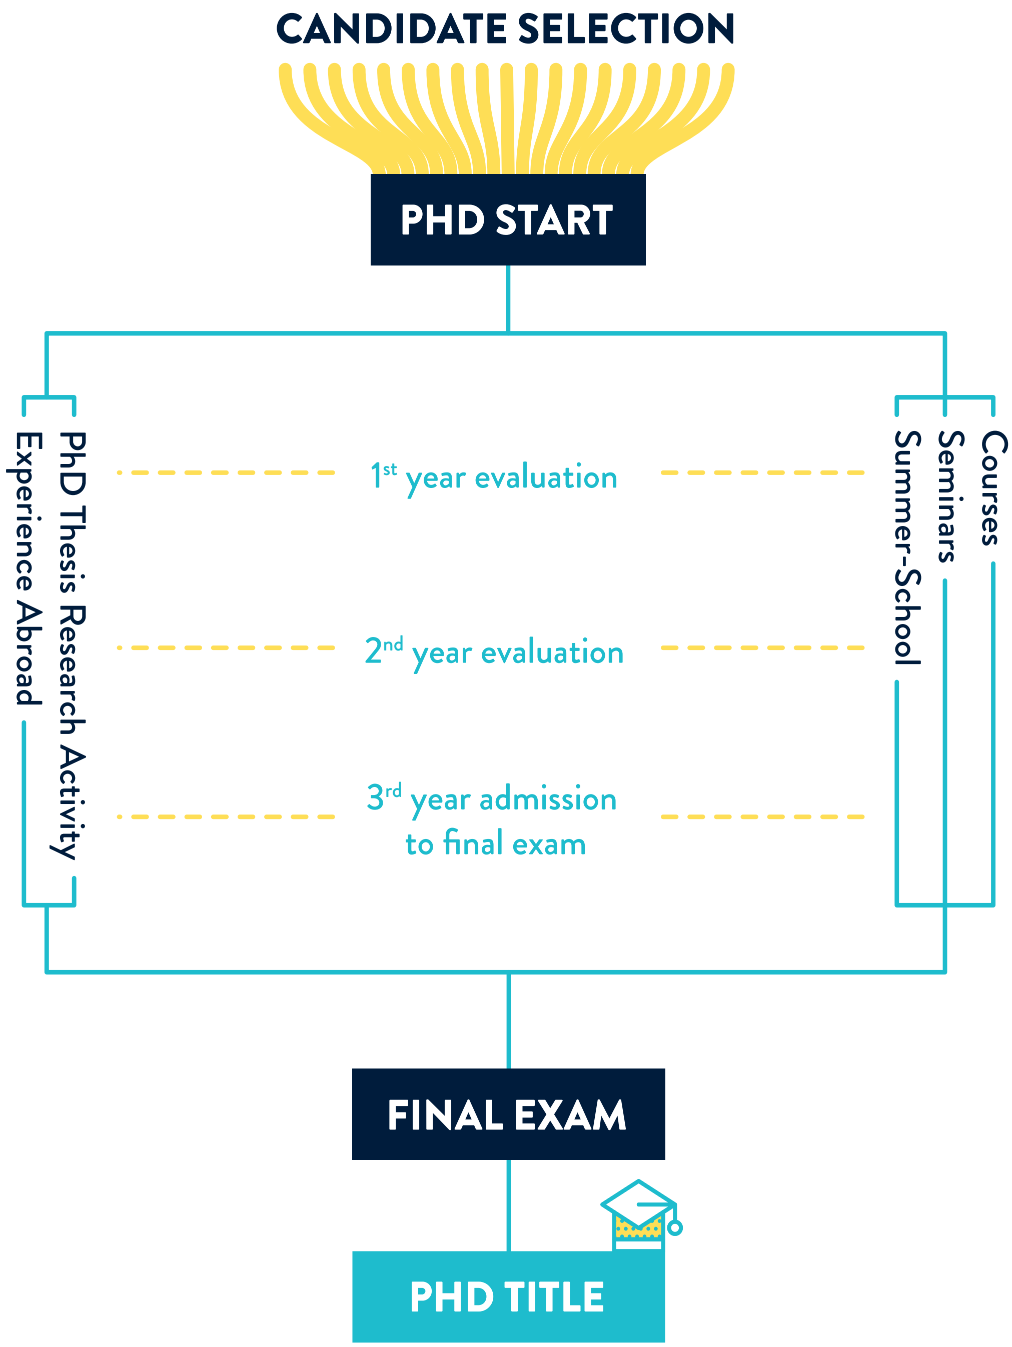

PHD IN IT

LIST OF TOPIC SCHOLARSHIPS

2018/2019

COMPUTER SCIENCE AND ENGINEERING

- Anomaly Detection in High-Dimensional and Evolving Data Streams
- Artificial Intelligence for Telecommunications with Quantum Computing
- Design of hardware accelerators for the internet of things
- Empathetic Conversational Agents
- Extracting Value from Scientific Data
- Multimodal and Reference-Based Anomaly Detection Algorithms
- Online evaluation and experimentation of UI personalization

ELECTRONICS

- 3D Printing of materials of relevance for the microelectronic industry
- Advanced ASICs for Scientific Instrumentation for Astrophysics
- Design of neuromorphic circuits with resistive memories
- Development of crosspoint memory arrays for neuromorphic computing
- Study and Development of Application Specific Integrated Circuits (ASICs) for High Resolution and High Speed X and Gamma Ray Spectroscopy

SYSTEMS AND CONTROL

- Condition Monitoring of the Built Environment with Multicopter Networks
- Systems of Tethered Multicopters: theoretical and experimental development

TELECOMMUNICATIONS

- Advanced Methods for Absolute and Relative Positioning in High-Performance Vehicles
- Advanced Techniques for the Mapping of Acoustic Sources
- Energy aware information transfer and coding solutions for green communications
- Mixed Signals Generic Testing in Photonic Integration
- Multilayer thin films for UV to LWIR operation

RESULTS ACHIEVED BY PHD STUDENTS

PUBLICATIONS

Publication of results at international level is encouraged by the PhD Program, since the presentation of the results to the national and international research communities provides an optimal forum for students to discuss research results, to get feedback, and to validate the achieved results.
An average of papers published for each PhD student at the end of their third year in the last three cycles is provided below.

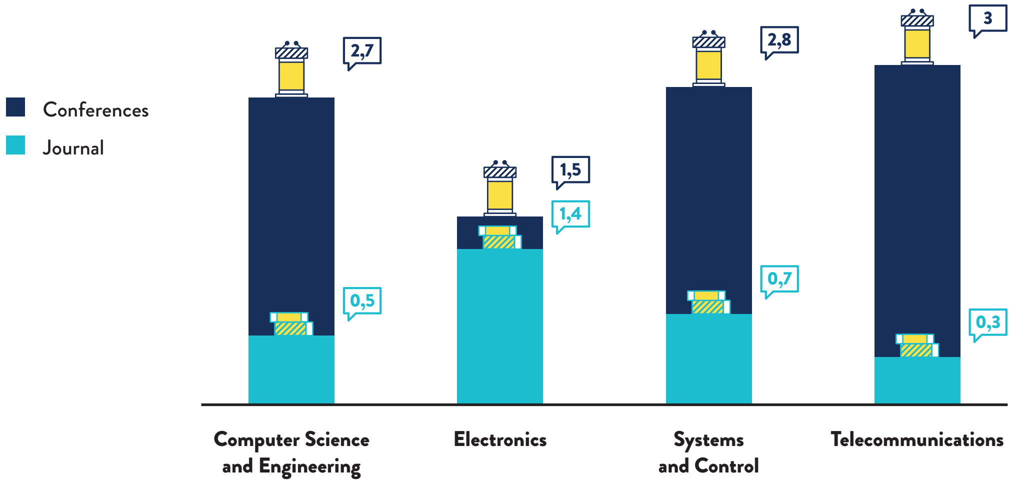

In addition to international journals and conferences, dissemination of results is also achieved through book chapters and publications.

The total number of publications by PhD students at the end of their third year in the last three cycles are listed in the following chart.

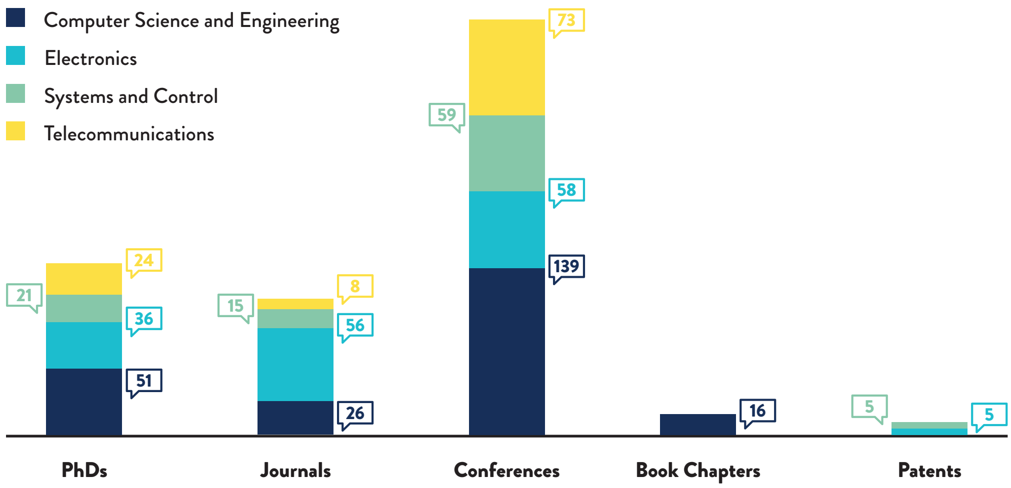

AWARDS TO PHD STUDENTS

Each year students receive awards for their scientific activities. In particular, two students from the PhD Program are awarded the Chorafas Foundation scholarship. The award is proposed by the PhD Board and evaluated by the Chorafas Foundation.
The awards received by the students in the last three years are listed below:

2017

- 17th International Conference on Web Engineering Best Student Paper Award - **Bernaschina Carlo**
- IEEE Computer Society Italy Section Chapter 2016 PhD Thesis Award - **Ashouri Amir Hossein**
- Maxeler Open Dataflow Design Competition 2017 - **Del Sozzo Emanuele**, **Pogliani Marcello**
- CAiSE Best Paper Award - **Marouan Mizmizi**
- Young Researcher Award of the International Hydropower Association (IHA) - **Schmitt Rafael**
- Chorafas Award - **Schmitt Rafael**, **Kaitoua Abdulrahman**
- 2017 Microsoft AI for Earth Azure Research Award - **Frajberg Darian**
- Innovation Design Contest 2017 - 3rd place - **Ragni Andrea**
- XILINX Open Hardware PhD Award 2017 - **Del Sozzo Emanuele**, **Rabozzi Marco**
- IFAC Young Author Prize - **Deori Luca**
- Premio Francesco Carassa 2017 - **Bernardini Alberto**

2018

- ACM CHI2018 PhD Student Best Paper Award - 3rd place - **Gelsomini Mirko**
- IEEE SSP 2018 Best Student Paper Award - **Brambilla Mattia**
- IEEE Inertial 2018 Best Student Paper Award - **Marra Cristiano Rocco**
- Outstanding Student Paper Award per Cristiano Marra - **Marra Cristiano Rocco**
- Chorafas Award - **De Nittis Giuseppe**, **Falsone Alessandro**
- RecSys Challenge 2018 – 2nd place - **Ferrari Dacrema Maurizio**
- 2018 IEEE Nuclear Science Symposium Best Student Paper Award - 2nd place - **Montagnani Giovanni Ludovico**
- 2018 IEEE Emilio Gatti and Franco Manfredi Best Ph.D. Thesis Award in Radiation Instrumentation - **Schembari Filippo**

2019

- Prof. Emilio Gatti Thesis Award issued by "Istituto Lombardo Accademia di Science e Lettere" - **Idham Hafizh**
- Young Academics Award of the Alpine Convention – 1st prize - **Zaniolo Marta**
- ECIO 2019 Best Student Poster Award - **Milanizadeh Maziyar**
- ACM International Conference on Computing Frontiers Best Poster Award 2019 - **Cherubin Stefano**
- 2019 Best Master Thesis Award "Il Quadrato della Radio" - **Brambilla Mattia**
- IFAC Young Author Award at the 15th IFAC LSS 2019 - **La Bella Alessio**
- IEEE ISCAS 2019 Best Paper Award - **Buonanno Luca, Montagnani Giovanni Ludovico**
- BalkanCom 2019 Best Student Paper Award and BalkanCom 2019 Best Paper Award - **Linsalata Francesco**
- 2019 CAiSE Ph.D. Award - **Meroni Giovanni**
- Intellectual Property Award for "ICT, Intelligenza Artificiale, IoT, Big Data, Logistica, Costruzioni" - **Pedretti Giacomo**
- Global Materialise Mimics Innovation Award issued during the ESB 2019 - **Caimi Alessandro**
- ACM Europe Council Best Paper Award - **Ferrari Dacrema Maurizio**
- Chorafas Award - **Bernardini Alberto, Cherniak Dmytro**

PHD THESES AWARDED

IN 2018 AND 2019

COMPUTER SCIENCE AND ENGINEERING

- **Affetti Lorenzo**, New Horizons for Stream Processing
- **Askarpour Mehrnoosh**, Safer-Hrc: A Methodology for Safety Assessment Through Formal Verification in Human-Robot Collaboration
- **Balduini Marco**, On the Continuous and Reactive Analysis of a Variety of Spatio-Temporal Data
- **Banfi Jacopo**, Multirobot Exploration of Communication-Restricted Environments
- **Bernaschina Carlo**, Tools, Semantics and Work-flows for Web and Mobile Model Driven Development
- **Bhatti Naveed Anwar**, System Support For Transiently-Powered Embedded Sensing Systems
- **Carrera Diego**, Learning and Adaptation to Detect Changes and Anomalies in High-Dimensional Data
- **Cenceschi Sonia**, Speech Analysis for Automatic Prosody Recognition
- **Cherubin Stefano**, Compiler-Assisted Dynamic Precision Tuning
- **Continella Andrea**, Defending from Financially-Motivated Software Abuses
- **De Nittis Giuseppe**, Patrolling Adversarial Environments Exploiting an Alarm System
- **Del Sozzo Emanuele**, On How to Effectively Target FPGAs from Domain Specific Tools
- **Deldjoo Yashar**, Video Recommendation by Exploiting the Multimedia Content
- **D'eramo Carlo**, On the Exploitation of Uncertainty to Improve Bellman Updates and Exploration in Reinforcement Learning
- **Di Federico Alessandro**, Compiler Techniques for Binary Analysis and Hardening
- **Fezzardi Pietro**, Discrepancy Analysis: a Methodology for Automated Bug Detection in Hardware Designs Generated with High-Level Synthesis
- **Filgueira Mendonça Danilo**, Self-Management of Geographically Distributed Infrastructures and Services
- **Gadioli Davide**, Dynamic Application Autotuning for Self-Aware Approximate Computing
- **Gelsomini Mirko**, Empowering Interactive Technologies for Children with Neuro-Developmental Disorders and their Caregivers
- **Geronazzo Angela**, Smart Buildings: A Methodological Approach to Data Design, Analysis and Exploitation
- **Gianniti Eugenio**, Performance Models, Design and Run Time Management of Big Data Applications
- **Guerriero Michele**, Model-Driven Engineering For Privacy-Aware Data-Intensive Applications
- **Imani Anita**, LIQDROID: a Middleware for Direct Interaction Between Multiple

Proximal Android Devices

- **Libutti Simone**, Multicore Resource Management: a horizontal perspective
- **Lopes Silva de Oliveira Ewerton**, Learning Models to Optimize the Player Experience in Robogames
- **Marconi Francesco**, Formal Verification of Timed Properties for Data-Intensive Applications
- **Meroni Giovanni**, Artifact-driven Business Process Monitoring
- **Natale Giuseppe**, On How to Design Optimized Spatial Architectures: from Iterative Stencils to Convolutional Neural Networks
- **Pagani Alessio**, Algorithms and Methods for the Design and Development of Intelligent, Context Aware and Sustainable Mobility Services
- **Paladino Stefano**, A Learning Approach for Pricing in e-Commerce Scenario
- **Peng Xuesong**, Monitoring Data Utilization in Fog Computing Environments
- **Perna Stefano**, Data-Driven Techniques for Knowledge Discovery in Regulomics
- **Pinciroli Riccardo**, Energy Efficiency in Large Data-Centers Using Performance Evaluation Techniques
- **Quarta Davide**, Embedded System Security: Attacks, Impacts & Defenses
- **Quattrocchi Giovanni**, Fast and Fine-Grained Elastic Resource Provisioning for Modern Software Systems
- **Rabozzi Marco**, CAOS: CAD as an Adaptive Open-platform Service for High Performance Reconfigurable Systems
- **Ramicic Mirza**, Perception as Behaviour Inducing Mechanism: a Reinforcement Learning Perspective
- **Riva Alessandro**, Development and Analysis of Algorithms for Information-Gathering in Autonomous Mobile Robotics
- **Sadeghi Mersedeh**, A Model-Centered Solution for Taming the Heterogeneity of Smart Devices
- **Scolari Alberto**, Optimizing Data-Intensive Applications for Modern Hardware Platforms
- **Tateo Davide**, Building Structured Hierarchical Agents
- **Zahmatkesh Shima**, On Relevant Query Answering over Streaming and Distributed Data

ELECTRONICS

- **Amirkhani Aidin**, Design of Analog Asics for X-Ray Detectors
- **Bellotti Giovanni**, Silicon Drift Detectors and Readout Electronics for High Throughput Spectroscopy Applications
- **Bricalli Alessandro**, Fabrication and Characterization of Resistive Switching Memory Devices for High-Density Storage and In-Memory Computing
- **Ceccarelli Francesco**, Development of Custom-Technology Single-Photon Avalanche Diode Arrays for High-Performance Applications
- **Cesarini Matteo**, Fully Printed Organic Imagers on Flexible Substrates for Large Area Applications and Novel Radiation Detectors
- **Cherniak Dmytro**, Digitally-Intensive Frequency Modulators for Mm-Wave Fmcw Radars
- **Cominelli Alessandro**, High-Speed, Low-Distortion Solutions for Time-Correlated Single Photon Counting Measurements
- **Cozzi Giulia**, Development of Scintillation Detectors Based on Silicon Photomultipliers for High-Energy Gamma-Ray Applications

- **Gandola Massimo**, Low-Noise, Low-Power Front-End Asics for High-Resolution X and Gamma Ray Spectroscopy for Radiation Semiconductor Detector
- **Giorgio Michele**, High-Resolution Direct-Written Field-Effect Transistors for High Frequency Applications
- **Grande Andrea**, Design of Ultra-Fast Front-End Electronics for New Pixel Detec-Tors for the European Xfel
- **Grimaldi Luigi**, Frequency Synthesizers Based on Digital Plls for Cellular Radio Applications
- **Guglielmi Emanuele**, Electronics Boosts Photonics: Detector and Electronic Design for Non-Invasive Monitoring and Control of Silicon Photonic Systems
- **Laudato Mario**, Study of Novel Devices for Crosspoint Memory and Neuromorphic Applications
- **Lusardi Nicola**, Advanced Methods, Techniques and Digital Architectures for High Performance Timing of Events
- **Marra Cristiano Rocco**, Time-Switched Frequency-Modulation for Low-Offset-Drift, Wide Range, Fully Integrated 3-Axis Mems Accelerometers
- **Milo Valerio**, Modeling and Simulation of Spiking Neural Networks with Resistive Switching Synapses
- **Minotti Paolo**, Towards Fully-Integrated Frequency-Modulated Mems Gyroscopes
- **Nicosia Gianluca**, Performance and Reliability Issues of Nand Flash Cells at the Transition from Planar to 3-D Array Architectures
- **Peronio Pietro**, Time-Correlated-Single-Photon-Counting Systems: Challenging the Limits
- **Portaluppi Davide**, Microelectronics and Instrumentation for Single-Photon Imaging
- **Resnati Davide**, Physical Modeling of Nanoscale Nand Flash Memory Reliability
- **Scuratti Francesca**, Printed Carbon Nanotubes Based Transistors: from Charge Transport Studies to Biosensing Applications
- **Tajfar Alireza**, High Speed and Intensity Laser Diode Drivers for Time-of-Flight 3D Ranging and Pico-Projectors
- **Vo Tuan Minh**, A Study of High-Performance Frequency Synthesizer Based on Bang-Bang Digital Phase-Locked Loop for Wireless Applications

SYSTEMS AND CONTROL

- **Amodio Alessandro**, Automatic Systems for Unsafe Lane Change Detection and Avoidance
- **Brankovic Aida**, Distributed Randomized Model Selection for Nonlinear Identification and Supervised Machine Learning
- **Busnelli Fabio**, Stability Control and Analysis of Two-Wheeled Vehicles out of Plane Dynamics
- **Colombo Tommaso**, Analysis and Design of Suspension Control Systems for Off-Highway Vehicles
- **Dao Le Anh**, Microgrids Energy Management with aHierarchical Distributed Model Predictive Control Approach
- **D'avico Luca**, Analysis and Design of Advanced Anti-Lock Braking Systems
- **Falsone Alessandro**, Distributed Decision Making with Application to Energy Systems
- **Farooqi Hafsa**, Design of Collaborative Eco-Drive Control Algorithms for Train Networks
- **Galluppi Olga**, Innovative Approaches to the Lateral Control Problem in Cars
- **Marelli Stefano**, Analysis and Development of Electrochemical Model-Based State Estimation Algorithms for Li-Ion Batteries

- **Mason Emanuele**, Beyond Full Rationality: Modeling Tradeoff Dynamics in Multi-Objective Water Management
- **Meddouri Soufiane**, Modeling and Control of an Autonomous Wind Energy Generator Based on a Saturated Squirrel Cage Induction Machine
- **Nicolis Davide**, A General Framework for Shared Control in Robot Teleoperation with Force and Visual Feedback
- **Pancerasa Mattia**, Data Analysis and Models for Long-Distance Bird Migration from Correlations on Ring Recoveries to Machine Learning on Geolocator Measurements
- **Parigi Polverini Matteo**, Novel Contributions to Robot Force Control for Industrial Manipulators
- **Pollok Alexander Josef**, Modelling and Control of Aircraft Environmental Control Systems
- **Rallo Gianmarco**, Robustness in Data-Driven Control: Theory and Automotive Applications
- **Rastegarpour Soroush**, Model Predictive Control Approaches for Energy Efficiency in Buildings with Heat Pump qnd Storage Systems
- **Recanati Francesca**, Sustainable Design and Management of Agroecosystems-Integrating Ecological Models and Optimization Techniques to Support Decision-Making
- **Roselli Federico**, Vehicle Dynamics Planning and Control for Safety and Comfort in Autonomous Cars
- **Sabatini Stefano**, An Autonomous Navigation Use Case for Last Mile Delivery in Urban Environments
- **Sakcak Basak**, Optimal Kinodynamic Planning for Autonomous Vehicles
- **Zhang Xinglong**, Hierarchical and Multilayer Control Structures Based on Mpc For Large-Scale Systems

TELECOMMUNICATIONS

- **Akgül Özgür Umut**, Real Time Trading of Mobile Resources in Beyond-5G Systems
- **Alvarez Villanueva Maria Antonieta**, Distributed Synchronization for Dense Wireless Network Systems
- **Aslani Hamid**, Printed Patch Antennas for Modern Wireless Communication Systems
- **Ayoub Omran**, Resource Management and Planning in Cloud-Enabled Optical Metro-Area Networks
- **Barzegar Hamid Reza**, Wireless Partial-Duplex Communications
- **Bernardini Alberto**, Advances in Wave Digital Modeling of Linear and Nonlinear Systems
- **Bondi Luca**, Data-Driven And Handcrafted Features For Forensics Analysis And Source Attribution
- **Combi Lorenzo**, Radio-relaying over Optical Fiber for Cloud Radio Access Networks
- **Dev Kapal**, Synthesis of Attenuation Time Series Across Fso Link for the Design of a Backhauling Solution for 5G Networks
- **Hmaity Ali**, Reliability Strategies for Next Generation Cloud Networks
- **Khosronejad Misagh**, Study of Level Measurement Pulse Radar Systems
- **Kumar Atul**, Synchronization and Performance Evaluation of Future Wireless Cellular System Based on the Use of New Multi-Carrier Transmission Techniques
- **Leanza Antonio**, Decorrelation Phenomena in a Geosynchronous Synthetic Aperture Radar: Theory, Techniques and Performance

- **Matera Andrea**, Interference Mitigation Techniques in Hybrid Wired-Wireless Communications Systems for Cloud Radio Access Networks with Analog Fronthauling
- **Memon Faisal Ahmed**, Silicon Oxycarbide Technology for Integrated Photonics
- **Mizmizi Marouan**, Network Sensing and Context Awareness
- **Morosi Jacopo**, Phase-Coded Brillouin Optical Correlation Domain Analysis for Fast And Highly-Resolved Distributed Strain and Temperature Monitoring
- **Oliveira Morais de Aguiar Douglas**, Reconfigurable Photonic Integrated Circuits For High Capacity Optical Networks
- **Peserico Nicola**, Integrated Optical Platform for Biosensor Applications
- **Setragno Francesco**, Feature-Based Analysis and Modelling of Violin Timbre
- **Shahid Saleem**, Design and Implementation of Microwave and Terahertz Material Characterization Mehtods
- **Sharma Navuday**, Increasing Capacity of Wireless Networks Through Aerial Base Stations
- **Shehata Mohamed Khaled**, Energy Efficiency and Survivability in 5G Centralized Access Networks
- **Tagliaferri Dario**, Visible Light Communications for Next Generation in-Flight Systems
- **Waqas Abi**, Stochastic Methods for Performance Prediction of Photonic Integrated Circuits

EMPLOYMENT DATA FOR PHD

The 2018 official employment survey of Politecnico di Milano, by the Career Service, targeter PhD graduates of 2015 and 2016.
PhD graduates in Information Technology at Politecnico di Milano were 103 in the two years considered. Total interviews: 80, accounting for 78% of the total number.

PhD graduates

103

Were employed 12 months after graduation

100%

Were employed before end of Phd

95%

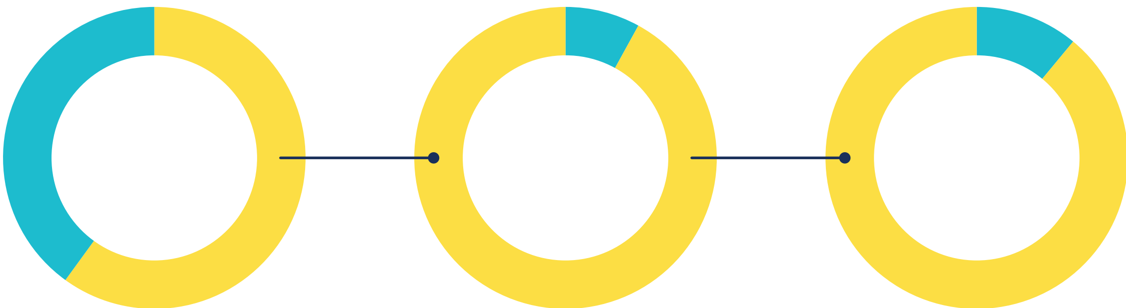

- Companies 60%
- Universities 40%

- Employed 92%
- Self-employed 8%

- Permanent contract 89%
- Self-employed 11%

Net monthly salary

€ 2.318

PhD is required for their job

78%

Would enroll again at Polimi

83%

ADVISORY BOARD

- Giorgio Ancona, *Atos*
- Matteo Bogana, *Cleafy*
- Mario Caironi, *IIT*
- Paolo Cederle, *Everis*
- Luigi Cicchese, *Concept Reply*
- Cristina Cremonesi, *The European Ambrosetti*
- Riccardo De Gaudenzi, *European Space Agency*
- Giuseppe Desoli, *STMicroelectronics*
- Alessandro Ferretti, *Tre-Altamira*
- Giuseppe Fogliazza, *MCE Srl*
- Bruno Garavelli, *Xnext s.r.l.*
- Sabino Illuzzi, *Prospera*
- Renato Marchi, *KPMG*
- Renato Lombardi, *Huawei Technologies*
- David Oliva Uribe, *Head Industrial Doctoral School, EIT Digital*
- Giorgio Parladori, *SM Optics srl*
- Francesco Prelz, *INFN*
- Enrico Ragaini, *ABB S.p.A.*
- Paolo Giuseppe Ravazzani, *CNR*
- Fabrizio Renzi, *IBM*
- Carlo Sandroni, *RSE S.p.A.*
- Massimo Valla, *TIM*
- Luisa Venturini, *Vodafone Italy*
- Stefano Verzura, *Huawei Technologies*
- Roberto Villa, *IBM Italy*

BOARD OF PROFESSORS

CHAIR OF THE PHD PROGRAM

» Barbara Pernici, barbara.pernici@polimi.it

AREA 1: COMPUTER SCIENCE AND ENGINEERING

» Cristina Silvano - Vice Chair Computer Science and Engineering, cristina.silvano@polimi.it
» Cesare Alippi
» Francesco Amigoni
» Florian Daniel
» Nicola Gatti
» Raffaela Mirandola
» Barbara Pernici
» Matteo Pradella
» Letizia Tanca

AREA 2: ELECTRONICS

» Angelo Geraci - *Vice Chair Electronics*, angelo.geraci@polimi.it
» Giuseppe Bertuccio
» Carlo Fiorini
» Ivan Rech
» Alessandro Sottocornola Spinelli

AREA 3: SYSTEMS AND CONTROL

» Luigi Piroddi - *Vice Chair Systems and Control*, luigi.piroddi@polimi.it
» Luca Bascetta
» Andrea Castelletti
» Fabio Dercole
» Lorenzo Fagiano
» Simone Garatti

AREA 4: TELECOMMUNICATIONS

- Matteo Cesana - *Vice Chair Telecommunications*, matteo.cesana@polimi.it
- Paolo Martelli
- Andrea Monti Guarnieri
- Carlo Riva
- Massimo Tornatore

REPRESENTATIVE FROM THE BIOENGINEERING AREA

- Giancarlo Ferrigno

THESES DEFENDED IN 2019

COMPUTER SCIENCE AND ENGINEERING

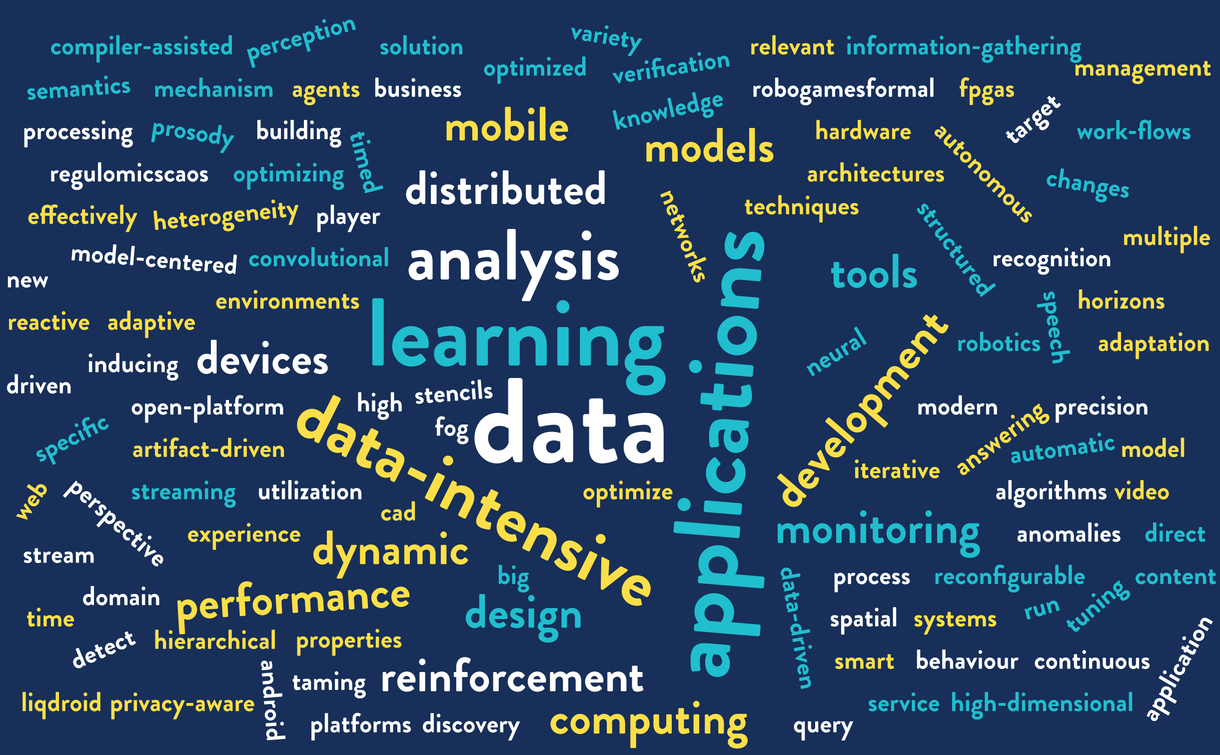

LORENZO AFFETTI

LEARNING MODELS TO OPTIMIZE THE PLAYER EXPERIENCE IN ROBOGAMES

Stream processing has gained tremendous attention over the last years and many Stream Processors (SPs) have been designed and engineered to cope with huge volumes of data coming at high velocity. Streams could contain stock options, user clicks in web applications, customer purchases in an e-commerce application, positions of robots in a warehouse, or temperature measurements from sensors. The common requirement for streaming applications is to process unbounded streams of elements and continuously compute queries like "what is the top purchased product?", or "what was the average temperature in the server room in the last second?" in order to take rapid compensating actions such as ordering a new stock of the top purchased product, or prevent fire in the server room. In order to continuously process huge amounts of elements and take real-time decisions, SPs exploit the computational power offered by multiple machines by distributing the computation and dividing data in shared-nothing partitions to avoid expensive data race management while processing. Stream processing is also a programming paradigm suited for designing novel event-driven applications with high throughput and low-latency requirements. Streams offer decoupling among the processing modules and, thus, enhance application modularity and composability. Indeed, SPs are playing a central role in the technology stacks of modern companies and they are covering more and more tasks that, in standard deployments, compete to other tools. The employment of one system instead of multiple ones reduces system integration complexity, communication latency, and facilitates application maintenance and modeling. Novel event-driven applications require a Database Management System (DBMS) for state management that is, indeed, embedded in the state of computation of the SP. However, due to its embedding, the DBMS suffers from some limitations such as the lack of multi-key transactions and consistent external querying. Eventually, their central role requires SPs to conform to a standardized execution semantics in order to improve their usability, interoperability, and interchangeability. This thesis takes a step towards SPs standardization through highlighting the discrepancies between them, and a step towards their integration with DBMSs by extending their computational model to deal with transactional computation. For SPs standardization, we use SECRET, a well recognized mathematical model to express their execution semantics, to model five distributed SPs that were developed after the introduction of SECRET itself and are today widely used in companies at the scale of Google, Twitter, and Netflix. We show that SECRET properly models a subset of the behavior of these systems and we shed light on the recent evolution of SPs by analyzing the elements that SECRET cannot fully capture. In order to decrease system integration overhead and to overcome the limitations of the current approaches for DBMS over SP, we enhance the capabilities of the SP with DBMS's ones by extending the SP computational model with transactional semantics: we develop a unified approach for multi-key transactions on the internal state of the SP, consistent external querying with respect to transactional operations on the state, and streaming data analysis. We implement TSpoon, a prototypal implementation of our extended model, as an extension to the open-source SP Apache Flink. We evaluate our prototype using synthetic workloads in various configurations to understand which metrics mostly impact its performance. Eventually, we evaluate a real use-case scenario and compare the results with the ones obtained from VoltDB, a commercial in-memory database known for its excellent level of performance: TSpoon outperforms VoltDB in the execution of multi-key transactions and proves to be a promising future direction for the integration of DBMSs and SPs.

MARCO BALDUINI

ON THE CONTINUOUS AND REACTIVE ANALYSIS OF A VARIETY OF SPATIO-TEMPORAL DATA

In recent years, an increasing number of situations call for reactive decisions making process based on a heterogeneous streaming data. In this context, the urban environment results particularly relevant, because there is a dense network of interactions between people and urban spaces that produces a great amount of spatio-temporal fast evolving data. Moreover, in a modern city there is a multitude of stakeholders who are interested in reactive decisions for urban planning, mobility management, tourism, etc. The growing usage of location-based social networks, and, in general, the diffusion of mobile devices improved the ability to create an accurate and up-to-date representation of reality (a.k.a. Digital footprint or Digital reflection or Digital twin). Five years ago, the state of the art was exploiting only a single data source either social media or mobile phones. However, better decisions can result from the analyses of multiple data sources simultaneously. Multiple heterogeneous data sources, and their simultaneous usage, of- fer a more accurate digital reflection of the reality. In this context, we investigate the problem of how to create an holistic conceptual model to represent multiple heterogeneous spatio-temporal data and how to develop a streaming computational model to enable reactive decisions. The main outcomes of this research are FraPPE conceptual model and RIVER streaming computational model with its implementations. FraPPE is a conceptual model, more precisely an ontology, that exploits digital image processing terms to model spatio-temporal data and to enable space, time, and content analysis. It uses image processing common terms to bridge the gap between the data engineer perspective and visual data analysis perspective. It does so to enable visual analytics on spatio-temporal data. During my PhD, we first formalize the spatial and temporal concepts in FraPPE 1.0, and, then, we add concepts related to the provenance and the content in FraPPE 2.0. We check the adherence of both versions of FraPPE to the five Tom Gruber's principles, and demonstrate the validity of the conceptual model in real world use cases. RIVER is a streaming computational inspired by two principles:(P1) everything is a data stream – a variety-proof stream processing engine must indifferently ingest data with different velocities from any sources and of any size –, and (P2) Continuous Ingestion – the data in input is continuously captured by the system and, once arrived, it is marked with an increasing timestamp. Most of the stream processing engines in the state of the art transform and adapt data at ingestion time. Contrariwise, RIVER is built around the idea of Lazy Transformation. So,a system that implements RIVER postpones data transformations until it can really benefits from them. Our hypothesis is that Lazy Transformation saves time and resources. RIVER relies on two main concepts: the Generic Data Stream (S<T>) and the Generic Time-Varying Collection (C<T>) and it proposes five different operators in order to ingest, process and emit data. The IN<T>operator is the entry point of the system, it takes an external data flow and injects the items into the system creating a new S<T>. The S2C<T>, C2C<T, T'> and C2S<T> operators in RIVER, inspired to the Continuous Query Language(CQL, the work on streaming data proposed by the Stanford DB Group) processing model, allows to move from S<T> to C<T> and vice-versa. The OUT<T> operator transform an S<T> into a new external data flow. Exploiting the Pipeline Definition Language (PDL) – our graphical language to abstract the operators' implementation complexity –, RIVER allows users to define computational plans, in the form of pipelines. In this thesis, we propose three different implementations of RIVER: Natron – a single-threaded vertically scalable implementation –, rvr@Spark and rvr@Hive – two horizontally scalable implementations based on distributed technologies (Spark and Hive). In order to prove the validity of the Lazy Transformation approach, we first evaluate Natron against our Streaming Linked Data engine that performs the data transformation at ingestion time. The result of this evaluation shows that Natron is cheaper – it consumes less resources in terms of memory and CPU load – and better approximates the correct answer under stress conditions. Moreover, we evaluate the cost effectiveness of Natron against rvr@Spark to prove that a distributed solution does not pay in all the situations. Indeed, in a mobile telco analysis, we observe that Natron is more cost-effective than rvr@Spark up to the scale of a nation. The results of those evaluations demonstrate the validity of the Lazy Transformation approach and confirm, in the stream processing engine field, that the a distributed solution does not pay at all scale. In order to prove the feasibility and the effectiveness of FraPPE and RIVER in enabling reactive decision-making processes on heterogeneous streaming spatio-temporal data, we present five real world use cases in Milan and Como. Moreover, during those case studies, we propose the data visualizations to different audiences (public users and stakeholders) in order to prove the guessability of our visual analytics interfaces. Finally, we reflect on limitations and state the future directions of this research work. In particular, those reflections involve the reasoning capabilities enabled by FraPPE, the future evaluations of RIVER against longer and more complex use cases and the evolution out the Pipeline Definition Language (PDL).

CARLO BERNASCHINA

TOOLS, SEMANTICS AND WORK-FLOWS FOR WEB AND MOBILE MODEL DRIVEN DEVELOPMENT

Web enabled mobile devices are becoming more an more ubiquitous in our lives. Application development for these devices opens newer and newer challenges. Model Driven Development was proposed as a solution able to reduce complexity and enhance productivity. This methodology was, and is still, not broadly adopted due to a proven, or perceived, high costs/advantages ratio making it difficult to reach a break-even point. The goal of the research presented in this thesis is to propose tools, semantics and work-flows aimed at reducing the costs of Model Driven Development, especially in the field of web and mobile applications. We will focus on tooling, by presenting an agile model transformation framework enabling the introduction of the Model Driven methodology in existing tools or the bootstrapping and rapid iterative development of new environments.
We present a formal semantics for the Interaction Flow Modeling Language, focused on web and mobile applications, having as objective a simple tool independent interpretation of IFML models enabling tools interoperability. We present an on-line tool for the rapid prototyping of web and mobile applications, showing how the proposed framework and semantics can be easily integrated together to produce a production ready Model Driven environment. We eventually present a Model and Text co-evolution work-flow which facilitates the interaction between code generators and human developers, by treating the application source code as the central artifact and the code generator as a virtual developer, i.e., yet another member of the team.
The experimental results show how the proposed methodology can reduce both the amount of work needed to obtain a production ready application and the level of expertise required in the process.

DIEGO CARRERA

LEARNING AND ADAPTATION TO DETECT CHANGES AND ANOMALIES IN HIGH-DIMENSIONAL DATA

Monitoring a datastream to detect whether the incoming data departs from normal conditions is a problem encountered in many important applications, ranging from quality control in industrial process to health monitoring. Many solutions in the literature adopt a model that describes normal data and check whether this model is not able to describe new data, detecting anomalous instances or permanent changes affecting the data-generating process. Pursuing this approach is challenging when data have high dimensions or feature complex structures (as in case of images or signals), and these are the settings we consider in this thesis.
We address this problem from two different perspectives. At first, we model data as realization of a random vector, i.e., we assume that data can be described by a smooth probability density function.
In these settings we focus on the change-detection problem, where the goal is to detect permanent changes affecting the data-generating process. In particular, we prove the detectability loss phenomenon, namely that performance of a popular change-detection algorithm that monitors the likelihood decreases when the data dimension increases. We also propose QuantTree, a novel algorithm to define histograms as density models for high dimensional data that are perfectly suit for change detection purposes. In fact, we prove that adopting QuantTree leads to an important property, i.e., that the distribution of any statistic computed over histograms generated by QuantTree does not depend on the distribution of the data-generating process. This enables non-parametric monitoring of multivariate datastreams with any statistic. Our experiments also show that combining several histograms computed by QuantTree into an ensemble can effectively mitigate the detectability loss phenomenon. In the second part we focus on data that feature complex structures, that cannot be described by a smooth probability density function. We adopt dictionaries yielding sparse representations to characterize normal data and propose a novel anomaly-detection algorithm that detects as anomalous any data that do not conform to the learned dictionary. To make our anomaly-detection algorithm effective in practical application, we propose two domain adaptation algorithms that adapt the anomaly detector when the process generating normal data changes. The proposed algorithms have been successfully tested in two real world applications: a quality inspection system monitoring the production of nanofibrous materials through the analysis of Scanning Electron Microscope (SEM) images, and ECG monitoring using wearable devices. Finally, we investigate convolutional sparse representations, translation invariant extensions of traditional sparse representations that are gaining much attention in the last few years. In particular our analysis focuses on image denoising and show that convolutional sparse representations outperform their traditional counterparts only when the image admits an extremely sparse representation.

SONIA CENCESCHI

SPEECH ANALYSIS FOR AUTOMATIC PROSODY RECOGNITION

This thesis presents a wide-ranging research work on prosody. Prosody is defined as the group of audio paralinguistic and suprasegmental clues involved in the communicative and understanding process of human speech. According to the main Universals in language, each Speech Act expresses common needs (e.g., talking about past or future events) similar for all humans, which are acoustically realized according to linguistic and phonotactics language related rules. At the same time, a spoken message can be uttered with a variable prosody because of countless factors as social context, emotions, intentions, rhetoric or spatial dislocation. This work starts proposing a new descriptive model in order to analyze prosody complexity in a structured and orderly manner, within which the sound of an utterance is considered as the final product of many exogenous and endogenous influences referring to the speaker. An Italian recited speech corpus and a psychoacoustic experiment were built in order to validate part of the model and to analyze the influence of semantics, phonotaxis and intonation on understanding processes. Results have been useful to defining the feature set to rely on the following parts of the work, regarding automatic recognition. Two neural network architectures have been developed, both of them regarding the Italian language. The first concerns the recognition of statements, questions and exclamations (using both textual and sound inputs), while the second identifies the presence of corrective focus into utterances (sound inputs only). A last section is focused on the semi-automatic characterization of prosody, laying the groundwork for further automatic recognition systems focused on prosodic skills. A monitoring protocol of expressivity and vocal qualities based on features extraction is then described, followed by practical applications to clinical, educational and forensics fields. The main contributions of this thesis are the definition of a new multi-dimensional conceptual model describing prosodic forms, two NNs based architectures for structures and corrective focus detection, two new audio/textual corpuses composed by recited and read speech used to feed NNs, and a proposal for the semiautomatic analysis of some aspects of prosody and expressiveness.

STEFANO CHERUBIN

COMPILER-ASSISTED DYNAMIC PRECISION TUNING

Given the current technology, approximating real numbers with finite-precision is unavoidable. Determining which finite-precision representation to exploit for each variable in the program is a difficult task. To face this problem, several precision mix solutions have been proposed so far in the state-of-the-art. However, the best precision mix configuration may vary at runtime along with input data. In this thesis we aim at suggesting two effective approaches to solve the precision tuning problem. The first approach follows the static precision tuning paradigm, i.e. it generates a single mixed precision version from the original code, which is designed to be used in place of the original version. Given the current technology, approximating real numbers with finite-precision is unavoidable. Determining which finite-precision representation to exploit for each variable in the program is a difficult task. To face this problem, several precision mix solutions have been proposed so far in the state-of-the-art. However, the best precision mix configuration may vary at runtime along with input data. In this thesis we aim at suggesting two effective approaches to solve the precision tuning problem. The first approach follows the static precision tuning paradigm, i.e. it generates a single mixed precision version from the original code, which is designed to be used in place of the original version. We later allow the possibility of changing the input conditions that may affect the best precision mix configuration. To solve this problem we propose a novel approach and a new toolchain that automitizes a large portion of this process. We present each component of the toolchain, and we provide guidelines to use them properly. We refer to this second approach as \emph{dynamic precision tuning}. We evaluate the static and the dynamic precision tuning solutions on a set of high performance computing and approximate computing benchmarks. We show how the dynamic precision tuning toolchain can be used -- under certain conditions -- also for static precision tuning. Our second toolchain is capable of achieving good results in terms of performance gain while maintaining acceptable precision loss threshold. In the future we aim at further improving this toolchain to extend its applicability to other use cases. Additionally, we highlight which improvements on the current toolchain may provide greater benefits on the quality of the output. Our proposed approach to static precision tuning exploits program profiling to understand which is the best precision mix configuration for the given accuracy requirements. We later allow the possibility of changing the input conditions that may affect the best precision mix configuration. To solve this problem we propose a novel approach and a new toolchain that automitizes a large portion of this process. We present each component of the toolchain, and we provide guidelines to use them properly. We refer to this second approach as dynamic precision tuning. We evaluate the static and the dynamic precision tuning solutions on a set of high performance computing and approximate computing benchmarks. We show how the dynamic precision tuning toolchain can be used - under certain conditions - also for static precision tuning. Our second toolchain is capable of achieving good results in terms of performance gain while maintaining acceptable precision loss threshold. In the future we aim at further improving this toolchain to extend its applicability to other use cases. Additionally, we highlight which improvements on the current toolchain may provide greater benefits on the quality of the output.

EMANUELE DEL SOZZO

ON HOW TO EFFECTIVELY TARGET FPGAS FROM DOMAIN SPECIFIC TOOLS

Heterogeneous System Architectures (HSAs) represent a promising solution to face the limitations of modern homogenous architectures, in terms of both performance and power efficiency. Indeed, thanks to the combination of hardware accelerators like GPUs, FPGAs, and dedicated ASICs, such systems are able to efficiently run performance demanding applications belonging to different application scenarios (like image and signal processing, linear algebra, computational biology, etc.) on the most suitable device for that domain. In order to fully take advantage of HSAs, in the last years new programming models and tools able to efficiently target such architectures, in terms of both final performance and productivity, emerged. Domain Specific Languages (DSLs) and Machine Learning (ML) frameworks are two significant examples. Both permit users to quickly and easily develop portable and efficient designs for multiple architectures. However, although DSLs and ML frameworks are highly effective in assisting users towards the generation of efficient designs for CPUs and GPUs, they still lack a concrete support for FPGAs. Indeed, even though FPGA toolchains have significantly improved and increased their features over the last years, the whole FPGA design process remains complex and the integration with high-productivity tools and languages is still limited. For these reasons, this research project focuses on the development of tools able to efficiently and easily target FPGAs from domain-specific scenarios. In particular, it consists in both a framework for the fast-prototyping and deployment of CNN accelerators on FPGA, and FROST, a unified backend to efficiently hardware-accelerate DSLs on FPGAs. On one hand, the goal of the CNN framework is to bridge the gap between high-productivity ML frameworks, like TensorFlow and Caffe, and FPGA design process. The framework automatizes the CNN implementation flow on FPGA, supports Caffe descriptions of the network, and provides a C++ library to design dataflow accelerators, as well as an integration with TensorFlow to train the network. On the other, starting from an algorithm described in one of the supported DSLs, FROST translates it into its Intermediate Representation (IR), performs a series of FPGA-oriented optimizations steps, and, finally, generates an optimized design suitable of FPGA tools. In order to better leverage the features of the FPGA and enhance the performance, FROST provides a high-level scheduling co-language the user can exploit to guide the optimizations to apply, as well as specify the architecture to implement. This allows to easily evaluate different hardware designs and choose the most suitable to the input algorithm.

YASHAR DELDJOO

VIDEO RECOMMENDATION BY EXPLOITING THE MULTIMEDIA CONTENT

VIDEO recordings are complex media types. For example, when we watch a movie, we can effortlessly register a lot of details conveyed to us (by the author) through different multimedia channels, in particular, the audio and visual channels. To date, the majority of content- based movie recommender systems (CBMRS) base their recommendations on metadata (e.g., editorial metadata such as genre or wisdom of the crowd such as user-generated tags) since they are human-generated and are as- sumed to cover the 'content semantics' of movies by a great degree. Multimedia features, on the other hand, provide the means to identify videos that 'look similar' or 'sound similar'. These discerning characteristics of heterogeneous feature sets meet users' differing information needs. In the context of this PhD thesis, methods for automatically extracting video-related information from the multimedia content (i.e., audio and visual channels) have been elaborated, implemented, and analyzed. Novel techniques have been developed as well as existing ones refined in order to extract useful information from the video content and incorporate them in recommendation systems. Different video recommendation tasks are solved using the extracted multimedia information under recommendation models based on content-based filtering (CBF) models and the ones based on combination of CBF and collaborative filtering (CF). As a branch of recommender systems, this thesis investigates a particular area in the design space of recommender system algorithm in which the generic recommender algorithm needs to be optimized in order to use a wealth of information encoded in the actual image and audio signals. The results and main findings of these assessments are reported via several offline studies or user-studies involving real users testing a prototype of developed movie recommender systems powered by multimedia content. The results are promising and show different scenarios in which multimedia content can be leveraged for successful video recommendation outperform- ing the alternatives, most notably in new-item settings.

CARLO D'ERAMO

ON THE EXPLOITATION OF UNCERTAINTY TO IMPROVE BELLMAN UPDATES AND EXPLORATION IN REINFORCEMENT LEARNING

The issue of sample efficiency always constituted a matter of concern in Reinforcement Learning (RL) research, where several works have been proposed to address the problem. It is historically well-known that this issue arises from the need of the agent to explore the environment it is moving in to improve its knowledge about it, and to exploit simultaneously the actions it considers to be the best to maximize its return, creating a trade-off known in RL as exploration-exploitation dilemma. The addressing of this trade-off is central and constitutes a measure of effectiveness of any algorithm available in literature. Moreover, the recent exponential growth of RL research, made possible by the comparable significant improvement in computational power, allowed researchers to extend the study of RL methodologies to high-dimensional problems that were unpractical before, opening the line of research that is now commonly known under the name of Deep Reinforcement Learning (DRL). However, the groundbreaking results that DRL is achieving are obtained at the cost of a huge amount of samples needed for learning, along with very large learning times usually in the order of days. One of the reasons why this is happening, besides the outstanding significance of the results that fundamentally poses the problems of the efficiency of these methodologies in the background, relies on the fact that often experiments are run in simulations in which the sample efficiency problem is not such an issue as in real applications. The purpose of this thesis is to study the previously described problems proposing novel methodologies that explicitly consider the concept of \textit{uncertainty} to speed up learning and improve its stability. Indeed, since a relevant goal of an RL agent is to reduce uncertainty about the environment in which it is moving, taking uncertainty explicitly into account can be intuitively an effective way of acting. This solution is not new in RL research, but there is still a lot of work that can be done in this direction and this thesis takes inspiration from the available literature on the subject extending it with novel significant improvements on the state of the art. In particular, the works included in this thesis can be grouped into two parts: one where uncertainty is used to improve the behavior of the Bellman equation and the other where it is used to improve exploration. The works belonging to the former group aim to address some of the problems of action-value estimation in the context of value-based RL, in particular in the estimate of the maximum operator involved in the famous optimal Bellman equation, and more generally in the estimate of all its components. On the other hand, the works belonging to the latter group study different methodologies to improve exploration by studying the use of Thompson Sampling in RL or by introducing a variant of the Bellman equation that incorporates an optimistic estimate of the action-value function to improve exploration according to the principle of Optimism in the Face of Uncertainty. All the works presented in this thesis are described, theoretically studied, and eventually, empirically evaluated on several RL problems. The obtained results highlight the benefits that the explicit exploitation of uncertainty in RL algorithms can provide; indeed, we show how in a large set of problems that have been chosen in order to highlight particular aspects we were interested in, e.g. exploration capabilities, our methods prove to be more stable and faster to learn than others available in the literature.

DANILO FILGUEIRA MENDONÇA

SELF-MANAGEMENT OF GEOGRAPHICALLY DISTRIBUTED INFRASTRUCTURES AND SERVICES

The paradigm of edge computing emerged in the last decade aiming to fill the gap between cloud data centres (accessible through multiple hops of networking) and the prosumers of information residing at the network edge. In edge computing, computing and storage resources are co-located with different kinds of infrastructures: with cellular infrastructures like base stations and aggregation sites; with core network components like ISP gateways; and with private infrastructures. Among the main goals of edge computing are the mitigation of network delay and the increase of bandwidth required by latency-sensitive and data-intensive applications hosted on mobile and Internet of Things (IoT) devices, including Autonomous Vehicles, Augmented/Virtual Reality, Mobile Multi-player Games, Natural Language Processing, Real-time Data Analytics, and Industry 4.0. The decentralised nature of the edge computing paradigm entails many challenges. First and foremost, the fine-grained distribution of edge nodes impose limitations to the capabilities offered by each node. Resources must be managed efficiently, and collaboration among surrogate nodes is paramount to enable more customers and services to be admitted into the system. Operational (Ops) aspects like the placement, deployment and scaling of edge-based services make automation and self-management properties first-class requirements. Last but not least, existing application and service models need to be adapted to cope with the characteristics of densely distributed infrastructure: heavyweight, monolithic applications may not fit into edge node resources or may fail to scale. In this thesis, we tackle the materialisation of edge computing from two main perspectives: architectural, in which we focus on the application and service models that enable the offloading of computation from latency-sensitive and data-intensive applications to edge nodes; and management, which handles the autonomic configuration, deployment, and scaling of services by geo distributed infrastructures. At the heart of our proposal is the paradigm of serverless computing and the Function-as-a-Service model. We leverage this alternative approach to cloud computing and propose a Serverless Architecture for Multi-Access Edge Computing. We then expand our contribution landscape with heterogeneous resources from mobile, edge, and cloud platforms, which we refer to as the Mobile-Edge-Cloud Continuum. To tackle the life-cycle of serverless functions deployed to the Continuum, we propose A3-E framework. A3-E moves away from centralised orchestration and management in favour of opportunistic, autonomic and decentralised provisioning of Function-as-a-Service to mobile applications with distinct requirements such as latency and battery consumption. We conclude our contributions with PAPS, a comprehensive framework that tackles the effective and efficient placement and scaling of serverless functions onto densely distributed edge nodes through multi-level self-management and control theory.

DAVIDE GADIOLI

DYNAMIC APPLICATION AUTOTUNING FOR SELF-AWARE APPROXIMATE COMPUTING

In the autonomic computing context, we perceive the system as an ensemble of autonomous elements capable of self-managing, where end-users define high-level goals and the system shall adapt to achieve the desired behaviour. This runtime adaptation creates several optimisation opportunities, especially if we consider approximate computing applications, where it is possible to trade off the result accuracy and the performance. Given the power consumption limit on modern systems, autonomic computing is an appealing approach to increase the computation efficiency. I divided this PhD thesis into three main sections.
The first section focuses on a dynamic autotuning framework, named mARGOt, which aims at enhancing the target application with an adaptation layer to provide self-optimisation capabilities at the production phase. In this context, the end-user might specify complex high-level requirements, and the proposed approach automatically tunes the application accordingly. The second section evaluates the mARGOt framework, by leveraging its features in two different scenarios. On the one hand, we evaluated the orthogonality between resource managers and application autotuning. On the other hand, we proposed an approach to enhance the application with a kernel-level compiler autotuning and adaptation layer in a seamless way for application developers. The third section focuses on two application case studies, showing how it is possible to significantly improve computation efficiency, by applying approximate computing techniques and by using mARGOt to manage them.

EUGENIO GIANNITI

PERFORMANCE MODELS, DESIGN AND RUN TIME MANAGEMENT OF BIG DATA APPLICATIONS

Nowadays the big data paradigm is consolidating its central position in the industry, as well as in society at large. Lots of applications, across disparate domains, operate on huge amounts of data and offer great advantages both for business and research. As data intensive applications (DIAs) gain more and more importance over time, it is fundamental for developers and maintainers to have the support of tools that enhance their efforts since early design stages and until run time. The present dissertation takes this perspective and addresses some pivotal issues with a quantitative approach, particularly in terms of deadline guarantees to ensure quality of service (QoS). Technically interesting scenarios, such as cloud deployments supporting a mix of heterogeneous applications, pose a series of challenges when it comes to predicting performance and exploiting this information for optimal design and management. Performance models, with their potential for what if analyses and informed design choices about DIAs, can be a major tool for both users and providers, yet they bring about a trade-off between accuracy and efficiency that may be tough to generally address. The picture is further complicated by the adoption of the cloud technology, which means that assessing operating costs in advance becomes harder, but also that the contention observed in data centers strongly affects big data applications' behavior. For all these reasons, ensuring QoS for novel DIAs is a difficult task that needs to be addressed in order to favor further development of the field. Over this background, the present dissertation takes two main routes towards facing such challenges. At first we describe and discuss a number of performance models based on various formalisms and techniques. Among these, there are both basic models aimed at predicting specific metrics, like response time or throughput, and more specialized extensions that target the impact on big data systems of some design decisions, e.g., privacy preserving mechanisms or cloud pricing models. On top of this, the proposed models are variously positioned across the spectrum between efficiency and accuracy, thus enabling different trade-offs depending on the main requirements at hand. This is relevant in the second main part of this dissertation, where performance prediction is at the core of some formulations for capacity allocation and cluster management. In order to obtain optimal solutions to these problems, in one case at design time and in the other at run time, we adopt both mathematical programming and several performance models, according to the different constraints on solving times and accuracy. More in detail, we propose performance models based on queueing networks (QNs), stochastic well formed nets (SWNs), and machine learning (ML). This variety is justified by the different uses of each methodology. ML provides algebraic formulas for execution times, which are perfectly fit to be added as constraints in our optimization problems' mathematical programming formulations, thus yielding initial solutions in closed form. Since ML can reliably provide accurate predictions only in regions properly explored during the training phase, the optimal solution is searched via a simulation-optimization procedure based on analytical models like QNs or SWNs, which in contrast are quite insensitive to the parameter range of evaluation, being devised from first principles. These kind of models boast relative errors below 10 % on average when predicting response times. In terms of optimization, first of all we consider the design time problem of capacity allocation in a cloud environment. The design space is explored via both ML and simulation techniques, so as to choose the best virtual machine type in the catalog offered by cloud providers and, subsequently, determine the minimum cost configuration that satisfies QoS constraints. We show also how this optimization approach was applied during the design phase of a tax fraud detection product developed by industrial partners, i.e., NETF Big Blu. Afterwards we also considered the run time issue of finding the minimum tardiness schedule for a set of jobs when the current workload exceeds predictions and the deployed capacity is not enough to ensure the agreed upon QoS. Thanks to the varied efficiency of performance models, it is possible to solve the design time problem in a matter of hours, whilst run time instances are solved within minutes, consistently with the different requirements.

MICHELE GUERRIERO

MODEL-DRIVEN ENGINEERING FOR PRIVACY-AWARE DATA-INTENSIVE APPLICATIONS

The pervasiveness of the modern digital ecosystem is leading to dramatic changes in our society and in the way software systems are designed. Nowadays smartphones, wearables, sensors and, in general, smart devices sensing the environment, allow to continuously collect huge volumes of data. Having the ability to deal with and to make sense of this data can give to companies and institutions important advantages. This is known as the “Big Data” phenomenon and is pushing significant investments in order to re-design software solutions in a more data-centric way. By leveraging several results from Model-Driven Engineering (MDE), DevOps, formal methods and data protection techniques, this thesis provides three research contributions along this direction. First, we present a model-driven approach for the development and prototyping of DIAs across different target platforms. Then, we complement our first result with a model-driven approach to simplify and automate the deployment of DIAs by means of the Infrastructure-as-Code paradigm. Finally, we present an approach for embedding privacy-awareness in DIAs to help simplifying the compliance with privacy requirements. Overall the evaluation of the three research contributions shows that, despite various limitations, the proposed approach is promising and reasonably achieves the research goals of this thesis.

ANITA IMANI

LIQDROID: A MIDDLEWARE FOR DIRECT INTERACTION BETWEEN MULTIPLE PROXIMAL ANDROID DEVICES

Nowadays, the speed of technology improvements in computing devices is very fast. The functionality and ease that using these different devices bring to us has influenced the way we are living and enabled us to benefit from them carrying out various daily tasks. But despite the improvements that have occurred in the field of mobile technology and the connection protocols, the current situation as regards multiple- device interaction techniques is still far behind what it could be, and these computing devices are still mostly working in isolation. Current multi-device interaction solutions enable the user to continue a task on another device, but their dependency on a specific set of devices and software artifacts (applications) on these devices limits their usage and forces the user to act as a bridge between the devices. This entails the user having to perform some preliminary and time-consuming steps to configure the next device on which to resume the task. While benefiting from the direct interaction between the multiple proximal devices, the user can interchangeably benefit from them in the execution of the different parts of the task he desires to carry out, exploiting the maximum potential that exists in these devices, to achieve better results and more user satisfaction. This dissertation covers the motivation, design, and development of a novel paradigm to support multiple-device direct interaction through distributing the Android execution.
The proposed solution to support this novel paradigm is a middleware infrastructure that will manage the direct interaction between a dynamic set of proximal devices to let the user seamlessly distribute the execution of a task between them, while also enhancing the process of the interaction and integration of these devices. The proposed middleware, which is called LIQDROID, will distribute the Android operating system between a set of proximal Android devices to create a bigger Android ecosystem. The proposed ecosystem transforms the current pattern of single-user single-device to a fully cooperative environment that will empower a user to start a task on his device and be able to interchangeably and collaboratively benefit from the potential that exists on the proximal devices or other users during its execution and reach his final goal with better results. More technically speaking, LIQDROID is an Android service which benefits from the features available within the Android framework to solve the challenges that already exist in multiple-device interaction to manage the execution of a distributed task, such as finding the best capable proximal Android device to perform a task, and synchronizing the state of the integrated devices and the data management. This will provide the required framework for developers to easily be able to distribute the execution of a task on proximal devices and be relieved of the underlying complexities, and instead put their focus on designing and developing more innovative distributed applications. A stable prototype of LIQDROID has already been implemented in Java, which is the main language of developing Android applications, and is ready to use. Also, a set of use case scenarios has been designed and considered in this thesis, based on the needs of real case scenarios. The developed versions of these use case scenarios in the shape of LIQDROID-compatible applications, along with LIQDROID, was tested on real devices to better evaluate and explore their strengths and weaknesses. The results were satisfying and also helped to apply the required changes in LIQDROID to improve it. These experiments have shown that LIQDROID introduces an innovative way of interacting between multiple proximal Android devices, which has the potential to introduce new and even more comprehensive features in the way of direct interaction between multiple proximal devices.

EWERTON LOPES SILVA de OLIVEIRA

LEARNING MODELS TO OPTIMIZE THE PLAYER EXPERIENCE IN ROBOGAMES

As technology progresses new game experiences emerges. Among these, a new type of game appears, where human players are involved in a physical activity against robotic agents. This type of games has been introduced as Physically Interactive Robogames (PIRG). In this work, we have developed methods and insights for modeling players in a PIRG environment with data from on-board sensors processed in real-time. This new type of game environment has as main characteristic the exploitation of the real world as environment (in both its dynamical, unstructured, and structured aspects), and of one or more real, physical, autonomous robots as game opponents or companions. The ultimate direction for PIRG is to obtain a robotic player purposefully aiming at maximizing human player entertainment. In our work, we provide a panorama of design for such robotic applications, advocating, in the process, the benefits of Machine Learning techniques to tackle the challenges. We present methods and insights for player modeling using Machine Learning techniques, as well as direction for future research to achieve full adaptation. Besides being an interesting field for testing approaches from Machine Learning, a PIRG scenario provides a challenging application for several other disciplines, among which: (general and specific) Artificial Intelligence, Statistics, Human-Robot Interaction, Robotics, Psychology, Design.

FRANCESCO MARCONI

FORMAL VERIFICATION OF TIMED PROPERTIES FOR DATA-INTENSIVE APPLICATIONS

In the past few years, cloud-based enterprise applications, leveraging the so-called data-intensive technologies, have emerged as pervasive solutions for modern computing systems. Their adoption has been motivated by the growing need for systems that are able to collect, process, analyze and store huge quantities of data coming from various sources (social media, sensors, bank transactions, etc.) in a reasonable time. Data-intensive applications (DIAs), taking advantage of those technologies, natively support horizontal scalability, and constitute a significant asset for the production of large-scale software systems. However, the adoption of data-intensive technologies in the small and medium enterprises (SMEs), constituting the vast majority of the European industry, is still slow for a number of reasons, such as the steep learning curve of the technologies, the lack of experience and resources to keep up with such innovation. The definition of methodologies and principles for good software design is, therefore, fundamental to support the development of DIAs. Non-functional (quality) requirements are an aspect of software design that is typically overlooked at design time and turns out to be crucial in later stages of development. Usually expressed in terms of Service Level Agreements (SLAs), if they are not met, further refinements of the applications are needed, resulting in additional costs. Design time quality analysis aims at detecting the presence of potential design flaws that could lead to later quality incidents, fostering the early detection of problems. Different approaches exist for the quality analysis of parallel distributed applications. Performance prediction is arguably the most common and is typically enacted on stochastic models by means of either analytical methods or simulation. A different approach is pursued by formal verification: it performs an exhaustive check on the model of the system to assess whether certain properties are satisfied by the modeled system. This thesis presents a unified model-driven approach for the formal analysis and verification of temporal properties for data-intensive applications. It addresses the computation of the two main classes of DIAs, i. e., streaming and batch processing, by proposing two formal models based on metric temporal logic and analyzable through state-of-the-art solvers. Both formalizations, albeit with relevant differences, capture the computational models of DIAs as DAGs, enriched with the most relevant quality aspects of the applications. We first describe timed counter networks, a formal model capturing the computation of Storm-like streaming applications. The model is devised by extending the CLTLoc metric temporal logic with positive discrete counters, and enables the analysis of properties concerning the growth of the queues of the nodes composing Storm topologies. The model can be automatically analyzed (with some limited modifications) by means of the Zot satisfiability checking tool, but it presents some undecidability issues, therefore we propose an additional check to assess the soundness of the analysis results. Then we present an analogous approach for the formalization of the deadline feasibility problem for Spark batch applications. The model is devised as well in terms of CLTLoc extended with counters, but describes a rather different scenario, in which application runs are finite and the goal of the analysis is to determine whether a specific deadline can be met by the designed application. We also discuss a model reduction strategy that exploits the properties of partially ordered sets to partition the DAG underlying the formal model. Next, we introduce D-VerT, a model-driven tool that allows designers to perform the formal analysis of streaming and batch applications by means of an automated toolchain, starting from suitably annotated UML diagrams. D-VerT embeds the models and analysis devised for Storm and Spark applications, and it automatically translates the UML design diagrams to the corresponding instances of the formal models, which are then fed to a state-of-the-art satisfiability solver. In this way, users with limited expertise of formal methods can benefit from the analysis without directly dealing with the specific formalisms. Experimental evaluation has been carried out over the two kinds of analyses performed by D-VerT. The queue boundedness analysis of Storm applications was shown to capture design flaws leading to memory saturation in some of the nodes of Storm topologies, and provided some important hints to designers. In the case of Spark applications, experiments show that the formal model reflects actual executions of the framework with good accuracy (error with respect to the actual execution time less than 10%). Moreover, the proposed optimization registered significant improvements in terms of time (up to

90%) and memory (up to 45%) needed to perform verification. Future works include further refinements of the formal models to improve their accuracy and performance. For example, the optimization strategy for DAG-based computations could be improved and possibly embedded in the decision procedures implemented by the solvers. We also consider the investigation of new properties and different technologies. Moreover, we plan to extend the experimental analysis to new use cases to have a more extensive evaluation. Finally, we are investigating the introduction of probabilities in our models to also carry out stochastic analysis.

GIOVANNI MERONI

ARTIFACT-DRIVEN BUSINESS PROCESS MONITORING

Traditionally, to monitor the execution of a business process, organizations rely on monitoring modules provided by Business Process Management Systems (BPMSs), which automate and keep track of the execution of processes. While the adoption of a BPMS to monitor a single-party, fully-automated business process is straightforward, the same cannot be said for multi-party processes heavily relying on manual activities. In fact, a BPMS requires explicit notifications to determine when activities that are not under its direct control are executed. This requires organizations to federate their BPMS, a complex task that has to be performed whenever a new organization participates in the process. Also, when activities are not automated, human operators are responsible for manually sending notifications to the BPMS, a task that disrupts the operators' work and, as such, is prone to be forgotten or postponed. To continuously and autonomously monitor multi-party processes involving non-automated activities, this thesis proposes a novel technique, named artifact-driven process monitoring. This technique exploits the Internet of Things (IoT) paradigm to make the physical objects participating in a process smart. Being equipped with sensors, a computing device, and a communication interface, such smart objects can then become self-aware of their own conditions and of the process they participate in, and exchange this information with the other smart objects and the involved organizations. This way, it is possible for the monitoring infrastructure to stay in close contact with the process, and to cross the boundaries of the organizations. To be aware of the process to monitor, instead of using activity-centric process models, usually adopted by BPMSs, smart objects rely on an extension of the Guard-Stage-Milestone (GSM) artifact-centric modeling language, named Extended-GSM (E-GSM). Normally, a BPMS expects the execution to rigidly adhere to the process model defined in advance. Therefore, whenever a deviation between the execution and the model is detected, a BPMS requires human intervention to resume process monitoring. E-GSM, on the other hand, treats the execution flow (i.e., dependencies among activities) in a descriptive rather than prescriptive way. Consequently, smart objects can detect violations during execution without interrupting the monitoring. Additionally, E-GSM can monitor if the physical objects evolve as expected while the process is executed. Finally, E-GSM provides constructs to determine, based on the conditions of the physical objects, when activities are started or ended. This thesis also presents an approach to determine to which extent smart objects are suited to monitor a particular process, given their sensing capabilities. To relieve process designers from learning the E-GSM notation, and to allow organizations to reuse preexisting process models, a method to instruct smart objects given Business Process Model and Notation (BPMN) collaboration diagrams is also presented. Finally, a prototype of an artifact-driven monitoring platform, named SMARTifact, is developed and tested against both historical and live sensor data.

GIUSEPPE NATALE

ON HOW TO DESIGN OPTIMIZED SPATIAL ARCHITECTURES: FROM ITERATIVE STENCILS TO CONVOLUTIONAL NEURAL NETWORKS

We are approaching the end of Moore's law and Dennard scaling, while demand for computing power increases constantly. Traditional processors are struggling to keep up with performance requirements, and both scientific research and industry are exploring reconfigurable architectures as an alternative. It is common knowledge that an increasing number of technology leaders, such as Microsoft, IBM, Intel, Google, Amazon to name a few, are currently exploring the employment of reconfigurable architectures as hardware accelerators. Fine grained inherent parallelism and low power consumption thanks to direct hardware execution are the key aspects that make reconfigurable architectures an attractive choice for High-Performance Computing (HPC). This thesis is focused on a set of algorithms sharing a similar computational pattern, namely iterative stencils and Convolutional Neural Networks (CNNs). The key computation for both algorithms consists in sliding a filter on the input data, computing new elements using a submatrix of the input. Iterative stencil algorithms are heavily employed in physics, numerical solvers and even finance, while CNNs are one of the recently developed deep learning algorithms, currently used in industry to perform image classification, analysis of video and even speech recognition and recommender systems. While High-Level Synthesis (HLS) capabilities have improved dramatically over the recent years, the synthesis tools have yet to reach the level of sophistication required to properly optimise these algorithms, and extract sufficient parallelism or generate highly scalable solution, or automate the acceleration creation and integration in a way that resembles the development experience that is in place for CPUs and GPUs. Indeed, the process of designing and deploying hardware accelerators for iterative stencils of CNNs is still a hard and complex task that requires expertise in FPGA programming and knowledge of hardware design tools. The complex dependencies that arise from the filtering conditions, the iterative nature of the algorihtms, and the low operational intensity, make current HLS solutions inadequate to really optimise the resulting implementations. This thesis objective is to improve with respect to the proposed solutions targeting all the presented challenges. For both the target algorithms, we designed optimized spatial architectures that are able to exploit different sources of parallelism of such algorithms, reduce the cost of data movements alleviating the burden on external memory and that can easily scale up to multi-FPGA systems. In particular, we propose spatial accelerators for iterative stencils, small CNNs and the features extraction stage of deep CNNs. The resulting designs, consisting of distributed architectures of independent elements communicating using read- and write-blocking FIFOs, can scale in size if enough resources are available, unfolding the algorithms computation in space. Such architectures can exploit different levels of parallelism offered by the target algorithms. Moreover, we address the scalability of the proposed solutions implementing the accelerators on very big Stacked Silicon Interconnect (SSI) FPGAs or even on custom-designed multi-FPGA systems. Finally, we provide automation flows and performance models that can automate completely or support the automation of the design process, starting from high-level specification of the algorithms, either being C/C++ for iterative stencils, or standard formats as the one implemented in Caffe for CNNs. The validation performed shows that for iterative stencils, the proposed solution is comparable with the state of the art for the single FPGA implementation, but we are able to perform substantially better on multi-FPGA, thanks to an approximately linear scaling in performance. Moreover, for deep CNNs our implementations for AlexNet and VGG-16 achieve a throughput of respectively 1.61 and 2.99 TOPS, where for VGG-16 we are the second-best implementation - by a narrow margin - available in the State of the Art, while for AlexNet we substantially outperform the previous work.

XUESONG PENG

MONITORING DATA UTILIZATION IN FOG COMPUTING ENVIRONMENTS

This thesis proposes a framework to improve monitoring data utilization in fog computing environments, reducing the data volume while maintaining the data quality and structure performance. Over the past few years, there has been a significant increase in the use of Monitoring Data for various domains such as IoT, cloud computing, smart spaces. In all the mentioned domains, applications use Monitoring Data as a critical pipeline to observe the target system (space) and to support the smart decision-making process. This work has been motivated by the increasing attention of monitoring data and the rising volume and dimensions of monitoring data, which brings the rapid grow of maintenance cost to the system owners in various aspects: energy, network bandwidth, storage, etc. In this work, we try to understand the situation and the environment of monitoring data utilization and how the monitoring data are collected and transmitted and modeled during its lifecycle. We propose a correlation-based approach to reduce the volume of monitoring data where some monitoring variables can predict some other variables based on their quantitative relation, which has been derived by analyses. We propose execution profiles to capture the resources demands of software modules and services to enable better management of computational resources. We propose a deployment methodology in fog computing environments, with hierarchical environment model and resources capacity constraints. We also offer a fault-tolerant scheme to ensure the data quality of raw sensor outputs in the monitoring system, to improve the performance of data utilization in the whole lifecycle. Combining the aspects above, we propose an adaptive monitoring data framework for fog computing environments. The framework learns the relationships between monitoring variables and reduces the monitoring dataset using data reduction services. It can respond to changes in the surrounding environment, keeping these relations updated. It distributes individual reduction services to the fog environment and optimizes the deployment of data reduction services. It also exploits the fault-tolerant scheme to the edge side of the monitoring system for reliable data collection.

STEFANO PERNA

DATA-DRIVEN TECHNIQUES FOR KNOWLEDGE DISCOVERY IN REGULOMICS

Novel technologies have led to exponentially increasing amounts of genomic data. However, while costs have been constantly reducing, modeling and analysis techniques have only just started to catch up in effectiveness and efficiency. Regulomics is a sub-field of genomics which studies the mechanics of gene expression regulation, i.e. how cells select and express different genes to respond to the different situations. Among those, Transcription Factors (TFs) are proteins that attach themselves the DNA of prokaryotic and eukaryotic organisms in highly specific Transcription Factor Binding Sites (TFBS), and modulate how accessible the surrounding DNA areas is by RNA transcription machinery. Such areas usually contain coding sequences of genes. For this reason, they are of great importance in regulomics. TF activity has been studied in isolation by various means, such as wet-lab experiments and computational methods, but the interplay of several TFs has not been studied as much. TF co-regulation is significantly harder to analyse directly, requiring novel computational methods. This thesis discusses a novel model aimed at predicting and classifying TF-TF interactions using a data-driven, model-based approach. The fundamental idea is that TFBS and coding sequences can be represented as a set of oriented, linear coordinates with features attached, and that the distance between binding sites in this coordinate system is an informative feature which can be used to predict TF-TF interactions. This approach relies on the properties of the distribution of genomic distances between matched, closest binding sites of potential interactors. To further refine this model, firstly protein-protein interaction (PPI) network data is mined to compute additional, independent features used in classification of TF-TF interactions, under the assumption that the more shared interactors two TFs have in the PPI network, the more likely it is that they are co-operating as opposed to competing for another partner; secondly, the number of detected copies of each TF at the relevant binding sites is used to infer whether the TFBS itself it highly bound or instead disrupted. The resulting classifiers are named TICA, NAUTICA and ESTETICA; the first two show good performance with respect both to reference databases and existing literature. Taken as a whole, they represent a powerful framework for inferring and classifying TF-TF interaction phenomena.

MARCO RABOZZI

CAOS: CAD AS AN ADAPTIVE OPEN-PLATFORM SERVICE FOR HIGH PERFORMANCE RECONFIGURABLE SYSTEMS

In current years we are assisting at a new era of computer architectures, in which the need for energy-efficiency is pushing towards hardware specialization and the adoption of heterogeneous systems. This trend is reflected in the High Performance Computing (HPC) domain that, in order to sustain the ever-increasing demand for performance and energy efficiency, started to embrace heterogeneity and to consider hardware accelerators such as Graphics Processing Units (GPUs), Field Programmable Gate Arrays (FPGAs) and dedicated Application-Specific Integrated Circuits (ASICs). Among the available solutions, FPGAs, thanks to their advancements, currently represent a very promising candidate, offering a compelling trade-off between efficiency and flexibility that is arguably the most beneficial. FPGA devices have also attained renewed interests in recent years as hardware accelerators within the cloud domain. Tangible examples of this are the Amazon EC2 F1 instances, which are compute instances equipped with Xilinx UltraScale+ FPGA boards. The possibility to access FPGAs as on demand resources is a key step towards the democratization of the technology and to expose it to a wide range of application domains. Despite the potential benefits given by embracing reconfigurable hardware both in the HPC and cloud contexts, we notice that one of the main limiting factor to the widespread adoption of FPGAs is complexity in programmability, as well as the effort required to port a pure software solution to an efficient hardware-software implementation targeting reconfigurable heterogeneous systems. The main objective of this dissertation is the development of CAD as an Adaptive Open-platform Service (CAOS), a platform able to guide the application developer in the implementation of efficient hardware-software solutions for high performance reconfigurable systems. The platform aims at assisting the designer starting from the high level analysis of the code, towards the definition of the functionalities to be accelerated on the reconfigurable nodes. Furthermore, the platform guides the user in the selection of a suitable architectural template for the FPGA devices and the final implementation of the system together with its runtime support. Finally, CAOS has been designed to facilitate the integration of external contributions and to foster research on the development of Computer Aided Design (CAD) tools for accelerating software applications on FPGA-based systems.

MIRZA RAMICIC

PERCEPTION AS BEHAVIOUR INDUCING MECHANISM: A REINFORCEMENT LEARNING PERSPECTIVE

Rapid advancement of machine learning make it possible to consider large amounts of data to learn from. In most of the implementations of reinforcement learning facing this type of data, approximation is obtained by neural networks and the process of drawing information from data is mediated by a short-term memory that stores the previous experiences for additional re-learning, to speed-up the learning process, mimicking what is done by people. In this work, we are proposing a range of novel computational approaches able to selectively filter the informational or cognitive load for the agent's short-term memory, thus simulating the attention mechanism characteristic also of human perception. Using the proposed attention filter block architecture, we were able to devise a variety of frameworks of agent's perception that are able to adapt to its environment by selecting the most suitable experiences. The adaptation also resulted in an emergence of different behavioural characteristics or traits among artificial learning agents.

ALESSANDRO RIVA

DEVELOPMENT AND ANALYSIS OF ALGORITHMS FOR INFORMATION-GATHERING IN AUTONOMOUS MOBILE ROBOTICS

The gathering of information is a problem that implicitly appears in a large number of autonomous mobile robotics tasks. Some examples are exploration, coverage, and surveillance, in which the abstract concept of information assumes different meanings, e.g., representation of parts of an unknown environment or the presence of entities or objects. In recent years, applications involving information gathering have received growing attention from both public institutions and private organizations. The development of proper algorithmic solutions, however, is still particularly challenging, being most of the problems involved computationally hard. In practice, several autonomous mobile robots adopt a layered architecture. Bottom levels include sensing and actuation, while the finding of a global plan to follow is demanded to higher levels. When a task requires to gather information by means of one or more autonomous mobile robots, the high-level plan often consists of a navigation plan, i.e., a sequence of locations (or poses) to orderly reach, where information is gathered. A navigation plan is usually computed on a high-level representation of the environment, in which most of the geometrical features and the robot dynamics are neglected, and that is commonly based on grids or graphs. On this abstract model, the information-gathering tasks can be formulated as optimization problems, in which an objective has to be minimized (e.g., the completion time or the energy consumption) and some application-dependent constraints must be satisfied. Unfortunately, most of the so-formulated problems are NP-hard and an effort has been made in the literature aimed at finding efficient sub-optimal algorithmic solutions. Although the approaches implemented can be broadly grouped in some standard classes (e.g., heuristics, greedy algorithms, linear programs relaxations), little is known about the quality of the navigation plans found. A typical tool to assess the performance of sub-optimal approaches is the analysis based on the approximation factor, that is, the ratio between the objective value of a sub-optimal solution and that of an optimal one. In this thesis, we address two different topics related to the gathering of information by means of one or more autonomous mobile robots. The tasks addressed in these two contexts are here formalized as high-level optimization problems, for which offline algorithmic solutions are developed and analyzed. The first topic is that of coverage of a known environment using tools mounted on the robots. In their classic formulation, coverage models assume that the tools are operated by independent actions of the robots (e.g., the sensing of the surrounding environment). However, the concept of coverage can be extended to encompass a larger number of tasks, in which actions may require the joint operations of two or more robots to cover some features of the environment. A first notable example is that of the measurement of the signal strength among pairs of locations. In this broader coverage framework, we address two problems of practical interest. (a) A robot that has to cover with a finite-range tool (e.g., a sensor or a brush) all the points of the free space of a given known environment. (b) A team of robots that has to jointly perform a given set of pairwise measurements in a given known environment. The second topic we consider is that of planning robot paths to reach given target locations. Differently from the typical path planning settings, we consider two scenarios in which communication constraints are imposed by practical application-dependent needs. These problems can arise in a number of information-gathering tasks, as robots deal with finite storage capabilities and it may be needed to flush data through a networked infrastructure, especially in long-term mission scenarios. Another critical aspect, emerging in multirobot tasks, is that of keeping robots connected in order to exchange information to enhance the decisions taken online. In this framework, the specific problems we address are the following. (c) A robot, moving from a start to a goal location, is required to gather data along its path (e.g., a video feed in a monitoring scenario). The robot has at its disposal only a limited amount of memory and hence gathered data has to be periodically delivered to a base station using transmission zones spread within the environment. (d) A team of robots moving in an environment must plan a set of start-goal joint paths which ensures global connectivity at each (discrete) time step, under some communication model. We assess the hardness of all the above problems by reduction from well-known hard problems. Once we have given strong evidence that optimal solutions are (in general) out of reach, we develop sub-optimal algorithmic solutions with particular attention to their efficiency, e.g., polynomial-time complexity. In some cases, an approximation factor guarantee is also devised. Finally, all the algorithms developed in the thesis are tested in simulation to provide an empirical evaluation of their performance. The results contained in this document can be considered a significant contribution to the field, as we provide the formalization of novel problems of practical interest, original analyses of the problems and their attribution to the proper complexity classes, polynomial-time algorithmic solutions, approximation factor results, and experiments in simulation. In a broader view, our achievements can provide useful insights for future analyses of similar (or related) problems, as well as be useful in the implementation and the deployment of efficient autonomous mobile robotics systems for specific applications.

MERSEDEH SADEGHI

A MODEL-CENTERED SOLUTION FOR TAMING THE HETEROGENEITY OF SMART DEVICES

Recent progress of Internet of Things (IoT) technologies has led to a competitive market and heterogeneity of communication protocols, a large diversity of smart device types, and a multitude of widely used Application Programming Interfaces (API) and standards. Even if, these advancements enhance the development of IoT domains including smart spaces, they intensify the interoperability challenge. In other words, the co-operability of smart devices is compromised by the numerous enabling technologies. Smart spaces are becoming more tightly bounded to the technology and standard proffered by their underlying frameworks. Moreover, the construction of pervasive environments could be subject to the thread of monopolistic smart spaces in which devices and technologies are exclusively harvested in the favor of giant companies. Thus, there is a pressing need for integration and possibly unified standardization. The current literature tackles this issue within bottom layers of IoT stack, and, predominantly toward the establishment of unified programming interface and communication means. Nevertheless, the harmonization of the interfaces to exchange data does not fully address the problem. In fact, a seamless integration and comprehensive interoperability do not accomplish if the exchanged data remain unstandardized and diversified. In this direction, although ontologies and shared vocabularies had promised the interoperability at the semantic level, yet, the upsurge of a competing domain-dependent and vendor-specific ontologies stirrers up the diversity trouble. This thesis aimed to provide a comprehensive solution for tackling the heterogeneity of smart devices and allow developers to integrate them seamlessly. It proposes a generic model, called TDeX, that encapsulates the peculiarity of a wide variety of objects and describes them in a unified manner. The TDeX that extends the W3C proposed TD (Thing Description), decouples devices from their enabling technology and envisions a device and technology agnostic modeling. Needless to remark that the smart devices are not isolated entities but they are essentially characterized in their relationship with other objects, human and the environment. Furthermore, their smartness becomes actually tangible for the end-user in the front-end application which visualizes and facilitates any interaction with the device. Consequently, to envisage a truly seamless integration, it is vital to feed the application with a multifaceted model covering the essential aspects. To this end, TDeX incorporates information about context and access rights, along with the elements to create device-agnostic Graphical User Interfaces (GUI) to interact with the devices. It thus paves the ground for additional contributions: (I) REST middleware infrastructure called M4HSD governs the interaction with the different devices by homogenous interfaces, harmonizes the diversified characteristic in the unified model and eventually cast them to the application as a proper instance of TDeX. (II) An innovative feature-oriented permission model for access management in IoT systems is proposed which is well integrate with TDeX. It enables M4HSD to offer a granular and context-aware access control mechanism to address the demanding privacy requirements of today smart space. (III) A model-driven solution for the automated GUI creation in front-end application is proposed. The TDeX model allows the application to visualize a smart device that automatically adapts to changes in the context of fruition and access rights.

ALBERTO SCOLARI

OPTIMIZING DATA-INTENSIVE APPLICATIONS FOR MODERN HARDWARE PLATFORMS

Data-intensive applications have become widespread in the years, especially in cloud-like environments. Among them, Data Analytics (DA) and Machine Learning (ML) applications are particularly important categories that deeply impacted business and science in the last decade, and are expected to have an even higher impact in the upcoming years. In the latest years, we also saw hardware platforms evolving along different directions to overcome the limitations of Dennard's scaling and the end of Moore's law. While heterogeneity is coming into play for several applications, Central Processing Units (CPUs) have also evolved towards a growing number of cores, specialized Single Instruction, Multiple Data (SIMD) units, high memory hierarchies and, in general, a more complex and diverse set of features. On the other side, also data-intensive applications became more complex, to the extent that a current ML model may comprise tens of diverse operators to compute a single prediction value, while taking in input data of multiple types like text, number vectors and images. Oftentimes these applications are structured as "data pipelines" and go through many steps like input parsing, data pre-processing, analysis and possibly loading the output to some "data sink" at the end. Mastering this complexity to achieve the best implementation and deployment of an application in a given setting (hardware platform, software stack, co-located applications, etc.) is becoming a key issue to achieve best usage of the infrastructure and cost-effectiveness. This problem is especially hard with heterogeneous platforms, whose hardware features may not suit some parts of an application to be accelerated or may require a long redesign effort. Here, where the inevitable complexity of applications determines the diversity of operations, CPUs are still central, as they provide the flexibility and the maturity to efficiently run most of these workloads with sufficient performance even for today's needs, while being easier to program than other architectures. Moreover, their general-purpose design naturally fits the diversity of data-intensive applications. This work explores the performance headroom that lies unused in modern CPUs with data-intensive applications. This headroom encompasses several dimensions, each one with specific problems and solutions. A first problem to solve is the performance isolation of co-located applications on the CPU, which has to take in account the sharing of the Last Level Cache (LLC). Here, we propose and evaluate a mechanism for partitioning the LLC that works on recent server-like CPUs and requires software-only modifications in the Operating System (OS), thus not impacting hardware nor applications. This solution proves to be effective with a diverse set of benchmarks and allows meeting Quality of Service (QoS) goals even in a contentious, hardly predictable environment. The second problem is the optimization of data-intensive applications, which can be composed of multiple, diverse computational kernels. This work explores the limitations of current solutions, revolving around a black-box approach: application kernels are individually optimized and run, disregarding their characteristics and their sequence along the data path; indeed, a simple case study shows that even a manual, naïve solution can achieve noticeable speedups with respect to the current state of the art. Building on these findings, we generalize them into a white-box approach for applications optimization: while applications should be written as sequences of high-level operators, as current development frameworks already do, this sequence should also be exposed to the system where these applications run. By looking at this structure during the deployment of the application, the system can optimize it in an end-to-end fashion, tailoring the implementation to the specific sequence of operators, to the hardware characteristics and to the overall system characteristics, and running it with the most appropriate settings; in this way the application can make the best use of the CPU and provide higher QoS. Such a re-thinking of current systems and frameworks towards a white-box also allows a cleaner support of heterogeneous accelerators. Indeed, the high-level description that we advocate allows the system to transparently map some operations to more specialized accelerators if need be. However, optimized solutions need to keep a sufficient degree of flexibility to cover the diversity of computational kernels. As an example, we explore a case study around Regular Expression (RE) matching, which is a ubiquitous kernel in data-intensive applications with limited performance on CPUs, and we propose an architecture that enhances previous work in terms of performance and flexibility, making it a good candidate for the integration with existing frameworks. Overall, this work proposes several solutions for the main issues around modern CPUs and data-intensive applications, breaking some common abstractions and advocating for an appropriate description level of those applications. The solutions proposed here leverage this level of description that enables various optimizations, providing novel guidelines in order to make the best use of the architecture. From this work, several research directions arise, especially around extending these abstractions and the related solutions to work with heterogeneous devices, whose usage for the masses calls for more automated optimization strategies and prediction models.

DAVIDE TATEO

BUILDING STRUCTURED HIERARCHICAL AGENTS

There is an increasing interest in Reinforcement Learning to solve new and more challenging problems. We are now able to solve moderately complex environments thanks to the advances in Policy Search methods and Deep Reinforcement Learning, even when using low-level data representations as images or raw sensor inputs. These advances have widened the set of application contexts in which machine learning techniques can be applied, bringing in the near future the application of these techniques in other emerging fields of research, such as robotics and unmanned autonomous vehicles. In these applications, autonomous agents are required to solve very complex tasks, using information taken from low-level sensors, in uncontrolled, dangerous, and unknown scenarios. However, many of these new methods suffer from major drawbacks: lack of theoretical results, even when based on sound theoretical frameworks, lack of interpretability of the learned behavior, instability of the learning process, domain knowledge not exploited systematically, extremely data hungry algorithms. The objective of this thesis is to address some of these problems and provide a set of tools to simplify the design of Reinforcement Learning agents, particularly when it comes to robotic systems that share some common characteristics. Most of these systems use continuous state and action variables that may need a fine-grained precision, making a good variety of deep learning approaches ineffective. They may exhibit different dynamics between different parts of the system, leading to a natural division based on different time scales, variable magnitudes, and abstraction levels. Finally, some of them are even difficult to formalize as a Reinforcement Learning task, making it difficult to define a reward function, while some human (or non-human) experts may be able to provide behavioral demonstrations. Based on these assumptions, we propose two approaches to improve the applicability of Reinforcement Learning techniques in these scenarios: hierarchical approaches to Reinforcement Learning, to exploit the structure of the problem, and Inverse Reinforcement Learning, which are a set of techniques able to extract the reward function i.e., the representation of the objective pursued by the agent, and the desired behavior from a set of experts' demonstrations. From these ideas follow the two major contributions of this work: a new Hierarchical Reinforcement Learning framework based on the Control Theory framework, which is particularly well-suited for robotic systems, and a family of Inverse Reinforcement Learning algorithms that are able to learn a suitable reward function for tasks (or subtasks) difficult to formalize as a reward function, particularly when demonstrations come from a set of different suboptimal experts. Our proposals make it possible to easily design a complex hierarchical control structure and learn the policy either by interacting directly with the environment or providing demonstrations for some subtasks or for the whole system.

SHIMA ZAHMATKESH

ON RELEVANT QUERY ANSWERING OVER STREAMING AND DISTRIBUTED DATA

Web applications that join streaming with distributed data to provide relevant answers are getting a growing attention in recent years. Answering in a timely fashion, i.e., reactively, is one of the most important performance indicators for those applications. The Semantic Web community showed that RDF Stream Processing (RSP) is an adequate framework to develop this type of applications. However, remaining reactive can be challenging, especially when the distributed data is slowly evolving, because accessing the distributed data can be highly time consuming as well as rate-limited. State-of-the-art work addresses this problem by proposing an architectural approach that keeps a local replica of the distributed data. The local replica progressively becomes stale if not updated to reflect the changes in the remote distributed data. For this reason, recently, the RSP community investigated maintenance policies of the local replica that guarantee reactiveness while maximizing the freshness of the replica. The investigated maintenance policies focus on a class of queries that join a data stream with a distributed data source. This thesis goes beyond the state of the art, focusing on finding the most relevant answers by continuously answering query over streaming and distributed data, while considering the reactiveness constraints imposed by the users. The contributions of this study are various maintenance policies, which are tailored for two classes of queries: i) queries that have to filter data in the distributed dataset before joining it with streaming data, and ii) top-k queries where the scoring function involves data that appears both in the streaming and the distributed datasets. The contributions of this doctoral thesis are advance policies that let RSP engines continuously answer the two classes of queries described above. In particular, the proposed policies focus on refreshing only the data in the replica that contributes to the relevancy of the results. For the class of queries that have to filter the distributed data, a new maintenance policy is proposed. Intuitively, the Filter Update Policy updates data which is likely to pass the filter condition and may affect the future evaluations. While the Filter Update Policy works for queries where the filter has high selectivities, other policies work better for low selectivity. To solve this problem, as the second contribution, a rank aggregation algorithm introduced to fairly consider the opinions of multiple policies simultaneously. In the next step, focusing on the class of top-k queries, the contribution is an extended top-k query evaluation which considers the join of streaming data with the distributed dataset. Keeping a local replica of the distributed dataset, two maintenance policies are proposed to approximately answer the continuous top-k query. The experimental evaluations empirically prove the ability of the proposed policies to guarantee reactiveness, while providing more accurate and relevant results than the state of the art.

ELECTRONICS

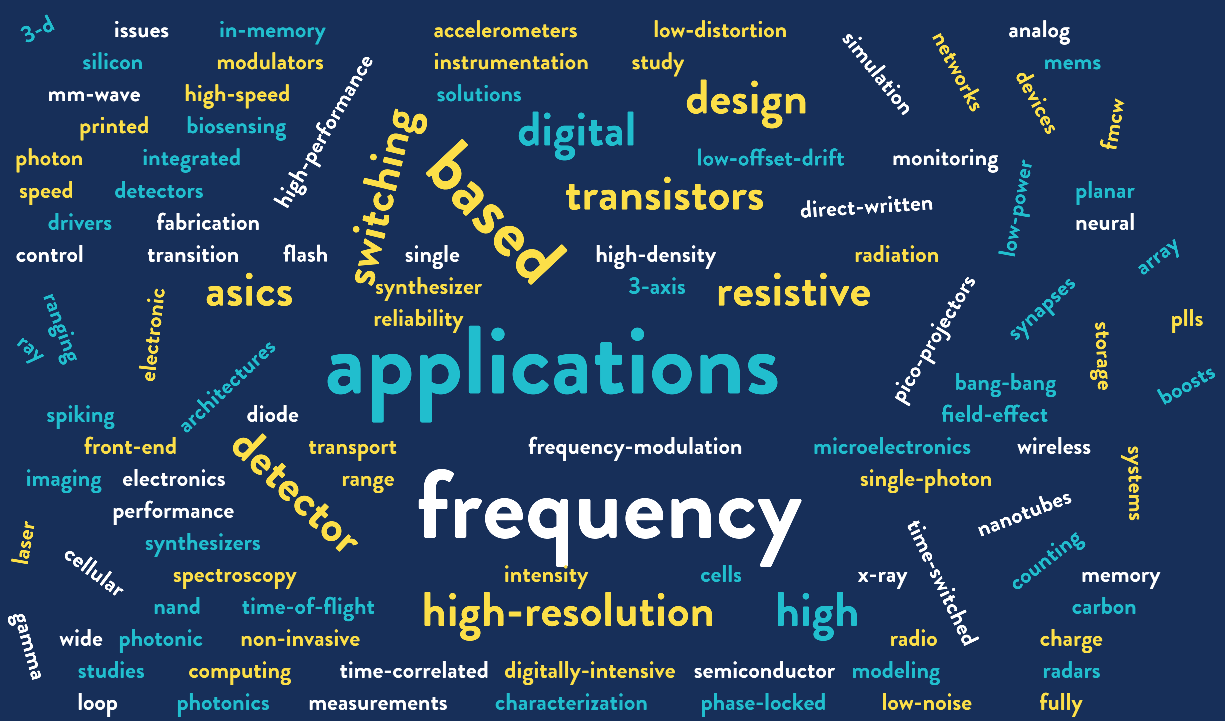

AIDIN AMIRKHANI

DESIGN OF ANALOG ASICS FOR X-RAY DETECTORS

This research project is focused on the development of readout ASICs for two main applications. The first part is mainly focused on the ASIC development for the SIDDHARTA experiment. The SIDDHARTA experiment is designed to investigate strong nuclear interactions using exotic atoms in the field of nuclear physics. Silicon Drift Detectors (SDDs) used in this experiment are arranged in arrays of 2×4 elements with total area of 612 mm2. At the final stage of SIDDHARTA experiment, 48 SDD arrays are needed to be utilized in a gantry structure to perform X-ray spectroscopy of exotic nuclei, like kaonic deuterium. Each single SDD unit in 2×4 formation of arrays is coupled to a charge sensitive preamplifier, namely CUBE, which is followed by shaping amplifier, and consequent analog and digital electronics that are all integrated on a custom developed multichannel chip called SFERA. In the first chapter of the thesis, we introduce SIDDHARTA experiment in details. After giving an insight about the main objective of the SIDDHARTA experiment and its requirements, main characteristics and working principles of SDDs will be discussed with a focus on SDDs that will be employed in SIDDHARTA Experiment. The chapter then continues with recalling basic concepts of detector signal processing as they will be used frequently in the following chapters. Finally, the charge sensitive preamplifier that will be coupled to SDDs in SIDDHARTA experiment will be introduced and discussed briefly. The second chapter introduces SFERA ASIC. SFERA (SDDs Front-End Readout ASIC), is a 16-channel ASIC that is composed of a 9th order semi-gaussian shaping filter with selectable peaking times of 0.5us, 1us, 2us, 3us, 4us and 6us as well as peak stretcher, pile-up rejection logic and different readout modalities. A dedicated digital logic, manages the data transfer to the downstream DAQ system, synchronously providing it both the "firing channels" addresses and related stretched pulse-peaks in the same time order the events are detected. This chapter then continues with the description of the detector characterization procedure as well as DAQ system used to characterize and troubleshoot detectors that will be used in final SIDDHARTA experiment. The third chapter focuses on the crosstalk and timing analyses made on SFERA and SIDDHARTA detection modules. In the SIDDHARTA experiment, in order to reject asynchronous background events, only X-rays detected within a time window opened by a kaon monitor trigger are selected. To test this detection operation, a beta source (90Sr) was used to provide a trigger and simultaneously excite fluorescence X-ray lines on a multi-element target. This chapter aims to investigate the anomalous behaviour, observed in both time and energy spectra of the SDD acquisition chain, when the detectors were exposed to radiation dominated by charged particles. The fourth chapter introduces the modifications made to SFERA ASIC to make it in-line with the experiment requirements. These modifications include but are not limited to increase of the ASIC time resolution as well as introducing new inhibit strategies that are vital for the final SIDDHARTA experiment. The second part of my research, is focused on the development of a 128-channel low-power ASIC for the readout of silicon microstrip detectors with high energy resolution and counting rate efficiency for diffractometry applications. In chapter five, we introduce the microstrip detector read-out chain and we will continue our discussion on the choice of the best filter for our application. Then the chapter continues describing the criteria considered for the choice of the filter family and its order. After this choice, count-rate linearity of the filter will be studied and statistically discussed. Later, different implementation topologies will be discussed. After deciding the filter topology, we will demonstrate on how to tune the corresponding filter components. By implementing the shaper, gain and peaking time spread will be studied through Monte Carlo analyses and by realizing the layout of the shaper, post-layout results and effects of parasitic capacitances will be discussed. In the last part of the chapter, we will discuss on the overall performance of the shaper after fabrication.

ALESSANDRO BRICALLI

FABRICATION AND CHARACTERIZATION OF RESISTIVE SWITCHING MEMORY DEVICES FOR HIGH-DENSITY STORAGE AND IN-MEMORY COMPUTING

Modern computing systems are facing fundamental challenges in processing large amounts of information efficiently, mainly due to the latency gap existing between the CPU and storage devices. Many approaches to solve this problem have been proposed, from the development of novel high-performance memories, to the implementation of non-Von Neumann computing paradigms such as neuromorphic computing and in-memory computing. Emerging non-volatile memories and devices based on 2D materials are promising candidates for this kind of application. In this work, the fabrication and characterization of novel resistive switching devices for high-performance memories and in-memory computing are discussed; moreover, the implementation of innovative in-memory computing systems based on such devices is addressed.

DMYTRO CHERNIAK

DIGITALLY-INTENSIVE FREQUENCY MODULATORS FOR MM-WAVE FMCW RADARS

The vast number of FMCW radar applications generates the demand for highly-linear, low-noise and reconfigurable fast chirp synthesizers implemented in high-volume deep sub-micron CMOS technologies. Conventional analog PLL-based chirp synthesizers realized in bipolar or BiCMOS technologies demonstrate an excellent phase noise performance, however, the modulation speed is typically limited by a narrow-bandwidth PLL. The emerging digital PLL-based synthesizers demonstrate prominent phase-noise and modulation-speed performance along with high degree of reconfigurability. The first part of the thesis derives the specification of the FMCW frequency synthesizer employing a new radar system model which was developed within the scope of this work. The proposed system model allows to analyze the impact of the chirp synthesizer impairments, such as phase noise and nonlinearity, on the performance of the FMCW radar system. The second part is focused on the analysis and design of two digital PLL prototypes in CMOS, which includes a novel pre-distortion scheme for DCO and DTC nonlinearity correction. The first implemented digital PLL-based FMCW modulator prototype is fabricated in 65-nm CMOS technology and demonstrates above-state-of-the-art performance of fast chirp synthesis. It is capable of maximum chirp slope of 173 MHz/μs and idle time of less than 200ns after an abrupt frequency step with no over or undershoot. The 23-GHz digital bang-bang PLL consuming 19.7 mA exhibits the phase noise of -100 dBc/Hz at 1MHz offset from the carrier and the worst case in-band fractional spur level is below -58 dBc. The second PLL prototype was developed and fabricated in 28 nm CMOS technology focusing on low-fractional spur operation. The novel digital pre-distortion scheme was applied to mitigate DTC nonlinearity in an 18-GHz digital bang-bang PLL which achieves in-band fractional spur level below -63dBc and exhibits the phase noise of -100 dBc/Hz at 1 MHz offset from the carrier.

ALESSANDRO COMINELLI

HIGH-SPEED, LOW-DISTORTION SOLUTIONS FOR TIME-CORRELATED SINGLE PHOTON COUNTING MEASUREMENTS

Nowadays, Time-Correlated Single Photon Counting (TCSPC) represents a key measurement technique in many scientific and industrial applications demanding for the acquisition of extremely fast and faint luminous signals with picosecond resolution. In particular, in life sciences time-resolved imaging by means of TCSPC is the enabling technology for several powerful analytical techniques, such as Fluorescence Lifetime Imaging Microscopy (FLIM), Förster Resonance Energy Transfer (FRET) or Fluorescence Lifetime Correlation Spectroscopy (FLCS). At the same time, it is becoming a gold standard in remote sensing applications, like LIght Detection And Ranging (LIDAR). In a typical TCSPC experiment, a sample is excited by means of a periodic laser source. Then, photons re-emitted by the sample are recorded to form a histogram, depending on their arrival times within the excitation period. In this way, after many periods the histogram represents a measurement of the average waveform of the luminous signal. Given most modern time-measurement circuits, TCSPC permits to reach a timing precision as low as 10 ps, that is well below the minimum resolution achievable by any analog recording technique. Unfortunately, the many advantages of TCSPC come along with a major drawback, that is a relatively long acquisition time. In particular, two effects concur in limiting the maximum measurement speed of a TCSPC acquisition channel. First of all, a conventional TCSPC system can detect only one photon per excitation cycle. As a consequence, if more than one photon impinges on the detector during a period, the reconstructed wave-form undergoes a distortion, which is known as classic pile-up. In order to avoid this issue, the intensity of the excitation source is typically adjusted to keep the average number of impinging photons in a period well below 1 (typical values range between 0.01 and 0.05). It follows that a relatively high number of excitation cycles is required to accumulate a statistically relevant number of events in the histogram. The second limit to the measurement speed is related to a relatively long dead time of both detector and time-measurement electronics, which typically ranges in the order of 100 ns. In this scenario, a TCSPC experiment can undergo a significant loss of events, thus leading to a further reduction of the measurement speed. In the last decade, TCSPC acquisition systems have been subject to a fast trend towards the parallelization of many independent channels in order to speed up the measure. On one hand, some multichannel modules based on discrete components are already available in the market, featuring the best in- class performance in terms of resolution and linearity, but the high power dissipation and the volume occupied by a single channel have limited the degree of parallelism to only 4 or 8 channels so far. On the other hand, the exploitation of CMOS technology has permitted the integration of hundreds and even thousands of independent channels on the same chip, including detectors, represented by Single Photon Avalanche Diodes (SPADs), and the whole acquisition and conversion electronics. Nevertheless, large arrays proposed so far with detectors and electronics integrated on the same chip suffer from a trade-off between number of channels and performance. In particular, the integration of both detectors and conversion electronics in the same pixel area has imposed tight constraints on power dissipation and area occupation of the electronics, limiting timing performance, both in terms of linearity and precision. At the same time, the exploitation of standard CMOS technologies has prevented designers from having the necessary degrees of freedom to pursue the best detector performance, in terms of photon detection efficiency, dark count rate and afterpulsing probability. Even worse, large multichannel systems are typically affected by a data-transfer bottleneck, which strongly limits the achievable measurement speed. In particular, the presence of a huge number of detectors can give rise to a considerably high data rate at the output of the system, which can easily reach 100 Gbit/s. Unfortunately, the real-time management of such a high rate demands for a huge bandwidth of the bus directed toward the external processor and for a considerable complexity of the system design. Instead, the maximum available transfer bandwidth is typically limited in the order of 10 Gbit/s. As a result, the efficient exploitation of a limited transfer bandwidth is, to date, one of the major challenges designers have to face to pursue the highest speed in TCSPC experiments. Recently, different readout architectures have been proposed in literature to cope with a limited transfer bandwidth, trying to maximize its exploitation under typical operating conditions. Nevertheless, solutions proposed so far are affected by relatively low efficiency and the measurement speed still lies well below the limit imposed by the saturation of the transfer rate towards the elaboration unit. The goal of this thesis work is to investigate novel approaches to speedup TCSPC measurements, avoiding at the same time any trade-off with performance. First of all, I deeply investigated the problem of pile-up distortion, which currently represents the major limitation to measurement speed in a single TCSPC acquisition channel. In this context, I propose a novel solution to keep pile-up distortion below a negligible value, paving the way to a remarkable increase of the excitation power, well above the classic pile-up limit, thus leading, in turn, to a significant speedup of TCSPC experiments. In particular, I theoretically demonstrated that negligible distortion (below 1%) is guaranteed if the dead time associated with the converter is kept below the dead time of the detector, and at the same time the detector dead time is matched to the duration of the excitation period. In this way, the speed of TCSPC experiments can be increased by a factor larger than 7.4, that is almost an order of magnitude, while providing negligible distortion regardless of the experimental conditions. It is worth noting that the proposed technique allows a single acquisition channel to reach a remarkable measurement speed, which can be achieved, to date, only using eight independent TCSPC measurement channels operating in parallel. In this scenario, my solution requires a considerably lower complexity of the system design and, even better, it can be easily extended to a multichannel approach to further increase the measurement speed. Moreover, a practical use of my solution is already feasible exploiting recently-proposed electronics, that is time-measurement circuits with negligible dead time and a SPAD coupled to a fast active quenching circuit, featuring a short and finely tunable dead time. In this work, I also investigated different approaches to address the trade-offs which currently affect large multichannel arrays, trying to maximize, at the same time, measurement speed and timing performance. In particular, I present a novel readout architecture, which has been conceived starting from a completely different perspective with respect to readout architectures proposed in the literature: a large detector array is shared with a limited set of high-performance time-measurement circuits, whose number is calculated starting from the maximum manageable data rate; then a smart routing logic has been designed to dynamically connect a large number of SPAD detectors to the external time-measurement electronics, in order to take full advantage of the available bus bandwidth. In addition, the proposed router-based architecture permits to exploit different technologies to design the various parts of the system, i.e. detectors, sensing electronics and time-measurement circuits, in order to optimize their performance. The core of the router-based architecture is a selection logic, whose task is to select a subset of the detectors carrying a valid signal during each excitation cycle, to connect them to the external converters. It is evident that a certain elaboration time is required to carry out this operation, so a low-jitter delay line has been designed to be integrated along with each pixel of the array, in order to preserve the timing information related to a photon detection, until the logic elaboration has been carried out. Then proposed delay line is able to provide a digitally-programmable delay up to 50 ns, while timing jitter is kept around 0.1% of the average delay, thus permitting excellent timing performance. At the same time, particular care has been devoted to the minimization of power dissipation and area occupation of the circuit, to make it compatible with the integration of a delay line for each pixel of a large array. Then, I designed a novel routing algorithm exploiting digital gates distributed in a tree structure, aimed at the future realization of a 32_32 array. The proposed algorithm is able to dynamically connect the array to five shared conversion channels operating at 80 MHz, thus providing an overall throughput up to 10.4 Gbit/s, including 2 bytes for the timing information and 10 bits to address the selected pixels within the array. In addition, the designed logic has a double advantage: it permits to minimize at the same time the elaboration time and the number of interconnections crossing the system, which is a major issue in dense multichannel arrays.

MASSIMO GANDOLA

LOW-NOISE, LOW-POWER FRONT-END ASICS FOR HIGH-RESOLUTION X AND GAMMA RAY SPECTROSCOPY FOR RADIATION SEMICONDUCTOR DETECTOR

The aim of this thesis is the study, the design and the characterization of two ASICs (named RIGEL and LYRA) employed in projects of the Agenzia Spaziale Italiana (ASI) and of the European Space Agency (ESA) regarding the study of deep space objects and other astrophysics applications. A dedicated readout electronics has been developed separately for each project in order to acquire and elaborate the signals coming from a SDD (Silicon Drift Detector) that detects X radiation. The design phase has been carried ùout carefully in order to satisfy the most important requirements for the ASICs: the low electronic noise and the reduced power consumption.

MICHELE GIORGIO

HIGH-RESOLUTION DIRECT-WRITTEN FIELD-EFFECT TRANSISTORS FOR HIGH FREQUENCY APPLICATIONS

Organic electronics is gaining more and more attention from scientific public and from industries because of his unique properties such as flexibility, transparency, and solution processability. High throughput manufacturing methods such as roll- to-roll coating and inkjet printing allow organic electronics to be suitable for new and diverse applications and to be an attractive viable way to make low-cost electronics. Before employing this technology for the production of complex circuits, the performance of the elementary block, i.e. the transistor, has to be optimized. The willingness to fabricate transistors capable of high-frequency operation is spurred by possible applications like high-resolution flexible displays or devices able to communicate via wireless. So far, a record frequency of transition of 27 MHz is achieved for transistors with lithographic contacts and evaporated semiconductors, while of 20 MHz for transistors fabricated without using masks in the production flow. Such record frequencies of operation were not measured directly but they were extrapolated due to bandwidth limitation, typically in the order of 1 MHz, present in the measurement setups used. Thanks to recent improvements in polymers charge carrier mobility, further enhancement of devices maximum operational frequency is in sight, and therefore it is not possible to rely further on such measurement methods for characterizing the upcoming organic high-frequency devices. In this thesis work, we resorted to Scattering parameters (S-parameters) for reliable determination of the frequency of transition without extrapolations. We report n-type FETs based on a solution-processed polymer semiconductor where the critical features have been realized by a large-area compatible direct-writing technique, allowing to obtain record frequencies of transition of 19 MHz in the case of n-type and 24 MHz for p-type polymer transistors. This is the first report of solution-processed organic FETs characterized with S-Parameters. We then made a step forward in the improvement of the OFETs maximum operational frequency by reducing the gate to source and gate to drain geometrical overlap to some hundreds of nanometers by fabricating also the gate contact by laser sintering. In this way we succeeded in the demonstration of the fastest n- and p-type organic transistors ever realized, showing transition frequency of 70 MHz and 100 MHz, respectively. Finally, in order to fabricate high-performance organic circuitry, we propose a method for self-assembled monolayers (SAMs) formation by using inkjet printing techniques, allowing to grow different SAMs on the desired contacts on the same substrate. We demonstrated both n and p unipolarized devices starting from a single ambipolar semiconductor uniformly bar-coated on the substrate, with mobilities higher than 1 cm^2/Vs in both cases.

LUIGI GRIMALDI

FREQUENCY SYNTHESIZERS BASED ON DIGITAL PLLS FOR CELLULAR RADIO APPLICATIONS

Current and future mobile communication standards are targeting a new 10x increase in data rate in the following 10 years, as it is required from 5G standard for instance. To meet these expectations, new and more complex modulation schemes are being standardized, and wider bandwidths at higher carrier frequencies are being employed. Modern transceivers are asked to manage a continuously increasing data rate while keeping high spectral efficiency at a restrained power consumption. Nowadays techniques as the beamforming and phased-arrays proved to be effective solutions to meet the stringent targets imposed by the standard, in terms of SNR of the signal chain and power efficiency. In this context, frequency generation circuits, in both transmitter and receiver side, are asked to reach lower phase noise and spur levels at lower power consumption, while operating in the range of several tenths of gigahertz. At the same time, the implementation of such circuits in new and more scaled CMOS processes imposes a radical change in the design methodology. In the last few years, digitally-assisted analog design, applied not only to frequency synthesis, has been proven to be effective in improving performance in more scaled CMOS nodes. In this scenario, the Ph.D. activity has been devoted to study and implement digital phase-locked loops (DPLLs) for future mobile communications standards, to get a deeper understanding of the adaptive filtering techniques employed in such circuits and to extend their application with the aim of improving the noise and spur-versus-power compromise. From this perspective the focus of this work is on the dissertation of digital techniques to improve the overall system performance to make it suitable for tight standard requirements. The digital approach makes it possible to exploit calibration techniques, working in the background of the PLL operation, to correct for the analog impairments, providing all the advantages of low-power operation, low area occupation, repeatability and portability to more scaled technology nodes. This work demonstrates the effectiveness of DTC-assisted digital PLL based on a single bit (bang-bang) phase detector when used as frequency synthesizer for modern transceiver architectures: the adoption of digital calibration algorithms leads to low-spurs and low-jitter whereas the utilization of the bang-bang phase detector, instead of a power-hungry multi-bit time-to-digital converter (TDC), is the key point to achieve low-power operation. Two test-chip in 65-nm CMOS technology, the first one in the sub-6GHz band at a center frequency of 3.7GHz, the second one in the mm-Wave band around 30GHz, have been implemented to show the validity and the performance of such approach.

EMANUELE GUGLIELMI

ELECTRONICS BOOSTS PHOTONICS: DETECTOR AND ELECTRONIC DESIGN FOR NON-INVASIVE MONITORING AND CONTROL OF SILICON PHOTONIC SYSTEMS

Electronics is an essential tool that can unlock the true potential of modern Silicon Photonic technologies, overcoming their limitations. My thesis contributes to the field of electronic-photonic integration, studying and improving the innovative CLIPP detector and developing the electronics to use it in novel scientific applications. Silicon Photonic technologies can achieve outstanding datarates, low losses, and power consumption, but require the closed-loop control of the optical devices, to compensate for their extreme sensitivity to fabrication tolerances and temperature fluctuations. The large-scale implementation of feedback control systems is halted by the inadequate state-of-the-art photo-detectors, that introduce losses to monitor the working point of the photonic circuits. The CLIPP, ContactLess Integrated Photonic Probe, is an innovative detector developed at Politecnico di Milano that overcomes these limitations by enabling non-invasive light monitoring in silicon photonic circuits, through an impedance measurement of the waveguide. My thesis explores the disruptive applications in Silicon Photonics that an innovative device like the CLIPP has unlocked. The work has mastered the electronic detection of the signal coming from the CLIPP, designing specific circuit implementations, in both standard discrete electronics and integrated CMOS technology. In particular, a new pseudo-resistor topology has been designed and tested to bias low-noise Trans-Impedance Amplifiers with capacitive feedback, achieving high linearity, wide bandwidth, three decades of resistance tuning range and dynamic compensation of the output offset. In order to improve the tiny signals available when measuring very low optical powers, this thesis has also addressed the sensor itself by demonstrating that more efficient designs can be achieved by exploiting the deep implantations at the same level of the waveguide, offered by active Silicon Photonics technologies. The new design achieved the best sensitivity ever measured with CLIPPs, while giving an insight into new aspects of the device to be studied in the future. The effectiveness of CLIPP-assisted circuit control was exploited for light path tracking, reconfiguration, and thermal crosstalk compensation on a switch fabric router, and for light mode unscrambling on a novel topology for Mode-Division Multiplexing, that requires non-invasive monitoring to avoid disrupting the orthogonality of the spatial modes. The results would not have been possible without advanced control strategies implemented on a reconfigurable FPGA-based electronic platform specifically conceived for Silicon Photonic applications.

CRISTIANO ROCCO MARRA

TIME-SWITCHED FREQUENCY-MODULATION FOR LOW-OFFSET-DRIFT, WIDE RANGE, FULLY INTEGRATED 3-AXIS MEMS ACCELEROMETERS

In this PhD thesis, a novel working principle for 3-axis frequency-modulated (FM) MEMS accelerometers is proposed. Indeed, the differential frequency readout is performed through the innovative time-switched approach: this methodology is based on a double sampling of the oscillation frequency of a single resonator, consecutively biased in two different configurations in time. The technique enables to avoid offset thermal drift contributions typical of differential resonant accelerometers based on two distinct resonators with unavoidable mismatch in the temperature coefficient of frequency. After a complete behavioral modeling of the conceived system, a compact tri-axis MEMS structure is designed and fabricated using an industrial process. The encouraging results obtained from the characterization of the sensor motivated the design of an integrated analog oscillator, in order to prove the possibility to keep the same key performance coupling the device with a low-power, small-footprint ASIC. To this aim, two different feedback oscillator loops are analyzed and fabricated, theoretically and experimentally identifying the optimum topology in terms of power-noise trade-off. Finally, the tri-axis MEMS-ASIC combo is coupled to an integrated frequency-to-digital converter, demonstrating the feasibility of a fully integrated, digital output, consumer grade tri-axis FM time-switched MEMS accelerometer. The designed system solves the trade-off between offset thermal drift and full-scale-range experienced by state-of-the-art capacitive AM accelerometers. The ratio between these two critical parameter results improved by more than an order of magnitude, without any post-acquisition digital temperature compensation. At the same time, the other key parameters (as resolution, power consumption and bandwidth) remain in line with consumer devices currently on the market. Thus, the presented approach represents a promising strategy to face tight requirements of next-generation applications as mixed-reality and pedestrian inertial navigation.

VALERIO MILO

MODELING AND SIMULATION OF SPIKING NEURAL NETWORKS WITH RESISTIVE SWITCHING SYNAPSES

Nowadays, functions such as the recognition of speech, image, and faces are becoming essential features in various applications, from smartphones to driver-less cars. However, the development of these cognitive systems offers critical challenges, in terms of accuracy, speed, energy efficiency, and cost. In this context, achieving brain-inspired circuits and algorithms seems a very promising pathway, given the unrivaled energy efficiency and computing performance of the human brain. Toward this goal, nanoscale devices capable of mimicking the biology of learning in silico, such as resistive switching memory (RRAM), are highly valuable. In this Ph.D. dissertation, the implementation of spike-timing dependent plasticity (STDP) and spike-rate dependent plasticity (SRDP) biological learning rules by RRAM-based synaptic structures is first described. Afterward, moving from device to network level, simulation and experimental results supporting the capability of two-layer feedforward spiking neural networks with RRAM-based synapses to achieve unsupervised learning of visual patterns are presented. Finally, the ability of Hopfield-type recurrent spiking neural networks with RRAM-based synapses to capture other fundamental brain-inspired primitives such as associative memory, pattern completion and error correction is also investigated by simulations.

GIANLUCA NICOSIA

PERFORMANCE AND RELIABILITY ISSUES OF NAND FLASH CELLS AT THE TRANSITION FROM PLANAR TO 3-D ARRAY ARCHITECTURES

NAND Flash memory technology is, nowadays, the most successful integrated solution for the nonvolatile storage of information. Its success stems from an uninterrupted increase in storage density, and, therefore, reduction of cost per gigabyte, that has continued during the past 20 years. Until around 2015 this was achieved thanks to scaling, i.e. the reduction of the feature size of each memory element. However, both physical limitations arising from the discrete nature of charge and matter and the soaring manufacturing costs due to process complexity have pushed the industry to a shift of paradigm. This is represented by the introduction of three-dimensional (3-D) NAND Flash arrays in which vertical stacking of many memory layers replaces cells miniaturization, thus allowing to break the traditional trade-off between the size of each memory element and storage density. In spite of the fact that the transition from a planar to a 3-D architecture has brought to a general improvement of the reliability of memory arrays, 3-D integration requires to implement a polycrystalline silicon channel in place of a monocrystalline one. The high density of spurious states at the polysilicon grain boundaries is responsible, first of all, for an increase in the resistivity of the channel material. Moreover, charge trapping/detrapping at these trap states represents an additional contribution to threshold voltage instabilities, negatively affecting array reliability. In this framework, this Ph.D. thesis aims at understanding the main reliability issues affecting NAND Flash technology at the transition from planar to 3-D arrays, highlighting the physical constrains that have halted horizontal scaling and discussing the new emerging issues that affect next generation 3-D NAND Flash arrays and that arise from the polycrystalline nature of the channel implemented in them.

DAVIDE PORTALUPPI

MICROELECTRONICS AND INSTRUMENTATION FOR SINGLE-PHOTON IMAGING

In recent years, an increasing number of applications have emerged which can significantly benefit from the ability of detecting fast and faint light signals, requiring sensitivity down to the single-photon level and sub-nanosecond resolution; these requirements are usually paired with the necessity of acquiring high frame rate, two- and three-dimensional movies of the scene under analysis. These applications range from industrial and automotive uses, such as object or obstacle recognition, road safety, depth-resolved ambient surveillance, to biomedical applications like fluorescence lifetime microscopy or time-resolved spectroscopy, and even to consumer applications such as gaming and gesture recognition. A number of technologies exist which can reach single-photon sensitivity; however, among those, SPADs (Single-Photon Avalanche Diodes) are the only solution capable of simultaneously achieving single-photon sensitivity, time-resolved detection and high frame rate, in a rugged and compact solution that can be realistically employed outside of a laboratory environment. As an example, Charge Coupled Devices (CCDs) can reach very high sensitivity, almost to the single photon, but suffer from very long integration times and require cooling systems to reduce their noise; electron multiplying or intensified CCDs can indeed detect single photons, but still require cooling and their bulky and delicate setups make them ill-suited for uses outside of a laboratory environment. Lastly, technologies such as Photomultiplier Tubes (PMTs) or Microchannel Plates (MCPs) also exist, but their vacuum-tube nature hinders integration with per-pixel electronics. SPAD arrays however are excellent candidates for such applications, due to their high detection efficiency, relatively low voltage operation, high timing resolution and compatibility with standard CMOS fabrication processes, thus allowing the development and exploitation of monolithic circuits integrating both sensor and processing electronics. This Ph.D. project aims at developing new CMOS SPAD imaging chips able to satisfy the numerous requirements that arise from such a varied set of applications: namely, high resolution (tens of picoseconds) time measurement capabilities, time-gated detection of incoming photons, simultaneous measurement of the photon counts and their time of arrival for reconstructing intensity and depth resolved images, and flexible readout systems to allow the end-user to tune the detector data- and frame-rate.

FRANCESCA SCURATTI

PRINTED CARBON NANOTUBES BASED TRANSISTORS: FROM CHARGE TRANSPORT STUDIES TO BIOSENSING APPLICATIONS

Solution-processable high mobility semiconductors, such as polymers and single-walled carbon nanotubes (s-SWCNTs), offer a concrete opportunity to develop high performance electronics integrable in novel large-area and lowcost technological applications, ranging from flexible and wearable devices to soft bioelectronics and biosensors. The possibility to formulate stable dispersions of carbon nanotubes by noncovalent functionalization through conjugated polymers allows the adoptin of solution-based printing deposition techniques, cheap and suitable for scalable fabrication processes. A deep knowledge of charge transport mechanisms in s-SWCNTs networks is crucial in order to address an efficient processing and realize high-performance devices and circuitry. Unfortunately, it is still unclear how the polydispersity of the semiconducting nanotubes and their interaction with the functionalizing polymer affect charge injection and transport in such networks. The interest in unraveling the dynamics of transport to enhance the devices' figures of merit is transversal to multiple fields, as carbon nanotubes networks can be employed also for bio sensing, thanks to their excellent stability in aqueous environments and biocompatibility, making them ideal candidates for developing a pervasive bridging technology for a variety of applications both in vitro and in vivo. This dissertation reports the investigation of charge transport dynamics in printed SWCNTs-based transistors and their application as cell-proliferation monitors. Charge accumulation and transport are first studied in monochiral and mixed networks of carbon nanotubes using charge-modulation spectroscopy, a technique ideally suited to study the nature and distribution of mobile carriers in working transistors. In agreement with previous experimental and theoretical studies, the mobile carrier distribution changes with applied gate voltage. Subsequently, several analytical tools were employed to assess the impact of network density and functionalizing polymer on charge transport mechanisms in inkjet printed networks of polymer-wrapped nanotubes, demonstrating the possibility to tune their ambipolar behavior and their charge transport efficiency by tailoring the printing process. Our findings provide useful insights on the mutual interplay between polymer and nanotubes, as well as on the dynamics of charge redistribution on the SWCNTs chiralities upon changes in network density. We then propose a practical application for such networks, introducing carbon nanotube based electrolyte-gated transistors as a novel tool to electrically monitor cell adherence and proliferation, able to operate consistently in aqueous environment and in conjunction with more established optical techniques. By establishing a direct relation between presence of adhered cells and electrical output, we demonstrate that carbon-based electrolytegated transistors constitute a promising technology to perform large-scale and cost-effective in vitro monitoring for toxicology and drug development applications.

ALIREZA TAJFAR

HIGH SPEED AND INTENSITY LASER DIODE DRIVERS FOR TIME-OF-FLIGHT 3D RANGING AND PICO-PROJECTORS

In the last decade, the ability to drive lasers with high frame-rate (higher than standard video-rate), with very faint illumination, has become more and more important in many fields, like ambient surveillance, road safety, identification of people and objects, biomedical imaging, studies on physics of materials as well as commercial applications such as gaming, laser-based projection and augmented virtual reality glasses. There is a growing interest in devices capable of projecting or acquiring high framerate 2D or 3D videos. This work presents the design of a 4-channel integrated laser diode driver (LDD) fabricated in 160 nm BCD technology for a low-power pico-projector application based on micro-electromechanical systems (MEMS) micro-mirrors and the design of a high-power and intensive single chip LDD in the same technology for direct time of flight (TOF) measurement in Light Detection and Ranging (LiDAR) application. The embedded 10-bit current DAC in pico-projector LDD known as video DAC, can produce 300MHz sharp current pulses with 1ns rise/fall time, less than 7% overshoot, 3-4ns settling time while the full-scale range (FSR) is also programmable through a static 10-bit current DAC, known as scale DAC, in the range of 160uA up to 160mA. These performances are guaranteed thanks to the novel active bootstrap presented in this work. TOF LDD, as the second part of this Ph.D. thesis, can produce a fully programmable sharp current pulses up to 20A with less than 1ns duration and a repetition rate of 40MHz. to reach such high performances, a new driving topology is proposed which has brought new challenges in the design that must be mitigated by innovative solutions. Chapter 1 illustrates a general introduction about both applications, pico-projection and time of flight measurements, focusing on the lase driving part and introducing well-known solutions, their drawbacks and possible improvements as would be explained in detail in this work. Chapter 2 presents a brief overview of the laser principles, specifically the semiconductor laser diode operation. This section is written to be more generic to cover both applications, but the main focus is on the ToF application as it is more sensitive to Laser behavior due to the high-frequency requirements. It gives a fundamental background about the laser operation and behavior according to the driving pulse, threshold current, temperature variation, and current pulse shapes. Chapter 3 focuses on the pico-projection LDD. The operation principle in system level and the laser driving role in such a system. The high-frequency compatibility, low power, high range FSR programmability, dynamic performances such as rise time, fall time, settling time are the challenges of the design that must be studied carefully. The solutions for such challenges end up with innovative solutions that are illustrated in detail within this chapter. The simulation and measurement results are shown in this chapter for each dedicated block. Chapter 4 focuses on the time of flight LDD. The dToF principle and its differences with respect to indirect ToF (iToF) are well described and its distinct design considerations such as optical pulse width, repetition rate, rise time, optical power are clarified as design challenges for such LDD. To improve the conventional laser driving approach, which is a limiting point to get further improvement in the dToF specs, a new driving topology is introduced. The new topology has some challenging drawbacks that require innovative solutions to deal with, such as, clamping circuit, PVT compensated monostable, eye-safety fault detection and out sync drift-less level shifter to drive TDC. In this chapter, the proposed solutions and their validations are described in detail and testified by means of simulations.

TUAN MINH VO

A STUDY OF HIGH-PERFORMANCE FREQUENCY SYNTHESIZER BASED ON DIGITAL BANG-BANG PHASE-LOCKED LOOP FOR WIRELESS APPLICATIONS

High-performance frequency synthesizer is a fundamental part of almost any modern wireless communication device, for example, used for coherent demodulation/modulation in wireless transceivers. The frequency synthesizer based on phase-locked loop (PLL) architecture, serving as a local oscillator in a transceiver, is indeed a negative feedback control system generating an output signal whose frequency is multiple of the reference signal frequency. The multiple can be an integer or a fractional number. Though fractional-N PLLs entail the key advantage of a finer frequency resolution, the noise-power figure-of-merit (FoM) of state-of-the-art integer-N PLLs is still better than in the fractional-N case. In addition, digital PLL synthesizers are taking over conventional analog ones, because of their benefits in terms of power consumption and area occupation in ultra-scaled CMOS technologies. The digital solution simplifies the design and, as this is portable to the next technology nodes, may potentially reduce the time-to-market. In this study, a type-II fractional-N digital PLL having a phase detector (PD) with only two output levels of 1 and −1 is of interest. Though this topology of fractional-N digital PLL has been demonstrated in practice being able to obtain a FoM close to the best of the integer-N ones, there is a lack of theoretical literature that explains in detail adopted techniques in the system. It is known that the one-bit (also known as bang-bang, BB) phase detector, when employed in frequency synthesizers, only acts like a linear element when phase (time) error at the PD input is dominated by random noise. Since the quantization noise of the digital Delta Sigma modulator (DSM) dithering the modulus control of the frequency divider is not white and much larger than other thermal random noise in the system, this noise may cause a limit cycle in the BB-PLLs. To address this issue, a digital/time converter (DTC) is placed in the feedback path between the divider and the BBPD. The control gain of the DTC is automatically adjusted in background by a calibration loop operating based on the principle of the leastmean-square (LMS) adaptive filter. The calibration loop helps the PLL to adapt to changes in the digitally-controlled oscillator (DCO) period as well as the DTC characteristics in practice. By far, in the presented publications, to guarantee a short convergence time, an at least second-order DSM is required when a fine frequency resolution is desired. This follows a large DTC range of twice of the DCO period in the system and this compromises jitter performance. In addition to the issue related to the quantization noise, the BB-PLLs also face an extremely long transient process when a large jump of the output frequency is required. Indeed, the frequency locking time of the BB-PLLs is shown to be inversely proportional to the values of the loop filter gains while these gains are very small to guarantee the loop stability. To solve the problem, a frequency-aid technique has been proposed. This technique is essentially based on exploiting digital ternary phase detectors (TPDs) to create multi filters in the feed-forward path of the BB-PLL. The outputs of the filters are used to tune a multi-bank DCO in corresponding order. The frequency-aid circuit is only triggered when the time error at the BBPD input is larger than a fixed value, i.e., the dead-zone of the TPD. In the published report, the lower bound of the dead-zone is large which may not be optimum for the frequency locking time. Furthermore, even when this technique is adopted, the frequency locking transient in the conventional system still requires relatively long time in cases the DCO tuning words are at their worst conditions. The objective of this thesis is firstly to give an insight, for the first time, into the behavior during the transient of the fractional-N digital BB-PLL in two separated cases, i.e., with the frequency aid technique in the first case and with the LMS calibration loop in the second one. In order to reach this goal, analysis is carried on in the time domain for the frequency aid technique. Based on the analyzed result, we propose a novel frequency aid technique to further improve the frequency locking speed. In the worst case of the frequency locking, the proposed technique reduces the transient time by a factor of 3.5. The LMS calibration loop is evaluated in accordance to the value of the fractional part of the frequency control word (fcw). The analysis, that is carried on in both the time-domain and the z-domain, not only gives the qualitative results but also quantity results in some particular cases. Moreover, two novel calibration schemes are proposed in order to use a smaller delay range DTC while keeping a short convergence time. At the same convergence time, the required DTC time range in the first proposed scheme is 0.57 times, and, the one in the second proposed scheme is 0.55 times

as of the one in the conventional schemes, respectively. All the analysis are verified by simulation based on accurate behavioral models. The models are built with real design parameters, and, designed for the BB-PLLs synthesizing an output frequency from 3.2 GHz to 4.0 GHz from a reference frequency of 52 MHz. The original contributions of the thesis are: a novel frequency aid technique; a detailed explanation to the convergence and the noise of the LMS calibration loop; and, two novel schemes for the DTC calibration loops. The results of study have been published in three articles on international conferences.

SYSTEMS AND CONTROL

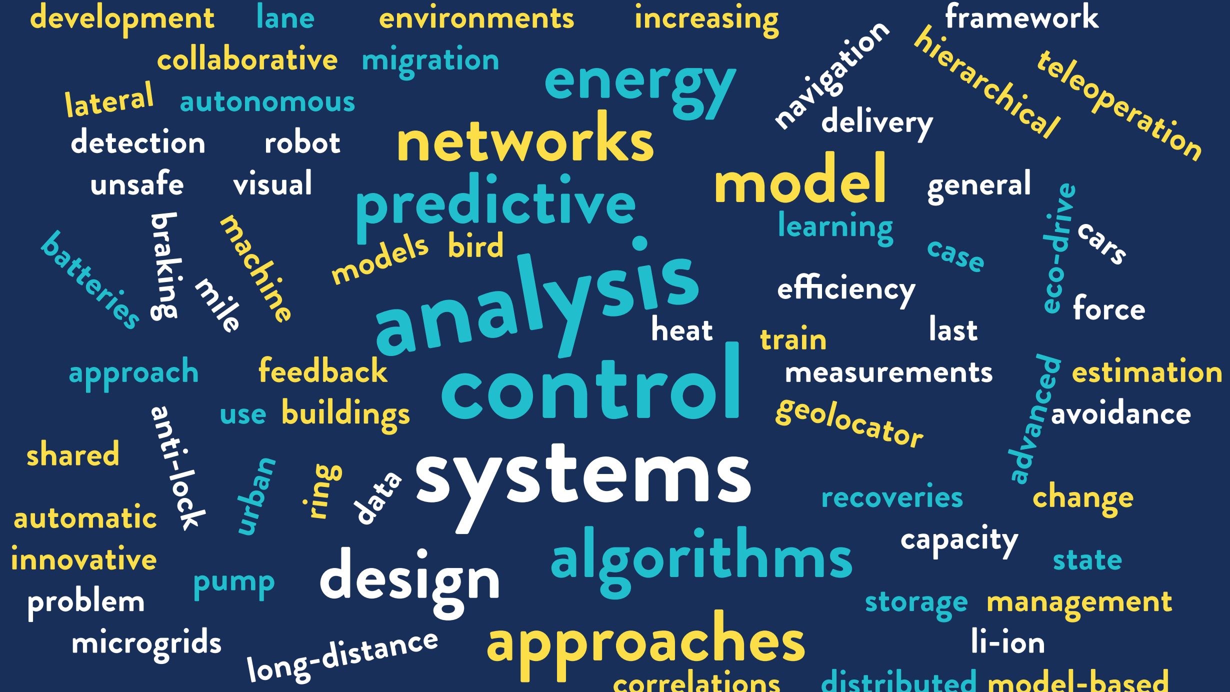

ALESSANDRO AMODIO

AUTOMATIC SYSTEMS FOR UNSAFE LANE CHANGE DETECTION AND AVOIDANCE

Road safety is currently recognized to be among the major societal issues, mainly due to road crashes being one of the major causes of death. In the past recent years, car manufacturers have been responding to this increasing need for safety by developing electronic systems called Advanced Driver Assistance Systems (ADAS), and increasingly deploying them on board of commercial vehicles. ADAS are systems designed for two- and four-wheeled vehicles which have the goal of increasing comfort and safety during everyday driving; in some cases, such systems are only asked to report information or alert the driver in dangerous situations, while in other cases they directly intervene on the vehicle dynamics to avoid potential hazards. This thesis proposes a composite and integrated system that helps the driver avoid safety hazards due to unsafe lane change maneuvers. The reasons that may lead to such maneuvers are manifold, and for this reason vehicle crashes are very often caused by unsafe lane change maneuvers. In order to avoid such dangerous situations, the proposed system is designed to perform to tasks: Prevention: The goal is to give alert signals to the driver to advise him about potentially dangerous situations that may lead to a lane change crash. This is done by checking the drunkenness state of the driver before the drive start, and by monitoring the vehicle surroundings during the drive. Intervention: The goal is to actively intervene to correct the vehicle trajectory in case a lane change maneuver is being performed in unsafe condition. The typical scenarios that are considered are twofold; the first is a vehicle drift towards the adjacent lane due to the driver losing control as a consequence of micro-sleep occurrence or drowsiness. The second is an intentional lane change performed during inattentive driving, without the driver noticing a potential danger in the adjacent lane. The thesis first introduces the ADAS systems and their main components, among which the sensors play a crucial role. For this reason, an extensive description of some of the main sensing technologies is given, namely stereo cameras, radar and lidar, and the results of on-road tests are presented to compare the different considered sensors. Then, the structure of the proposed system is presented, being composed by a Warning Module, in charge of performing the Prevention task, and the Control Module, which addresses the Intervention task. A Supervisor is presented that regulates activation and deactivation of the Control Module, together with experimental results to validate to proposed methodologies. In particular, the Supervisor is shown to allow timely activation of the controller, without providing excessive false alarms. The Warning Module is composed by two blocks; the first is a system for revealing the drunkenness condition of a driver by analyzing the dynamics of its Pupillary Light Reflex; a light stimulus is applied to the subject's eye, and the pupil constriction is recorded through cameras. A method is presented to extract the pupil's diameter profile from video sequences, from which a set of features is computed; such features are used by a classifier to discriminate between "sober" and "drunk" states of the subject. Experimental results obtained with different classification techniques on a database experimentally acquired reveal a good classification performance. The second block implements a lane change warning function, which is in charge of monitoring the vehicle's rear surroundings and give alert signals to the driver in case an object is detected in the adjacent lane in the vehicle's blind spot or approaching at high speed. Extensive validation is performed on an instrumented motorbike to show the effectiveness of the proposed solution. The Control Module is composed by a vehicle's lateral position control loop, where the actuation variable is the differential torque at the wheels. A linear controller in the PID form is tuned within the H-infinity framework, with anti-windup and gain scheduling schemes. Different actuation and activation strategies are proposed, and results of experimental and simulation tests show the effectiveness of the control system, which is able to stabilize the system and correct the vehicle trajectory in each of the considered cases of interest. In the Appendix, a vehicle's lateral position estimation method is proposed which makes use of a set of magnetic sensors mounted on the road surface. Upon transit of a target vehicle, each sensor records a magnetic signature for a different vehicle section; such signatures are then compared with a database of previously acquired signatures of known vehicles transited at known position. A classifier then identifies the unknown vehicle's model and its lateral transit position: the results obtained on real experimental data show the effectiveness of the algorithm, with an evaluation of the sucess rate.

LE ANH DAO

MICROGRIDS ENERGY MANAGEMENT WITH A HIERARCHICAL DISTRIBUTED MODEL PREDICTIVE CONTROL APPROACH

The dissertation addresses the problem of management and coordination of energy resources in a typical microgrid, including smart buildings as flexible loads, energy storages and renewables. The overall goal is to provide a comprehensive and innovative framework to maximize the overall benefit, while still accounting for possible requests to change the load profile coming from the grid and leaving every single building or user to balance between servicing those requests and satisfying his own comfort levels. The user involvement in the decision-making process is granted by a management and control solution exploiting an innovative consensus-based distributed model predictive control approach with coordination. In addition, a hierarchical structure is proposed to integrate the distributed MPC user-side with the microgrid control, also implemented with an MPC technique. The proposed overall approach has been implemented and tested in several experiments in the laboratory facility for distributed energy systems (Smart RUE) at National Technical University of Athens - NTUA, Athens, Greece. Simulation analysis and results complement the testing, showing the accuracy and the potential of the method, also from the perspective of implementation. On the other hand, special attention was also put into the work of estimating the production of Renewable Energy Resources (RESs) in general or photovoltaic in this dissertation. Specifically, short-term and medium-term predictions of the day-ahead generated power (GP) of a photovoltaic plant using predicted regional solar radiation (SR) are concerned. In this work, different predictors are developed, which are then combined with the weather forecast service using ensemble methods. Afterwards, under a similar ensemble framework, the impact of the accuracy in the prediction of meteorological variables on the quality of the GP prediction is evaluated. The validation of the approach is performed by using a pilot PV plant and several meteorological stations situated in Northern Italy.

LUCA D'AVICO

ANALYSIS AND DESIGN OF ADVANCED ANTI-LOCK BRAKING SYSTEMS

Safety is one of the major concerns on vehicles and one of the most safety-critical situation is the braking manoeuvre. The focus of this work is the analysis, design and experimental validation of anti-lock braking systems on aircrafts and bicycles. Since the early 1900s, anti-skids have been used on aircrafts based upon industrial practice and, up to now, with almost no variation of the control logic. A formal description of commercial anti-skids is provided in this thesis and the experimental comparison between the designed and the commercial one is reported as well. Innovative anti-skids have been proposed and proved to improve the performances of current solutions. Safety devices are an important features of vehicles and more often part of the mandatory equipment. However there is a physiological gap in time between the formation of well established mobility trends and the introduction of safety regulations. Bicycles are increasing popularity in most of the urban areas but not all the cyclists are expert enough to avoid dangerous situations due to the lack of appropriate safety devices. For these reasons, an anti-lock braking system for bicycles has been designed, implemented and tested on different road conditions (asphalt, clay ground, gravel/sand) showing the improvement in terms of safety without compromising the braking performance.

HAFSA FAROOQI

DESIGN OF COLLABORATIVE ECO-DRIVE CONTROL ALGORITHMS FOR TRAIN NETWORKS

In today's age, green transportation remains one of the most important topics of research. The main goal is to promote vehicle technologies and driving styles which are energy efficient and environment friendly. In this thesis, the main focus is on the Energy Efficient Train Control (EETC) or eco-driving strategies of railways. For this purpose, two main research paths have been explored. The first research direction is associated with a single train control problem, where the control problem is to find the best driving strategy for the train to go from one stop to another, given an optimal timetable. EETC strategies can be either fully automated (ATO) or serve as an advisory system to the driver (DAS) for the purpose of assisting drivers in following an energy efficient driving style. For this purpose, three control strategies using Model Predictive Control (MPC) have been presented. In the first two strategies, shrinking horizon techniques have been combined with input parametrization approaches to reduce the computational burden of the control problem and to realize the nonlinear integer programming control problem which arises in the DAS scenario, while the third strategy is based on switching MPC with receding horizon. All the strategies have been tested on the official simulation tool CITHEL of our industrial partner Alstom, and the obtained results in comparison with the existing techniques have proven to be more energy efficient. The second research direction falls under the paradigm of collaborative eco-drive control strategies, involving multiple trains belonging to a substation network. The main aim is to use the energy regenerated by the braking trains through collaboration among the trains connected and active in the network. In this case, three strategies to decide the collaborative law have been presented along with the extensions from the single train control strategies presented in the first part of the thesis. For the design of collaborative laws, techniques such as manual supervision, substation modeling and dissension based adaptive laws with concept similar to Markov chains have been used. The strategies have been validated with simulation examples. Finally, comparisons of energy efficiency with and without collaboration have been presented, which show the advantage of using the developed collaborated laws as compared to no collaboration.

OLGA GALLUPPI

INNOVATIVE APPROACHES TO THE LATERAL CONTROL PROBLEM IN CARS

Cars represent the primary means of passenger transport in the world and road accidents are one of the mainspring of premature human death. Although, a decrease in the number of road deaths is observed throughout the last years. Technology advances in vehicle dynamics control are considered one of the reasons for this trend: the automotive industry is transforming, by becoming software intensive rather than mechanically intensive. Road vehicles are nowadays very complex systems, composed of several subsystem which are often interacting among each other. Although each module is responsible for a specific function, all of the subsystems influence the vehicle dynamic behaviour and, perhaps more importantly, the quality of the driving experience which is perceived by the driver and passengers, in terms of both performance and safety. The research studies of this Thesis describe innovative approaches to the lateral control problem in cars. The derived argumentations are human-oriented: they assort and embrace relevant perspectives on the lateral control question entailing sensibility over passengers safety, over the attainment of high-performance dynamical car behaviour, but also over the ecological burden. Throughout the discussion, multiple aspects which comprehend different layers of the problem are investigated and integrated. Placing in nowadays original core research on the subject, some of the studies are re-intepreted in an autonomous driving framework, though most of them mainly being conceived for the operation in all driving settings. Contemplated theorethical and applicative advances are endorsed by experimental and simulation studies. Concerned vehicle investigations include hybrid vehicles eco-routing, electrification of driving users, data-driven MIMO nonlinear control steering input solution, sideslip angle and longitudinal speed estimation, semi-active suspension control.

STEFANO MARELLI

ANALYSIS AND DEVELOPMENT OF ELECTROCHEMICAL MODEL-BASED STATE ESTIMATION ALGORITHMS FOR LI-ION BATTERIES

Lithium ion (Li-ion) batteries are the most widely adopted technology for electric mobility and consumer electronics, thanks to their ability to store and deliver electric energy more efficiently and effectively than other chemistries. However, costs, performance limits and safety concerns are aspects that still require investments and research efforts. Since fully electric or Battery Electric Vehicles (BEVs) are still seen as a costly solution by consumers, less onerous solutions like Hybrid Electric Vehicles (HEVs) and Plug-in Hybrid Electric Vehicles (PHEVs) are more appealing choices. In these vehicles the battery pack is compact, because it is not the main or the only source of energy onboard, thus entailing less initial cost for the batteries. This puts even more attention on battery performances: a smaller battery is typically required to stand higher powers, compared to its size, than those required to larger BEVs batteries. Given the non-negligible costs and the need for compact yet high-performance solutions, batteries need to be exploited to their limit; a conservative approach would be too costly both for companies and customers. Unfortunately, Li-ion batteries are chemically unstable systems, that require Battery Management Systems (BMSs) to be operated safely and efficiently. The BMS continuously monitors and controls the battery states, such as: temperature, current, voltage, amount of remaining energy, and battery degradation. Many of these states cannot be directly measured; one of the key functions of the BMS is therefore to provide an estimate of these states. The more accurate this estimate is, the closer the battery can be exploited to its fundamental limits, which allows for an efficient and cost-effective utilization. Accurate state estimation and physical insights into cells behavior are enabled by electrochemical models. In the present thesis, two physics-based electrochemical and thermal models are implemented with efficient formulations, namely a Single Particle and Thermal Model (SPTM) and a Pseudo 2-Dimensional Thermal (P2DT) model. These models describe the dynamics of lithium concentrations and temperatures at different levels of detail. The parameters of the former model are identified experimentally on a commercial Li-ion cell. The latter model is implemented by solving the algebraic constraints on the states via a model decomposition and a control-oriented coupling of the equations. The SPTM is used to develop a Sliding Mode Observer, to estimate lithium concentration inside the model particles from the measured terminal voltage. An analytic computation of the gain matrix allows to enforce mass conservation in the cell; an extremely efficient, yet robust observer is obtained. The State of Charge (SoC) estimation error converges to less than 2.5%. Also, this model is used to design a backstepping observer, which is another computationally efficient solution. Thanks to the inclusion of bulk thermal dynamics, this observer is validated with experimental tests, showing less than 2% SoC estimation error; importantly, local concentrations and bulk temperature estimates are provided by the observer. The P2DT model is first used, without thermal dynamics, to develop an Unscented Kalman Filter (UKF), for which observability issues are solved via a soft-constraint on total lithium mass, and computational burden is reduced by more than a factor 3 with a parallel computation implementation. This approach gives SoC estimation errors of less than 5% in realistic conditions, and local concentration errors of less than 3%. With the inclusion of distributed thermal dynamics, a Dual Unscented Kalman Filter (DUKF) is designed, again with a soft-constraint on lithium mass and a parallel implementation. This structure not only allows to estimate the SoC with less than 1.5% error and the local concentration with less than 4% error, but also to estimate the temperature in any point of the cell with an error of only 0.2°C, under currents as large as 50C and noisy measured voltage.

DAVIDE NICOLIS

A GENERAL FRAMEWORK FOR SHARED CONTROL IN ROBOT TELEOPERATION WITH FORCE AND VISUAL FEEDBACK

Since the ‘50s, robot teleoperation has been employed in a variety of applications where a human user is required to operate from a distance a robotic device, often a robot manipulator. The use of telerobotics is often motivated by the inaccessibility of the environment where the task must be performed, caused by hostile conditions, such as on-orbit maintenance, and decommissioning of hazardous materials, or simply by the different scale of the workspace in robot-assisted surgery. Currently, the topic of interaction between user and robotic devices has been receiving increasing attention from the research community and the industry. As teleoperation applications and platforms grow more complex, the employed control framework should be able to relieve the user of some of the burden caused by operating such devices, establishing a sort of shared control. This work aims at proposing a comprehensive control framework for teleoperation systems comprising robot manipulators. At a local lower control level, sliding mode control theory is employed to achieve a prescribed system behavior, by robustly shaping master and slave manipulators impedances irrespective of uncertainties. An outer hierarchical optimization layer considers control and motion constraints. To help and guide the operator, the specification of hard and soft virtual fixtures is tackled at this level, with virtual force feedback rendered through the analysis of the dual solution of the optimization. A stability analysis of the overall control scheme in presence of variable communication delays during contact is performed by relying on small gain theorem and absolute stability requirements, which provide clear tuning guidelines for master and slave robot control parameters. Furthermore, optical feedback by means of visual servoing is integrated and experimentally validated on a teleoperated dual-arm platform. The proposed controller helps the user in navigating cluttered environments and keep a line of sight with its target by completely avoiding occlusions, reducing the operator workload required to complete a reaching task. Finally, machine learning techniques are employed to infer the user intention and predict his/her motion to actively assist in task execution and reduce fatigue.

MATTIA PANCERASA

DATA ANALYSIS AND MODELS FOR LONG-DISTANCE BIRD MIGRATION. FROM CORRELATIONS ON RING RECOVERIES TO MACHINE LEARNING ON GEOLOCATOR MEASUREMENTS

Climate is one of the fundamental shapers of ecosystems, thus its ongoing changes deeply influence the behavior, distribution and dynamics of plant and animal populations. Migratory birds are among the species most affected by this phenomenon, as they need to fine-tune their phenology according to the climatic conditions of their breeding and wintering areas. To investigate how and to what extent alterations of climate regimes may determine key changes in the movement ecology of migratory birds, a detailed knowledge of their staging sites, a trustable reconstruction of their migration routes and of the time schedules of their journeys is very necessary. The classic methods for studying migration, such as bird ringing, can now be complemented by new technologies, such as GPS loggers or light level geolocators, that allow to record proxies of organisms' positions throughout their routes. Focusing on a model species, the barn swallow (Hirundo rustica), in this work we first developed a method to investigate the occurrence of climatic connections between the African wintering and European breeding areas of this migratory passerine bird: surprisingly significant correlations between the average temperatures in the wintering and breeding locations of individuals emerged at the precise weeks of individuals' spring migration. Correlations have high significance only in the proximity of barn swallow wintering sites and if the temperature series refer to the precise weeks of migration. Second, we reconstructed migratory routes of 88 barn swallows using the measurements provided by light level geolocators, verifying the repeatability of the estimation method we used. The results obtained allowed us to identify four groups of individuals, as well as a possible effect of the year of migration on many indicators of the migration schedules obtained from the reconstructed routes. Third, using the routes data as reconstructed in the second research step, we have automated a long manual phase of data pre-processing by implementing filters based on Machine Learning algorithms. The migratory routes reconstructed using the automated pre-processing are completely comparable with those obtained from the manual selection of geolocator data. The work confirms that models based on data gathered with ICT devices may be helpful tools to let us gain insights on the influence of environmental and climate changes on species and ecosystems connectivity.

SOROUSH RASTEGARPOUR

MODEL PREDICTIVE CONTROL APPROACHES FOR ENERGY EFFICIENCY IN BUILDINGS WITH HEAT PUMP AND STORAGE SYSTEMS

Air-to-water heat pump is one of the most common and energy efficient way for a building heating system, particularly floor-heating plants. One way to further improve its effectiveness is to control the heat pump exploiting the load dependency of its coefficient of performance, and exploit it in the control decision in a predictive manner, anticipating its effect on the building temperature. This dissertation studies the impact of using different types of energy storages integrated with a heat pump for energy efficiency in radiant-floor buildings. In particular, on one hand, the performance of the building energy resources management system is improved through the application of model predictive control to better anticipate the effects of disturbances and real-time pricing together with following the modular structure of the system under control. To this end, the load side and heating system are decoupled through a 3-element mixing valve, which enforces a fixed water flow rate in the building pipelines. Hence, the building temperature control is executed by a linear model predictive control, which in turn is able to exchange the building information with the heating system controller. On the contrary, there is a variable action of the mixing valve, which enforces a variable circulated water flow rate within the tank. In this case, the optimization problem is more complex than in literature due to the variable circulation water flow rate within the tank layers, which gives rise to a nonlinear model. Therefore, nonlinear model predictive control is used to deal with many physical constraints and nonlinear problems. Alternatively, linear time-varying MPC is also used, based on successive linearizations around a reference trajectory. The other goal of the dissertation is to analyze the advantages and disadvantages of those MPC techniques for temperature control in radiant-floor buildings. Moreover, a robustness analysis has been conducted, showing the impact of the heat pump efficiency on the control performance.

STEFANO SABATINI

AN AUTONOMOUS NAVIGATION USE CASE FOR LAST MILE DELIVERY IN URBAN ENVIRONMENTS

Autonomous mobile robots are changing many industrial fields. Up to now their use has been limited to indoor confined applications (e.g. warehouse logistics) or outdoor applications in isolated areas (e.g. agriculture applications). In the last couple of years, due to the boom of the e-commerce, a great interest has been raised around the possibility of using robots to deliver parcels navigating city's sidewalks. Despite the enormous recent scientific and technological progress, urban environments still represent a challenge for robot autonomous navigation. In the first place, it is clear the need of a robust sidewalk level localization solution that is able to cope with GPS inaccuracies typical of urban environments. Furthermore, to navigate urban sidewalks, robots need to be extremely maneuverable and able to handle even the worst sidewalk surface conditions without getting stuck. This research aims to address these challenges by developing and experimentally validating a complete autonomous navigation solution capable of navigating urban sidewalks. An autonomous mobile robot named YAPE (Your Autonomous Pony Express) is designed from scratch based on a two wheeled self-balancing configuration. This design, although more difficult to control compared to other robots, guarantees great maneuverability and the flexibility to handle sidewalk irregularities. A method to map large scale urban areas is also developed. Using a map-based localization system, sidewalk level localization is proved even in GPS degraded and denied environments. The complete autonomous navigation system has been extensively tested in real-world situations on Milan urban sidewalks.

TELECOMMUNICATIONS

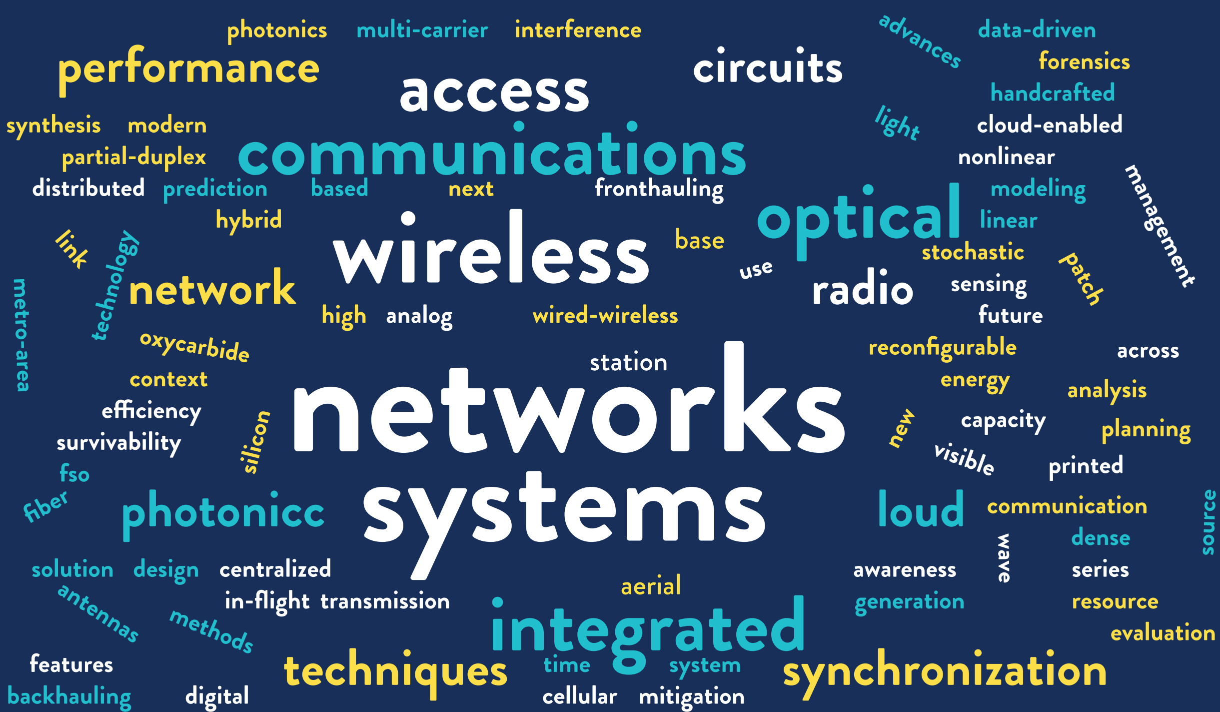

OMRAN AYOUB

RESOURCE MANAGEMENT AND PLANNING IN CLOUD-ENABLED OPTICAL METRO-AREA NETWORKS

The Internet is experiencing an exponential increase in terms of number of users, data traffic, connected devices and latency-stringent services. It is foreseen that the introduction of new technologies, like 5G, will boost this growth even further. To cope with this growth, future communication networks are required to provide unprecedented performance in terms of enhanced throughput, increased coverage, reduced latency and power consumption. Moreover, these networks must be continuously evolved, by resorting to novel technologies and architectural solutions, to meet such requirements. A promising solution consists in enhancing nodes at the network edge, taking advantage of Network Function Virtualization and Cloud Computing, with cloud capabilities, i.e., with storage and computing capabilities, such that services can be terminated locally. In this context, a strategic planning of the deployment of cloud-enabled edge nodes is needed to efficiently enhance the network performance. In this thesis, we investigate the deployment of cloud-enabled edge nodes and propose novel strategies for improved network resource management. We focus on optical metro-area networks, which are currently evolving from a rigid ring-based aggregation infrastructure to a composite cloud-network ecosystem where novel 5G cloud-based services can be implemented and supported. Specifically, we consider Video-on-Demand (VoD) content delivery service, which poses the strongest pressure on the network infrastructure being responsible for roughly 70% of the global Internet traffic, and assume a case study where edge-nodes (e.g., caches) host and terminate VoD services. First, we examine energy and techno-economic aspects of cache deployments and then propose integer linear programming formulations for efficient VoD content caching and distribution. In addition, we propose simulation-based optimization strategies and derive an analytical model to optimally plan the deployment of caches in hierarchical metro-area networks considering emerging network architectures such as fixed mobile convergence and filterless optical networks. Based on the obtained results, we provide a framework to identify the optimal cache deployment which minimizes the overall network resource occupation due to VoD delivery considering network and service characteristics. Moreover, complementary to the deployment of services, we focus on moving (i.e., migrating) services hosted on Virtual Machines (VMs) between cloud-enabled edge-nodes and data centers. Specifically, we focus on online VMs migration due to its advantages such as allowing the migration of a service from one data center to another with minimal service interruption. However, online VMs migration consumes high amount of network resources. In this context, we propose efficient routing and bandwidth assignment algorithms and integer linear programming formulations aimed at minimizing the overall network resource consumption due to online VMs migration.

MARIA ANTONIETA ALVAREZ VILLANUEVA

DISTRIBUTED SYNCHRONIZATION FOR DENSE WIRELESS NETWORK SYSTEMS

Distributed synchronization on dense wireless network constitutes a new regime for network synchronization, mainly due to the challenges involved on the inter-connectivity of large number of devices for management and coordination of the network. This inter-connectivity enables various services that are envisioned for Internet of Things (IoT), such as smart meters, smart traffic, public safety, medical metering, smart homes, etc. IoT paradigm embeds the inter-communication of massive number of devices that is challenging in terms of scalability, sustainability and improved efficiency. The inter-connectivity of massive number of devices is an open issue, and it is expected to be part of 5G ecosystem. Therefore, time and carrier frequency synchronization are critical aspects to be considered to guarantee the proper network communication. In this context, there are some characteristics to be taken into account for a proper synchronization solution in dense networks. On PHY-level context, the multiple connectivity for heterogeneous communication devices require scalability, and the constraint consumption of energy due to the capacity of the devices, requires a fast convergence and optimum processing. The proposed synchronization methodology considers the following statements: i) to allow scalability of the network; ii) to allow fast synchronization of the whole network; and iii) to mitigate power consumption. The network scenario consists of a dense and not coordinated wireless connected nodes without any external agent as a reference, where the internode distance is small (compared to the bandwidth) to neglect any propagation delay. Each node is equipped with a local free-running reference that skews from the others, and local control is by changing timing offset (TO) and carrier frequency offset (CFO) on each node independently of the others. The periodic synchronization allocates specific signatures in data communication to exchange the synchronization state by every node so that the network iteratively reaches a global convergence. In this Doctoral Dissertation, a distributed synchronization algorithm based on consensus paradigms is proposed, that enables the network to reach asymptotically a global convergence based on the exchange of a common beacon (i.e., the same synchronization beacon is used by all nodes in the network) with features that enable a fast and accurate timing and carrier frequency synchronization. Each node broadcasts the same signature that superimpose (collide) with the others and compound signal of multiple collisions represents a reference signature that embeds the TO and CFO reference for the entire network. Once the network reaches a convergence, the frames are aligned giving the start of time-slot. Contrary to conventional synchronization methods, the feature of the proposed distributed synchronization algorithm is that the collision of signatures does not impair the synchronization, rather it is used by the receiving node as ensemble reference to enable its synchronization. The design of a unique synchronization signature (beacon) largely simplifies the setting and it is based on chirplike sequences with good correlation properties (e.g. Zadoff-Chu sequences used on LTE), that allows the join estimation of TO and CFO errors. The accuracy of the correlator-based estimator is analyzed, and the impact of stochastic perturbations product of the oscillator's instability and estimator's error is studied to evaluate the convergence condition of the proposed distributed synchronization algorithm. The proposed distributed synchronization algorithm is implemented on a hardware demonstrator, based on software-defined radios programmed in GNU radio, showing the algorithms ability to decouple the TO and CFO estimate and it is analyzed the convergence time of TO and CFO synchronization. Then, the optimization of the distributed synchronization is carried out based on i) an optimal duplexing strategy, and ii) an optimal synchronization protocol. In context of dense inter-connected networks with oscillators affected by drifts, two synchronization approaches are compared based on: collision-avoidance and collision of signals, to investigate the impact of the network scalability by comparing the convergence time and synchronization dispersion error. Finally, the impact of synchronization is evaluated for a resource management scheme (spectrum allocation) in Device-to-Device (D2D) communications. The selection of the resources to be added or released in the allocation is performed by minimizing the boundary extension of the time-frequency (TF) spectrum region, this criteria avoids fragmented region allocations with large boundary areas that could increases the cross-interference due to TF jitter.

HAMID ASLANI

PRINTED PATCH ANTENNAS FOR MODERN WIRELESS COMMUNICATION SYSTEMS

The increasing demand for higher bandwidth and higher speed wireless communication motivates the exploration of modern wireless communication. Ultra-wide band (UWB) technology is one of the most promising solutions for future communication systems due to the high data rate and excellent immunity to multi-path interference. Also, The IEEE 802.11ad and IEEE 802.11ay operating on 60 GHz mmWave are the two most expected wireless local area network (WLAN) technologies for ultra-high-speed communications. The 802.11ad standard (WiGig) provides throughput speeds of multi-Gb/s covering tens of meters by offering a wide beamforming channel in 60GHz ISM band channel. In the proposed work, we will focus on both the above-mentioned technologies. As a first approach, a novel wide band microstrip patch antenna (MPA) configurations that can be used for UWB applications with enhanced performance is designed and discussed. The impedance bandwidth of the proposed antennas has been enhanced by using techniques such as patch slotting, as an improved technique. An impedance bandwidth between 5 -11 GHz has been achieved. The antenna has been examined in terms of the return loss, peak gain and radiation characteristics. We also present an efficient microstrip patch antenna (MPA) with superstrate Technique. The antenna configuration can be used for UWB applications. Based on the interference problem in modern wireless communication systems, the proposed work has been extended in order to design an efficient UWB antennas with band notch characteristics at centre frequency of 5.2 GHz. The Simulations were performed using different EM software such as Ansys HFSS which uses the Finite Difference Time Domain (FDTD), and Finite Element Method (FEM). We follow the strategy as the first step by designing patch antenna suitable for Ultra Wide Band (UWB) bandwidth range and the next step, an anisotropic filter ring resonator imposed and etched on the UWB antenna patch to produce band rejection properties. Several investigations and analysis of printed antenna arrays with high-gain characteristic are performed for mm-waves wireless systems. The structure presented based on microstrip arrays antennas, the proposed antenna is designed for 60GHz high gain backhaul transceiver systems. Several array antenna structures have been also proposed to cover 360 degree. In order to analyse and validate the antenna performance designs, several numerical analyses have been done by using electromagnetic software simulators.

HAMID REZA BARZEGAR

WIRELESS PARTIAL-DUPLEX COMMUNICATIONS

Wireless communication exploits electromagnetic signals that are broadcast by a transmitter device acting as a sender or as an intermediate relay. The communication between two devices occurs when the receiver device captures these signals creating a wireless communication bridge with the sender device. Wireless communications express technologies in many different system families, such as Wireless Personal Area Networks (WPAN), Wireless Local Area Networks (WLAN), Metropolitan Area Networks (MAN), mobile cellular systems, satellite communication, etc. So far wireless systems cannot transmit and receive on the same frequency band at the same time (Full-Duplex wireless communications) due to the strong Self-Interference (SI). The duplexing of transmission and reception must be done via either frequency division, i.e. Frequency Division Duplex (FDD), or time division, i.e. Time Division Duplex (TDD). In fact, Full-Duplex (FD) wireless communication means facing a high amount of SI that should be canceled. The SI (or echo) canceler is indeed one of the most critical components and key-factors of wireless FD communication systems: in this context, it offers the potential to complement and sustain two-way communications and, in general, Self-Interference Cancellation Schemes (SICS) are expected to have a tremendous impact on the implementation of two-way communications. Wireless FD potentially has the ability to make double the spectral efficiency with respect to traditional Half-Duplex (HD) operation. This phenomenon happens by exploiting the same wireless resource to transmit and receive at the same time, with the cost of a large power difference between its own transmitted signal, i.e. SI from its own transmitted signal, and the low-power received signal from the other transceiver. The SI can be reduced gradually, thanks to the different Self-Interference Cancellation (SIC) techniques at radio frequency and baseband levels. However, this potential still suffers from the residual SI and the resulting short distance range of communication. This research activity has been focused on the analysis and design of new schemes for wireless communication based on the reuse of the same channel for both communication directions (FD) in order to increase link performance in terms of transmission range or capacity. The proposed solutions are intended to relax the challenging constraints of the self-interference cancelers; in addition they provide, especially in multicarrier systems, an additional flexibility degree in the resource allocation strategies in presence of multiple channels and interfering devices. Therefore, in order to overcome these issues, this research introduced a new scheme, named as Partial-Duplex (PD) approach: this solution consists of a communication link with the capability of supporting the connection in both directions at the same time in a portion of the bandwidth and with a frequency division for the two communication directions (for example uplink and downlink) in the rest of the band. The rationale behind this approach is to limit the level of SI finding a compromise between HD and FD transmission and relaxing consequently the constraints on the echo canceller design in order to increase the distance range between transmitter and receiver. Equivalently, PD transmission aims to increase the overall bidirectional system rate w.r.t. an equivalent HD system, relaxing at the same time the high SIC requirements that practical FD systems have to provide. This hybrid transmission method between classical HD and FD is considered for point-to-point single carrier and multi-carrier (Orthogonal Frequency Division Multiplexing (OFDM)) links experiencing Additive White Gaussian Noise (AWGN) channels and also random fading, flat or frequency selective. In the first part of the study, the regions of SIC performance where PD systems outperform HD ones in terms of achievable spectral efficiency are analyzed by deriving the analytical distributions of the spectral efficiency gain regions in presence of random frequency-selective fading, for different strategies in the selection of FD sub-carriers in PD schemes; therefore, it is investigated the potential of this hybrid method, highlighting the role of the different parameters involved and the peculiarities of this flexible system design approach. Nevertheless, we have wondered also how to see this approach in the upcoming next generation of wireless networks Fifth Generation of mobile communications (5G), where some common trends and promising technologies have been already identified. In addition to capacity gain and improvement of performance, which are expected from network densification, device-to-device communication, and small cells, we can mention an increased spectrum sharing and integration as well as spectrum enhanced carrier aggregation and other advanced wireless communication technologies like massive Multiple-Input Multiple-Output (MIMO) and resource management based on machine learning. As an example, massive Machine Type Communication (mMTC) is one of the use cases of 5G and it is expected to increase significantly both the number of connected devices to the network and the traffic. In these cases, a proper use of wireless PD communication can improve the overall throughput of the mobile network. In this context, the second part of this study has been focused on the channel encoding for PD. One of the best candidates for efficient channel coding and error correction is the class of Low Density Parity Check (LDPC) codes. We have identified that, in Partial Duplex Communication (PDC), the receiver faces a mixture of high and low bit (symbol) Signal-to-Noise Ratio (SNR) in the same codeword, a condition associated with a scheme in which part of the bits/sub-carriers is subject to FD interference. LDPC codes have shown, also w.r.t. turbo and polar codes, the best performance for PD and we have studied how to optimize further their performance in this specific context. The Digital Video Broadcasting- Second Edition (DVB-S2) was one of the standards that includes the use of LDPC. These codes were implemented for long codewords and we proposed to create LDPC codewords with arbitrary length, derived from the original DVB-S2 but maintaining the same parity check matrix structure. Then, the application of this class of LDPC codes, derived from DVB-S2 to PD communication can be improved by a specific allocation techniques of the high and low SNR bits, suited to PD communications. The main innovative results of this activity are related to the fact that, in the Partial Duplex Scheme (PDS), part of the band is transmitted in FD and the rest in HD and, consequently, some transmitted bits (in single carrier) and sub-carriers (in OFDM) will be characterized by high SNR and the others by low SNR according to a pattern which is known, a-priori, by the system. Therefore, combining properly the patterns of these high and low SNR bits affects the coding performance of the system in a way that depends also on the parity check matrix structure of the code used in the transmission. In addition, in order to validate further the results and according to upcoming 5G standardization, this channel encoding procedure has been applied also to the 5G encoding schemes recently considered by 3rd Generation Partnership Project (3GPP). Mainly two types of channel coding are adopted in 5G, Polar and LDPC codes. Therefore, in this part of the study, we have investigated and compared our approach in the context of 3GPP standard for 5G and Long Term Evolution (LTE), showing the performance of encoded PD. Results have turned out to be really promising for a specific class of LDPC codes, when, as in PD, we can exploit the a-priori knowledge of high and low SNR bits in the transmitted codeword with a specific bit allocation.

ALBERTO BERNARDINI

ADVANCES IN WAVE DIGITAL MODELING OF LINEAR AND NONLINEAR SYSTEMS

This thesis presents my contribution to the recent evolution of modeling and implementation techniques of linear and nonlinear physical systems in the Wave Digital (WD) domain. The overarching goal of WD methods is to build digital implementations of analog systems, which are able to emulate the behavior of their analog counterpart in an efficient and accurate fashion. Though such methods usually focus on the WD modeling of analog audio circuits; the methodologies that we address in this thesis are general enough as to be applicable to whatever physical system that can be described by an equivalent electric circuit, which includes any system that can be thought of as a port-wise interconnection of lumped physical elements. Such systems, in fact, are fully described by a "dual" pair of Kirchhoff variables per port (e.g. a current-voltage pair, a force-velocity pair, etc.), and the circuit connecting such elements implements the topological relations that exist between them (e.g. Kirchhoff laws, laws of dynamics/fluid-dynamics, etc.). The possibility of describing systems through electrical equivalents has relevant implications not only in the field of numerical simulation of physical phenomena, but also in the field of digital signal processing, as it allows us to model different kinds of processing structures in a unified fashion and to easily manage the energetic properties of their input-output signals. However, digitally implementing nonlinear circuits in the Kirchhoff domain is not straightforward, because dual variables (currents and voltages) are related by implicit equations which make computability very hard, as instantaneous dependencies between input and output signals cannot be eliminated. Spice-like software, based on the Modified Nodal Analysis (MNA) framework, is not always suitable for realizing efficient and interactive digital applications, mainly because it requires the use of iterative methods for solving multi-dimensional systems of equations. WD Filters (WDFs) are a very attractive alternative. During the seventies, Alfred Fettweis introduced WDFs as a special category of digital filters based on a lumped discretization of reference analog circuits. A WDF is created by port-wise consideration of a reference circuit, i.e., decomposition into one-port and multi-port circuit elements, a linear transformation of Kirchhoff variables to wave signals (incident and reflected waves) with the introduction of a free parameter per port, called reference port resistance, and a discretization of reactive elements via the bilinear transform. Linear circuit elements, such as resistors, real sources, capacitors and inductors, can be described through wave mappings without instantaneous reflections, as they can be all "adapted" exploiting the mentioned free parameter; in such a way that local delay-free loops (implicit relations between port variables) are eliminated. Series and parallel topological connections between the elements are implemented using scattering topological junctions called "adaptors", which impose special adaptation conditions to eliminate global delay-free loops and ensure computability. It follows that WDFs, as opposed to approaches based on the MNA, allow us to model separately the topology and the elements of the reference circuit. Moreover, WDFs are characterized by stability, accuracy, pseudo-passivity, possibility of parameterization, modularity and low computational complexity, making many real-time interactive applications easy to be realized. Most classical WD structures can be implemented in an explicit fashion, using binary connection trees, whose leaves are linear one-ports, nodes are 3-port adaptors and the root may be a nonlinear element. However, WDFs are also characterized by important limitations. The first main weakness of state-of-the-art WDFs is that WD structures, characterized by explicit input-output relations, can contain only one non-linear element, as nonlinear elements cannot be adapted. In fact, the presence of multiple nonlinear elements might affect computability, which characterizes classical linear WDFs, as delay-free loops arise. As a second main limitation of traditional WDFs, up to three years ago, there were no systematic methods for modeling connection networks which embed non-reciprocal linear multi-ports, such as nullors or controlled sources. Finally, very few studies were presented on the use of discretization methods alternative to the bilinear transform and potentially varying step-size in WD structures. This thesis presents various techniques to overcome the aforementioned limitations. After a review of the state of the art on WD modeling of lumped systems up to 2015, this thesis begins with redefining wave signals, first by offering a unified definition that accommodates the existing ones (those based on one free parameter per port), then by introducing a new class of waves with two free parameters per port, which leads to WD structures that exhibit a doubled number of degrees of freedom, called "biparametric" WDFs. The second part discusses how to implement nonlinear one-port and multi-port elements in the WD domain; finding a compromise between accuracy, efficiency and robustness. In particular, it shows how scalar nonlinearities can be represented in canonical piecewise-linear form in the WD domain; it presents a technique based on the Lambert function to make a class of exponential nonlinearities, e.g., diodes or Bipolar Junction Transistors, explicit in the WD domain; and it provides an in depth discussion on the modeling of nonlinear 3-terminal devices with various examples of applications on circuits containing transistors. The third part focuses on the description of connection networks as WD adaptors. In particular, this part discusses how to model arbitrary reciprocal and non-reciprocal junctions in WD structures based on various definitions of wave signals and how to compute waves reflected from the scattering junctions in an efficient fashion. The fourth part focuses on the cases in which we cannot do without iterative solvers. In particular, it introduces a novel relaxation method based on WD principles, developed for implementing circuits with multiple nonlinearities. The method is called Scattering Iterative Method (SIM). A proven theorem guarantees that SIM converges when applied to circuits with an arbitrary number of nonlinearities characterized by monotonic voltage-current characteristics. As an example of application of this method, I show that power curves of large Photovoltaic (PV) arrays (constituted of thousands of nonlinear PV units) with arbitrary topologies can be obtained far more efficiently using the proposed SIM than using MNA-based approaches, (SIM is at least 30 times faster). I also show that SIM is highly parallelizable and suitable for the implementation of time-varying circuits. Moreover, a strategy for deriving WD models of dynamic elements based on arbitrary linear multi-step discretization methods with potentially adaptive time-step size is here introduced and discussed. This strategy allows us to describe capacitors and inductors using time-varying Thévenin or Norton equivalents, and proves particularly useful in conjunction with SIM for implementing dynamic circuits with multiple nonlinearities. These results also motivated the simulation of some audio circuits (e.g., a dynamic ring modulator), which proved SIM to be a promising implementation method that is also employable for virtual analog applications, as it is characterized by high efficiency, parallelizability and inherent capability of handling time-varying elements. The combination of the new techniques for modeling arbitrary reciprocal and non-reciprocal connection networks with SIM is promising for the future development of a general purpose simulation program which might be more efficient than Spice-like software, as far as the time-domain analysis of arbitrary nonlinear circuits is concerned. As a further example of application of WD principles, the last part of the thesis discusses a novel approach for implementing a class of differential beamformers using WDFs.

LUCA BONDI

DATA-DRIVEN AND HANDCRAFTED FEATURES FOR FORENSICS ANALYSIS AND SOURCE ATTRIBUTION

The communication power associated with visual content makes digital images a powerful and effective tool to deliver messages, spread ideas, and prove facts. Smartphones, digital cameras, and camcorders are becoming more affordable every day, and thus constitute a rapid and convenient way of capturing and sharing photos quickly and inexpensively. The increasing diversity of brands, models, and devices makes the creation of new visual contents easier every day, while the ever-growing access to social network and picture sharing platforms poses a set of challenges, from the diffusion of illegal content to copyright infringement. The wide availability and ease of use of image manipulation software makes the process of altering an image simple and fast. This could severely reduce the trustworthiness of digital images for users, legal courts, and police investigators. The fake news phenomenon is a well-known and widespread example of the malicious use of digital pictures and manipulation software. Modified images done with precision are used to create false proofs for made-up stories, exploiting the often unquestionable trust with which readers take in visual content. In this thesis we face several challenges related to the analysis of digital images. A first step in assessing image authenticity, and tracing an image back to its origins, consists in determining which device shot a specific picture. State-of-the-art techniques based on Photo Response Non-Uniformity (PRNU) prove to be very effective in determining the specific sensor that shot a picture. However, given the highly increasing number of devices, a full-range search over all the existing devices is impractical and time consuming. One of the ways to reduce the search space is to first find the camera model that took a picture, then test the image under analysis against the devices from the same camera model. In this thesis we present the first data-driven method designed to learn camera model features directly from a collection of images, showing how modern deep-learning techniques based on Convolutional Neural Networks (CNN) can be adapted to multimedia forensics tasks. When it comes to a large-scale search of picture-device matches based on PRNU, at least two challenges arise: time and storage space constraints. To address such challenges, the forensics community explored a series of techniques to compress PRNU fingerprints and residuals. In order to reduce storage space requirements, while lowering the computational complexity, we introduce two techniques to address PRNU compression, by exploiting classical signal processing analysis and data reduction techniques. While determining the origin of a digital image is important to solve copyright infringement cases, digital images can be locally altered by adding, removing, or modifying objects with the goal of changing the semantics of the image. We present how to exploit the features learned with a CNN trained for camera model identification with the goal of detecting and localizing tampered regions within an image. Under both device identification and camera model identification perspectives, we study a set of possible antiforensics attacks tailored at anonymizing an image to prevent the correct identification of its origin. This allows us to understand the limitations and weaknesses of the proposed camera model and device identification techniques. Finally, we leverage the knowledge and skills acquired in mixing together handcrafted signal processing and data-driven methods in two different forensics applications: Laser Printer Attribution and Single versus Double JPEG Detection. In both scenarios the key to tackle the forensics task at hand is fusing together a proper signal pre-processing technique with a carefully designed data-driven system.

LORENZO COMBI

RADIO-RELAYING OVER OPTICAL FIBER FOR CLOUD RADIO ACCESS NETWORKS

The focus of this Thesis is on analog signal transport and processing for future mobile radio access networks, and its content can be broadly divided into four macro-areas. Analog optical PWM is suitable for analog fronthauling as it combines the best features of digital and analog transmission. Indeed, the 2-level waveform involves a relaxation of the linearity requirements, still allowing for an analog signal transport that avoids the bandwidth expansion of digital fronthauling. PWM is experimentally validated for different optical architectures. Besides the conventional optical network layout (that can be summarized in the cascade of a laser, fiber and photodiode), here an innovative architecture is proposed, based on reflection. In this reflective PWM the transmitter (at the remote antennas) is equipped with a modulated reflective semiconductor optical amplifier (RSOA) that receives, modulates and reflects back a continuous waveform (CW) from an optical source located at the receiver (at the centralized baseband unit, BBU). The modulated reflected signal which is the uplink fronthauling is polarization-separated from the CW seeding signal.PWM analog fronthauling effectively avoids the bandwidth expansion related to digitization of fronthaul signals in state-of-the-art systems, but massive antenna arrays and larger signal bandwidth may call for additional capacity of the fronthaul link. Mode division multiplexing is considered in this Thesis to provide an additional multiplexing dimension and, to ensure the cost-effectiveness of the proposed solution, modal multiplexing is all-optical and passive. To cope with the arising intermodal interference, the mapping between radio resources and optical fronthauling resources is studied and it is shown that appropriate resource assignment can almost overcome the limitations introduced by intermodal interference. In particular, the resources available on the two channels (wireless link antennas-users and wired fronthaul link) are defined in a space-frequency domain, the space being defined by the antenna array for the radio signal and by the propagation modes for the optical fronthaul. Massive antenna arrays and larger signal bandwidth are key features of millimeter wave radio communication, and hybrid beamforming has been proved in literature to provide a valid mean to overcome the channel limitations (mainly the increased path loss) while keeping a reasonable hardware complexity and energy consumption. In this Thesis, the integration of hybrid beamforming and analog fronthauling is tackled. In particular, realistic analog signal processing techniques in the optical domain are considered in order to provide wideband analog processing via tunable delay lines. Tunability is shown to be essential in dealing with a fast-changing radio environment and with time division multiplexing (TDM) of users, but the speed of typical optical tunable elements in achieving a time-varying delay response is not enough. This causes a transient in system performance with degradation of the first transmitted symbol of the TDM frame. To counteract this effect, this Thesis proposes a technology-aware scheduling of time division multiplexed users, together with a digital precompensation of the transient in delay response. Tuning the analog processing at the remote antennas requires a signaling channel parallel to the fronthauling. The proposed solution is based on joint A-RoF fronthaul transmission of the radio signals paired with a binary polarization shift keying (PolSK). This provides a low-cost, low-complexity parallel channel between the BBU and the RRH that can be used, e.g., to control the analog beamformer at the RRH and to feed back CSI from the radio channel.

KAPAL DEV

SYNTHESIS OF ATTENUATION TIME SERIES ACROSS FSO LINK FOR THE DESIGN OF A BACKHAULING SOLUTION FOR 5G BEYOND NETWORKS

FREE Space Optical (FSO) communication is one of the most widely researched technology due to its very interesting characteristics such as high data rate, free license, no electromagnetic interference, light weight, small volume, secure due to narrow laser beam, portability, and low power consumption. However, some limitations due to weather dependency occur which include scattering and absorption, caused by gases and various hydrometers, fog being the most challenging issue which majorly impact on the performance of the FSO link. From the application perspective, we know that in order to provide 5G technology opportunity to users to utilize extremely large bandwidth need the use of a dense network with mini base stations at short range (few 100 meters) which could be connected in two ways i.e. wired or wireless. Solving backhaul connectivity is critical before any 5G small cell deployments can scale up. There is no way to even consider adding wired backhaul drops to thousands of sites in an urban environment which will be expensive require more time and physical efforts. In this respect, FSO links can handle the scale, they are easy to deploy, very large capacity and represent an economic solution over a distance of 100 meters. It is essential that operators planning for high density small cell deployment seriously consider FSO as an option before 5G rollouts begin. In order to use FSO for backhauling, our proposed work is very useful in providing the methodology to generate synthetic attenuation values of signal fading and its characteristics over different low visibility conditions even before the deployment. This research work is directed to propose a procedure for a synthetic attenuation time series synthesizer for low visibility events along terrestrial free space optical links useful for 5G backhauling networks design. To this aim, preprocessing for correct use of data had been done which includes identifying the low visibility events with optimal approach, bias removal on average and event basis, and identifying the best time integration value for all three databases. Later, it is demonstrated that visibility is suitable to derive attenuation time series on a slow time sampling but for fast time sampling, a general statistical technique is proposed. After doing an extensive literature, we selected a procedure for the development of a time series synthesizer among the ones proposed for mm wave and modified the procedure with the introduction of visibility time series at the input instead of using large set of attenuation time series. Also, fast variations obtained through the statistical technique will be superposed to slow fading which made the predicted synthetic attenuation time series more accurate. Finally, after executing detailed step by step procedure, synthetic attenuation time series is constituted by a synthetic component (Fast Fading) each superposed to a component obtained by manipulating measured data (Slow Fading). Large database of measured data collected at Politecnico di Milano and in other two experimental sites of the Europe are considered to validate the proposed work. We tested our procedure on all the events from the different sites and compared synthetic time series attenuation with the measured one on a statistical basis. This testing is based on number of indicators: mean value, standard deviation and RMSE. Overall, we found an excellent result as 88.88% events have 1 dB/km and 0.5 dB/km in terms of difference in mean and standard deviation respectively and RMSE difference is within the 1.5 when considering that it represents the comparison of actual measurements on event basis with a statistical model based on the data from two sites. The performance of the proposed work is excellent in reproducing moderate visibility conditions in the presence of fog (with 88.88% accuracy), rain (100% accuracy) is found. In the case of Milan and Prague, the full model of generating the synthetic attenuation series is applicable and it gives very good results. Unfortunately, considering Milesovka, it works only on generating the slow fading attenuation time series because the measurements were of low sampling rate which doesn't allow us to identify fast variations. The proposed procedure can be in principle applied to any location provided visibility time series (very simple and inexpensive to collect) are available where the optical link is required to set up.

ATUL KUMAR

SYNCHRONIZATION AND PERFORMANCE EVALUATION OF FUTURE WIRELESS CELLULAR SYSTEM BASED ON THE USE OF NEW MULTI-CARRIER TRANSMISSION TECHNIQUES

In recent years, the evolution of cellular technology and connectivity has led to a major revolution in the wireless communication industry, with differnt major application scenarios i.e. enhanced mobile broadband (eMBB), massive machine type communication (mMTC), ultra-reliable and low latency communications (URLLC), and vehicle-to-everything (eV2X). In order to, target mMTC applications, the important scenario considered here is Internet of Things (IoT). Physical layer (PHY) requirements for mMTC application are: low latency; large number of connected devices; high reliability; asynchronous transmission; and dynamic fragmented spectrum access. Current (4th generation) cellular transmission schemes, which are based on Long- Term Evolution (LTE) and LTE-Advanced (LTE-A), employ orthogonal frequency division multiplexing (OFDM) waveform in PHY design. As is well known, synchronization represents one of the most challenging issues and plays a major role in PHY design and its impact can also be higher in future wireless cellular system, where the use of non-proportional sub-carrier spacing will be required to accommodate the necessary bandwidth. Motivated by this, I studied and realized this Thesis work that concerns with "Synchronization and Performance Evaluation of Future Wireless Cellular System Based on the Use of New Multi-Carrier Transmission Techniques". With this aim, I started studying about the two categories of multi-carrier waveforms: legacy OFDM waveforms and new waveforms for future wireless cellular systems. For the legacy OFDM waveforms, I studied mainly about the cyclic prefix (CP)- based OFDM system and also an OFDM system with an advanced transformation tool, known as discrete fractional Fourier transform (DFrFT). DFrFT-based OFDM system is analogous to the conventional OFDM one with the difference that DFT and inverse DFT (IDFT) are replaced by DFrFT and inverse DFrFT (IDFrFT), respectively. DFrFT is a generalization of the ordinary DFT with an DFrFT angle parameter (\alpha). Moreover, I started analyzing the effect of synchronization issues, in particular, the problem of evaluating performance in the presence of symbol timing offset (STO) and carrier fre-quency offset (CFO). Here, we consider the chirp-based DFrFT method for both CFO and STO estimation. From the demonstrated results of mean square error (MSE) performance of CFO and STO estimations, it is realized that CFO and STO estimators with chirp-based DFrFT method performs better than other methods available in literature in case of the transmission over multipath fading channel. However, in the presence of oscillator drifts and time-varying Doppler shifts, residual CFO and STO are still present in the received signal after the application of synchronization algorithms. In order to improve the robustness to residual CFO and STO, an OFDM system based on DFrFT is considered. Therefore, it would be better to analyse the analytical symbol error probability (SEP) performance of the DFrFT-based OFDM by considering residual CFO and STO together. This is demonstrated by quantifying its effect in an analytical expression of the term responsible for introducing inter carrier interference (ICI) and inter symbol interference (ISI), which is one of the novel contributions of this Thesis. The results of the SEP performance of the DFrFT-based OFDM system in the presence of residual synchronization errors, i.e., CFO and STO, demonstrate that the performance of DFrFT-based OFDM system depends on the DFrFT angle parameter ( ) and, by properly selecting its optimal value opt, DFrFT-based OFDM system performs better than the one based on the DFT. Therefore, the calculation of optimal DFrFT angle parameter is important for better performance of DFrFT-based OFDM system. In this Thesis, we also derived an analytical expression to calculate the optimal value of opt based on CFO and STO values. This is another novel contribution of this Thesis. Finally, we also considered more practical aspects related to the realization of a Software Defined Radio (SDR) system which is used to implement hardware co-simulation of multi-carrier transmission techniques. We have considered field-programmable gate array (FPGA)-in-the-Loop (FIL) co-simulation of receiver with equalization of DFrFT-based OFDM system transmission over a frequency selective Rayleigh fading channel in presence of CFO. My simulation results clearly demonstrate that the FPGA implementation of a DFrFT-based OFDM system in presence of CFO has the same performance as that obtained from Monte Carlo simulation. Also, the performance is validated with the fixed-point model of DFrFT-based OFDM. The approach described in the thesis constitutes an efficient way to convert the floating point model into a fixed-point one to be run in an FPGA and then verifies its correctness through FIL co-simulation. However, as it is well known, CP-based OFDM uses fixed set of waveform parameters, including sub-carrier spacing and length of CP that, are uniformly applied across the entire system bandwidth. Due to the lack of flexibility in supporting mixed services with different waveform parameters within one carrier, which is a key requirement in the PHY design of future cellular network. Additionally, high out-of-band (OOB) emission in frequency-domain is introduced by the time-domain rectangular pulse shaping filter. Also, OFDM signal with one CP per symbol may have a prohibitive low spectral efficiency and also increase the OFDM symbol duration. For this reason, new waveforms that are able to support variable and customizable pulse shaping filters in order to achieve a better trade-off between time-domain and frequency-domain localization is one of the research priorities. With this aim, generalized frequency division multiplexing (GFDM) is one of the proposed multi-carrier waveforms for future wireless cellular systems, which is based on the use of circular filtering at sub-carrier level. Compared to CP-based OFDM, the important advantages of GFDM consist in a reduction of OOB emission, achieved by means of filtering at sub-carrier level, in an increase of spectral efficiency, obtained through the introduction of tail biting, which makes the length of the CP independent from that of pulse shaping filter. Moreover, the flexible frame structure of GFDM is achieved, by changing the number of time slots and sub-carriers in a frame, covering both conventional OFDM and DFT-spread OFDM. (DFT-s-OFDM), which results in complete backward compatibility with LTE. Motivated with the GFDM, we focus on the integration of GFDM in the timefrequency grid of LTE system and then analyse the impact of “Better than Nyquist” pulse shaping filters on OOB emission and symbol error rate (SER). Moreover, we also consider the concept of wavelet for better time-frequency localization of pulse shaping filters by using the Meyer auxiliary function. After the impact of pulse shaping filter, for an efficient implementation of receiver in time-domain, a relationship between GFDM signal and discrete Gabor transform (DGT) is investigated for reducing the complexity. After that, we implement DGT-based GFDM when the synthesis function, i.e. pulse shaping filter, and the analysis function, i.e. receiving filter, satisfy the Wexler-Raz identity. Choosing functions that satisfy the Wexler-Raz condition allows for optimal symbol-by-symbol detection after DGT-based receiver in case of ideal channel. However, when transmission takes place over a frequency selective channel, symbol-by-symbol detection of sub-symbols is no longer optimal due to intersub- symbol interference (ISSI) generated by sub-symbols transmitted over the same sub-carrier, to improve the performance maximum likelihood detection (MLD) is implemented as a optimal detector for all the sub-symbols on the same sub-carrier, which is another novel contribution of this thesis. Finally, we derive exact SEP expressions for GFDM waveform in the presence of CFO in AWGN channel and frequency selective Rayleigh fading channel. The analytical expressions of SEP are derived when matched filtering is implemented at the receiver for BPSK, QPSK, and 16-QAM in case of AWGN channel and for BPSK only in case of frequency selective Rayleigh fading channel which is another contribution. Monte Carlo simulations are presented to demonstrate the exactness of the derived SEP expressions.

ANDREA MATERA

INTERFERENCE MITIGATION TECHNIQUES IN HYBRID WIRED-WIRELESS COMMUNICATIONS SYSTEMS FOR CLOUD RADIO ACCESS NETWORKS WITH ANALOG FRONTHAULING

Radio Access Network (C-RAN) is an attractive solution to handle the huge number of user devices and antennas that are expected to populate next generation (5G and beyond) communication networks. C-RAN is already adopted in current (4G) mobile networks, in which BaseBand Units (BBUs) and Remote Antenna Units (RAUs) exchange In-phase and Quadrature (I/Q) digitized signals over the so-called FrontHaul (FH) link. However, the expected increase in radio frequency bandwidth calls into question the effectiveness of digital I/Q stream. Among the numerous FH functional splits that have been proposed over the last years, mostly based on digital FH, C-RAN based on analog fronthauling is considered a low-cost and bandwidth efficient solution. In analog C-RAN, the RAUs directly relay the radio signals to/from the BBU, thus bypassing any bandwidth expansion due to digitization, reducing latency, and providing synchronization among multiple RAUs. In particular, C-RAN with analog FH based on copper cables, namely Analog Radio-over-Copper (A-RoC), is a suitable solution to enhance indoor wireless coverage into buildings and enterprises. In this thesis, A-RoC is considered for multiple-antennas RAUs and multiple twisted-pairs copper cables, e.g., Local Area Networks (LAN) cables, thus leading to the Analog MIMO RoC paradigm. In the proposed A-MIMO-RoC architecture, radio signals from multiple-antennas are opportunely mapped over the copper cables multiplexed both in space dimension, given by the multiple twisted-pairs, and in frequency dimension, given by the bandwidth of each twisted pair. Such all-analog radio-cable resource mapping enables the full exploitation of the transport bandwidth capability of copper cables. The A-MIMO-RoC system design requires the optimization over the cascade of two different MIMO channels, i.e., the radio and cable channels, each of which with different characteristics and constraints. As a first step, signal processing techniques have been investigated for these two channels, separately. In particular, on the one hand, the focus was on non-linear precoding techniques for next generation xDSL systems, namely G.fast; on the other hand, interference mitigation techniques have been designed for wireless systems with focus on multi-operator scenarios, and for optical wireless systems, i.e., Visible Light Communications (VLC). Merging the knowledge acquired from this preliminary step, the core of the thesis discusses the design and optimization of the Analog-MIMO-RoC architecture, which is complicated by the mutual interaction between the wireless and wired communication channels. This thesis provides extensive numerical analysis showing the benefits of the proposed A-MIMO-RoC architecture for 5G indoor networks. As a conclusive step, the real world implementation of the proposed architecture has been demonstrated by developing a first A-MIMO-RoC prototype, which was able to prove experimentally and for the first time the potentials of the proposed analog C-RAN architecture for FH indoor applications.

FAISAL AHMED MEMON

SILICON OXYCARBIDE TECHNOLOGY FOR INTEGRATED PHOTONICS

Materials and their characteristics are the backbone of any technological field. The advancement of a technology demands novel materials with appealing properties. Integrated photonics is considered as emerging technology of 21st century. The fundamental block of integrated photonics technology is a waveguide that propagates photons. The waveguides can be used to design devices with enhanced functionality for telecommunication, optical signal processing, control and sensing. Unlike microelectronics, several material platforms in addition to silicon are being used in integrated photonics. The major passive waveguide platforms in integrated photonics are based on dielectrics. The well-established dielectric material platforms include doped-silica, silicon oxynitride and silicon nitride. These platforms have their own advantages and disadvantages and none of them can fulfill the needs of photonics applications such as CMOS compatibility, high index reconfigurability and low losses. The primary theme of this doctoral dissertation is to develop a dielectric platform that is CMOS compatible, widely index tunable, low loss, efficiently reconfigurable and can be integrated with typical dielectrics. To this aim, silicon oxycarbide (SiOC) a novel class of glass compounds has been exploited. Reactive RF magnetron sputtering was employed to deposit a system of silicon oxycarbide thin films over a wide composition range and large refractive index window from silica to amorphous silicon carbide. The films properties were investigated in greater detail to assess its potential for micro-photonic device fabrication. Further to advance the development of the platform, medium-to-high contrast photonic waveguides in silicon oxycarbide system were realized employing microfabrication process. The classical characteristics of the waveguides have been measured in the commercial telecom window to reveal its transparency and potential for photonic applications. Silicon oxycarbides have been investigated for possible application in reconfigurable photonic integrated systems. The record high thermooptic effect in silicon oxycarbides has been discovered that is one order of magnitude larger than typical dielectric platforms. As a further exploitation, integration of silicon oxycarbide with conventional dielectrics resulted in power efficient phase actuators that is a great achievement. Within the scope of this work, we have been successful in developing a versatile platform with appealing characteristics of refractive index tunability providing low losses in the telecom wavelength range and efficient reconfigurability that was not possible with other typical dielectric platforms.

MAROUAN MIZMIZI

NETWORK SENSING AND CONTEXT AWARENESS

Position information is one of the core element for context awareness and pervasive systems. In recent years, with the increase of the demand for location-based services (LBS), there is more and more interest in localization systems. In the outdoor environment, location information is obtained from the Global Positioning System (GPS), while for indoor environment, intense research is ongoing both form academic in industry words to achieve a global solution. In this thesis, we focus on positioning and navigation techniques, for indoor, outdoor, and molecular communication environment. The main achievements in the indoor scenario are related to the fingerprinting-based localization, which is one of the most popular techniques for Indoor Positioning System (IPS). In particular, we have developed an analytical model to predict the performance at the design time. Therefore it is possible to tune the system parameters without the time consuming and expensive practical implementation and testing methods. A more important result is the introduction of the Binary Fingerprinting (BFP) technique, where the quantization bits of the Received Signal Strength Indicator (RSSI) is reduced till the binary level. The obtained Radio-Map and the online fingerprints are binary, therefore it is possible to apply concepts from binary coding theory to analyze and estimate the user position. Numerical and experimental evidence has shown the benefits of the proposed technique, in terms of computational complexity and scalability. Moreover, we extended the concept of the BFP beyond the mere positioning purpose, introducing the Binary Space Labeling concept, which can serve as a network access and resource sharing strategy with less signaling and control w.r.t state of the art techniques. For the outdoor environment, the Vulnerable Road User (VRU) protection is presented. It is a location-based service in the network, where Road users are localized and tracked by the network with the goal of preventing road accidents. Within 5GCAR project, VRU protection will be demonstrated live on June 27, 2019, in Paris. Another localization application studied in this work is related to Molecular Communication (MC). The study of communication systems based on the exchange of molecules is currently a fertile research direction in communications, encouraged by novel practical tools to engineer such systems from living substrates. Position estimation in these systems, especially when applied to the pervasive monitoring and control of biochemical environments, e.g., intra-body, is a critical functionality. In this context, two distance estimation methods are introduced. The first estimator is based on the measure of the peak attenuation, while the second is based on the measure of the widening, respectively, of a diffusing pulse of emitted molecules. Their performance is evaluated through simulations in terms of normalized Root Mean-Squared Error (RMSE) by varying several parameters, such as the transmitter-receiver distance, the number of released molecules, and the diffusion coefficient. The use of a variable-length smoothing filter is also proposed to further improve the performance of the two considered methods. The simulation results have shown that a positioning algorithm is feasible in the bio-environment, with unlimited practical applications.

DOUGLAS OLIVEIRA MORAIS DE AGUIAR

RECONFIGURABLE PHOTONIC INTEGRATED CIRCUITS FOR HIGH CAPACITY OPTICAL NETWORKS

In this thesis, the main advances in control of complex and flexible Photonic Integrated Circuits (PICs) are presented, and a novel control technique applied to a reconfigurable PIC is demonstrated. It involves a complex optical chip with an array of Microring Resonators (MRRs) that are tuned and locked to a reference optical channel. Targeting the application on high capacity optical networks, the requirements of such networks are analyzed and an evaluation on how they can be fulfilled by integrated photonics technology is done. Then, two integrated photonics platforms, Silicon Photonics (SiP) and Silicon Oxynitride (SiON), are analyzed and their key features and applications highlighted. The required tasks to fabricate, test, operate and calibrate optical chips are discussed. Next, the key building blocks required to assemble a complete solution making use of integrated photonics technology are addressed. State-of-the-art optical devices are designed and analyzed in the level required to use them to construct a telecom graded reconfigurable optical filter. Such device is then designed, realized, characterized and operated with the fabrication of the optical chips being done in a commercial SiP foundry run. The automated tuning and locking of a hitless silicon MRR filter, exploiting a novel channel labeling scheme is demonstrated. Hitless tuning with more than 30 dB isolation is achieved. A dynamic channel reconfiguration making use of the SiP PIC is demonstrated by measuring the Bit Error Rate (BER) of a channel while the system is reconfigured. It is shown that the already established channel was not affected by the newly added channel. Moreover, a technique to measure in-band Optical Signal to Noise Ratio (OSNR) of optical channels also making use of the labeling technique is demonstrated. All the demonstrated techniques, architectures and strategies are key enabling technologies for the applications of PICs in add-drop reconfigurable optical node architectures.

NAVUDAY SHARMA

INCREASING CAPACITY OF WIRELESS NETWORKS THROUGH AERIAL BASE STATIONS

In the past decade, due to immense high speed data and wider connectivity requirements, the cellular technologies have been continuously evolving leading to a major revolution in telecommunication industry. Currently, under IMT 2020, commonly known as 5th Generation of Cellular Technology, the targeted applications have been broadly classified as: enhanced mobile broadband (eMBB), massive machine type communication (mMTC), ultra-reliable and low latency communications (uRLLC), vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2X). In order to address the increasing data requirements, particularly, during flash crowds such as concerts, rallies, festivals, sport events etc, where many people gathered around in an area, use data services such as video streaming, photo sharing, video calls etc, higher densification of terrestrial network architecture is gaining immense importance in the name of Ultra-Dense Networks (UDNs). The data traffic increases manifold due to development of high resolution and big screen smart devices such as phones, tablets, laptops with 4K resolution, which eventually demands for higher data. Recently, Unmanned Aerial Vehicles (UAVs), generally known as drones were started to be investigated to provide data service to the users in suburban areas. Such projects were initially started by Google and Facebook to develop solar powered drones providing internet services. Later, companies such as Qualcomm, China mobile and Nokia started to investigate on other aspects such as controlling the drones through LTE base stations, cooperative communication among the drone network etc. Successful handovers were reported by Qualcomm with zero link failure in autonomous drone control through LTE network. Similar project was conducted by Ericsson and China mobile with deployment of prototype in field trial. Also, with techniques such as swarm optimization and collision avoidance algorithms, the drone network could be deployed in the city environment as well, as reported by Nokia. With such interest from industry and academia, research on UAVs acting as Aerial Base Stations (ABS) attained immediate attention. Therefore, this thesis mainly discusses about several technological aspects pertaining to such a system. Although, there are many research directions to the development of an ABS network, here we address certain directions. The work in this thesis, initially starts with the Air-to-Ground (A2G) Channel Modeling. There are many A2G channel models existing in the literature but since the work done here is to provide cellular network by UAVs, Low Altitude Aerial Platforms (LAPs) were preferred up to the altitude of 2000 m. The existing channel models mostly deal with commercial or military aircrafts, which fly at very high speeds near to subsonic or supersonic ranges. Such high speeds are not preferred for an ABS network to avoid frequent handovers. Also, some researchers have argued about the effect of Doppler shift at higher speeds which would finally lower the performance of the system. However, from the analysis and simulations provided in this thesis, effect of Doppler was not much observed with A2G channel and implementation of 5G waveforms. Due to unavailability of channel parameters for LAPs, measurements were performed using a radio propagation simulator for different environments: Suburban, Urban and Urban High Rise. These environments were generic since they were developed using ITU-R parameters for the simulator. Therefore, results from this simulator could be applied to practical environments with small degree of inaccuracy. Further, the work on cell coverage, capacity and interference analysis was conducted with simulation results obtained from the ray tracing simulator and verification of these results were done by performing analytical analysis and obtaining closed-form expressions. The graphs plotted using these expressions, matched with the graphs obtained from simulations with same simulation parameters. Later, in this thesis a new system was proposed as an optimal replacement to UDNs, to support the flash crowds. This system was termed as Ultra-Dense Cloud Drone Network (UDCDN). This system is advantageous as it offers reduction in Total Cost of Ownership (TCO) as compared to UDNs. Also, UDCDN is on-demand deployment system, i.e. it is deployed by the mobile operator only when required based on the data traffic information obtained from the cloud. Further, in the thesis, work has been done on implementing parameters for Ultra-Reliable Low Latency Communication (uRLLC) of 5G Physical Layer (PHY) on the ABS network to provide reliable and faster connectivity for ground users. Symbol Error Rate (SER) improvements were seen when uRLLC was implemented for A2G channel with Generalized Frequency Division Multiplexing (GFDM) modulation.

Further, results were also provided by introducing Carrier Frequency Offset (CFO) in the system model. Following the above work, other aspects were also considered in the thesis. Apart from primarily ABS network, a heterogeneous network (HetNet), consisting of multi-tier drone and terrestrial network, is seen as a more feasible options for the present cellular network. Therefore, further this thesis discusses about three major aspects which optimize the multi-tier ABS network: survivability, coverage and mobility laws to avoid issues related to inter-cell interference, frequent handovers, power deficiency etc. Later, work has also been done on interference alignment (IA) for maximizing the sum rate. However, IA demands for independent channels to provide better efficiency but at LAPs obtaining independent channels seems improbable. Therefore, such study serves better for a High Altitude Aerial Platform (HAP). An optimal receiver separation distance was also defined for the system. Another set of work has been performed in this thesis, which formed as a minor project apart from the major one reported throughout. The aim of this project was to develop an Internet of Thing (IoT) based Health and Usage Monitoring Systems (HUMS) for a helicopter. HUMS is an integrated recording and monitoring system that includes sensors, data acquisition technology and software algorithms (both on-board and ground-based) that are provided as a unit with the goals of reducing maintenance costs and improving safety and availability. For this system, the goal was to deploy various sensors all throughout the aircraft for monitoring different avionics components health and lifecycle. This data was to be send to a common gateway, known as a Transmission Data Concentrator (TDC) for data aggregations and processing. This processed data was sent to the cloud using cellular network when the aircraft is in the range of terrestrial network, otherwise using the satellite network.

MOHAMED KHALED SHEHATA

ENERGY EFFICIENCY AND SURVIVABILITY IN 5G CENTRALIZED ACCESS NETWORKS

The continuous demand for better wireless data services in terms of very high data rates (typically of Gbps order), extremely low latency, and significant improvement in users' perceived Quality-of-Service, has triggered the research on the fifth generation (5G) wireless systems that are expected to be deployed beyond 2020. Maintaining the current network architecture will lead to an unsustainable network-cost increase as well as to a dramatic expansion in the network power consumption. Hence, minimization of network cost and energy consumption have become a necessity for mobile network operators. In order to do so, the network infrastructure has to evolve from the old static architecture towards a more scalable, dynamic and agile one by resorting to novel technologies and architectural solutions to improve cost and energy efficiency. Centralized Radio Access Network (C-RAN) is a promising mobile network architecture designed to support the requirements of future 5G mobile networks. C-RAN is a new mobile access network architecture where base stations are splitted among BaseBand Units (BBU) and Remote Radio Heads (RRH) and BBUs are centralized into single physical location, with the consequent introduction of the new "fronthaul" transport network. In C-RAN, the "centralization" of baseband units enables substantial savings of computational resources and significant power savings. On the other hand, the deployment of C-RAN requires high capacity and imposes strict latency requirements on the fronthaul transport-network. To address these issues, various alternative architectures, known as "RAN functional splits", have been introduced to relax these strict fronthaul requirements. In this thesis, we investigate the opportunities enabled by C-RAN. First we provide a quick survey on the C-RAN stat of art. Then, we model the computational savings (what we called multiplexing gain) enabled by C-RAN under four different functional splits. Furthermore, we show the cost savings arises from centralization. To estimate the power savings -resulting from reduction in the computational resources- for the various cases, we identify the main power consumption contributors in a BS and provide a power consumption model for the different RAN split options. Following this centralization savings, we design a survivable C-RAN against BBU pool and link failures. We propose three different approaches for the survivable BBU pool placement problem and traffic routing in C-RAN deployment over a 5G optical aggregation network. We formalize different protection scenarios as Integer linear Programming (ILPs) problems. The ILPs objectives are to minimize the number of BBU pools, the number of used wavelengths and the baseband processing computational resources. Finally, we design survivable C-RAN based on shared path protection schemes with the objective to minimize the number of BBU pools and the number of used wavelengths. The results show the cost and energy advantages of CRAN with respect to classical RANs, due to "centralization" of BBUs into a few sites. Moreover, the results give insights on the latency and the transport network capacity on the BBU placement problem.

DARIO TAGLIAFERRI

VISIBLE LIGHT COMMUNICATIONS FOR NEXT GENERATION IN-FLIGHT SYSTEMS

In recent years, the amount of flights have undergone a rapid increase and at the same time also the number of passengers complaining for the lack of Internet connectivity during the flight is increased as well. Although some airline companies currently provide an Internet service based on both satellite communications (mostly in the Ku band) to ensure the aircrafts' coverage and on on-board cognitive Wi-Fi systems to interface with the users' terminals, the capacity is limited by the large geographical coverage of the geostationary satellites. In fact, the huge number of served airplanes forces the overall bandwidth to be subdivided in several small portions which are insufficient to serve all the passengers. Moreover, the potentials of cognitive-type radio frequency systems are greatly limited by the need to not generate spurious interference to the avionic instrumentation. This thesis work aims at investigating the possibility of providing the passengers with a radio frequency-free Internet access through optical wireless communications, in particular by using the visible light emitted by the LED reading lamps as a carrier for the downlink and the infrared light for the uplink. The wireless optical channel is studied in detail both in terms of impulse response and signal-to-interference-plus-noise ratio, characterizing its peculiarities. Moreover, optimized digital transmission techniques are tested to best suit with on-board scenarios. An innovative high-capacity backbone network is proposed to interconnect the LED reading lamps one to each other, reducing as much as possible the extra weight and limiting the energy consumption. Finally, the development of a visible light optical transceiver with its relative preliminary tests is reported.

ABI WAQAS

STOCHASTIC METHODS FOR PERFORMANCE PREDICTION OF PHOTONIC INTEGRATED CIRCUITS

Integrated photonic is the field where optical devices (such as couplers, interferometers, and so on) can be incorporated on a chip by the use of dielectric waveguides. Technology progress in integrated photonic has resulted in the implementation of complex photonic circuits combining many functions on a single chip, significant production volumes and reduced fabrication costs. While standard fabrication technologies are an essential condition for the commercial exploitation of photonic, they still have to face an unavoidable reality of uncertainties. As photonic devices are much longer when compare to wavelength, a slight variation in device geometry can cause a dramatic phase error, especially for devices based on interferometers such as Mach-Zehnder and micro-ring resonators. Each fabrication run is subject to several possible variations (waveguide width or height deviation, improper gap opening, change in material composition and surface roughness) that may eventually cause a fabrication deviations that reduce the yield to low levels to be economically sustainable. This means that several realized devices, which are designed to be nominally the same, differ from one to another due to variations of fabrication process. As a result, the device response is no longer considered as deterministic but is more suitably interpreted as a stochastic process and the analysis of photonic circuit is incomplete without the inclusion of stochastic analysis due to the presence of unavoidable uncertainties. To obtain a design of a photonic circuit with high yield and reduce performance variation, it is essential to consider such uncertainties in process design kits and to isolate the most critical parameters of the circuits and to estimate and reduce the cost of post-fabrication correction of the process variability. Therefore, statistical data and efficient statistical tools to include this data to predict the statistical behaviour of the final circuit are becoming fundamental instruments in photonic. In this doctoral dissertation, we focus on exploration and development of statistical tools to predict the performance of the integrated photonic circuit in the presence of unavoidable fabrication uncertainties. To address the problem of post-fabrication correction and to isolate the most critical parameters of the circuit, advanced statistical methods were studied and applied. More specifically, we focus on the application of two sensitivity analysis techniques, namely Elementary Effect Test and Variance-Based Sensitivity Analysis to investigate the behaviour of a photonic circuit under fabrication uncertainties, with the aim to identify the most critical parameters affecting the circuits' performance. We demonstrate, for the first time, the possibility to use the results of these methods to reduce the power consumption for the mitigation of statistical variation of circuits' parameters and increase the yield. To include information on the effect of fabrication uncertainties in each building block, we propose a Building-Block-based Generalized Polynomial Chaos method (BB-gPC) by properly exploiting stochastic collocation and Galerkin methods to realize a completely novel class of device models for the preparation of stochastic process design kits. Using these proposed stochastic macro-models that inherently convey stochastic information, only a single deterministic simulation is required to compute the statistical features of any arbitrary photonic circuit, without the need of running a large number of time-consuming circuit simulations thereby dramatically improving simulation efficiency. The BB models, in the form of transmission or scattering matrices, are circuit independent and can be stored and replace the original deterministic macro-model of the building blocks in the process design kit. The new matrices can hence be combined according to the building blocks connections to derive with a single run of the deterministic circuit simulator the stochastic behaviour of any circuit. The stochastic properties of the BB can reflect for example the foundry technological process and they are embedded in the PDK and should not be recalculated for every circuit. In the last part of work, applications of the statistical methods (sensitivity analysis and generalized polynomial chaos) are discussed and demonstrated by means of relevant photonic circuit examples. Beside pure analysis purposes, we demonstrate the exploitation of sensitivity analysis results to estimate the cost of a post-fabrication correction of the process variability, for example in terms of energy dissipated for the compensation of the uncertainties, relating this to the maximum yield that can be achieved by the process. Next, for the first time, the use of generalized polynomial chaos technique is demonstrated for stochastic modelling of the photonic circuit from experimental data. The statistical properties of a circuit are correctly estimated from a reduced number of experimental characterizations, saving costly and time-consuming measurements whilst achieving good accuracy comparable to those obtained by Monte Carlo. This approach should help to test and to tune in future mass production technological processes. Lastly, the effect of spatial correlation on the yield of photonic devices is discussed and efficient method that includes the treatment of spatial correlation is presented.

Novembre 2019

For more information

DEIB: www.deib.polimi.it
IT PhD Program: https://dottoratoit.deib.polimi.it/
PhD secretariat: phd-inf@polimi.it
PhD School: http://www.dottorato.polimi.it/