


ERCIM



NEWS

Special theme:

QUANTUM Technology

Also in this issue's Research and Innovation Section:

**SYNAPSE: An Integrated Platform for Cyber Risk
and Resilience Management**

Editorial Information

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of about 2,000 printed copies and is also available online at <https://ercim-news.ercim.eu>.

ERCIM News is published by ERCIM EEIG
BP 93, F-06902 Sophia Antipolis Cedex, France
+33 4 9238 5010, contact@ercim.eu
Director: Dominique Hazaël-Massieux,
ISSN 0926-4981

Contributions

Contributions should be submitted to the local editor of your country

Copyright notice

All authors, as identified in each article, retain copyright of their work. ERCIM News is licensed under a Creative Commons Attribution 4.0 International License (CC-BY).

Advertising

For current advertising rates and conditions, see
<https://ercim-news.ercim.eu/> or contact peter.kunz@ercim.eu

ERCIM News online edition

<https://ercim-news.ercim.eu>

Subscription

Subscribe to ERCIM News by sending an email to
en-subscriptions@ercim.eu

Editorial Board:

Central editor: Peter Kunz, ERCIM office (peter.kunz@ercim.eu)

Special Theme guest editors:

- Miroslav Dobsicek (RISE)
- Sebastian Raubitzek (SBA Research)
- Filippo Vella (CNR-ICAR)

Local Editors:

- Ferran Argelaguet, Inria, France (ferran.argelaguet@inria.fr)
- Andras Benczur, SZTAKI, Hungary (benczur@info.ilab.sztaki.hu)
- José Borbinha, Univ. of Technology Lisboa, Portugal (jlb@ist.utl.pt)
- Monica Divitini, NTNU, Norway (divitini@ntnu.no)
- Marie-Claire Forgue, ERCIM/W3C (mcf@w3.org)
- Lida Harami, ICS-FORTH, Greece (lida@ics.forth.gr)
- Athanasios Kalogeras, ISI, Greece (kalogeras@isi.gr)
- Georgia Kapitsaki, Univ. of Cyprus, Cyprus (gkapi@cs.ucy.ac.cy)
- Annette Kik, CWI, The Netherlands (Annette.Kik@cwi.nl)
- Alexander Nouak, Fraunhofer-Gesellschaft, Germany (alexander.nouak@iuk.fraunhofer.de)
- Laura Panizo, University of Malaga (laurapanizo@uma.es)
- Erwin Schoitsch, AIT, Austria (erwin.schoitsch@ait.ac.at)
- Thomas Tamisier, LIST, Luxembourg (thomas.tamisier@list.lu)
- Maurice ter Beek, CNR-ISTI, Italy (maurice.terbeek@isti.cnr.it)

Image by Gerd Altmann from Pixabay

Contents

JOINT ERCIM ACTIONS

- 4 **European Software Community Launches Strategic Research and Innovation Agenda for the Future of Software**
- 5 **Margot Herin Wins the 2026 ERCIM Cor Baayen Early Career Researcher Award**
- 6 **Report on the 31st International Conference on Formal Methods for Industrial Critical Systems**
by Maurice ter Beek (CNR-ISTI)

SPECIAL THEME

- 7 **Quantum Technology – Introduction to the Special Theme**
by Miroslav Dobsicek (RISE), Sebastian Raubitzek (SBA Research) and Filippo Vella (CNR-ICAR)
- 10 **Quantum Optimization for Constraint-Aware Scheduling**
by José A. Tirado-Domínguez, Eladio Gutiérrez, and Oscar Plata (University of Málaga)
- 11 **Towards Practical Quantum Multi-Objective Optimisation**
by Zakaria Abdelmoiz Dahi (University of Lille, Inria, CNRS), Gabriel Luque and Francisco Chicano (University of Malaga)
- 13 **Quantum Neuromorphic Computing**
by Jacopo Settino, Andrea Vinci and Carlo Mastroianni (CNR-ICAR)
- 15 **Hard Is Not Enough: Where Quantum Optimization May Matter**
by Pascal Halffmann (Fraunhofer Institute for Industrial Mathematics ITWM)
- 16 **The Bottleneck is not Cryptography: Understanding the Costs of Quantum-Safe Communications**
by Ruben Rios and José A. Montenegro (University of Málaga)
- 18 **Austria's Quantum-Communication Infrastructure Testbed**
by Hannes Hübel, Manuela Kos and Sebastian Ramacher (AIT Austrian Institute of Technology)
- 20 **How the Laws of Physics Help to Safeguard our Digital Future**
by Florian Kanitschar and Christoph Pacher (AIT Austrian Institute of Technology)

- 22 Planning and Sizing National Quantum Communication Networks**
by Sebastian Raubitzeck, Rebecca Klingbeil (SBA Research) and Werner Strasser (fragmentiX Storage Solutions)
- 24 From QKD to the Quantum Internet: Building a Full-Stack Quantum Network Infrastructure**
by Claudio Cicconetti (CNR-IIT) and Alessandro Zavatta (CNR-INO)
- 26 Forecasting Chaotic Dynamics with Hybrid Quantum Machine Learning**
by Luca Salatino, Francesco D'Amore and Andrea Giordano (CNR-ICAR)
- 27 Towards Quantum-Inspired Reactive Robots**
by Giovanni Pilato (CNR-ICAR), Antonio Chella (Università degli Studi di Palermo) and Salvatore Gaglio (Università degli Studi di Palermo and CNR-ICAR)
- 29 Quantum Machine Learning for Malware Analysis**
by Giovanni Ciaramella (IMT Lucca and CNR-IIT), Fabio Martinelli (CNR-ICAR) and Francesco Mercaldo (University of Molise)
- 30 Quantum Optimization in Systems Control**
by George T. Stamatiou and Kostas Magoutis (FORTH-ICS and University of Crete)
- 32 Designing how Quantum Neural Networks Represent Data**
by Smita Chakraborty (RISE)
- 33 From Quantum Randomness to Financial Simulation: Synthetic Data and Quantum Monte Carlo**
by Enrico Barbierato and Nicoleta Mihalachi (Catholic University of the Sacred Heart)
- 35 Genetic Programming for Quantum Oracle Optimization**
by Javier Zayas-Gallardo (University of Extremadura), Francisco Chicano (University of Málaga) and Juan Manuel Murillo (University of Extremadura)
- 37 Beyond a Single Measurement: Training Reusable Quantum Models**
by Erik Källman (RISE)
- 38 Quantum Logic Programming in Prolog**
by Giovanni Pilato (CNR-ICAR), Salvatore Gaglio (Università degli Studi di Palermo and CNR-ICAR) and Filippo Vella (CNR-ICAR)
- 39 The Quantum Software Crisis: A Reminder of Classical Software Engineering**
by Enrique Moguel, Jaime Alvarado-Valiente, and Javier Romero-Alvarez (Universidad de Extremadura)
- 41 Towards the Design of Service-Oriented Hybrid Quantum-Classical Systems**
by Álvaro M. Aparicio-Morales, Ernesto Pimentel (University of Málaga) and Jose Garcia-Alonso (Universidad de Extremadura)

- 43 Quantum-Inspired Tensor Networks for Counting Software Configurations**
by Daniel-Jesus Munoz and Lidia Fuentes (Universidad de Málaga)
- 45 Quantum Retina: A Quantum Eye on the Brain and Other Things**
by Göran Wendin (RISE)
- 46 Formal Methods Move into Quantum Computing**
by Kim Guldstrand Larsen, Christian Schilling (Aalborg University), Mirco Tribastone (IMT School for Advanced Studies Lucca), Max Tschaikowski and Paolo Zuliani (Sapienza University of Rome)
- 48 Cold Light in Orbit: A Photonic Path to Quantum Computing in Space**
by Andreas M. Hein and Carl Shneider (University of Luxembourg)

RESEARCH AND INNOVATION

- 50 Crete Looks to the Stars: SPACE-Crete Secures European Excellence Funding**
by Panagiotis Tsakalides (FORTH-ICS) and Vassilis Charmandaris (FORTH)
- 52 SYNAPSE: An Integrated Platform for Cyber Risk and Resilience Management**
by Massimo Cossentino, Giuseppe Papuzzo, Vishal Sonth and Fabio Martinelli (ICAR-CNR)
- 54 OpenSwarm: Coordinated Smart Devices for a Sustainable Future**
by Giulia Petrarulo, Filip Maksimovic and Thomas Watteyne (Inria)

ANNOUNCEMENTS

- 56 Dagstuhl Seminars and Perspectives Workshops**
- 56 ERCIM Forum Beyond Compliance on Digital Ethics in Research**
- 57 Training in Digital Ethics**
- 57 Van Wijngaarden Awards 2026 Honour Pioneers in Mathematics and Computer Science**
- 57 ERC Starting Grant to Advance Reliable Quantum Error Correction**
- 58 ERCIM “Alain Bensoussan” Fellowship Programme**
- 59 Showcase your EU Project in ERCIM News**

NEXT ISSUE

ERCIM News 147, January 2027
Special theme: Data Spaces

European Software Community Launches Strategic Research and Innovation Agenda for the Future of Software

European researchers, research organisations and companies have presented a joint Strategic Research and Innovation Agenda (SRIA) for the Future of Software in Brussels. The agenda sets out the software capabilities Europe needs to strengthen in the coming years to help shape its digital future.

The SRIA was developed through the Software Research Working Groups of Informatics Europe and ERCIM, with contributions from researchers, research organisations, industry and other stakeholders from across the European software ecosystem.

At the heart of the agenda is a clear message: “Software is the foundation of Europe’s digital future. Let’s make sure Europe has the capabilities to shape it.”

The agenda identifies four key areas in which Europe needs to strengthen its software capabilities:

- trustworthy, secure and compliant software systems;
- large-scale software systems that can evolve sustainably;
- sovereign and resilient digital infrastructure;
- human and organisational capabilities.

These priorities are considered through five cross-cutting perspectives: societal value, sustainability, skills, governance and innovation.

Bringing together research, industry and policy

The launch event brought together representatives from research, industry and European policy organisations. Contributions to the programme included Siemens, Canon, Ericsson, Eiffel Software, Software Heritage, HiPEAC and the French Secrétariat général pour l’investissement (SGPI).

Following the launch, the SRIA was formally presented at the European Commission to Thibaut Kleiner, Director for Future Networks at the Directorate-General for Communications Networks, Content and Technology (DG CONNECT). His area of responsibility includes software, cloud and other digital infrastructures.

A starting point for further European action

The presentation of the SRIA marks the beginning of a broader dialogue on how Europe can strengthen its software capabilities through future investments in research, innovation, infrastructure and skills.

The initiative was led by Prof. Dr. Tanja Vos, Professor at the Universitat Politècnica de València and the Open Universiteit in the Netherlands, and Chair of the Software Research Working Group of Informatics Europe, together



Figure 1: Participants at the launch of the Strategic Research and Innovation Agenda for the Future of Software in Brussels.



Tanja Vos and Tijs van der Storm presenting the SRIA to Thibaut Kleiner at the European Commission.

with Prof. Dr. Tijs van der Storm, Senior Researcher at CWI, Professor at the University of Amsterdam, and Chair of the Software Engineering Working Group of ERCIM.

Both are affiliated with VERSEN, the Dutch association for software engineering research, from which the initiative to develop the agenda at European level originated. Vos and Van der Storm coordinated the joint efforts of the European software research communities to develop the strategic agenda.

The Strategic Research and Innovation Agenda for the Future of Software is available on the ERCIM website: <https://www.ercim.eu/publications/strategic-reports>

Margot Herin Wins the 2026 ERCIM Cor Baayen Early Career Researcher Award

Margot Herin, currently a postdoctoral researcher at the Instituto de Telecomunicações in Lisbon, has been named the winner of the 2026 Cor Baayen Early Career Researcher Award. The annual ERCIM award recognises outstanding early-career researchers in computer science and applied mathematics.

Margot Herin works at the intersection of machine learning and decision theory, with the aim of developing artificial intelligence systems that are both more transparent and more rational. Her research focuses in particular on preference learning: methods that learn how decisions or choices are made while preserving desirable rationality properties and offering interpretable models whose parameters and computations have a clear meaning.

Herin received her PhD in 2025 from Sorbonne University in Paris for a thesis entitled Learning preference models: a marriage between decision theory and machine learning, supervised by Patrice Perny and Nataliya Sokolovska. The thesis combines mathematical foundations with extensive experimental validation on real and synthetic data. It has already received significant recognition, including the Gilles Kahn 2025 PhD Award of the French Computer Science Society, under the patronage of the French Academy of Sciences, and the 2026 PhD Award of the CNRS Operations Research and Decision research group.

Her work also demonstrates the potential of preference-learning methods in applications beyond core AI research. In an interdisciplinary project involving researchers in law and computer science, the methods were applied to predicting compensatory damages in divorce cases. She has also collaborated on uncertainty quantification and sensitivity analysis, developing new indices for identifying influential variables in numerical models of physical phenomena, which have found applications, for instance, in environmental monitoring.

More recently, her research has expanded towards large language models, including methods for aligning LLMs with diverse preferences and optimisation techniques for steering their generation processes.

Herin's work has appeared at leading international AI conferences including IJCAI, AAAI and UAI, as well as in journals such as Theory and Decision, Annals of Mathematics and Artificial Intelligence and SIAM/ASA Journal on Uncertainty Quantification. She received an IJCAI 2024 Distinguished Paper Award.

The 2026 selection committee unanimously selected Margot Herin from seven finalists. She was nominated by Jean-Frédéric Gerbeau of Inria, France.



The Cor Baayen Early Career Researcher Award was established in 1995 and is named after ERCIM's first President. It recognises outstanding scientific promise and achievement in informatics and applied mathematics and carries a prize of €5,000.

More information about the ERCIM Cor Baayen Early Career Researcher Award:

<https://www.ercim.eu/people-careers/cor-baayen-award>

2026 Cor Baayen Early Career Researcher Award

Winner:

Margot Herin, nominated by Jean-Frédéric Gerbeau, Inria, France

Finalists

- Adrián Segura Ortiz, nominated by University of Málaga, Spain
- Giovanni Puccetti, nominated by CNR, Italy
- Betül Bayrak Çik, nominated by NTNU, Norway
- Delia Rico Marchena, nominated by University of Málaga, Spain
- Marco De Vincenzi, nominated by CNR, Italy
- Corentin Bisot, nominated by Inria, France

Selection Committee

The Selection Committee was composed of Gabriel David (INESC-TEC), Monica Divitini (NTNU – chair of the ERCIM Human Capital Task Group), Georgia Kapitsaki (University of Cyprus), Bruno Lévy (Inria), Theodore Patkos (ICS-FORTH), Ernesto Pimentel (University of Malaga), and Fabrizio Sebastiani (ISTI-CNR).

The decision was unanimous.

Report on the 31st International Conference on Formal Methods for Industrial Critical Systems

by Maurice ter Beek (CNR-ISTI)

The annual conference of the ERCIM Working Group on Formal Methods for Industrial Critical Systems, FMICS, is the key conference at the intersection of industrial applications and formal methods. This year, the participants met in Liverpool, UK, on 3–4 September 2026.

The aim of FMICS is to provide a forum for researchers and practitioners interested in the development and application of formal methods in industry. It brings together scientists and engineers to exchange experience in their industrial use and promotes research and development to improve formal methods and tools for industrial applications. FMICS 2026 encouraged submissions combining formal methods with industrial practice, including their application in standardisation and industrial forums.

The conference was chaired by Kristin Yvonne Rozier (Iowa State University, USA) and Peter Gorm Larsen (Aarhus University, Denmark) and organised by the CONFEST 2026 general chairs David Purser and Patrick Totzke (University of Liverpool, UK) and their team as a co-located event, alongside CONCUR and QEST+FORMATS, as well as a number of pre- and post-conference workshops. FMICS 2026 overall attracted about 30 participants with larger audiences for the keynotes.

The international program committee (PC), with 25 members from 13 different countries, received 34 abstract submissions, resulting in 30 full-paper submissions from 17 different countries, and decided to accept 8 papers after a rigorous reviewing process. The program moreover included four excellent invited talks. Representing NASA JSC, USA, and sponsored by ERCIM, Dr. Julia Badger delivered a keynote talk on “Application of Formal Methods to Design and Verification of Autonomous Space Systems.” Dr. Colin O’Halloran, CTO of D-RiSQ Ltd, UK, and Honorary Visiting Professor at the

University of York, UK, delivered a keynote talk on “Reducing Time to Market through Formal Methods.” In conjunction with CONFEST, Prof. Ezio Bartocci of TU Wien, Austria, delivered a keynote talk on “Reasoning About Probabilistic Loops, Moment by Moment” and Prof. Ichiro Hasuo of the National Institute of Informatics (NII) in Tokyo, Japan, delivered a keynote talk on “Category Theory for Fast Model Checking Algorithms.”

Following a tradition established over the years, Springer and EASST (European Association for the Study of Science and Technology) sponsored the FMICS Best Paper Awards. This year, the PC selected the contributions “Exploiting the Layout of a Railway Interlocking System for Path Reliability Evaluation” by Alessandro Fantechi, Gloria Gori and Jacopo Zecchi (University of Florence, Italy) and “Reasoning about Concurrent Loops and Recursion with Rely-Guarantee Rules” by Ian J. Hayes and Larissa Meinicke (The University of Queensland, Australia) and Cliff Jones (University of Newcastle, UK).

The PC chairs of FMICS 2027 are Rosemary Monahan (Maynooth University, Ireland) and Michael Leuschel (Heinrich-Heine-University Düsseldorf, Germany). FMICS 2027 will take place in Athens, co-located with the 28th International Symposium on Formal Methods, which is organised by the general chairs Anna Philippou (University of Cyprus) and Petros Stefanias (National Technical University of Athens, Greece) and their team on 18–22 October 2027.

Links:

FMICS 2026 conference website:

<https://confest-2026.github.io/fmics/>

ERCIM WG FMICS: <https://fmics.inria.fr/>

Reference:

P.G. Larsen and K.Y. Rozier (eds.), “Formal Methods for Industrial Critical Systems”, Proc. of the 31st Int. Conf. on Formal Methods for Industrial Critical Systems (FMICS’26), Liverpool, UK, 3–4 September 2026, LNCS, vol. 16967, Springer, 2026.

<https://link.springer.com/book/9783032368485>

Please contact:

Maurice ter Beek, CNR-ISTI, Pisa, Italy
maurice.terbeek@isti.cnr.it



Alessandro Fantechi and Gloria Gori received their FMICS 2026 Best Paper Award from the PC chairs Kristin Yvonne Rozier and Peter Gorm Larsen.



Larissa Meinicke received her FMICS 2026 Best Paper Award from the PC chairs Kristin Yvonne Rozier and Peter Gorm Larsen.

Introduction to the Special Theme

Quantum Technology

by Miroslav Dobsicek (RISE), Sebastian Raubitsek (SBA Research) and Filippo Vella (CNR-ICAR)

In recent years, major advances in quantum hardware, algorithms and software have created great opportunities for scientific discovery and industrial innovation.

What can be observed is that quantum technology is entering a new phase where systems are becoming the central focus. Progress is no longer defined only by better qubits, individual algorithms or isolated demonstrations, but increasingly by whether quantum components can be integrated with classical computing, software, networks, sensing platforms and real applications. This transition, from quantum components to quantum systems, is changing both the research questions and the criteria by which progress should be judged.

State of the art

When ERCIM News last devoted a Special Theme to quantum computing, in January 2022, the field was still largely organised around individual building blocks: improving quantum hardware, developing algorithms and software stacks, identifying promising applications, and asking when quantum advantage might become possible [1]. These questions remain central, but four years later a broader transition is becoming visible.

Quantum technology is entering a systems phase. Progress is defined not only by better qubits, individual algorithms or isolated demonstrations, but increasingly by whether quantum components can be integrated with classical computing, software, networks, sensing platforms and real applications. This transition from quantum components to quantum systems is changing both the research questions and the criteria by which progress should be judged.

Quantum processors are increasingly considered as specialised accelerators connected to HPC and classical computing resources rather than as stand-alone machines. Quantum algorithms are embedded into classical optimisation, machine-learning, simulation and control workflows. Quantum communication is moving from individual links towards networks that must coexist with conventional telecommunications, key management and post-quantum cryptography. Quantum software research is expanding from circuit construction towards abstraction, interoperability, verification and composability. In sensing, longer-term visions increasingly combine sensing, communication and information processing within a single architecture.

This system perspective also changes how progress should be evaluated. Quantum algorithms need strong classical baselines, quantum processors must be assessed as parts of complete workflows, and quantum key distribution (QKD), quantum software and sensing technologies must function

within operational infrastructures. European policy and infrastructure increasingly reflect this broader view. The Quantum Europe Strategy adopted in 2025 connects research and innovation with infrastructure, ecosystem development, skills and emerging applications [2], while EuroHPC is integrating several quantum computers with European supercomputing resources [3].

The contributions in this Special Theme reflect the same transition. They do not point towards a single route to quantum advantage or one dominant technology. Instead, they show how quantum effects, devices and algorithms may become components of larger computational, communication and sensing systems, and what new scientific and engineering problems arise when they do.

Main challenges

The familiar physical limitations of quantum technology remain: noise, restricted scale, limited circuit depth and the cost of fault tolerance constrain quantum computing, while communication and sensing face their own technological limitations. But system integration adds another layer of difficulty.

One central challenge is identifying where quantum resources genuinely add value. Computational hardness alone does not imply quantum advantage. Classical optimisation, simulation and machine learning benefit from mature algorithms, specialised hardware and decades of engineering. Credible quantum applications therefore require strong classical comparisons and a clear understanding of which part of a workflow might plausibly benefit from quantum processing.

A second challenge is integration. Most near- and medium-term applications are likely to be hybrid. Quantum processors must interact with CPUs, GPUs, HPC systems, classical optimisation loops, networks and domain-specific software. Performance therefore becomes an end-to-end system property: latency, communication, resource scheduling, data encoding and orchestration may matter as much as quantum execution itself.

A third challenge is software and infrastructure maturity. Quantum software remains strongly influenced by particular devices, providers and software development kits (SDKs). Quantum communication likewise needs more than secure physical-layer protocols: practical systems require interoperability, reliability, key management and integration with existing networks. Across the field, it is increasingly important to distinguish clearly between demonstrated capability, plausible research direction and longer-term vision.

Future directions

These challenges point towards a field increasingly defined by co-design. Hardware, algorithms, software and applications cannot be optimised independently. Quantum processors will increasingly be treated as specialised components within larger computing systems; quantum communication will combine quantum and classical security mechanisms;

sensing architectures may integrate local information processing; and software abstractions must hide unnecessary hardware complexity without losing the information required for meaningful resource decisions.

Shared infrastructure will be particularly important because it allows these complete systems to be studied. HPC-integrated quantum computers, operational QKD testbeds, open software tools and network experimentation platforms provide environments in which assumptions about performance, scalability and interoperability can be tested.

Contributions in this issue

The contributions to this Special Theme span a wide range of topics, from foundational algorithms and simulation to quantum networks, communication and security, hybrid quantum-classical systems, quantum software engineering, sensing, and longer-term technology visions.

Several articles address the fundamental question of where quantum computation may provide useful capabilities. **Halffmann (p. 15)** argues that computational hardness alone is not enough to identify promising quantum-optimisation applications and instead emphasises problem structure and comparison with mature classical solvers. **Tirado-Domínguez and colleagues (p. 10)** introduce a constraint-aware scheduling framework that separates optimisation objectives from constraints within a hybrid QAOA formulation. **Dahi and colleagues (p. 11)** address multi-objective optimization, combining distributed execution, hardware-aware circuit design and classical assistance.

Quantum machine learning raises similar questions about representation and structure. **Chakraborty (p. 32)** focuses on the embedding of classical data into quantum states and discusses trainability, inductive bias and quantum variants of dropout as elements of a more resource-aware design methodology. **Settino and colleagues (p. 13)** explore quantum neuromorphic computing, where the dynamics, memory and dissipation of the physical quantum systems themselves become computational resources. **Muñoz and Fuentes (p. 43)** demonstrate the reverse flow of ideas: tensor-network methods originating in quantum physics are used on classical GPU hardware for exact counting of software configurations. These contributions illustrate how quantum-algorithm research is increasingly concerned with identifying useful structure rather than simply mapping hard problems onto quantum hardware.

The articles on quantum communication and security show particularly clearly the movement from components towards systems. **Hübel, Kos and Ramacher (p. 18)** report operational experience from Austria's QCI-CAT infrastructure, where QKD has been combined with post-quantum cryptography, key-management systems and governmental and medical applications. **Raubitzek and colleagues (p. 22)** ask what scaling towards national QKD infrastructure would require in terms of nodes, fibre and equipment.

Cicconetti and Zavatta (p. 24) describe an Italian inter-regional quantum-network testbed initially based on commercial

QKD but designed to evolve towards quantum repeaters and Quantum Internet experimentation. **Kanitschar and Pacher (p. 20)** present continuous-variable QKD and its potential compatibility with existing optical-telecommunication technology. **Rios and Montenegro (p. 16)** approach quantum-safe communication from the complementary perspective of post-quantum cryptography, showing through end-to-end benchmarking that deployment bottlenecks can arise as much from protocols, software and networks as from cryptographic primitives themselves. Together, these contributions move the discussion from individual security mechanisms towards complete quantum-safe communication infrastructures.

The system perspective is equally apparent in hybrid quantum-classical applications. **Stamatiou and Magoutis (p. 30)** embed quantum annealers within model-predictive control loops and show that surrounding communication and infrastructure can dominate the quantum-computation latency. **Salatino and colleagues (p. 26)** divide chaotic forecasting between a quantum reservoir providing nonlinear information processing and a classical recurrent component providing memory.

Barbierato and Mihalachi (p. 33) examine quantum generative models and quantum Monte Carlo methods in financial applications. **Ciaramella and colleagues (p. 29)** investigate quantum machine learning for malware analysis, including cases where classical baselines remain more accurate or dramatically faster – an important result in a field where realistic comparison is essential. **Pilato and colleagues (p. 27)** explore hybrid and quantum-inspired representations for reactive robotics. Collectively, these articles reinforce the view that useful near-term quantum technologies are more naturally understood as components inside classical workflows than as replacements for complete classical systems.

As these systems become more complex, quantum software engineering becomes increasingly important. **Alvarado-Valiente and colleagues (p. 39)** revisit the historical software crisis and argue that engineering practices should mature alongside quantum hardware rather than after it. **Aparicio-Morales and colleagues (p. 41)** explore service-oriented architectures for exposing, selecting and orchestrating heterogeneous quantum capabilities through higher-level interfaces.

Zayas-Gallardo and colleagues (p. 35) use grammar-based genetic programming to automate the synthesis and optimisation of quantum oracles. **Källman (p. 37)** examines composability from another direction: when the output of one quantum model becomes the input of another, preserving a single observable is insufficient; what matters is the reduced quantum state exposed at the component boundary. **Pilato and colleagues (p. 38)** investigate quantum logic programming in Prolog as a higher-level route from logical descriptions towards quantum operations. **Larsen and colleagues (p. 46)** transfer established formal-methods techniques, including model reduction, model checking and satisfiability solving, to quantum-circuit simulation and verification. Together, these contributions move quantum software beyond the programming of individual circuits towards components that can be abstracted, composed, optimised and verified.

The final contributions extend the systems perspective towards sensing and longer-term architectures. **Wendin (p. 45)** presents the Quantum Retina, a vision in which arrays of quantum sensors are progressively integrated with information processing and communication, ultimately treating sensing and computation as parts of one architecture. **Hein and Shneider (p. 48)** examine cryogenic photonic computing in space: primarily a classical photonic architecture in the nearer term, but also a possible infrastructure pathway towards future photonic quantum processing and networking.

Taken together, the contributions in this issue show a field becoming broader, more concrete and more engineering-driven. Progress in quantum technology still depends critically on better devices and deeper understanding of quantum phenomena, but useful quantum technology will not emerge from the quantum component alone. It will depend on how that component interacts with algorithms, classical computing, software, networks, sensors and applications.

The transition from quantum components to quantum systems may therefore be one of the defining challenges of the coming years. The articles collected here approach that transition from very different directions: improving algorithms and abstractions, integrating quantum technologies with existing infrastructure, identifying bottlenecks and limitations, and exploring architectures whose enabling components are only beginning to emerge. Together, they illustrate an important change in the field: the question is increasingly not simply what quantum technology can do, but how quantum technology can become a useful part of the systems we build.

References

- [1] S. Ali and S. Selstø, “Quantum Computing – Introduction to the Special Theme,” ERCIM News, No. 128, January 2022.
- [2] European Commission, Quantum Europe Strategy: Quantum Europe in a Changing World, COM(2025) 363, 2 July 2025.
- [3] EuroHPC Joint Undertaking, “EuroHPC JU Opens Access to its Quantum Computers,” 25 June 2026. The access programme initially included superconducting, photonic and trapped-ion systems and forms part of EuroHPC’s explicit strategy of integrating quantum computers with European HPC infrastructure.

Please contact:

Miroslav Dobsicek, RISE, Sweden
miroslav.dobsicek@ri.se

Sebastian Raubitzek, SBA Research, Austria
SRaubitzek2@sba-research.org

Filippo Vella, CNR-ICAR, Italy
filippo.vella@icar.cnr.it

ERCIM Working Group on Quantum Technologies

The ERCIM Working Group on Quantum Technologies (QT) brings together researchers working across quantum computing, communication, sensing, hardware and software. Its aim is to strengthen European collaboration in quantum technologies by connecting expertise across research institutions and by building bridges between academia and industry.

The group covers a broad range of topics, including quantum algorithms and optimisation, quantum machine learning, quantum-computing architectures and error correction, quantum communication and networking, quantum key distribution, sensing and metrology, quantum software engineering, hybrid quantum–classical systems, scalability and system integration. Particular attention is also given to reproducibility and benchmarking, including common baselines, datasets and reference models.

Beyond research, the Working Group provides a forum for networking and knowledge exchange. It encourages cooperation between research groups, supports the preparation of joint project proposals, promotes collaboration with industry and contributes to workshops, training activities and other events. An important goal is also to support young researchers and engineers and to help other scientific communities understand where quantum technologies may become relevant to their own work.

Participating ERCIM institutions currently include CNR (Italy), FORTH and ISI (Greece), Fraunhofer (Germany), INESC (Portugal), Inria (France), LIST (Luxembourg), NTNU (Norway), RISE (Sweden), SBA Research (Austria), UCY (Cyprus) and Universidad de Málaga (Spain), together with non-ERCIM participants King’s College London (United Kingdom), University of Galway (Ireland) and Fujitsu (Spain).

The group is led by Francisco Chicano (Universidad de Málaga) and Carlo Mastroianni (CNR), supported by a steering committee representing several ERCIM organisations.

The Working Group also played an important role in bringing together the community behind this ERCIM News Special Theme on Quantum Technology, helping to identify contributions and coordinate the scientific scope of the issue.

The Working Group is open to researchers and practitioners beyond ERCIM member organisations, and interested experts are invited to participate and contribute to its activities.

More information:

<https://ercim-qt-wg.icar.cnr.it/>

Quantum Optimization for Constraint-Aware Scheduling

by José A. Tirado-Domínguez, Eladio Gutiérrez, and Oscar Plata (University of Málaga)

Could quantum computing help address complex scheduling problems? We explore this question by proposing a structured problem representation that isolates constraints from optimization objectives. Our framework improves solution quality using fewer circuit layers than conventional implementations.

In the field of high-performance computing, task scheduling remains one of the most important optimization challenges. Across domains as diverse as air traffic control, hospital operating rooms, and future 6G communications networks, the challenge is to schedule activities with fixed durations within strict time windows while making efficient use of limited resources and avoiding overlaps. The combination of these temporal and resource constraints gives rise to problems that are often classified as NP-hard, with complexity growing exponentially as the number of tasks increases, making them difficult to solve using traditional classical methods.

Quantum computing offers a promising avenue for tackling complex optimization problems. While current Noisy Intermediate-Scale Quantum (NISQ) devices are still limited in scale and highly sensitive to noise, they provide an ideal platform for exploring new computational paradigms. In this context, we have developed QTIS (Quantum Time Interval Scheduler) [1], a hybrid quantum–classical framework designed to manage time intervals while exploring new approaches to resource allocation and constraint-aware scheduling.

The Hybrid Approach: A Bridge to the Future

At the core of QTIS lies a variant of the Quantum Approximate Optimization Algorithm (QAOA) [2], one of the most promising optimization techniques in the current NISQ era. Inspired by the principles of quantum annealing, QAOA builds a solution through a sequence of parameterized quantum operations that alternate between encoding the optimization problem and exploring new candidate states. The parameters governing these operations are iteratively tuned by a classical optimizer, which acts as a “coach” that guides the quantum circuit toward increasingly better solutions. This creates a hybrid quantum–classical workflow capable of tackling complex combinatorial problems.

What Makes QTIS Different? Divide and Conquer

Most conventional quantum scheduling approaches encode both the optimization objective and the scheduling constraints within a single mathematical formulation known as the problem Hamiltonian, which guides the quantum optimization process. QTIS takes a different approach.

Rather than treating this formulation as a single block, QTIS decomposes it into separate terms, each responsible for a specific aspect of the scheduling problem:

- The objective Hamiltonian: focused exclusively on what we want to achieve, namely selecting the best possible set of tasks.
- The constraint Hamiltonian: a guardian that ensures the rules of the problem are respected by preventing tasks assigned to the same resource from overlapping in time.
- The mixer operator: the quantum mechanism that enables the exploration of different candidate solutions across the search space. Unlike the complex problem-specific mixers proposed in [3], we employ the standard mixer to reduce circuit overhead.

This separation is more than a mathematical convenience. By assigning independent parameters to the objective and con-




Figure 1: Overview of QTIS (Quantum Time Interval Scheduler). The framework separates optimization objectives and scheduling constraints into independent Hamiltonians (HP and HC), combined with a standard mixer (HB) within a hybrid quantum-classical optimization loop. Conflict detection guides the search toward valid, high-quality schedules.

straint components, QTIS can balance optimization goals and scheduling constraints more effectively than conventional QAOA implementations. Experimental results show that QTIS consistently outperforms standard QAOA when using circuits of the same depth (see Figure 1). In most cases, QTIS with 10 layers achieves better solutions than a conventional QAOA implementation using 15 layers, demonstrating that a more structured problem representation can improve solution quality while reducing quantum resource requirements.

The Damselfly Gate: Detecting Scheduling Conflicts

One of the greatest challenges in scheduling is determining whether two tasks assigned to the same resource overlap in time. To address this issue, QTIS introduces an ancilla-assisted quantum subcircuit that explicitly incorporates conflict detection into the optimization process.

The framework supports two alternative approaches. In the Full Quantum variant, task overlaps are estimated directly within a quantum circuit by encoding temporal relationships into quantum states through parameterized rotations. Ancillary qubits are then used to mark potential conflicts between pairs of tasks. In the hybrid quantum–classical variant, overlap detection is performed classically and the resulting conflict information is transferred to the quantum circuit, providing a more deterministic alternative for noisy quantum hardware.

At the heart of the quantum approach lies a custom structure, which we call the Damselfly gate. Whenever an overlap is detected, the corresponding ancilla qubit activates a penalty in the constraint component, increasing the cost of invalid schedules. In this way, QTIS is naturally guided away from conflicting task assignments while converging on the optimal schedule.

A Smarter Starting Point

The quality of the solutions produced by QTIS depends not only on the quantum circuit itself, but also on how its parameters are initialized and optimized. To explore this aspect, we evaluated three optimization strategies: conventional parameter minimization, a progressive approach known as T-QAOA, and our newly developed method, Homotopy-Transfer QAOA (HT-QAOA).

Inspired by the principles of quantum annealing and homotopy optimization, HT-QAOA first optimizes a simple quantum circuit and then uses the resulting parameters as a guide for circuits with a larger number of layers. By transferring information from a simpler optimization stage to a more complex one, the algorithm starts its search from a more informed configuration.

This strategy improves the chances of finding high-quality schedules without increasing the computational cost. In the QTIS experiments, HT-QAOA achieved competitive performance while maintaining execution times comparable to those of standard optimization methods.

Looking Ahead

QTIS demonstrates how quantum computing can already contribute to solving real-world scheduling problems, even within the limitations of current NISQ hardware. By combining a novel problem decomposition, ancilla-assisted conflict detection through the Damselfly gate, and advanced optimization strategies such as HT-QAOA, the framework provides a new

perspective on constraint-aware scheduling. Although practical quantum advantage remains a long-term goal, QTIS highlights how today’s quantum devices can serve as a valuable testbed for developing the optimization technologies of tomorrow.

References:

- [1] J. A. Tirado-Domínguez et al., “QTIS: A QAOA-based quantum time interval scheduler,” *Future Gener. Comput. Syst.*, 2026.
- [2] E. Farhi et al., “A quantum approximate optimization algorithm”, 2014.
- [3] S. Hadfield et al., “From the quantum approximate optimization algorithm to a quantum alternating operator ansatz”, *Algorithms*, 12(2), 34, 2019.

Please contact:

José Antonio Tirado Domínguez
University of Malaga, Spain
jatirado@uma.es

Towards Practical Quantum Multi-Objective Optimisation

by Zakaria Abdelmoiz Dahi (University of Lille, Inria, CNRS), Gabriel Luque and Francisco Chicano (University of Málaga)

Quantum computing has the potential to revolutionise optimisation, yet most existing approaches focus on problems with a single objective. Our work bridges this gap by developing scalable quantum algorithms capable of tackling realistic optimisation problems involving multiple conflicting objectives.

Optimisation is widely regarded as one of the application domains where quantum computing could have the greatest impact. This field, grounded in mathematics and engineering, aims to find solutions that maximise (or minimise) a given quality metric, known as an objective function, which indicates the quality of a solution (e.g. determining where to place electric vehicle charging stations to satisfy as much of the city’s demand as possible). Single-objective problems involve one objective at a time. However, many practical problems involve several, often conflicting, objectives that must be optimised simultaneously. This category of problems, known as multi-objective problems, is quite complex and widespread (e.g. deciding where to build parking facilities to maximise city coverage while minimising construction costs). Rather than seeking a single optimum, multi-objective optimisation aims to identify a set of trade-off solutions, known as the Pareto front.

Many optimisation problems can naturally be formulated using binary decision variables (e.g. deciding whether to place or not an electric charging station in a candidate spot). Such problems, known as pseudo-Boolean problems, have received




Figure 1: Axes of research.

considerable attention in quantum single-objective optimisation. In contrast, quantum multi-objective pseudo-Boolean optimisation has received comparatively little attention. As a result, the applicability of quantum optimisation remains limited for many real-world problems involving multiple conflicting objectives. Devising multi-objective quantum algorithms poses several challenges from both the algorithmic and hardware perspectives. From an algorithmic perspective, integrating search mechanisms that capture the objectives' conflicting relationship, as well as solving problems with complex features are particularly challenging (even in the classical realm). From a hardware point of view, today's quantum machines remain limited in terms of qubit counts, gate fidelity, and robustness to errors.

The limitations of current quantum hardware, together with the complexity of pseudo-Boolean multi-objective optimisation problems, raise challenging research questions on how to design optimisers that (i) can efficiently tackle complex problems, (ii) scale despite the limitations of current quantum hardware, (iii) are resilient to quantum errors, and (iv) are practically implementable (see Figure 1). Thus, our efforts have focused on four axes: efficiency, scalability, feasibility, and applicability.

To address the scalability challenge, we first proposed a distributed optimisation framework [1]. The approach exploits high-performance computing to execute multiple quantum optimisation processes in parallel, each exploring different regions of the objective space through eigensolver-based quantum optimisers. This distributed strategy enables the approximation of multiple trade-off solutions while overcoming the limited computational capacity of individual quantum processors. Building upon this framework, we further improved the proposal's optimisation capabilities in [1,2]. The initial scalarisation-based approach was extended with a reference-based formulation (known as Techebycheff scalarisation), which guides the search towards different regions of the objective space and allows the optimiser to approximate much more complex Pareto fronts than conventional weighted-sum techniques.

Beyond scalability and efficiency, we also focused on improving the feasibility and practical applicability of quantum multi-objective optimisation. In [2], we exploited algorithmic properties to generate more compact quantum circuits, reducing the number of quantum operations and therefore the exposure to noise sources such as relaxation and dephasing. Finally, in [3], we proposed a machine learning-assisted distributed eigensolver that significantly reduces the amount of quantum-classical interaction required by NISQ algorithms (see Figure 2). This decreases communication overhead, execution costs, and error accumulation, while making the overall optimisation process more practical for current quantum hardware. Together, these contributions form a scalable quantum optimisation framework that combines distributed execution, improved Pareto front exploration, hardware-aware circuit design, and machine learning-assisted optimisation, enabling a broad class of pseudo-Boolean multi-objective problems to be tackled on current NISQ platforms.

Our research on quantum pseudo-Boolean multi-objective optimisation combines theoretical advances with extensive experimental validation. Our proposals [1-3] have been validated using the largest collection of quantum hardware platforms and benchmark instances reported so far, including the largest




Figure 2: Proposed approach.

IBM gate-based quantum computers and D-Wave quantum annealers, benchmark instances with up to one thousand variables, and problems of varying sizes and complexities. This is worth highlighting because our research pursues two complementary goals: advancing the theoretical foundations of quantum multi-objective optimisation while demonstrating its practical applicability. We believe both are essential stepping stones towards future advances in quantum multi-objective or general optimisation. As quantum hardware continues to evolve, these techniques are expected to scale to increasingly larger and more complex optimisation problems. We hope that this work contributes to bringing quantum multi-objective optimisation closer to real-world scientific and industrial applications.

References:

- [1] Z.A. Dahi, et al, “Scalable Quantum Approximate Optimiser for Pseudo-Boolean Multi-objective Optimisation”, in: Proc. of the 18th Int. Conf. on Parallel Problem Solving from Nature (PPSN’24), Springer LNCS vol 15151, 2024.
https://doi.org/10.1007/978-3-031-70085-9_17
- [2] Z.A Dahi, et al, “Scalable quantum Trotterised-vs-continuous annealing for pseudo-Boolean multi-objective optimisation”, Future Generation Computer Systems 183 2026, 108568.
<https://doi.org/10.1016/j.future.2026.108568>
- [3] Z.A. Dahi, F. Chicano, and G. Luque, “Tightening the Calculation Gap in Quantum Multi-Objective Pseudo-Boolean Optimisation”, in Proc. of the 19th Int. Conf. on Parallel Problem Solving From Nature (PPSN’26), Springer. https://doi.org/10.1007/978-3-032-36226-1_21

Please contact:

Zakaria Abdelmoiz Dahi
Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189
CRISTAL, France
abdelmoiz-zakaria.dahi@inria.fr

Gabriel Luque and Francisco Chicano
ITIS Software, University of Málaga, Spain
gluque@uma.es, chicano@uma.es

Quantum Neuromorphic Computing

by Jacopo Settino, Andrea Vinci and Carlo Mastroianni
(CNR-ICAR)

Neuromorphic computing shows how the physical dynamics of a device can become part of the computation itself. Quantum systems offer a different but potentially complementary opportunity. This perspective considers how quantum dynamics, interactions, memory, and dissipation might open new directions for neuromorphic computing, beyond simply running neural-network models on quantum processors.

In recent years, advances in artificial intelligence have been accompanied by a rapid increase in the computational resources required to train and run increasingly complex models. Training, in particular, involves repeatedly processing vast amounts of data and iteratively updating growing numbers of parameters, with substantial costs in energy, time, and computing infrastructure. This makes it increasingly clear that the future of AI will depend not only on new algorithms, but also on rethinking the hardware on which they run. In conventional architectures, memory and processing units are physically separated, and the continuous movement of data between them contributes significantly to computational and energy costs.

Neuromorphic computing offers one possible architectural response. Inspired by organizational principles found in biological neural systems [1], neuromorphic architectures can distribute memory and processing across the same physical network and operate in an event-driven manner, activating computation only when required. More broadly, they aim for a close integration between computational function and physical substrate, exploiting properties that arise from the system’s internal dynamics. From this perspective, hardware is no longer merely the substrate on which an algorithm is executed: its physical dynamics can become part of the computation itself.

Quantum neuromorphic computing

A useful conceptual parallel can be drawn with quantum computing. In a quantum device, information is encoded in the physical state of the system, whose evolution can itself be controlled to process information. Conventional quantum computing is largely organized around programmable sequences of quantum gates. A neuromorphic perspective, however, suggests a different question: can the collective dynamics of a quantum system, including its interactions, temporal response, and memory, be directly used to perform computational tasks?

This idea is central to quantum neuromorphic computing (QNC), an emerging field that seeks to combine principles of neuromorphic computation with resources available in quantum physical systems [2]. The more distinctively neuromorphic direction is to make the dynamics and physical organization of the quantum device part of the computational architecture itself.

One of the clearest early examples of this dynamics-based approach is represented by Quantum Reservoir Computing

(QRC). In reservoir computing, the internal dynamics of a complex system transform input information into a high-dimensional representation, while training is largely confined to a simple linear output layer.

In QRC, this transformation is performed by the evolution of a quantum system, whose state is represented by a vector in a Hilbert space whose dimension grows exponentially with the number of qubits, making the system's dynamical properties an integral part of information processing.

Recent hybrid architectures further illustrate how computational functions such as memory can be distributed between quantum dynamics and classical post-processing. For example, memory augmentation through classical processing of quantum measurements can reduce experimentally demanding requirements while retaining the dynamical processing performed by the quantum system [3].

Beyond the reservoir: designing quantum-neuromorphic architectures

A broader perspective is to move from using a quantum system as a reservoir to designing it explicitly to implement specific neuromorphic functions. In such devices, the physical structure of an interacting quantum network would itself become part of the computational architecture.

For example, the topology of interactions could determine how information propagates through the system and which collective modes are excited. Different relaxation timescales, engineered through controlled coupling to the environment, could provide forms of temporal memory by making the present response depend on recent inputs. Controlled dissipation could thus become part of the design space rather than simply a process to be minimized.

This points towards quantum networks in which connectivity, timescales, dissipation, and measurement points are co-designed for a target task (Figure 1). The goal would no longer be only to exploit whatever complex dynamics a device happens to provide, as in a generic reservoir, but to construct dynamics with desired computational properties: memory, temporal selectivity, collective response, and distributed processing.

Several emerging concepts already point in this direction. Quantum memristors introduce history-dependent responses at the device level, while quantum spiking-neuron and leaky-integrate-and-fire models explore ways of translating neuromorphic primitives such as integration, leakage, thresholding, and spiking into quantum systems. These approaches are not yet a mature architectural framework. Rather, they illustrate a possible evolution of the field: from quantum systems that merely provide complex dynamics to quantum elements deliberately




Figure 1: Concept of quantum neuromorphic computing. Local inputs drive an interacting quantum network whose topology, dissipation and temporal memory are designed so that its physical dynamics perform a desired computational task.

designed to implement specific neuromorphic primitives, which could eventually be combined into larger networks.

The decisive question is therefore not whether a quantum-neuromorphic device can be built, but what it could do better or fundamentally differently than the best classical neuromorphic architectures. A possible advantage might arise when the input itself is quantum and can be processed without first being converted into a classical representation, or when, for a specific task, the functions naturally generated by the quantum dynamics provide an effective representation of the relevant input-output relation, without requiring the explicit construction of complex functions by hand.

The long-term potential of quantum neuromorphic computing lies in bringing together two complementary ideas: the quantum view of information as a physical resource and the neuromorphic view of computation as an emergent property of dynamics, memory, and connectivity.

References:

- [1] D. Kudithipudi et al., “Neuromorphic computing at scale,” *Nature* 637, 801–812 (2025). DOI: 10.1038/s41586-024-08253-8.
- [2] D. Marković and J. Grollier, “Quantum neuromorphic computing,” *Applied Physics Letters* 117, 150501 (2020). DOI: 10.1063/5.0020014.
- [3] J. Settino et al., “Memory-augmented hybrid quantum reservoir computing,” *Physical Review Applied* 24, 024019 (2025). DOI: 10.1103/wzww-7rk2.

Please contact:

Jacopo Settino
ICAR-CNR, Italy
jacopo.settino@icar.cnr.it

Hard Is Not Enough: Where Quantum Optimization May Matter

by Pascal Halffmann (Fraunhofer Institute for Industrial Mathematics ITWM)

Hardness alone does not identify promising quantum applications. Classical optimization has spent decades exploiting problem structure; thus, quantum methods must compete with far more than brute force. Where, then, might quantum optimization genuinely contribute, and what kinds of problems offer a plausible path to real-world utility or even advantage?

Quantum optimization is booming

Quantum optimization has become a highly active area of quantum computing. A search of arXiv's Quantum Physics category identified 182 preprints related to quantum optimization that were published this year until September 1 [L1]. The appeal is clear: optimization problems are ubiquitous, economically important, and often computationally hard. To date, no quantum approach has demonstrated a practical advantage over classical methods in solution time or quality. Some critics therefore question whether such an advantage is attainable at all.

What theory actually promises

Aaronson's "law of conservation of weirdness" expresses the intuition that exponential quantum speedups require unusual problem structure. The related, unproven Aaronson–Ambainis conjecture concerns the black-box/query model stating that a quantum algorithm's acceptance probability ("yes"- vs. "no"-answer) can be approximated classically, to fixed additive accuracy, on most inputs with only polynomial overhead. If true, it rules out generic exponential query speedups for decision

problems on typical unstructured inputs but does not prohibit quantum advantage: it does not establish efficient classical simulation of arbitrary quantum computations and polynomial advantages for structured, search or approximation problems remain possible.

The best-known quantum optimization methods do not overturn this cautious conclusion. For the Quantum Approximate Optimization Algorithm (QAOA), the workhorse of gate-based quantum optimization, no general scaling advantage is known. Quantum annealing can be exponentially faster on a constructed oracle problem, yet no advantage is established for practical NP-hard instances. Grover search, and its minimum-finding extension, offers a proven quadratic advantage over exhaustive enumeration, not over structure-exploiting classical algorithms.

More promising routes exploit structure directly. Quantum backtracking and branch-and-bound provably accelerate structured search trees. The recently published algorithm based on Decoded Quantum Interferometry (DQI) turns Fourier-structured optimization into decoding. For distinct artificial problems, it is superpolynomially faster than known classical algorithms [1]. While it has been beaten for Max-k-XORSAT, DQI has challenged classical solvers given the right problem structure.

An end-to-end reality check

We ran a small benchmark on the symmetric traveling salesperson problem (TSP). We compared the classical optimization solver SCIP on both standard and Quadratic Unconstrained Binary Optimization (QUBO) formulations (SCIP-TSP, SCIP-QUBO), a warm-started QUBO-based QAOA (Penalty QAOA), and a QAOA with permutation-preserving four-qubit exchange mixers (Swap-QAOA) using IBM quantum hardware (ibm_miami with 120 qubits). The details and results are given in Figure 1.




Figure 1: Benchmarking results for solution quality (left) and time-to-solution (right) for TSP instances with {5, 6, 7, 8, 10, 20, 30, 40, 50} cities (five per instance size). The classical solvers have a time limit of 60 seconds. SCIP-QUBO and QAOA stop at eight cities. The QAOA solvers have depth $p=1$ and use COBYLA. In each iteration, 1,024 shots are evaluated, with at most 25 evaluations; the last evaluation contains 8,192 shots. For the solution quality, we provide the gap to the optimal objective function value for both the best feasible solution (if found) and the best infeasible solution (given by its QUBO value).

The comparison is unsurprising but instructive. Both classical methods outperform the quantum approaches in both quality and solution time. SCIP on the standard formulation performs best with an optimal solution for 50 cities in under one minute. The QAOA solvers need minutes to terminate; both barely find feasible solutions. Apart from warm starting and custom mixers, our implementations remain close to vanilla QAOA. Thus, performance could be improved, for example, by using deeper circuits and error mitigation. Yet this reinforces the practical lesson. Mature classical solvers are close to plug-and-play; competitive quantum performance requires carefully engineering each step of the pipeline.

Where quantum optimization may matter

Does the TSP result make quantum optimization a dead end? No; it changes where we should look. One route is modular: use quantum tree search, amplitude estimation, or quantum optimization inside decomposition, while classical software retains modeling and coordination. The comparison is then against one bottleneck, not an entire solver. The second route targets problems that are richer in structure and solution demand than the regular NP-hard combinatorial optimization problem: multiobjective, robust or stochastic, bilevel, or simulation-based optimization requiring exponentially many outputs or scenario evaluations. For such problems, classical tooling is also less mature.

Multiobjective optimization seeks a potentially exponentially large set of Pareto-optimal trade-offs. In [2], we adapted an existing multiobjective variant of QAOA and incorporated it into a classical multiobjective heuristic, where the outer loop takes over QAOA parameter training. In our experiments using the NSGA-II heuristic, we demonstrated a 10% performance gain in terms of hypervolume and increased stability. Additionally, we were able to increase coverage by up to 40% with little to no hypervolume loss. This shows a promising co-design of quantum sampling and classical search.

In robust optimization, a solution must work across a set of parameter scenarios. In [3], we used the stochastic output of quantum hardware in two approaches applied to use cases from the energy sector with uncertain energy demand: for the unit commitment problem, we solved its deterministic variant with a quantum annealer and ranked near-optimal samples by robustness, producing feasible schedules close to the robust optimum. Further, we trained a QAOA on the expected-value instance of the electric vehicle charging problem and reused the parameters on the scenario-specific circuits. On a simulator, we were able to recover the robust-optimal schedules. These proof-of-concept results suggest a credible role for quantum devices: generating structured candidate sets inside harder problems.

Outlook

The directions outlined above are promising, but none offers a free lunch for quantum advantage. Progress will require tightly integrated and optimized quantum-classical workflows, continued hardware improvements, rigorous theoretical guarantees, and convincing empirical evidence. Yet quantum optimization still rests on only a handful of algorithmic paradigms. Combining insights from quantum physics with classical opti-

mization therefore leaves substantial room for paths toward quantum advantage.

Links:

[L1] <https://arxiv.org/search/advanced> (query: quant-ph; title or abstract contains the phrase “quantum optimization” in either American or British spelling; first announced January 1 to September 1, 2026; cross-lists included)

References:

- [1] S. P. Jordan et al., “Optimization by decoded quantum interferometry,” *Nature*, vol. 646, pp. 831–836, 2025. doi:10.1038/s41586-025-09527-5
- [2] I. Turkalj et al., “Enhancing variational quantum algorithms for multicriteria optimization,” *Quantum Inf. Process.*, vol. 25, Art. no. 209, 2026. doi:10.1007/s11128-026-05232-y.
- [3] P. Halfmann et al., “Harnessing inferior solutions for superior outcomes: obtaining robust solutions from quantum algorithms,” in *Proc. Genetic Evol. Comput. Conf. Companion (GECCO ‘24)*, 2024, pp. 1950–1953. doi:10.1145/3638530.3664160.

Please contact:

Pascal Halfmann

Fraunhofer Institute for Industrial Mathematics ITWM,
Germany

pascal.halfmann@itwm.fraunhofer.de

The Bottleneck is not Cryptography: Understanding the Costs of Quantum-Safe Communications

by Ruben Rios and José A. Montenegro (University of Málaga)

Post-quantum cryptography is no longer a theoretical exercise. As quantum-safe algorithms move towards large-scale deployment, recent studies by NICS Lab in the context of the PQSecNg project show that the main challenges extend far beyond the cryptographic primitives themselves, revealing hidden bottlenecks across protocols, software stacks and network infrastructures.

The transition to post-quantum cryptography (PQC) marks one of the most significant transformations in the recent history of cybersecurity. For decades, the security of Internet communications has relied on public-key primitives such as RSA and elliptic-curve cryptography. However, sufficiently powerful quantum computers are expected to eventually break these mechanisms, compelling governments, industries, and infrastructure operators worldwide to prepare for a quantum-safe future.

Today, the discussion is no longer centred on whether this transition will occur, but on how it can be achieved efficiently and at scale, ensuring cryptoagility — the ability of systems to rapidly adapt cryptographic mechanisms as needed.

Our research has focused on understanding the practical implications of deploying PQC in real communication systems. While considerable effort has been devoted to designing and standardising quantum-resistant algorithms, much less attention has been paid to the impact of integrating them into the layered infrastructure that supports modern Internet services.

The starting point was the development of a framework for evaluating post-quantum Transport Layer Security (TLS), the protocol that provides the security foundation for most Internet communications. Existing studies often focused on isolated cryptographic operations, providing useful but incomplete insights. In contrast, real-world deployments involve complex interactions between libraries, protocols, and services, as well as variable network conditions. To capture these effects, our evaluation framework was built to measure end-to-end performance under realistic deployments.

Using this framework, we benchmarked a wide range of TLS configurations during the handshake phase [1], covering different post-quantum primitives for key establishment (KEM) and authentication (signatures). The results demonstrated that quantum-safe communications are already technically viable. However, not all configurations exhibit the same deployment costs. While some configurations introduced only modest overhead, others required considerably larger handshake mes-

sages and created additional operational challenges that make them unsuitable even for simple deployments.

While TLS remains the cornerstone of secure web communications, newer transport protocols are rapidly gaining ground. The scope of this work was therefore extended to QUIC — a transport protocol increasingly adopted by web browsers, cloud providers, and large-scale Internet platforms — which was specifically designed to reduce latency and improve connection establishment, making it an increasingly relevant deployment target for post-quantum cryptography.

A comparative analysis of post-quantum deployments in TLS and QUIC revealed that protocol behaviour becomes considerably more complex when quantum-safe algorithms are introduced [2]. The benefits associated with modern transport protocols may be undermined by larger cryptographic payloads and modified handshake procedures. These findings highlight the importance of evaluating complete protocol ecosystems rather than studying cryptographic algorithms in isolation.

Perhaps the most important insight emerging from our research is that the bottleneck is not necessarily the PQC primitive. Our latest work [3] adopted a broader perspective and investigated the deployment of PQC across different layers of the communication stack — including both KEM and signature primitives, mutually authenticated TLS, and TLS-based VPN — and across different hardware architectures — Intel and ARM. The results revealed that protocol interactions, software architectures and network conditions often contribute as much as, or even more than, the cryptographic primitives




Figure 1: Effect of replacing classical cryptographic primitives with post-quantum primitives across different network layers and microprocessor architectures. The results show that deployment overhead depends not only on the cryptographic algorithm itself, but also on protocol behaviour, software architecture, hardware platform and network conditions.

themselves to the overall deployment cost. A summary of results is shown in Figure 1.

In other words, replacing a classical primitive with a post-quantum alternative is only one part of a much larger engineering challenge. This observation is particularly relevant in the current transition toward quantum-safe communications. Understanding where bottlenecks appear and how they propagate through different layers is therefore essential for achieving efficient quantum-safe systems.

As Europe continues to invest in quantum technologies and cybersecurity resilience, practical deployment research will become increasingly important. Building secure communication infrastructures for the quantum era requires not only strong cryptographic primitives but also a deep understanding of how these primitives interact with network protocols, software, and hardware architectures.

Taken together, our studies suggest an encouraging but nuanced message. The quantum-safe Internet is already within reach. However, the path towards successful large-scale adoption will depend on rigorous benchmarking, interoperability testing and system-level optimisation efforts that extend well beyond cryptographic design.

The future quantum-safe Internet will not be defined solely by cryptography. It will be defined by our ability to efficiently integrate post-quantum cryptography into the protocols, platforms, and network infrastructure on which our society depends.

Links:

- [1] <https://www.nics.uma.es>
- [2] <https://github.com/montenegro-montes>

References:

- [1] J.A. Montenegro, R. Rios, J. Lopez-Cerezo, “A Performance Evaluation Framework for Post-Quantum TLS”, *Future Generation Computer Systems*, 2025.
- [2] J.A. Montenegro, R. Rios, J. Bonilla, “Comparative Analysis of Post-Quantum Handshake Performance in QUIC and TLS Protocols”, *Computer Networks*, 2025.
- [3] R. Rios, J.A. Montenegro, “Post-Quantum Cryptography: A Multi-layer Bottleneck Analysis”, *International Journal of Information Security*, 2026.

Please contact:

Ruben Rios and José A. Montenegro
E.T.S.I. Informática, Spain
{ruben.rdp, jmmontes}@uma.es

Austria's Quantum-Communication Infrastructure Testbed

by Hannes Hübel, Manuela Kos and Sebastian Ramacher
(AIT Austrian Institute of Technology)

The project QCI-CAT is advancing Austria's contribution to Europe's quantum-secure communication infrastructure by bringing quantum key distribution and post-quantum cryptography from research into operational networks and real-world applications. Austrian research, industry and government partners have deployed long-distance QKD networks and demonstrated quantum-secure governmental and medical use cases laying the foundations for scalable, resilient and quantum-safe communications across Europe.

As Europe moves towards a continent-wide quantum communication infrastructure, the Austrian QCI-CAT project [L1] has taken an important step from experimental quantum networks towards demonstrating real-world use cases. The project operated quantum key distribution (QKD) networks connecting sites across Austria, demonstrated quantum-secure applications with governmental and medical stakeholders, developed training activities for future users and specialists, and advanced research towards future quantum repeater networks.

A central achievement was the deployment of three complementary network environments (Figure 1): an inner-city network in Vienna, a connection between Vienna and Graz, and a long-distance proof-of-concept network extending approximately 400 km from Vienna to St. Johann im Pongau. The deployments brought together QKD equipment from multiple European suppliers, trusted nodes, classical networking infrastructure and quantum-enabled encryption systems.

The Vienna–Graz connection linked the Medical University of Graz with the Vienna network over approximately 150 km. The longer Vienna–St. Johann link connected the Austrian Ministry of Defence in Vienna with a backup data centre of the Republic of Austria. Because of the distance involved, the latter network required several intermediate links along a route passing through St. Pölten, Amstetten, Linz and Salzburg. All QKD devices procured by QCI-CAT were brought together for this demonstration, making it an important test of long-distance terrestrial QKD operation.

An important aspect of QCI-CAT was the integration of quantum and classical security technologies rather than treating QKD as an isolated technology. A key management system developed by AIT [L2] managed key material and supported hybrid QKD/post-quantum cryptography (PQC) keys. Point-to-point hardware encryptors and the software-based Arnika VPN [L3] provided QKD/PQC-hybrid protection at the data layer.

The deployed networks became platforms for demonstrating applications in operational environments [1]. In the governmental use case, representatives of Austrian ministries participated in demonstrations combining QKD with secret sharing




Figure 1: QKD-secured networks deployed in Austria as part of QCI-CAT combining hybrid QKD/PQC key management, the Arnika VPN solution, and governmental and medical applications.

and secure video conferencing. A quantum-secured video call and simultaneous data exchange demonstrated how different security technologies can be combined to protect sensitive governmental communications.

A second application addressed highly sensitive medical information. Pathogen genome sequence data are increasingly important information assets for public health, while their sensitivity and dual-use potential impose confidentiality requirements. QCI-CAT demonstrated secure transmission and decentralized storage of genomic and experimental data between the Medical University of Graz and Austrian governmental sites. The approach combined QKD-generated keys with a secret-sharing architecture in which information was divided between independent storage locations. During the demonstration, genomic data and associated files were transferred and subsequently reconstructed, with integrity checks confirming that the recovered data matched the original data. The experiment illustrated how quantum-secure communications can form part of a broader security architecture rather than simply providing a point-to-point cryptographic link.

Operating QKD networks also revealed challenges that will matter for future large-scale infrastructures. Intermittent instabilities required remote reboots and, in some cases, physical power cycling or intervention by equipment manufacturers. This experience highlights that deploying a future European quantum communication infrastructure will depend not only

on QKD performance, but also on hardware reliability, software resilience, interoperability and maintainability.

QCI-CAT therefore also investigated what scaling from experimental networks to national infrastructure would entail with respect to capital and operational expenditures [2]. Network planning and cost-estimation work examined a terrestrial Austrian QKD backbone for critical infrastructure. This work included abstract network planning and Monte Carlo simulations to estimate network length and was subsequently extended to all EU Member States.

Developing infrastructure must be accompanied by developing skills. QCI-CAT established training curricula for engineers as well as stakeholders and decision-makers, implemented online courses [L4] using a Moodle-based learning environment, and complemented these with in-person training. More than 50 participants registered for the online courses and over 160 attended in-person training activities, substantially exceeding the project's original training target. Feedback also led to a compact quantum-communication course designed specifically for employees entering the sector from related technical fields.

Alongside deployment, QCI-CAT continued research towards technologies that may underpin future generations of quantum networks. For instance, a new rack-mounted trapped-ion node with an integrated optical cavity was developed. Entanglement between trapped-ion systems was established using telecom-

wavelength photons transmitted over a 50 km optical-fibre link. The experiment represented the first trapped-ion entanglement over such fibre distances using a two-photon technique robust against path-length fluctuations [3]. QCI-CAT also contributed to the design and implementation of FAEST [L5], a post-quantum digital signature scheme and a 3rd round candidate in NIST's call for additional signature schemes. FAEST builds its security on well-established cryptographic primitives, notably AES and SHA-3, and constructs signatures as non-interactive zero-knowledge proofs of knowledge of an AES key. These contributions complement the project's QKD activities by advancing post-quantum cryptographic mechanisms that can protect authentication mechanisms against quantum adversaries.

QCI-CAT's results illustrate the transition now taking place in European quantum communications: from individual technologies and laboratory experiments towards integrated networks, operational applications and an emerging ecosystem of users, engineers and infrastructure providers. Just as importantly, the project has identified practical challenges that must be addressed for production deployments. These lessons provide a foundation for Austria's contribution to the European Quantum Communication Infrastructure (EuroQCI) and for the next generation of interconnected quantum-secure infrastructure across Europe.

Links:

- [L1] <https://qci-cat.at>
- [L2] <https://qkd-kms.ait.ac.at/>
- [L3] <https://github.com/arnika-project/arnika>
- [L4] <https://qurios.academy/>
- [L5] <https://faest.info>

References:

- [1] C. Döberl, W. Eibner, H. Hübel, M. Kos, F. Kutschera, S. Laschet, A. Neuhold, S. Ramacher, P. Stanzer, and W. Strasser, "Implementation of a hybrid QKD/PQC network and applications," in Proc. IEEE Int. Conf. Quantum Commun., Netw., Comput. (QCNC), 2026. <https://www.ieee-qcnc.org/2026/>
- [2] S. Raubitzek, W. Strasser, S. Ramacher, T. Lebeth, A. Neuhold, and C. Pacher, "Towards national quantum communication in Europe: Planning and sizing terrestrial QKD networks," arXiv preprint arXiv:2604.06764, 2026. <https://arxiv.org/abs/2604.06764>
- [3] V. Krutyanskiy, M. Canteri, M. Meraner, V. Krcmarsky, and B. Lanyon, "Multimode ion-photon entanglement over 101 kilometers," PRX Quantum, vol. 5, Art. no. 020308, 2024. <https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.5.020308>

Please contact:

Sebastian Ramacher
AIT Austrian Institute of Technology, Austria
sebastian.ramacher@ait.ac.at

How the Laws of Physics Help to Safeguard our Digital Future

by Florian Kanitschar and Christoph Pacher (AIT Austrian Institute of Technology)

From bank transactions to medical records – our digital lives depend on secrets that may not stay secure for long. Discover how cutting-edge research collaborations are turning quantum mechanics into a guardian for our most sensitive data – and how laser pulses instead of single photons can ensure that high-performance solutions remain compatible with existing telecom infrastructure.

When exchanging confidential information, we trust that the transmitted data remains secret. This trust is founded on cryptographic encryption that uses a key to render the information inaccessible during transmission. That way, only the recipient of the message who holds the appropriate key can decrypt the original message. Thus, one key (pun intended) challenge is to securely distribute the key from the sender to the recipient.

However, today's most established methods – Diffie-Hellman key exchange (DHKE) and its elliptic curve analog ECDH – rest on a risky assumption: that certain mathematical problems, like the discrete logarithm problem, are practically unsolvable. Since the 1990s, when Peter Shor discovered his famous algorithm, we have known that a sufficiently large quantum computer (QC) could solve these problems efficiently. One recent analysis [L1] predicts that RSA-2048, an encryption standard that is harder to attack with QCs than ECDH, could be broken by 2030. Even more alarming is the "harvest now, decrypt later" strategy, where attackers copy and store encrypted communications today, intending to decrypt them later with more advanced QCs. Corporate secrets, health data, and critical infrastructure information are therefore already at risk today.

Post-Quantum Cryptography (PQC) has already provided NIST-standardized algorithms that are currently not known to be efficiently breakable with QCs and that should be deployed now. A fundamentally different response to the same threat comes from quantum physics itself: Quantum Key Distribution (QKD). Instead of relying on unproven mathematical conjectures, QKD exploits the fundamental properties of quantum mechanics: unknown quantum states cannot be perfectly copied and measuring them unavoidably disturbs them. These properties enable key distribution protocols whose security is rooted in physical laws rather than mathematical assumptions.

While the most-studied QKD protocols rely on costly single-photon measurements, Continuous-Variable (CV) QKD measures the quadratures of faint laser pulses using homodyne detection – a technology well-known from optical data communication. This makes it possible to leverage existing telecom technology and infrastructure. The QKD protocol steps are illustrated in Figure 1 and are detailed in the following.

In each round, Alice randomly chooses and prepares one out of four coherent states in her lab and transmits them over an insecure channel that is under the control of an eavesdropper (Eve). Bob receives the states and performs homodyne detection. Based on Bob's measurement results, Alice and Bob perform statistical tests on a random subset of the detected states. The test results are used to infer how much information about the key Eve might have gained. The remaining detected states are mapped to key symbols. This is followed by information reconciliation, where Alice and Bob correct errors in their respective shares of the key. Finally, during privacy amplification, they apply a randomly chosen hash function to their corrected raw keys. This step removes Eve's partial knowledge about their shared key at the cost of shortening the key. If successful, they hold a shared key that is both secret and identical for Alice and Bob. An illustration of the protocol steps is shown in Figure 1.

For years, however, the security analysis and the implementation of CV protocols relied on idealized assumptions. Recent breakthroughs have progressively closed the gap between theory and practice. Notably, the security of CV-QKD protocols using discrete modulation (DM) constellations, like Quadrature Phase Shift Keying (QPSK), has advanced considerably. Within the last two years, studies have moved these insights from theory to the lab: researchers from the Technical University of Denmark (DTU) and the Austrian Institute of Technology (AIT) implemented a full CV-QKD protocol with QPSK modulation over a 20 km fiber link at telecom wavelengths [1].

The system used advanced signal processing and simple optical modules to enable secure key exchange between sender (Alice) and receiver (Bob). By integrating hardware functions into the digital signal processing (DSP) module, the optical system was significantly simplified. The digital postprocessing of the recorded data was also fully implemented, including the correction of errors that occurred during transmission or measurement. This was followed by privacy amplification, a classical routine that hashes the recorded and corrected raw key to produce a shorter but private secret key. This work demonstrated composable secure keys that can be safely used in any larger ("composed") cryptographic context.

While standard lab-based QKD systems often rely on bulky tabletop setups, miniaturization is critical for large-scale deployment. In collaboration with the National University of Singapore, AIT demonstrated that high performance and practicality can coexist. The team developed a fully integrated on-chip quantum transmitter and receiver operating at 40 Gbaud at room temperature, featuring high-performance homodyne detectors [2]. For a QPSK CV-QKD protocol, the system achieved a record key rate of 1.2 Gbit/s with composable security.

The standard setting of QKD is typically limited to exactly two communicating parties, whereas modern telecom networks often involve multiple users. A recent study [3] extends the security of DM CV-QKD protocols to point-to-multipoint networks, addressing various trust scenarios. Applied to a three-party passive-optical network with a distance of 10 km be-




Figure 1: Illustration of the QPSK CV-QKD protocol, showing quantum-state preparation by Alice, transmission over the insecure quantum channel, measurement by Bob, statistical testing, reconciliation and privacy amplification.

tween the central node and each user, this work validates the practical compatibility of CV-QKD with multi-user settings, demonstrating scalability for deployment in real-world networks.

These results underscore the rapid progress toward practical, high-performance QKD systems. By combining theoretical rigor with engineering ingenuity, researchers are paving the way for a scalable, cost-effective quantum secure communication infrastructure. As threats advance, QKD stands out as a future-proof solution, with DM CV-QKD offering robust security without sacrificing compatibility with existing infrastructure. This transition from laboratory to real-world applications marks an important step towards ensuring that the confidentiality of our digital communications remains intact in the quantum computing era.

Links:

[L1] <https://postquantum.com/q-day/q-day-y2q-rsa-broken-2030/>

References:

- [1] A.A.E. Hajomer, et al., “Experimental composable key distribution using discrete-modulated continuous variable quantum cryptography”, *Light: Science&Applications*, 14 (1), 255, 2025.
- [2] S.Q. Ng, et al., “Gigabit-rate quantum key distribution on integrated photonic chips”, *Optica*, vol. 13 (6), pp. 1043-1055, 2026.
- [3] F. Kanitschar, et al., “Composable Continuous-Variable Multi-User QKD with Discrete Modulation: Theory and Implementation”, *Optica*, 13 (8), pp. 1563-1580, 2026.

Please contact:

Florian Kanitschar
AIT – Austrian Institute of Technology, Austria
florian.kanitschar@ait.ac.at

Christoph Pacher
AIT – Austrian Institute of Technology, Austria
christoph.pacher@ait.ac.at

Planning and Sizing National Quantum Communication Networks

by Sebastian Raubitze, Rebecca Klingbeil (SBA Research) and Werner Strasser (fragmentiX Storage Solutions)

The European Quantum Communication Infrastructure (EuroQCI) [L1] sets a strategic goal for secure communication across the EU, yet the size and shape of the national networks needed to achieve it remain unspecified. We present a reproducible method to estimate the number of nodes, total fibre length, and equipment required for national terrestrial quantum key distribution (QKD) networks, using Austria as a reference and scaling the results to other Member States.

EuroQCI targets secure communication for security-relevant sites across all Member States. Its technical core is QKD, which enables key generation and secure distribution. The security of this distribution relies on physical principles, as measuring a qubit in the wrong basis disturbs its state and yields a random result. Still, a basic planning question remains open: how large must a national QKD network actually be? It is not yet clear how many nodes each country needs, how densely they must be connected, or what total fibre length is required.

This work addresses that gap with a computational planning-level method based on a small set of stated assumptions. Endpoints represent sites where key material must be available, such as ministries and operators of critical infrastructure. They are connected by direct QKD links wherever the distance allows, and most sites host two or three links so that they do not depend on a single connection for their key supply. All distances are computed as geodesic distances and multiplied by a constant detour factor of 1.5 to approximate real fibre routing along roads and existing corridors.

A central constraint is QKD’s limited transmission distance over optical fibre. In the BB84 protocol, with photon sources at the endpoints, practical links operate over roughly 40–100 km per span, with a planning target below 75 km for stable operation. Entanglement-based QKD, with a central photon source, can reach up to twice that much [1]. To remain protocol-agnostic and to guarantee a workable key rate, we assume spans of around 50 km and use trusted repeater nodes to bridge longer distances, where a photon source for an entanglement-based QKD protocol could be placed instead.

Monte Carlo simulation is used to sample through possible network shapes. Each run constructs a network that gives every site its prescribed number of connections while preferring short links, and then places trusted repeater nodes on the longest spans first. Running many realisations produces stable estimates for the total fibre length, the distribution of span lengths, and the number of required components. The full implementation is openly available [L2].

Austria serves as the reference case, with 250 endpoints and 50 trusted repeater nodes; the choice is also convenient in




Figure 1: Representative simulation of a national terrestrial QKD network for Austria, with 250 QKD endpoints (blue) and 50 trusted repeater nodes (red). Endpoints cluster around major cities, and trusted repeater nodes are placed along the longest links so that individual QKD spans stay within a feasible range. The same method is applied to every EU Member State, with endpoint counts scaled by population and trusted repeater node counts by land area.

practice, since one of the first metropolitan QKD networks was demonstrated in Vienna [2]. Across 1000 simulation runs, the route-adjusted total fibre length is about 8,600 km and varies only slightly between runs, which shows that network size is driven mainly by national geography and the fixed endpoint count rather than by random placement. After trusted-node insertion, the mean span length is about 23 km and the maximum about 52 km after route correction, both within a practical range for QKD. The reference case also fixes the equipment counts: 750 QKD devices in total, together with one key management system and one hardware security module per participating site [3].

Austria is then used as a baseline for the other EU Member States, following two rules. The number of endpoints grows mainly with population, because a larger population implies more ministries, agencies, and operators of essential services that need key material. The number of trusted repeater nodes grows mainly with land area, because more territory means longer inter-regional links that must be split into feasible spans. Under these rules, a country with a large population but a small area, such as the Netherlands, needs many endpoints but few trusted nodes, while a country with moderate population but large territory, such as Finland, needs fewer endpoints but many trusted nodes. Large countries such as Germany and France require increases in both (Figure 1). For states with islands or remote territories, the simulation is restricted to the mainland, so that the terrestrial comparison is not distorted by unrealistic maritime links.

The results are first-order planning references rather than deployment designs. They indicate that trusted repeater nodes,

key management systems, and hardware-protected key handling are central components of a national QKD architecture rather than secondary details. The next steps are to replace the synthetic endpoints with real inventories of protected sites, to use measured fibre corridors instead of a uniform detour factor, and to add cross-border links, so that national estimates can be combined into the pan-European network that EuroQCI ultimately requires. This work was supported by the Digital Europe Programme under project number 101091642 (QCI-CAT).

Links:

- [L1] <https://tinyurl.com/i-euroqci>
- [L2] https://github.com/Raubkatz/QCI_Simulation
- [L3] <https://arxiv.org/abs/2604.06764>

References:

- [1] S. P. Neumann et al., “Continuous entanglement distribution over a transnational 248 km fiber link”, *Nat. Commun.* 13, 6134, 2022.
- [2] M. Peev et al., “The SECOQC quantum key distribution network in Vienna”, *New J. Phys.* 11(7), 075001, 2009.
- [3] P. James et al., “Key management systems for large-scale quantum key distribution networks”, *Proc. ARES*, 2023.

Please contact:

Sebastian Raubitsek
SBA Research, Vienna, Austria
SRaubitzek2@sba-research.org

From QKD to the Quantum Internet: Building a Full-Stack Quantum Network Infrastructure

by Claudio Cicconetti (CNR-IIT) and Alessandro Zavatta (CNR-INO)

Quantum Key Distribution is already moving out of the laboratory, but today's QKD networks are only a first step towards a full Quantum Internet. CNR is building an inter-regional quantum network across Pisa, Florence and Rome to explore how commercial QKD can be integrated into telecom infrastructure today — and how the same testbed can evolve towards quantum repeaters and networked quantum computing tomorrow.

Quantum communication technologies make it possible to transmit quantum states, or qubits in quantum information terminology, over a distance, using photons as carriers and encoding information in degrees of freedom such as space, time,

or polarization. Quantum communications have matured rapidly in recent years, fostering experimentation worldwide and enabling the commercialization of products for specific applications. The opportunity is huge: unlike their classical counterparts, arbitrary unknown quantum states cannot be perfectly copied, while distant quantum systems can be entangled, opening the door to new applications connecting quantum computers and sensors across a future Quantum Internet.

Unfortunately, what makes qubits powerful also makes them fragile: quantum states cannot simply be amplified, buffered, or retransmitted as classical information can, invalidating many principles underlying today’s network protocols and devices. The research community has long debated which architectures and technologies will best support quantum networking, for instance within the Quantum Internet Research Group of the IRTF [L1], but much of this work has so far remained theoretical.

At the Consiglio Nazionale delle Ricerche (CNR), the largest public research body in Italy, we are addressing future quantum networks from a different perspective. Through a collaboration between the Institute of Informatics and Telematics (IIT) and the National Institute of Optics (INO), we are deploying an inter-regional quantum network over standard telecom fibre




Figure 1: Quantum Key Distribution (QKD) allows two parties—traditionally called Alice and Bob, as illustrated in the bottom-left box—to establish a shared secret key by preparing and measuring qubits according to a defined protocol. CNR is developing an inter-regional QKD network connecting Pisa, Florence, and Rome over telecom fibre links. The infrastructure is designed to support experimentation on the integration of QKD with existing communication and computing systems, while providing a path towards more advanced technologies, including quantum repeaters and, ultimately, the Quantum Internet.

links as a long-term infrastructure for research and experimentation. The testbed extends from Pisa (CNR-IIT) through Florence (CNR-INO) to Rome, where the CNR headquarters are located, with intermediate sites to overcome the distance limitations of current quantum communication technologies.

Compared with existing quantum communication infrastructures, our testbed will combine three distinctive features: inter-regional scale, the use of off-the-shelf commercial technologies operating under realistic operational conditions, and support for full-stack Quantum Internet experimentation beyond the physical layer.

Following a bottom-up, technology-driven approach, we start from the most mature quantum communication technology currently available: Quantum Key Distribution (QKD). QKD comprises protocols that allow two parties, traditionally called Alice and Bob, to establish a shared secret key by exchanging quantum states and performing authenticated classical post-processing [1]. In prepare-and-measure protocols, Alice prepares qubits in different bases, while Bob measures them using randomly selected bases. Figure 1 illustrates the basic QKD process and how it fits into the planned CNR quantum network infrastructure.

The CNR infrastructure initially consists of several QKD links interconnecting the three main sites and other stations along the route. End-to-end secret keys will be established through intermediate trusted nodes. Each trusted node establishes keys with its neighbours and participates in relaying key material between the end points. It is “trusted” because compromising such a node may also compromise the end-to-end key.

From this initial stage, we expect the infrastructure to evolve along two complementary directions: technology integration in the short term and more fundamental research in the longer term.

From a technology perspective, significant effort has been devoted to QKD devices, which are now commercially available from multiple vendors, but their integration with existing computing, storage, and communication technologies remains a practical challenge. We therefore plan to make the infrastructure available within SLICES [L2], the ESFRI research infrastructure devoted to medium- and long-term evolutions of the Internet. It will also become a core component of the national AI-PHOQUS project, integrating three key digital technologies: quantum, Internet, and AI.

This will allow experimenters from across Europe to investigate how commercial QKD devices can be integrated into telecom infrastructures through proof-of-concept scenarios involving, among others, 5G/6G mobile networks, the edge-to-cloud continuum, and IoT and vehicular systems. The infrastructure will also serve as a sandbox in which application developers can test emerging QKD-enabled applications under realistic conditions, moving beyond point-to-point key provisioning towards services such as secure storage, financial applications, authentication, and access control for sensitive data.

From a research perspective, we plan to upgrade the infrastructure as more advanced technologies become available. A first major step will be the possible replacement of trusted nodes with quantum repeaters: intermediate devices that allow entanglement to be distributed between distant end nodes. Unlike trusted-node QKD networks, a repeater-based architecture does not require intermediate stations to access the end users’ secret key material. Combined with appropriate quantum communication and verification protocols, this will enable secure communication across networks whose intermediate stations need not be trusted. This evolution will benefit from the QuRE (Quantum Repeaters) project, funded by the Italian Ministry of University and Research.

Quantum repeaters will also enable quantum processors and memories at different sites to be interconnected over long distances. The infrastructure could thus evolve into a true long-distance quantum network, allowing experiments to move beyond individual laboratories (e.g., [2]) and fostering new software and applications for the Quantum Internet. This will complement activities within the Quantum Internet Alliance [L3], where CNR is already involved in benchmarking quantum applications and systems.

A further strategic evolution will be the creation of a free-space quantum communication link across the Florence metropolitan area. Initially, it will provide a metropolitan testbed for quantum communication over free-space channels. In the longer term, it will be upgraded to operate as a ground station for satellite quantum communications, bridging terrestrial quantum networks and future space-based systems and contributing to integrated global quantum communication and sensor networks.

Links:

[L1]: <https://www.irtf.org/qirg.html>

[L2]: <https://www.slices-ri.eu/>

[L3]: <https://quantuminternetalliance.org/>

References:

- [1] H.-K. Lo, M. Curty, and K. Tamaki, “Secure quantum key distribution,” *Nature Photonics*, vol. 8, pp. 595–604, 2014, doi: 10.1038/nphoton.2014.149.
- [2] C. Delle Donne et al., “An operating system for executing applications on quantum network nodes,” *Nature*, vol. 639, no. 8054, pp. 321–328, Mar. 2025, doi: 10.1038/s41586-025-08704-w.

Please contact:

Claudio Cicconetti
CNR-IIT, Italy
c.cicconetti@iit.cnr.it

Forecasting Chaotic Dynamics with Hybrid Quantum Machine Learning

by Luca Salatino, Francesco D'Amore and Andrea Giordano (CNR-ICAR)

Quantum computing is often associated with future large-scale machines capable of solving problems beyond the reach of classical computers. However, useful applications can already emerge from today's small quantum processors. Researchers from the Italian National Research Council (CNR-ICAR) and the University of Calabria have developed a hybrid quantum-classical machine learning framework that significantly improves predictions of how chaotic systems evolve. This opens new opportunities for scientific computing and quantum-enhanced artificial intelligence.

Artificial intelligence has become an indispensable tool for analysing and predicting the behaviour of complex dynamical systems. Applications such as weather forecasting, fluid dynamics, energy management and digital twins all rely on the ability to anticipate the future evolution of systems whose behaviour is often highly nonlinear. Unfortunately, many of these systems are chaotic: small uncertainties in their current state rapidly amplify over time, making accurate long-term prediction one of the most challenging problems in computational science.

Quantum computing has recently emerged as a promising technology for scientific machine learning. Rather than replacing classical computers, current quantum devices are expected to complement them by exploiting physical processes that are

difficult to reproduce efficiently using conventional hardware. Among the most promising approaches is Quantum Reservoir Computing (QRC), a machine learning paradigm in which the natural dynamics of a quantum system are used to transform input data into a representation suitable for prediction tasks, while only a simple output layer is trained.

Although QRC has attracted considerable interest, its practical implementation faces important challenges. Existing approaches often require the quantum system not only to process incoming information but also to preserve memory of previous inputs. This dual role places considerable demands on present-day quantum hardware, where noise, decoherence and repeated measurements inevitably limit the amount of information that can be retained during the computation.

To address this limitation, researchers from the Italian National Research Council (CNR-ICAR) [L1] and the Physics Department of the University of Calabria [L2] developed a new Hybrid Quantum Reservoir Computing (HQRC) architecture that combines the complementary strengths of quantum and classical computing. Instead of assigning every computational task to the quantum processor, the proposed framework separates nonlinear information processing from memory. The quantum system generates a rich set of features from the input data through its natural dynamics, while a lightweight classical recurrent layer retains the temporal information required for forecasting.

This division of responsibilities considerably simplifies the role of the quantum processor while preserving its computational advantages. The first version of the HQRC framework [1] demonstrated that this hybrid strategy significantly improves the capability of quantum reservoirs to process temporal information. By transferring memory to the classical component, the architecture becomes more robust and better suited to the capabilities of current noisy intermediate-scale quantum (NISQ) devices, which contain only a limited number of qubits but are already available in several experimental platforms.

Building upon these results, the researchers extended the framework to address a far more demanding class of problems: the prediction of multidimensional chaotic systems. Unlike standard benchmark tasks involving a single time series, many real physical systems evolve through the interaction of several coupled variables. The new architecture was therefore designed to encode multidimensional inputs [2] directly into the quantum reservoir, allowing the quantum dynamics to capture the correlations among all components of the system simultaneously.

A second important innovation is to sample the quantum system at multiple points during its evolution, an approach known as temporal multiplexing. Each observation provides complementary information about the ongoing quantum dynamics, enriching the representation generated by the same quantum processor without requiring additional qubits. This strategy substantially increases the amount of information available to the learning algorithm while remaining fully compatible with existing quantum hardware.




Figure 1: Evolution of the Lorenz-63 system (solid line), compared with the signals predicted by the QRC algorithm (dashed line). The best prediction extends up to approximately 500 time steps, corresponding to about 13 Lyapunov times (t/LT), as indicated by the vertical dashed green line.

The extended HQRC architecture was evaluated on two well-known models of chaotic dynamics. The first is the Lorenz-63 system (Figure 1), a classical benchmark that introduced the concept of deterministic chaos and remains one of the most widely used models for testing forecasting algorithms. The second is a reduced-order model derived from the two-dimensional Navier-Stokes equations through a Galerkin projection. Although this simplified model consists of only five interacting modes, it preserves the essential nonlinear mechanisms responsible for the transition to turbulence and therefore represents a realistic and challenging forecasting problem. Numerical experiments show that the proposed framework successfully reconstructs the evolution of both systems over significantly longer time intervals than previous quantum reservoir implementations [3]. In particular, temporal multiplexing plays a crucial role by enabling the quantum reservoir to extract a richer description of the underlying dynamics. Perhaps the most significant outcome of this research is that these results are obtained using quantum reservoirs composed of only a handful of qubits. Rather than waiting for large fault-tolerant quantum computers, the proposed approach demonstrates that useful scientific machine learning applications can already be developed on today's NISQ technologies by combining quantum and classical resources in an efficient way.

The work illustrates a broader trend in quantum computing, where hybrid algorithms are increasingly replacing the idea of purely quantum solutions. By allowing each component to perform the task for which it is best suited, hybrid architectures offer a practical route to exploiting quantum technologies in realistic scientific applications. Future developments will focus on implementing the proposed framework on real quantum processors and extending it to increasingly complex dynamical systems. Beyond chaotic forecasting, the same methodology may contribute to new quantum-enhanced tools for modelling, monitoring and controlling complex physical processes in science and engineering.

Links:

[L1] <https://www.icar.cnr.it>
[L2] <https://fisica.unical.it>

References:

- [1] J. Settino et al., "Memory-augmented hybrid quantum reservoir computing," *Phys. Rev. Appl.*, vol. 24, no. 2, p. 024019, Aug. 2025.
- [2] L. Salatino et al., "Forecasting Low-Dimensional Turbulence via Multi-Dimensional Hybrid Quantum Reservoir Computing," *Chaos*, vol. 36, no. 8, p. 083110, Aug. 2026.
- [3] O. Ahmed, F. Tennie, and L. Magri, "Prediction of chaotic dynamics and extreme events: A recurrence-free quantum reservoir computing approach," *Phys. Rev. Research*, vol. 6, no. 4, p. 043082, Nov. 2024.

Please contact:

Luca Salatino
CNR-ICAR, Italy
luca.salatino@icar.cnr.it

Towards Quantum-Inspired Reactive Robots

by Giovanni Pilato (CNR-ICAR), Antonio Chella (Università degli Studi di Palermo) and Salvatore Gaglio (Università degli Studi di Palermo and CNR-ICAR)

Researchers from the University of Palermo and CNR-ICAR are investigating how quantum computing concepts can support action selection in reactive robots. Their work proposes a hybrid classical-quantum architecture in which sensor data are mapped into a Hilbert-space representation and used to select robot actions under uncertainty.

Reactive robots are designed to act fast. They sense the world, evaluate the current situation, and select an action without the long deliberation typical of planning-based systems. This makes them suitable for dynamic environments, where it is important to have quick responses. However, such agents may face uncertainty when sensor readings are noisy, when perceptual situations are ambiguous, or when the robot's current sensor readings are near a decision threshold, so that small noise or small variations could change the selected action.

A research team from the University of Palermo and the Institute of High-Performance Computing and Networking of the Italian National Research Council (CNR-ICAR) is exploring whether concepts from quantum computing can provide a useful representational framework for this type of robotic decision-making. The work has led to a hybrid classical-quantum architecture for reactive action selection, developed by Antonio Chella, Salvatore Gaglio, Giovanni Pilato and Filippo Vella [1].

The central idea is to take the robot's sensor readings and encode them as a classical vector, which is then mapped into a quantum feature space. In that space, typical situations or rule-defined regions are used as references associated with possible actions. By comparing the current sensor state with these references, the system can decide which action is most appropriate. This provides a bridge between classical robot sensing and Hilbert-space representations, where similarity, uncertainty, and ambiguous cases can be modelled in a natural way (Figure 1).

The work is motivated by the following question: how can a reactive robot select one action when the perceived situation is not completely clear? Classical condition-action rules are usually deterministic: if a condition is satisfied, the corresponding action is executed. In realistic settings, however, sensor values may be close to several decision regions at once. The proposed architecture uses quantum-inspired representations to model this uncertainty before projecting the result back to a definite classical command [1].

The research is framed in relation to teleo-reactive robot control. In the classical teleo-reactive formalism, behaviour is specified through prioritised condition-action rules that are continuously re-evaluated. The quantum component proposed in this work focuses on the perceptual evaluation and action-selection layer, where conditions can be represented as regions or reference states in a Hilbert space.




Figure 1: Simplified hybrid classical-quantum workflow for reactive robot action selection. Classical sensor readings are mapped into a quantum representation that supports action selection, while the resulting command executed by the robot remains classical.

Two concrete implementations were studied. The first uses a quantum kernel based on ZZ feature maps, combined with a classical Support Vector Classifier. In this setting, the quantum feature map transforms the sensor values into a quantum feature space, while the final classification is performed by a classical machine-learning model. The second implementation is a dedicated quantum circuit for rule-based action selection. In this circuit, front and left sensor readings are encoded into condition qubits. Their activations are propagated to action qubits, an oracle-diffuser block inspired by Grover’s search increases the probability of valid single-action configurations, and a final measurement selects a classical action.

The architecture was evaluated using a public wall-following robot dataset containing ultrasonic sensor readings collected from a SCITOS G5 robot performing a wall-following task in an indoor environment. In the simplified setting considered by the team, the robot’s behaviour is determined from two sensor readings: the front distance and the left distance. These values define rule-based regions corresponding to actions such as moving forward, turning sharply right, turning slightly right, or turning slightly left.

The experimental evaluation was deliberately framed as an offline proxy for reactive decision-making. The aim was to test whether a hybrid classical-quantum representation can support the classification of perceptual states and illustrate how measurement-based action selection could be embedded in a future robotic control architecture.

The results show that the proposed framework can support both data-driven classification through quantum kernels and explicit measurement-based action selection through a quantum circuit. A classical Support Vector Classifier using a radial basis function (RBF) kernel, a standard method for measuring similarity between data points, was included as a baseline. The quantum-kernel results were then compared using single and replicated ZZ feature-maps, which encode classical data into quantum states. The rule-based quantum circuit was also evaluated through a confusion matrix, showing that most actions were selected correctly in the simplified two-sensor setting.

The work is exploratory, but it identifies a clear research direction: using quantum representations inside embodied AI systems. Future work will move beyond offline tests and evaluate the approach in closed-loop robotic platforms or simulators,

where perception and action continuously influence each other. The model will also be extended to richer perceptual and proprioceptive inputs, and to learnable quantum circuits. The goal is to retain the flexibility of data-driven models while keeping an explicit mechanism for measurement-based action selection.

In the longer term, quantum-inspired robotic architectures may offer new ways to represent uncertainty, contextual similarity and action ambiguity in autonomous systems. While practical quantum robotics is still at an early stage, this work is a step toward understanding how classical robotic control and quantum computational models can interact within a coherent architecture for reactive behaviour.

References:

- [1] A. Chella, et al., “An Architecture for a Quantum Teleo-Reactive Robot”, *Entropy*, 2026.
- [2] A. Chella, et al., “A Quantum Planner for Robot Motion”, *Mathematics*, vol. 10, no. 14, Art. no. 2475, 2022.
- [3] A. Chella, et al., “Quantum planning for swarm robotics”, *Robotics and Autonomous Systems*, vol. 161, 104362, 2023.

Please contact:

Giovanni Pilato
CNR-ICAR, Italy
giovanni.pilato@icar.cnr.it

Quantum Machine Learning for Malware Analysis

by Giovanni Ciaramella (IMT Lucca and CNR-IIT), Fabio Martinelli (CNR-ICAR) and Francesco Mercaldo (University of Molise)

Quantum computing has emerged as a critical domain in recent cybersecurity research. This article explores the use of quantum machine learning for malware analysis in Windows and Android environments. We also consider explainability in the hybrid quantum approach.

Quantum computing is an emerging computational paradigm that exploits the principles of quantum mechanics to process and analyze information. Unlike classical computers, which use bits represented as 0 or 1, quantum computers use qubits, which can exist in a superposition of multiple states. Quantum phenomena such as superposition, entanglement, and interference enable quantum systems to perform certain types of computations in fundamentally different ways from classical machines. Currently, this computational paradigm is intersecting directly with cybersecurity, a domain under constant pressure from increasingly frequent and sophisticated cyberattacks. Machine Learning has become a mainstay for automated threat identification. According to the PTR 25-27 project P2.01 (Cybersecurity) objectives, in this paper we discuss several approaches using quantum machine learning for malware analysis across mobile and desktop environments (Figure 1).

For instance, in [1] we propose a hybrid quantum-classical method tailored for Android malware identification, by integrating a single quantum layer with standard classical convolutional neural network layers. To assess performance, we benchmarked this hybrid network against established classical architectures, including LeNet, AlexNet, a standard Convolutional Neural Network (CNN), and VGG16. Using a balanced dataset comprising 8,446 Android applications, divided between verified malicious files and legitimate software, the hybrid setup

achieved a competitive classification accuracy of 0.91, compared to 0.95 for the classical VGG16 baseline.

Despite these encouraging recognition rates, the experimental results highlighted significant practical trade-offs. Because contemporary classical simulators and early-stage quantum hardware impose severe input layer size limitations, our hybrid model was restricted to processing flattened inputs at a 25×1 resolution. In contrast, the fully classical VGG16 network operated on higher-resolution 100×3 spatial representations. Computational resource demands were similarly asymmetric: training the hybrid network required over seven hours on our setup, whereas the top-performing classical VGG16 model converged in roughly six minutes.

To explore how far quantum representations can be pushed without relying on heavy classical feature extraction, in [2] we propose a quantum circuit composed of two-qubit unitaries that can encode the label of any n-bit Boolean function. The network employs two main stages: a Parameterized Quantum Circuit (PQC) layer on a single qubit, which handles parameter initialization and quantum training, followed by a dense classical layer that performs final classification, analogous to conventional deep networks.

When evaluated under identical conditions, this fully quantum-simulated model achieved a lower accuracy of 0.413. The performance decline directly illustrates the spatial and hardware bottlenecks inherent to current quantum implementations. Even after compressing image data down to a minimal 4×1 input resolution, generating the corresponding circuit required allocating 17 qubits. Scaling such models to larger input resolutions remains heavily constrained by the classical hardware overhead needed to emulate complex multi-qubit states.

Because operational security deployment requires transparent decision-making rather than opaque black-box predictions, we integrated Explainable Artificial Intelligence (XAI) techniques directly into our research framework. In our initial hybrid evaluations, we incorporated the Gradient-weighted Class Activation Mapping (Grad-CAM) algorithm. Grad-CAM generates visual heatmaps highlighting the spatial regions in an




Figure 1: Conceptual workflow for multi-platform malware analysis using hybrid classical-quantum models and Explainable AI.

input representation that most strongly influence the network's classification outputs.

In reference [3], we deepened the explainability analysis by directly comparing the proposed hybrid approach against LeNet on an expanded dataset of 14,226 malware samples spanning 25 distinct malware families, alongside 1,832 trusted applications. To ensure a controlled evaluation, both networks were trained using identical hyperparameters and a standardized 32×1 input size. Under these uniform conditions, the hybrid network achieved an accuracy of 0.887, compared with 0.830 for LeNet. Furthermore, we expanded our interpretability toolkit by testing two advanced diagnostic algorithms: Grad-CAM++, which utilizes positive partial derivatives from the final convolutional layer's feature maps to deliver refined multi-instance visual explanations, and Score-CAM. This alternative visualization technique operates independently of gradient calculations.

Building upon these findings, in future work we will consider designing multi-layer fully quantum networks to reduce reliance on classical convolutional layers, with the aim of improving malware-detection accuracy and the explainability of quantum machine learning.

References:

- [1] G. Ciaramella et al., "Introducing quantum computing in mobile malware detection," in Proc. 17th Int. Conf. Availability, Rel. Secur. (ARES), Aug. 2022, pp. 1–8.
- [2] F. Mercaldo et al., "Towards explainable quantum machine learning for mobile malware detection and classification," Appl. Sci., vol. 12, no. 23, Art. no. 12025, Nov. 2022.
- [3] G. Ciaramella et al., "Exploring quantum machine learning for explainable malware detection," in Proc. IEEE Int. Joint Conf. Neural Netw. (IJCNN), Jun. 2023, pp. 1–6.

Please contact:

Giovanni Ciaramella
IMT Lucca, Italy and CNR-IIT, Italy
giovanni.ciaramella@imtlucca.it,
giovanni.ciaramella@iit.cnr.it

Francesco Mercaldo
University of Molise, Italy
francesco.mercaldo@unimol.it

Quantum Optimization in Systems Control

by George T. Stamatiou and Kostas Magoutis (FORTH-ICS and University of Crete)

Adaptive computer systems must decide their next action in a repeated fashion and against a time deadline. Optimization lies at the heart of that decision-making process and limits how fast they can react. Researchers at FORTH-ICS and the University of Crete used a quantum annealer inside that loop across different systems, from a queueing model to a running web server to an inverted pendulum. The controllers were shown to be operational; what still stands in the way, however, is not the quantum computation itself but everything around it.

Systems that adapt themselves, such as a data center reallocating resources, a web server resizing its worker pool, a robotic system correcting its balance, share a common theme: they must choose an action now, on the basis of where the system is heading, and each must choose before the next control period elapses. Model Predictive Control (MPC) is the established approach of making such choices in an optimal way. It is based on a mathematical model of the system, predicts its behavior over a future horizon, and selects the sequence of inputs that best balances close tracking of a reference value against the cost of acting. Only the first input is applied and at the next step the calculation is repeated based on the new system state. It is the prediction step that makes MPC proactive rather than merely reactive, while the optimization step lets it respect hard operational constraints; re-solving at every step keeps the controller responsive to model error and unforeseen disturbances.

The cost of the optimization problem that sits inside the MPC control loop grows with the length of the prediction horizon and with the number of discrete choices available at each step. Longer horizons are expected to lead to more stable control and also make the search space combinatorial. Engineers usually respond by simplifying the model, shortening the horizon, or solving the problem offline in advance, each of which buys speed at the expense of control quality. For distributed systems aiming at control periods of tens of milliseconds, that is an expensive trade-off.

Quantum annealing suggests a different route. A quantum annealer is a special-purpose device that searches for the lowest-energy configuration of a system of coupled qubits, exploiting quantum effects, such as quantum tunneling, to explore a rugged energy landscape. If the control cost function is reformulated so that its minimum corresponds to that lowest-energy configuration, then the annealer becomes a solver for the optimization problem in the control loop. However, the reformulation is not trivial: the control inputs must be encoded in binary form and the objective must be expressed as a quadratic function over those binary variables, which is the form accepted by quantum machines. This new form is a Quadratic Unconstrained Binary Optimization (QUBO) problem. Once the translation is completed, however, the same recipe may be applied to any system whose dynamics can be modeled.




Figure 1: Model predictive control with a quantum optimizer. The controller predicts the behaviour of the computing system over a future horizon, compares the prediction against a reference, and hands the resulting cost function (reformulated as a QUBO problem) to the optimizer. The first input of the returned optimal sequence is applied to the system, and the loop repeats at the next control step.

Figure 1 shows the resulting workflow, with a quantum computer (or quantum annealer) as the optimizer of an otherwise conventional control loop.

We have tested this workflow in three case studies of deliberately different character. The first was a simulated tandem queueing system, an idealized model of a computer system, in which the controller regulated end-to-end response time by tuning a fixed buffer size [1]. The second was a running Apache HTTP server under a realistic synthetic workload, where the controller held processor utilization at a reference by adjusting the number of concurrent workers, while accounting for the substantial cost of the server's own reconfiguration [2]. The third was the inverted pendulum on a cart, a classic benchmark in control engineering, chosen because it is fast, physically demanding and small enough to fit today's hardware [3]. In each case the optimization was solved on D-Wave quantum annealers and compared against classical alternatives: exact solvers, simulated annealing, and hybrid quantum-classical solvers.

The picture that emerges is encouraging and sobering, in roughly equal measure. Quantum annealing-based controllers do achieve control. They stabilize the queueing system, they track utilization references on a server in real-time, and they manage to keep the pendulum upright. Solution quality is competitive with the classical baselines, though with noticeably greater run-to-run variation. The advantage also grows with difficulty: as the prediction horizon increases and the optimization becomes genuinely hard, the annealer's solution times compare increasingly well against simulated annealing.

The honest accounting, however, is not about the quantum annealing itself. Decomposing the wall-clock time of a control step shows that the quantum computation is a small fraction of the budget. What dominates is the round trip to a cloud-hosted machine and the waiting in its queue, together with the preprocessing needed to lay the problem out on the physical qubit lattice. That layout step, which maps a densely connected prob-

lem onto a physical chip of limited connectivity, costs tens of milliseconds per step, and the network contribution varies with the geographical position of the quantum devices, which are repeatedly accessed remotely. None of these costs is fundamental to quantum annealing. They are properties of how state-of-the-art quantum machines are currently accessed.

That is where the interesting engineering lies. Precomputing the problem layout, placing quantum resources close to the systems they control, and reserving quantum capacity rather than queueing for it would each remove a large factor from the latency budget, while the hardware itself continues to improve in qubit count, connectivity and error-correction. Work in the QUADS project [L1], funded by the Hellenic Foundation for Research and Innovation (HFRI) and carried out at the

Information Systems Laboratory of FORTH-ICS [L2], aims to follow and extend this line of work towards adaptive systems at different scales, leveraging different quantum hardware.

Taken together, these case studies show that quantum annealing can be integrated into model predictive control across very different systems. While there is more work to be done to broaden the applicability of quantum optimization, an equally important challenge lies with the surrounding latency caused by remote access, queueing and problem embedding. Reducing these overheads will be essential if quantum optimisers are to become practically competitive in real-time control.

Links:

[L1] <https://quads-project.eu/>

[L2] <https://www.ics.forth.gr/isl>

References:

- [1] G. T. Stamatiou, K. Magoutis, "Quantum-Enhanced Control of a Tandem Queue System", *Performance Evaluation Methodologies and Tools (VALUETOOLS 2023)*, LNCS 539, pp. 99–114, 2024.
- [2] E. Papageorgiou, G. T. Stamatiou, K. Magoutis, "Model Predictive Control of Internet Servers using Quantum Annealing", *IEEE Int. Conf. on Autonomic Computing and Self-Organizing Systems (ACSOS)*, 2024.
- [3] G. T. Stamatiou, K. Magoutis, "Quantum-Optimized Control of the Inverted-Pendulum System", *IEEE Int. Conf. on Quantum Software (QSW)*, 2026.

Please contact:

George T. Stamatiou and Kostas Magoutis
FORTH-ICS and University of Crete, Greece
stamatiou@ics.forth.gr
magoutis@csd.uoc.gr

Designing how Quantum Neural Networks Represent Data

by Smita Chakraborty (RISE)

Quantum computers could sharpen machine learning, but only if we get one step right: how everyday data such as images, audio and tables is “embedded” into a quantum system. Understanding feature representation means building a predictive theory of embedding design, including a quantum version of dropout, a classical technique used to reduce overfitting, to determine which choices actually improve learning for a given dataset and computing budget.

Machine learning has transformed how we analyse data, and researchers now ask whether quantum computers could push it further. The idea at the heart of quantum machine learning [L1, L2] is simple: take ordinary data and “embed” it into the state of a quantum system, where a vast, high-dimensional space of possibilities opens up for a model to work in. This embedding step, which turns a photograph or an audio clip into a pattern of quantum states through a so-called feature map, is where the hoped-for advantage lives, and where most of the hard questions now sit.


A powerful unifying insight is that many quantum learning models are, mathematically, kernel methods: the model compares how similar two embedded data points are, and makes decisions from those comparisons. Early work [1] showed that data could be classified inside these quantum feature spaces, and later results proved that, for carefully designed problems, quantum embeddings can offer advantages no classical computer could match. That promise, however, has become far more nuanced. A widely cited “power of data” study [2] showed that classical models, once given enough example data, can quietly match quantum ones on many realistic tasks, and introduced refined projected quantum kernels to regain a genuine edge and generalise better.

Two obstacles now dominate the field. The first is trainability. Quantum models are tuned by adjusting parameters to reduce error, but in large and highly entangled circuits the signal guiding those adjustments can vanish exponentially because of barren plateaux, where the training landscape becomes almost perfectly flat and learning stalls. The second is inductive bias: an embedding quietly encodes assumptions about what patterns matter, and if those assumptions are misaligned with the real data, the model generalises poorly no matter how long it is trained. The community’s most promising answers involve structured embeddings, designs that respect inherent data symmetries and avoid needless complexity, so that models stay trainable and well-behaved on today’s noisy hardware.

This is where a quantum version of a familiar classical trick becomes interesting and is currently being investigated. In ordinary neural networks, dropout randomly switches off parts of the network during training, which prevents over-reliance on any single feature and improves performance on unseen data.

Recent work has begun to translate this idea to quantum circuits, randomly dropping quantum operations, or temporarily setting aside and measuring individual qubits, to reduce overfitting and, in some cases, to soften barren plateaux [3]. This is a genuinely appealing direction for real datasets. For images, dropout-like schemes could stop a model fixating on irrelevant pixels. For audio, they could encourage robustness to background noise and pitch shifts, while for tabular data, they could curb over-reliance on a handful of correlated columns.

But the evidence so far is fragmentary, a mix of small demonstrations and specific circuit designs, and no clear method yet tells us which form of quantum dropout suits which kind of data, or how strongly to apply it. Nor is it clear how quantum dropout should interact with the embedding itself: the same drop that stabilises training on one feature map may erase exactly the structure that made another useful. This gap deserves far more thorough and systematic investigation across realistic, multi-dimensional datasets.



Inner products of embedded states define the quantum kernel $\kappa(x, x') = |\langle \phi(x) | \phi(x') \rangle|^2$ - the object a “quantum dropout” step regularises during training.

Figure 1: From data to decision in quantum machine learning. Classical data (an image, an audio clip or a table) is turned into a quantum state by a feature map, or embedding.

Crucially, none of this happens on quantum hardware alone. Quantum machine learning today is largely hybrid – the quantum circuit proposes an embedding, while a classical computer, very often a high-performance GPU, handles the optimisation, the gradient estimates and the simulation of the quantum system itself. Simulating quantum states is exponentially expensive, so memory and energy consumption climb steeply as circuits grow.

Techniques such as quantum dropout are attractive here for a second reason: by thinning circuits and shortening training, they could cut the number of costly circuit evaluations and the associated GPU workload. Understanding embeddings therefore cannot be separated from the compute budget they demand. As models scale toward problem sizes that classical simulation can no longer reach, the same efficiency pressures will shape which experiments are feasible on real quantum processors at all. The right question is therefore not only “does this embedding help?” but “does it help enough to justify its cost in qubits, runs and energy?”

The open challenge here is the absence of a predictive, unifying account linking embedding design to real-world performance. Most positive results are existence proofs or small synthetic demonstrations. Most negative results are generic warnings. What is missing is a design theory that, given a particular dataset and a fixed hardware and energy budget, predicts which embeddings, and which regularisation tricks such as quantum dropout, will actually help. Building that theory, together with the scalable constructions needed to test it, is the aim of ongoing and upcoming projects at RISE. Achieving this would be a landmark result: moving quantum machine learning from intriguing possibility toward dependable, resource-aware practice.

Links:

[L1] <https://pennylane.ai/qml/>

[L2] <https://www.tensorflow.org/quantum>

References:

- [1] V. Havlíček, et al., “Supervised learning with quantum-enhanced feature spaces”, *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
- [2] H.-Y. Huang, et al., “Power of data in quantum machine learning”, *Nat. Commun.*, vol. 12, no. 1, Art. no. 2631, 2021.
- [3] F. Scala, et al., “A general approach to dropout in quantum neural networks”, *Adv. Quantum Technol.*, vol. 8, no. 12, Art. no. 2300220, 2025.

Please contact:

Smita Chakraborty
RISE, Sweden
smita.chakraborty@ri.se

From Quantum Randomness to Financial Simulation: Synthetic Data and Quantum Monte Carlo

by Enrico Barbierato and Nicoleta Mihalachi (Catholic University of the Sacred Heart)

The article connects the probabilistic foundations of quantum mechanics with practical computational opportunities in finance, while emphasizing the likely role of hybrid quantum–classical approaches.

Quantum computing inherits some conceptual difficulties from quantum mechanics. Quantum advantage depends on the ability to manipulate superpositions, create correlations through entanglement, and exploit interference so that useful computational outcomes are favoured while others are suppressed. At the same time, quantum mechanics imposes important constraints: the internal quantum state cannot be inspected freely, and measurement reveals only limited information. Many quantum algorithms must therefore be executed repeatedly to estimate the probabilities associated with their possible outcomes.

This combination of controlled evolution and probabilistic observation becomes particularly interesting in finance, where many computational problems involve probability distributions. Two examples illustrate this connection particularly well: the generation of synthetic financial data and Monte Carlo estimation (see Figure 1).

Quantum randomness and synthetic financial data

Randomness is essential to simulation and generative modelling. Classical computers normally rely on pseudorandom numbers produced by deterministic algorithms, whereas quantum measurement can provide randomness of genuinely physical origin. Yet randomness alone is not sufficient to generate meaningful financial data, and useful synthetic time series must reproduce statistical distributions, temporal dependencies, volatility patterns and, ideally, rare market events.

Quantum generative models address this more demanding task by using quantum states and circuits to represent and sample probability distributions. Orlandi et al. [1], for example, investigated a hybrid Quantum Wasserstein Generative Adversarial Network with Gradient Penalty for generating synthetic S&P 500 data. Their architecture combines a quantum generator with a classical discriminator, illustrating how a quantum component can be embedded within a conventional machine-learning workflow rather than replacing it entirely.

Synthetic-data generation therefore provides one example of a broader idea: quantum computation may be useful when financial applications require efficient manipulation or sampling of complex probability distributions.

Quantum Monte Carlo in finance

Probability plays an even more fundamental role in Monte Carlo simulation. Many problems in quantitative finance can be expressed as the estimation of an average value over a large number of possible future scenarios. The price of a derivative is a typical example: its value depends on the possible future evolution of one or more market variables and on the payoff associated with each possible outcome.

For simple financial instruments, analytical solutions may exist. For path-dependent derivatives, portfolios involving several assets, complex interest-rate models or credit-risk calculations, analytical solutions often become unavailable or impractical. Monte Carlo simulation then provides a remarkably flexible alternative. A large number of possible scenarios is generated, the relevant financial quantity is evaluated in each scenario, and the results are averaged.

The strength of classical Monte Carlo lies in its generality. Its convergence rate is largely independent of the dimensionality of the problem, which makes it particularly attractive for high-dimensional financial models. However, its weakness is that convergence is relatively slow. To reduce the statistical error by a factor of ten, approximately one hundred times as many samples are required. High accuracy can therefore become computationally expensive, especially when generating and evaluating a single scenario already involves a complex financial model.

Quantum Amplitude Estimation, one of the central techniques behind Quantum Monte Carlo methods [2], can encode infor-

mation about probabilities or payoffs into a quantum state and use interference to estimate the quantity of interest. Under ideal conditions, this leads to a quadratic improvement in the number of evaluations needed to achieve a given accuracy.

The importance of this improvement is easier to appreciate by considering what Monte Carlo simulation actually does. A classical algorithm obtains greater accuracy essentially by generating more samples, while a quantum algorithm instead exploits the structure of quantum amplitudes to extract information about the expected value more efficiently. The potential advantage therefore does not arise from generating random numbers faster, but rather from a different method of processing probabilistic information.

Derivative pricing is one of the most natural applications. European options provide simple examples, but the potential becomes more interesting for products whose payoff depends on an entire price trajectory, on several correlated assets, or on complex market conditions. In these cases, classical Monte Carlo may require very large numbers of simulations, while Quantum Amplitude Estimation offers, at least theoretically, a more favourable scaling with the required precision.

The same principle extends beyond pricing. Value-at-Risk estimates the loss threshold associated with a specified probability, while Conditional Value-at-Risk focuses on the expected loss once that threshold has been exceeded. Both depend on the distribution of future financial outcomes and can require extensive simulation, particularly for large portfolios. Credit-risk calculations, counterparty exposure and scenario-based




Figure 1: Overview of the main quantum approaches discussed for financial simulation, linking synthetic-data generation and Quantum Monte Carlo within a hybrid quantum-classical workflow.

stress testing similarly involve repeated evaluation across many possible future states.

The theoretical advantage, however, depends on the assumption that the relevant probability distribution and financial payoff can be encoded efficiently into a quantum circuit. Preparing these states may itself be expensive. Current quantum hardware also remains affected by noise, limited circuit depth, and restricted numbers of usable qubits. The quadratic improvement of the algorithm concerns computational complexity under suitable assumptions; it does not guarantee that present-day quantum processors will outperform highly optimised classical Monte Carlo implementations. These limitations reinforce the case for hybrid computation.

References:

- [1] F. Orlandi, E. Barbierato and A. Gatti, “Enhancing Financial Time Series Prediction with Quantum-Enhanced Synthetic Data Generation: A Case Study on the S&P 500 Using a Quantum Wasserstein Generative Adversarial Network Approach with a Gradient Penalty,” *Electronics*, vol. 13, no. 11, 2158, 2024.
- [2] P. Rebentrost, B. Gupta, and T. R. Bromley, “Quantum computational finance: Monte Carlo pricing of financial derivatives”, *Physical Review A*, 98(2), 022321, 2018.

Please contact:

Enrico Barbierato
Catholic University of the Sacred Heart, Italy
enrico.barbierato@unicatt.it

Genetic Programming for Quantum Oracle Optimization

by Javier Zayas-Gallardo (Quercus SEG, University of Extremadura), Francisco Chicano (ITIS Software, University of Málaga) and Juan Manuel Murillo (Quercus SEG, University of Extremadura)

Designing quantum oracles by hand becomes unfeasible as problem complexity grows. We introduce a grammar-based genetic programming framework that automatically generates and repairs diagonal oracle circuits, minimising circuit depth. Across standard, composite, and SAT instances, the resulting oracles match or outperform existing literature and naïve baselines, in some cases by more than 80%.

Quantum algorithms often rely on oracles: subroutines that encode a logical condition, such as “is this number prime?” or “does this item satisfy the search criteria?”, so a quantum algorithm can act on it. Oracles allow quantum algorithms to evaluate many possibilities at once and are essential for classification, filtering, and optimization tasks. Grover’s algorithm is an example: it enables unstructured searches in sublinear time through a quantum oracle that marks desired states via phase inversion.

Designing these oracles, however, is far from straightforward. A major challenge lies in their compositionality: combining multiple logical conditions into a single oracle is far from trivial, since encoding conjunctions, disjunctions, or nested predicates does not compose linearly in circuit depth. Manual design becomes unfeasible as logical complexity or qubit count increases, and automatically synthesized oracles often exhibit substantially higher depth than manually optimized ones, limiting their practical viability in the current Noisy Intermediate-Scale Quantum (NISQ) era.

To see why, consider how quantum computing represents and manipulates information. A state in a quantum computer is represented by a complex-valued vector, and transformations occur through unitary, reversible gates (linear functions) such as Pauli X, Y, Z, Hadamard, and multi-qubit gates like CNOT and MCZ. Diagonal gates (Z, MCZ) apply phase shifts without altering measurement probabilities, precisely the mechanism algorithms like Grover’s use to mark desired states. This constraint, that oracles must be diagonal in the computational basis, is part of what makes automatic synthesis so difficult.

A suitable way to specify and generate quantum oracles is to use a formal grammar that constrains how elementary gates (e.g., X, Y, etc.) may be combined. Generic compositions of standard gates do not guarantee that the circuit is diagonal. For this reason, we define a grammar whose production rules enforce that any expression it generates corresponds to a diagonal operator, ensuring the resulting oracle performs phase-only transformations on the computational basis states. Then, we propose a Grammar-Based Genetic Programming algorithm [1] to explore the space of possible quantum oracle circuits (the implementation can be found at [L1]). Full details of the algorithm are available in [2]. Genetic Programming is a nature-inspired approach that draws on the principles of biological evolution, such as natural selection and reproduction, to automatically evolve computer programs or, in this case, circuits, gradually improving them over successive generations. In Genetic Programming, a set of circuits, also called individuals, is evolved using random variation operators, such as recombination (Figure 1), which takes two circuits and produces a new one with features of the original two, and mutation, which introduces a change in the circuit.

Because the circuit-generation process is random, the resulting circuit does not always behave exactly like the target operator it is meant to reproduce. For some input states, it may assign the wrong phase instead of the correct one. To address this, a repair step compares the generated circuit against the target, state by state, and identifies exactly where the two disagree. Extra gates are added to the circuit to correct just those specific cases, leaving everything else untouched. This ensures that every generated circuit implements exactly the target operator, guaranteeing functional correctness and greatly improving the efficiency of the evolutionary search.

Finally, once the circuit has been repaired, the objective function is computed over the circuit. The metric to optimize is the circuit depth, formally defined as the length of the longest path from the input to the output of the circuit, where gates that act on disjoint sets of qubits can be executed in parallel and, thus, occupy the same layer. An example of a generated circuit can be seen in Figure 2.




Figure 1: Example of the recombination (crossover) operator applied to two parent circuits. A subtree/fragment from each parent is exchanged to produce a new offspring circuit that combines features from both.

The proposed framework was evaluated on a diverse collection of quantum oracles of up to 6 qubits, ranging from simple arithmetic conditions to more complex combinations of logical constraints and arbitrary Boolean formulas. The resulting oracles were compared with those obtained in the work by Sánchez et al. [3]. Overall, the results show that the framework consistently produces quantum circuits whose depth is equal to or lower than that of established approaches. The improvements are particularly noticeable when several logical conditions must be encoded within the same oracle, where the method is able to exploit shared structure instead of treating each condition independently. Even for highly unstructured Boolean formulas, where little regularity can be leveraged, the approach still generates compact circuits that outperform straightforward construction techniques. These findings suggest that the framework provides a practical and versatile way to automatically synthesise efficient quantum oracles across a broad range of applications.

While the current results are encouraging, the proposed approach has some limitations that we are now facing. In partic-




Figure 2: An example of a circuit generated by our proposal, showing a composite oracle that mixes the conditions x is prime, $x < 15$, and $x > 6$.

ular, the repair mechanism does not scale well with the number of qubits, due to its enumerative nature. We are now exploring a symbolic treatment of the repairs that helps us to scale the approach to more than 15 qubits, which is a reasonable size for producing predicates over integers. We are also considering the inclusion of ancilla qubits in our evolutionary framework to reduce the depth of the generated circuits.

Overall, the results show that grammar-based genetic programming can provide an effective and scalable way to automate the design of compact, functionally correct quantum oracles, bringing automatic oracle synthesis closer to practical use in the NISQ era.

Link:

[L1] <https://doi.org/10.5281/zenodo.19218736>

References:

- [1] R. I. McKay, et al., “Grammar-based genetic programming: A survey,” *Genetic Programming and Evolvable Machines*, vol. 11, no. 3, pp. 365–396, 2010.
- [2] J. Zayas-Gallardo, et al., “Generation of optimised quantum oracles using grammar-based genetic programming,” in *Parallel Problem Solving from Nature – PPSN XIX*, G. Iacca et al., Eds., *Lecture Notes in Computer Science*, vol. 16987, pp. 303–318, Springer, 2027, doi: 10.1007/978-3-032-36226-1_19.
- [3] J. Sanchez-Rivero, et al., “Automatic generation of efficient oracles: The less-than case,” *Journal of Systems and Software*, vol. 219, Art. no. 112203, 2025.

Please contact:

Javier Zayas Gallardo
University of Extremadura, Spain
javierzg@unex.es

Beyond a Single Measurement: Training Reusable Quantum Models

by Erik Källman (RISE Research Institutes of Sweden)

Quantum machine-learning models are usually trained to match a measured number. When a model's output becomes the input of another quantum model, that is not enough. Work at RISE, built on the open-source Shim library, shows what a reusable and composable quantum model must actually preserve, and tests whether an optimiser can learn it.

Quantum machine learning rarely runs one model in isolation. One model prepares a quantum state; a later model, or a later quantum channel, consumes part of that output. Making such composed and hybrid computations reliable and useful is one of RISE's contributions to QSTAR, the Norwegian national Centre for Quantum Computing and Applications, which works on quantum algorithms and software [L2]. RISE participates as a Swedish project member, and this reduced-state work is one of our early contributions on the software side. It answers a question that sounds simple: when the output of one quantum model becomes the input of another, what must training preserve?

Many quantum models are trained against a single number, the expectation value of one measurement. That is enough when the number is the final output. It is not enough when a downstream quantum model receives the whole state. The smallest counterexample needs one qubit. Two states can give the same fifty-fifty result for one measurement and still be driven to opposite, certain outcomes by a single downstream gate. Agreement on the trained number then survives no further processing.

The right target is the state carried by the wires that cross the boundary between the two models — the exposed reduced state (see Figure 1). Two upstream models that expose the same state are interchangeable: any downstream channel receives the same input from either, so no later measurement can tell them apart. However, when the states are only close, trace distance, a measure of how distinguishable two quantum states are, sets the tolerance, and a basic law of quantum information [1] — that no physical processing can make two states more distinguishable — carries that tolerance through to every later measurement. For exact interchangeability, nothing less than equality will do: the next model might only measure the state it receives, and if the two states differ, some measurement can tell them apart.

Shim, an open-source Rust library developed at RISE [L1], lets us both state and test this in ways previously not possible. It holds circuits, states, density matrices and quantum channels together with category-theoretic interfaces around them [2]. You name which wires pass to the next model and discard the rest; the library forms the reduced state, applies the downstream channel, and reads out the result. Substitution then has a formal meaning, pulling a later measurement back through the downstream channel [3]. Shim turns this exposed-state criterion into something computable.

The value of the category theory approach is that it names the interface as an object: say which system crosses the boundary, and the upstream model is a map that ends there while the downstream model is a map that begins there. Replacement then follows from the fact that the same map on the same input gives the same output, instead of requiring a separate argument for each downstream circuit. The view also fixes what to compare and where: the state on that boundary, read after the discard. And because the guarantee sits at the interface, it composes: a bound proved at one boundary still holds when the model is dropped into a longer chain of models.




Figure 1: The exposed state is the only part of the upstream output that crosses the boundary into the downstream channel; the rest is discarded. The downstream model sees only the exposed state. One upstream model can replace another when their exposed states match; if the states differ by at most epsilon in trace distance, no later measurement can differ by more than epsilon.

We then asked a separate, empirical question: can an optimiser actually learn the exposed state? A four-qubit model with three Ry-CNOT layers exposes two qubits and discards two. Five target states were fixed and then relearned from six random starts under two losses, one matching the full exposed state, the other matching only the single measured number. All sixty runs behaved as designed. Every scalar-trained run reached its target number, to within nine parts in a hundred thousand, yet its exposed states stayed far apart, up to a trace distance of 0.75. Training on the exposed state cut that spread to between 0.001 and 0.016, a median ninety-nine-fold improvement. A downstream check sent both trained states of every pair through the same fixed channel and a measurement, and no pair ever exceeded the predicted bound.

Four of five cases met the threshold of 0.01 that had been fixed in advance; one reached 0.016, and several runs used the full optimisation budget. That keeps two claims apart: it is proven that the exposed state defines what a reusable quantum model must preserve, while whether a given optimiser can preserve it in practice remains an open question, with a qualified answer. The distinction is the point. As quantum toolchains grow and hybrid pipelines start to interoperate, components will be swapped in and out, and what swapping means needs exactly this criterion, with a number attached to how close is close enough.

This is a deterministic, exact-simulator study of one small model family, meant as a baseline. Later experiments move to shot noise, real hardware, and larger exposed systems. With this result at hand, later studies can also investigate different model topologies and how trace distance affects their output, which is a fundamental aspect in quantum machine-learning model building.

Links:

- [L1] <https://github.com/erikkallman/shim>
 [L2] <https://tinyurl.com/qstar-kick-off>

References:

- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2010.
 [2] B. Coecke and A. Kissinger, *Picturing Quantum Processes*, Cambridge University Press, 2017.
 [3] E. D'Hondt and P. Panangaden, "Quantum weakest preconditions", *Math. Struct. Comput. Sci.*, vol. 16, no. 3, pp. 429-451, 2006.

Please contact:

Erik Källman
 RISE Research Institutes of Sweden
erik.kallman@ri.se

Quantum Logic Programming in Prolog

by Giovanni Pilato (CNR-ICAR), Salvatore Gaglio (Università degli Studi di Palermo and CNR-ICAR) and Filippo Vella (CNR-ICAR)

Programming quantum systems still requires working with low-level circuit operations. Researchers at CNR-ICAR and the University of Palermo are exploring a higher-level alternative: Quantum Logic Programming in Prolog (QLPP), which translates logical formulas into sequences of quantum gates. Although quantum computing is progressing from theoretical promise toward increasingly powerful hardware platforms, programming quantum systems remains a challenge. QLPP aims to address this by providing a high-level approach for researchers who are used to rule-based and logic-based descriptions of problems.

Quantum logic is a field of study that starts from the logical structures inherent in quantum mechanics and moves beyond the framework of classical logic. Its formal origins can be traced to the 1936 article "The logic of quantum mechanics" by Birkhoff and von Neumann, where the authors identify the underlying logical structures of physical theories, such as quantum mechanics, that differ from classical logic.

Quantum logical connectives are declared as Prolog operators and associated with quantum gates such as NOT, CNOT and Toffoli. The syntactic structure of a formula is then used to build a quantum circuit layer by layer.

In QLPP, the structure of the Prolog formula determines the circuit, while each logical connective is replaced by its corresponding quantum gate.

The relevance of non-classical logic in quantum theory stems from the way propositions about physical quantities are represented mathematically. These representations differ from those used in classical physics and give rise to non-classical logical models. Non-classical logics form a family of approaches that arise from the abstract structures of quantum theory.

An example is provided by Dalla Chiara et al. [3], who introduced a logic for quantum computation implemented by quantum circuits. Sentences in the logical language are composed of atomic symbols corresponding to quantum registers, and




Figure 1: Quantum circuit implementing a robot-movement rule on a 2×2 map. Qubit x represents the x -coordinate, y the y -coordinate, and mov the movement.

quantum computational Boolean connectives such as $\neg$ (NOT), $\wedge$ (AND), $\vee$ (inclusive OR), and $\oplus$ (exclusive OR). These connectives are implemented using X quantum gates for NOT, Toffoli gates for AND, a mapping of inclusive OR to AND via the De Morgan laws, and controlled NOT gates for the exclusive OR. To the classical Boolean operators are also added specific quantum computational connectives, such as $\sqrt{\neg}$ and $\sqrt{\text{id}}$.

The quantum Boolean computation is represented and analysed using a syntactic tree. The generic formula α is represented with a sequence of levels from the original formula to its elementary parts. Starting from the original formula α , which forms the bottom level, the next levels are obtained by removing the principal connective and repeating all the atomic sentences of the current level.

Each element of the syntactic tree can be associated with a meaning through the so-called holistic map. This map is a density operator, implemented through partial trace [2]. This operator is useful for interpreting formulas because its value depends on the entire formula rather than on a single symbol. This is a distinctive feature of the formalism that is difficult to reproduce using a classical representation.

Quantum logic for robot planning

Logic-based systems described through Boolean representations and assertions can describe both system evolution and target functions. System states are represented at the bit-level and evolve according to predefined update rules.

A similar method can be applied to all systems described by production rules. These rules are often used to characterise problems in a way that resembles human reasoning. They are also connected with Markov models, which in turn can be related to Universal Turing Machines.

The representation and evolution through rules that can be translated into quantum operators are well suited to logic programming, particularly through the Prolog language, which can be used to implement a computational logic formulation.

A simple example is robot movement on a small grid. The robot's current position and its possible movement are represented by Boolean variables, and the corresponding update rules can be expressed as logical formulas in Prolog. QLPP then translates these formulas automatically into a sequence of quantum gates (Figure 1). In this example, the circuit uses controlled-NOT, NOT and SWAP operations. In our case study [2], the resulting circuit is subsequently combined with Grover's algorithm, a quantum search method, to identify movements that lead the robot to a target position.

One advantage of using Prolog for this formulation is its declarative nature: the programmer can describe the problem in logical terms, while QLPP generates the corresponding sequence of quantum gates. The benefit at this stage is therefore a higher level of abstraction rather than a demonstrated computational speed-up over existing quantum programming approaches.

It is very suitable for symbolic AI, theorem proving and constraint satisfaction problems. Pattern matching can support symbolic manipulation and translation from one domain to another,

facilitating the creation of a quantum solution given the definition of the problem with classical formulas.

References:

- [1] A. Chella, et al., "An architecture for a quantum teleo-reactive robot", Entropy, 2026.
- [2] G. Pilato, et al., "Towards quantum logic programming", in Proc. 2025 IEEE Conf. Pervasive Intell. Comput. (PCom), 2025, pp. 267–272.
- [3] M. L. Dalla Chiara, et al., "Quantum Computation and Logic: How Quantum Computers Have Inspired Logical Investigations", Springer, 2018.

Please contact:

Giovanni Pilato
CNR-ICAR, Italy
giovanni.pilato@icar.cnr.it

The Quantum Software Crisis: A Reminder of Classical Software Engineering

by Enrique Moguel, Jaime Alvarado-Valiente, and Javier Romero-Alvarez (Universidad de Extremadura)

If we wait until quantum hardware becomes mature before thinking seriously about Quantum Software Engineering (QSE), we risk repeating the same mistakes. The software crisis showed what happens when hardware evolves faster than software engineering. This article revisits the classic software crisis and emphasises that learning from the past is essential for developing the methodologies and abstractions needed to prevent a future quantum software crisis.

Over the past century, quantum mechanics has transformed our understanding of the physical world and has driven revolutionary advances in disciplines ranging from chemistry and telecommunications to microelectronics and computing. Today, these same principles are giving rise to quantum computers, devices that exploit phenomena such as superposition and entanglement to process information in ways that are fundamentally different from classical computers. Their unique capabilities promise not only to solve certain problems dramatically faster, but also to tackle computational challenges that are currently beyond the reach of even the most powerful supercomputers [1].

However, revolutionary hardware alone is never enough. Without equally mature software engineering practices, powerful computers become difficult to program, expensive to maintain, and ultimately inaccessible to most developers.

Learning from the Past to Build the Future

During the first decades of classical computing, computers were expensive, enormous, and highly unreliable. Early sys-




Figure 1: The evolution of computer technology, from the first classical computers to modern quantum computers. Illustrating the relationship between the evolution of hardware and the maturity of software engineering.

tems based on vacuum tubes suffered frequent failures, while software was written using machine-specific instructions and tightly coupled to the hardware. As a result, programs were rarely portable and often had to be rewritten for each new machine. Ironically, as hardware advanced, software development became increasingly complex rather than easier.

Edsger Dijkstra famously summarized the situation: “The major cause of the software crisis is that the machines have become several orders of magnitude more powerful!” [2].

The problem was not simply that programs became larger. Developers lacked the necessary capabilities to handle increasing complexity. The response eventually became known as Software Engineering. In this context, instead of viewing programming as an individual craft, software development evolved into a disciplined engineering process supported by modular programming, object-oriented design, software architectures, and testing methodologies, among others.

Following the Same Path in Quantum Computing

Figure 1 compares the historical evolution of consumer computers, supercomputers and quantum computers. The current state of quantum computing closely resembles the early stages of classical computing. Both paradigms have faced remarkably similar challenges: (1) the first generations of classical computers were expensive, scarce, and unreliable; today’s quantum computers are equally costly, limited in availability, and highly susceptible to noise. (2) Classical software was tightly coupled to specific hardware platforms, making applications difficult to port; similarly, quantum software remains strongly dependent on vendor-specific hardware, SDKs, and execution models. (3) Programming early classical computers required highly specialized expertise and considerable effort; likewise, developing quantum circuits and hybrid quantum-

classical applications still demands scarce expertise and significant development time. (4) Finally, access to early mainframes was constrained by limited computing resources and shared infrastructures; similarly, quantum developers continue to face restricted access to QPUs and long execution queues.

Classical computing eventually overcame these limitations through advances in Software Engineering. Given the rapid progress of quantum hardware in recent years, quantum computing is expected to follow a similar evolution, where QSE will play a key role in enabling mature, portable, and widely accessible quantum systems [L1].

Beyond Classical Thinking

History shows that abstractions transformed classical computing. High-level programming languages replaced assembly language. Functions, objects, modules, components, and services enabled programmers to reason about increasingly complex systems without having to constantly consider processor instructions. Figure 1 also illustrates this historical progression, showing how advances in hardware have consistently been accompanied by corresponding advances in software engineering methodologies. Quantum computing must follow a similar path and needs similar abstractions to become a mature software platform rather than a collection of isolated experiments.

Developers require programming constructs that naturally express quantum algorithms while hiding unnecessary hardware details, since they currently manipulate quantum circuits directly, rather than reasoning at higher conceptual levels. They also need software architectures, reusable components, debugging techniques, testing methodologies and design patterns. Equally important, they need quantum equivalents of the fundamental building blocks that support iterative software development.

Looking at this trend, classical software engineering methodologies might seem directly applicable to quantum computing. However, they are based on the Von Neumann model [L2], where programs manipulate data to produce deterministic results. Quantum computing instead evolves through probability amplitudes, with measurement collapsing the state into a single outcome, making familiar concepts such as variable inspection or loop evaluation impossible or fundamentally different.

This is where one of the greatest challenges facing QSE arises [3]. The goal is not simply to provide new libraries, but to raise the level of abstraction so that developers can design applications without reasoning directly in terms of qubits. This requires introducing quantum-specific software constructs, such as reusable quantum data types, composable oracles, high-level programming primitives, and design patterns for hybrid quantum-classical applications, that encapsulate complex circuit implementations while preserving quantum behavior.

Conclusion

The history of classical computing offers a clear warning. Software crises do not arise because hardware is weak, they arise because hardware advances faster than our ability to engineer software. Today, quantum computing risks following the same trajectory.

The next generation of software engineers needs new abstractions and software engineering fundamentals that incorporate the unique characteristics of quantum computing, enabling them to “think in quantum”, while preserving the best lessons learned over more than sixty years of classical software engineering. If we wait for quantum hardware to mature and become standardized before tackling these software challenges, it will already be too late to avoid a new crisis in quantum software.

This work was funded by grants PDC2025-165096-C31 and PID2024-155693NB-C41 from MICIU/AEI/10.13039/501100011033, co-funded by the European Regional Development Fund (ERDF/EU).

Links:

[L1] <https://www.ibm.com/quantum/hardware>
[L2] <https://tinyurl.com/modelodvn>

References:

- [1] M. J. Klein, “Max Planck and the beginnings of the quantum theory”, *Archive for History of Exact Sciences*, vol. 1, no 5, pp. 459-479, 1961.
- [2] E. W. Dijkstra, “The humble programmer”, *Comm. of the ACM*, vol. 15, no. 10, pp. 859-866, 1972.
- [3] J. M. Murillo, J. Garcia-Alonso, et al., “Quantum software engineering: Roadmap and challenges ahead”, *ACM Transactions on Software Engineering and Methodology*, vol. 34, no. 5, pp. 1-48, 2025.

Please contact:

Jaime Alvarado-Valiente
Universidad de Extremadura, Spain
jaimeav@unex.es

Javier Romero-Alvarez
Universidad de Extremadura, Cáceres, Spain
jromero@unex.es

Towards the Design of Service-Oriented Hybrid Quantum-Classical Systems

by Álvaro M. Aparicio-Morales, Ernesto Pimentel (University of Málaga) and José Garcia-Alonso (Universidad de Extremadura)

Quantum computing is unlikely to operate in isolation. Its practical impact will depend on its integration with existing software ecosystems. This article explores how Service-Oriented Architecture can help expose, compose and manage quantum capabilities as services, and why classical service-oriented principles must be reconsidered in the quantum domain.

Quantum computing promises to provide new computational capabilities for problems that remain challenging for current systems. Its integration into existing software ecosystems is therefore becoming increasingly relevant, giving rise to hybrid quantum-classical systems. However, current quantum computing is largely based on remote access to resources offered by different providers and deployed over heterogeneous infrastructures. In this context, Service-Oriented Architecture (SOA) provides a suitable perspective by promoting the exposure, composition and consumption of quantum capabilities through well-defined interfaces, contracts and access mechanisms. This vision builds on Service-Oriented Computing, where services have enabled interoperable, flexible and reusable distributed systems [1]. Nevertheless, transferring these principles to the quantum domain is not straightforward: provider heterogeneity, hardware dependence, result variability and probabilistic execution introduce new tensions into abstraction, interoperability, service composition and quality-of-service management, which must be reconsidered for hybrid quantum-classical systems [2]. This article explores how SOA principles can be applied and extended to address these challenges in hybrid quantum-classical software systems, as summarised in Figure 1.

Abstraction

In general, reasoning about complex systems requires an adequate level of abstraction. In SOA, this is achieved by representing software capabilities as services that encapsulate business logic behind well-defined interfaces, promoting loose coupling, technological independence, reuse and scalability. Transferring this idea to the quantum domain means enabling applications to consume quantum capabilities without becoming tightly coupled to a provider, SDK or device. As quantum hardware is mostly accessed through cloud-based or managed environments, service-level abstractions are essential for its integration into broader software systems. Achieving this level of abstraction still requires progress in the hardware and software infrastructure needed to offer quantum services with a maturity comparable to that of classical services [1].

Abstraction should also extend to the hardware layer. Today, properties such as the number of qubits, connectivity, native




Figure 1: Service-oriented challenges for hybrid quantum-classical systems. The figure summarises three main challenges that arise when applying Service-Oriented Architecture (SOA) principles to hybrid quantum-classical systems. Abstraction concerns the representation of quantum capabilities as services, while accounting for hardware-aware properties and application-level confidence or feasibility thresholds. Interoperability, composition and orchestration refer to the need for provider-agnostic access, common APIs and contracts, and mechanisms to coordinate hybrid workflows, including techniques such as circuit cutting. Quality of Service (QoS) highlights the need to manage classical service properties, such as latency, cost and availability, together with quantum-specific factors such as noise, result variability, calibration status and adaptive monitoring and selection.

gate sets, noise levels and calibration status remain fundamental when deploying circuits on specific devices [1]. From a service-oriented perspective, however, these characteristics should not be managed directly by the consuming application. Instead, they could be encapsulated in higher-level service descriptions associated with the computational capability offered. Thus, rather than reasoning in terms of individual qubits or provider-specific constraints, an application could express requirements related to feasibility, expected result quality, cost, waiting time or a minimum confidence threshold. The service infrastructure would then select, configure and adapt the most suitable quantum resource.

Interoperability, Composition and Orchestration

Interoperability is another fundamental aspect of SOA. Here, it should not be understood only as the ability to execute the same service across platforms, but as the ability to integrate heterogeneous services through common interfaces, contracts and access mechanisms. This is particularly relevant in quantum computing, where the ecosystem comprises multiple providers, technologies, SDKs, programming models and devices.

At present, important limitations still prevent quantum capabilities from being truly provider-agnostic [2]. Hybrid applications may become coupled to a specific platform due to the programming language, APIs, circuit representation format, native gates or backend constraints. From a service-oriented perspective, the challenge lies in defining mechanisms that allow quantum services to be replaced, combined or reused without redesigning the whole application [3].

Moreover, quantum services will not operate in isolation. Hybrid applications must coordinate classical and quantum services within broader workflows, including data preparation, resource selection, quantum execution, post-processing and validation. Consequently, composition and orchestration become essential to build applications from independent services while transparently managing provider selection, execution

changes, resource availability and adaptation to failures [3]. This also invites us to rethink granularity: instead of treating a quantum algorithm as a monolithic circuit, service-oriented solutions may decompose quantum capabilities into smaller, reusable units. Techniques such as circuit cutting may support this vision, although models, tools and middleware are still needed to connect these low-level techniques with service-oriented orchestration.

Quality of Service (QoS) Criteria

Service-oriented applications must satisfy Quality of Service (QoS) requirements, including cost, response time, availability, reliability, scalability and resilience. In many domains, these properties are part of service-level agreements, whose strictness depends on the application context, becoming particularly demanding in critical environments [3].

The integration of quantum resources into service-oriented applications therefore requires QoS to be reconsidered in the hybrid quantum-classical context. Architectural decisions must support the selection of the most suitable execution resource, whether classical, quantum or hybrid, and among alternative services. These decisions may affect deployment, provider selection, execution configuration, the use of simulators or real QPUs, and dynamic adaptation when conditions change.

This problem is intensified by quantum-specific factors. Remote access, execution queues, provider heterogeneity, device noise, result variability and calibration changes introduce uncertainty that must be incorporated into QoS management. Monitoring latency or cost is not sufficient; quantum-specific properties such as expected result quality, execution error or confidence from a given number of shots must also be considered. Therefore, autonomous monitoring, selection and adaptation mechanisms are needed to verify requirements, detect quality degradation and activate strategies such as backend changes, provider selection, parameter adjustment or service redeployment.

Conclusion

Service-Oriented Architecture provides a suitable foundation for integrating quantum computing into current software systems, but this integration cannot be treated as a simple extension of classical models. Provider heterogeneity, hardware dependence, probabilistic behaviour and dynamic resource constraints require abstraction, interoperability, service composition and QoS management to be reconsidered. Progress will require new infrastructures, description models, orchestration mechanisms and adaptation strategies. Achieving practical quantum advantage will depend not only on quantum hardware and algorithms, but also on their effective integration with classical systems through extended service-oriented principles.

References:

- [1] J. García-Alonso et al., “Rethinking Services in the Quantum Age: The SOQ Paradigm,” *ACM Trans. Softw. Eng. Methodol.*, 2026, doi: 10.1145/3805807
- [2] Á. M. Aparicio-Morales, J. Garcia-Alonso, J. Cámara, and J. M. Murillo, “Architecting Hybrid Quantum-Classical Software Systems: Exploration of the Design Trade-off Space with Quantitative Guarantees,” *arXiv preprint arXiv:2606.24260*, 2026, doi: 10.48550/arXiv.2606.24260.
- [3] J. M. Murillo et al., “Quantum Software Engineering: Roadmap and Challenges Ahead,” *ACM Trans. Softw. Eng. Methodol.*, vol. 34, no. 5, Art. no. 154, pp. 1–48, May 2025, doi: 10.1145/3712002.

Please contact:

Álvaro Manuel Aparicio Morales
ITIS Software, University of Málaga, Spain
alvaroapamor@uma.es

Quantum-Inspired Tensor Networks for Counting Software Configurations

by Daniel-Jesus Munoz and Lidia Fuentes (Universidad de Málaga)

Industrial variability models routinely reach thousands of options, and counting their valid configurations governs almost every other analysis performed on them. Search-based counters cope well until they do not. We are designing a contraction engine that counts by algebra rather than by search.

A Software Product Line (SPL) is a family of systems built from shared assets, with their differences made explicit. A feature model organises those differences as a tree of features plus cross-tree constraints, and a configuration is valid when its selected features satisfy all of them. Determining whether such a configuration exists is known as the Boolean satisfiability (SAT) problem; #SAT, or model counting, goes a step

further by asking how many valid configurations exist. In practice, that number lets engineers draw test configurations uniformly at random, measure how often each feature appears across the valid products, and spot options that no valid product can ever include.

Regarding who needs those counts, the stakeholders are concrete: automotive suppliers, operating-system distributors, and industrial control vendors whose certification obligations make the count auditable rather than interesting. While the question sounds arithmetic, it is #P-complete, the counting counterpart of NP-completeness, and the configuration spaces at hand are colossal.

Nevertheless, the usual motivation for physically inspired hardware is not the one the evidence supports. Boolean analysis of industrial feature models is empirically easy, and counters built on knowledge compilation count most industrial models quickly, although a few systems remain on which every counter tested so far times out. Substrates that sample return low-energy solutions without certificates, and annealing hardware pays twice, since its sparse qubit graph forces minor embedding, whose qubit overhead and chain breaks have so far kept densely constrained models well short of industrial size. Gate-based formulations are elegant, but are restricted to sizes no practitioner would call industrial: in [3], the authors report at most 16 features. Work tackling exact counting is rare.

Aiming to preserve the exactness counting demands, we plan to build on tensor networks, the formalism devised for weakly entangled quantum states and now used to simulate quantum circuits classically. Each clause becomes a small tensor, joining two legs means summing over a shared index, and contracting the whole network yields the number of models [1]. No bond dimension is truncated. The contraction is therefore exact as mathematics, whatever precision stores its intermediate tensors; the arithmetic required to represent such large values is a separate design consideration.

Concretely, exact tensor-network counting is not new [1]; what is missing is its adaptation to realistic variability models, which is what we are designing. The pipeline has four stages: (1) a reader from variability formats to a hypergraph, (2) a contraction-order search [L1], (3) contraction on the GPU [L2] with slicing bounded by video memory, and (4) an accumulation scheme sized to the result. Figure 1 walks those stages, which plug into flamapy [L3] as a backend.

While quantum devices remain noisy, with few qubits and short coherence, the same algebra already runs on a graphics card. By dispensing with annealing hardware, we also avoid the need to accommodate a specific device topology: the treewidth, the main structural determinant of intermediate tensor size, is a property of the model, while memory, bandwidth, contraction order and slicing decide what is feasible in practice. Amplitude estimation, a quantum counter with a quadratic guarantee, would be the successor once circuit depth allows.

As for evaluation, the target operation is exact configuration counting and the corpus comes from public variability repositories and from Kconfig models of the Linux kernel, which run beyond 14,000 options. The baselines are the established exact counters d4, GANAK, sharpSAT and ddnnfe, and the dimen-

sions measured are wall-clock seconds, peak memory in gigabytes, and joules per completed count, read from hardware counters, not estimated. Metaheuristics do not count exactly.

Then, the structure of real feature models has to be measured first, because contraction cost grows exponentially with the treewidth of the primal graph, the main structural bound on the intermediate tensors. Our hypothesis, no more than that, puts the useful regime below a treewidth of 30 and the hopeless one above 45; off-the-shelf heuristics allow us to bound it over real models in a weekend. On the other hand, to our knowledge no published study reports them, so a negative result would settle a question answered today by intuition.

We are aware that the approach rests on a structural property we have not yet measured, and that numeric features, whose translation to logic no published approach covers completely [2], inflate any encoding through bit-blasting. That inflation, however, may be structured. At least, sums and comparisons admit bit-blasted encodings of low width, whereas products

between decision variables are typically far costlier, so which arithmetic fragment stays tractable is an encoding question we intend to measure.

Continuing past that measurement, the next step parameterises the engine by semiring, so one kernel serves counting, family-wide reliability estimates and minimum-cost queries. The formalism is quantum, the hardware classical. Delivered as an open artefact in the community's standard exchange format, it would matter most to smaller European vendors shipping configurable software under certification duties without large compute budgets. Our central hypothesis is that feature models, built by people to stay comprehensible, carry the hierarchy and modularity that keep treewidth low, and that is the bet this line rests on.

Configuration counting underpins how configurable software, from the Linux kernel to automotive control units, is tested, analysed and certified, yet some industrial models still defeat every existing counter. We propose to take tensor networks, a formalism from quantum physics, and run them on classical GPUs, so that counting becomes an exact algebraic contraction rather than a search. The result would be a contraction engine that plugs into existing variability tools, together with a measurement of whether real feature models are structured enough for it to work.

This work is supported by the projects SAVIA PID2024-159945NB-I00 (co-financed by FEDER funds) and DUNE DGP_PIDI_2024_00092.

Links:

- [L1] <https://github.com/jcmgray/cotengra>
 [L2] <https://developer.nvidia.com/cuquantum-sdk>
 [L3] <https://github.com/flamapy/flamapy>

References:

- [1] S. Kourtis et al., “Fast counting with tensor networks”, *SciPost Physics*, vol. 7, no. 5, art. 060, 2019. <https://doi.org/10.21468/SciPostPhys.7.5.060>
 [2] D. Fernandez-Amoros et al., “Pragmatic random sampling of Kconfig-based systems: A unified approach”, *Journal of Systems and Software*, vol. 230, art. 112577, 2025. <https://doi.org/10.1016/j.jss.2025.112577>
 [3] J. Ammermann et al., “Quantum Solution for Configuration Selection and Prioritization”, in *Proceedings of the 5th ACM/IEEE International Workshop on Quantum Software Engineering (Q-SE 2024)*, ACM, New York, NY, USA, 2024, pp. 21–28. <https://doi.org/10.1145/3643667.3648221>

Please contact:

Daniel-Jesus Munoz
 ITIS Software, Universidad de Málaga, Spain
dm@uma.es




Figure 1: The four-stage counting pipeline, from feature model to configuration count, and the memory wall that governs it. Intermediate tensor size grows exponentially with the treewidth of the primal graph, which makes that parameter the primary model-side determinant, while device memory and bandwidth, contraction order and slicing decide practical feasibility. Sizes are memory estimates for the intermediates, not the precision of the final count.

Quantum Retina: A Quantum Eye on the Brain and Other Things

by Göran Wendin (RISE)

The biological retina is an intelligent sensor because it uses complex neural circuits to compute, filter and pre-process data locally before sending meaningful compressed data to the brain via the optic nerve. The artificial Quantum Retina is based on the same concept: a grid of quantum sensors coherently coupled to a quantum neural network that sends pre-processed compressed quantum data to a powerful quantum processor. The concept is realistic and feasible in a long-term perspective and provides a blueprint for stepwise integration of quantum sensors, communication and information processing into an intelligent imaging system.

The overarching goal is to demonstrate quantum-enhanced sensing modalities for applications such as brain imaging, cardiovascular diagnostics, and cancer detection, with systematic benchmarking against established clinical technologies [1]. The specific goal is to develop an intelligent quantum sensor - a Quantum Retina - that coherently integrates and entangles a quantum sensor array with a quantum data processing device. It will consist of a grid of quantum sensors coherently connected to a quantum backplane that processes quantum data from the sensors. The quantum processor backplane may ultimately be a trainable analogue quantum artificial neural network. A vision of a “quantum computer at the tip of a scanning

probe - SPQ” was formulated already by G. Wendin in 2017 emphasizing coherent quantum data transfer from sensor to processor. This is now a hot topic [2,3] and the rapid development of quantum technologies makes it possible to experiment with real applications.

The current project started in the spring of 2026 and involves the Emerging Technologies [L1] and Computer Science divisions [L2] at RISE. It provides a platform for broad collaboration to build systems integrating quantum sensors with both classical and quantum communication and computing. It also paves the way for in-depth collaborations between institutes and academia in Sweden and beyond.

A fully functional Quantum Retina will avoid classical-quantum conversion and quantum memory problems and will provide advantages over conventional experiments in which quantum states are measured and pre-processed, and the resulting classical data processed with classical, or even quantum, computers. One goal is to develop quantum-enhanced sensing modalities for applications such as brain imaging [1] to eventually pave the way for quantum-coherent brain-machine-computer interfaces.

Particularly interesting as elementary sensors are electronic spin qubits in NV centres (NVC) in diamond, quantum dots in nanotubes, and molecules. The grid of quantum sensors will be directly connected to a layer of qubits, potentially a quantum neural network (QNN), functioning as a quantum pre-processor. In principle, the sensors and the processing qubits could be of the same kind, or different types coherently interfaced together - time will show what works best. Of particular interest will be to apply quantum reservoir computing and machine learning to a coherently coupled QNN to analyse time-depen-




Figure 1: Intelligent quantum sensor-based on spin qubits. Conceptual image generated with Microsoft Copilot.

dent signals driving the quantum sensor array, providing pre-processed classical or quantum data communicated to an external powerful information processor.

The groundbreaking impact of the Quantum Retina builds on the potential for “quantum advantage” that the project can develop in tune with the current and future evolution of quantum technologies and AI. In the initial stage, quantum data collected by quantum sensors: NVC, quantum dots, molecules, and photonic sensors will be read out and processed classically. Later, adding and measuring a coherently connected qubit network will provide pre-processed compact classical data representing an intelligent quantum sensor. Finally, the system will be miniaturised to a compact Quantum Retina and integrated in measuring systems, e.g. forming a grid of Quantum Retinas in a brain scanning helmet. (Figure 1).

The Quantum Retina promises to provide clear Quantum Advantage. Huang et al. (Science, 2022) investigated quantum advantage in learning from experiments that processes quantum data with a quantum computer and proved that quantum machines with access to quantum data can learn from exponentially fewer experiments than the number required by conventional experiments. Exponential advantage was shown for predicting properties of physical systems, performing quantum principal component analysis, and learning about physical dynamics. 10 years from now, Quantum Retina types of projects will allow us to scan the brain, map in detail its functional network of hubs and circuits, identify problematic functional circuits, localise the “faulty” or “poorly programmed” physical neural circuits.

Quantum technologies will be at the core of this project at all levels: sensing, communication, simulation and computing. It will progress in time with the increasing maturity of quantum components, quantum communication and computing systems. AI will be integrated at all levels.

These types of projects also have a big overlap with distributed quantum computing - quantum computers connected through a quantum internet - achieving this in a local setting, with potential for useful results ahead of a working system of distributed quantum computers. On the way to that, the advanced sensing system in the form of a Quantum Retina will be groundbreaking. In a larger IoT network perspective, the Quantum Retina is an Edge quantum processor that is connected by a quantum internet to central quantum computers. And on a smaller scale, an intelligent quantum pixel in a quantum-coherent sensor for brain scanning and medical diagnostics.

Finally, printed electronics for quantum technologies is a topic of great interest and importance. It is at the focus of ongoing research at RISE [L3] and will provide a platform for experiments toward intelligent quantum sensing and imaging systems.

Links:

[L1] <https://www.ri.se/en/quantum-technologies>

[L2] <https://www.ri.se/en/quantum-technology/expertise/quantum-computing-and-simulation>

[L3] <https://www.ri.se/en/advanced-electronics/printed-electronics/expertise/electrochromic-systems-and-displays>

References:

- [1] J. Waldthaler, I. Comarovschii, and D. Lundqvist, “Magnetoencephalography-based prediction of longitudinal symptom progression in Parkinson’s disease,” *npj Parkinson’s Disease*, vol. 12, Art. no. 29, 2026.
- [2] R. R. Allen, F. Machado, I. L. Chuang, H.-Y. Huang, and S. Choi, “Quantum computing enhanced sensing,” *arXiv preprint arXiv:2501.07625*, 2025.
- [3] S. A. Khan, et al., “Quantum computational sensing using quantum signal processing, quantum neural networks, and Hamiltonian engineering,” *npj Quantum Information*, vol. 12, Art. no. 119, 2026.
Also available: *arXiv:2609.35016v1*.

Please contact:

Göran Wendin, RISE, Sweden
goran.wendin@ri.se

Formal Methods Move into Quantum Computing

by Kim Guldstrand Larsen, Christian Schilling (Aalborg University), Mirco Tribastone (IMT School for Advanced Studies Lucca), Max Tschaikowski and Paolo Zuliani (Sapienza University of Rome)

Techniques developed to analyze classical software are finding a new role in quantum computing. Our recent work shows how model reduction, model checking and satisfiability solving can help simulate quantum programs and verify quantum communication protocols.

Quantum computing is often discussed as a hardware race: more qubits, lower error rates and longer coherence times. But useful quantum computers will also need a mature software stack. Programmers must be able to describe algorithms at a convenient level, simulate them before execution and gain confidence that the resulting software does what it is supposed to do.

Quantum software makes this particularly challenging. A classical program can be paused and inspected. A quantum state, by contrast, cannot simply be read without changing what happens next, and access to real quantum hardware remains limited. This makes conventional testing difficult.

This is where formal methods can help. Formal methods are mathematical and algorithmic techniques for analyzing software and systems, developed over decades for areas such as embedded systems, communication protocols and safety-critical software. Our vision is to transfer these techniques between research communities and bring them into quantum computing rather than starting the quantum software toolbox from scratch.

This direction is central to the Danish ODAQS project, which brings together Aarhus University, Aalborg University and the quantum-software company Kvantify. ODAQS develops a

software stack spanning compilation, simulation and verification, with computational chemistry as a major application area. The vision is shared by the Italian project QUMOBA at the Sapienza University of Rome, which develops scalable methods for verifying quantum circuits by combining model reduction with barrier certificates.

The broader agenda also includes CLASSIQUE, the Center for Classical Communication in the Quantum Era at Aalborg University. CLASSIQUE is funded by the Danish National Research Foundation with about EUR 12 million. Its research includes reliability and verification for systems combining classical and quantum communication.

Building a community around this transfer is equally important. In 2025, we initiated the Workshop on Formal Methods in Quantum Computing (FMQC) [L1]. It brings together researchers from quantum computing and formal methods and encourages the exchange of methods between the two communities. A second edition was held in Lisbon in 2026, and a third is planned for 2027 in Copenhagen.

What can this transfer look like in practice? Three recent results illustrate the idea. Figure 1 summarizes this technology transfer from formal methods to quantum computing.

First, a quantum system with only a modest number of qubits already has an enormous classical state space, so simulation quickly becomes expensive. In [1], we adapted a formal-methods technique called bisimulation, traditionally used to identify states that behave in the same way. The method searches for a smaller representation of a quantum computation while preserving the information that matters for the result. Combined with efficient data structures called decision diagrams, this produced speed-ups of several orders of magnitude on quantum-circuit benchmarks, including established algorithms for search and optimization. For Grover’s search algorithm, the dynamics can even be reduced to a two-state description while preserving the relevant outcome.

Second, formal methods can help with quantum communication. In [2], we extended the verification tool UPPAAL, originally developed for classical real-time systems, so that it can analyze protocols containing quantum states, operations and




Figure 1: Technology transfer from formal methods to quantum computing. Techniques such as model reduction, model checking, and satisfiability solvers support the simulation and verification of quantum programs and communication protocols. This research direction is pursued in the ODAQS, QUMOBA and CLASSIQUE projects and is supported by the FMQC workshop series, which brings together the formal-methods and quantum-computing communities.

measurements. The approach was applied to a two-way quantum communication protocol together with entanglement distillation. UPPAAL can exhaustively check ideal, noiseless versions and use so-called statistical model checking when realistic timing and noise are included. This matters because future quantum networks will combine classical control, timing constraints and fragile quantum resources.

Third, we can borrow automatic reasoning technology from software verification. In [3], the tool symQV translates the question “does this quantum program satisfy its specification?” into mathematical constraints handled by a so-called satisfiability modulo theories (SMT) solver. Such solvers are widely used in formal verification and model checking.

These examples use different techniques, but the message is the same. Quantum computing creates genuinely new challenges, yet many surrounding software problems - reducing complex systems, checking protocol behavior and proving program correctness - have close classical relatives. Decades of research in formal methods therefore provide a rich source of ideas, algorithms and mature tools.

The next step is to make these techniques part of everyday quantum software engineering. ODAQS, QUMOB and CLASSIQUE pursue this goal, while FMQC provides a meeting place for the communities whose ideas need to come together. The ambition is not to make quantum computing classical. It is to make quantum technology benefit from what computer science already knows about building software that is efficient, predictable and trustworthy.

Links:

[L1] <https://fmqc-workshop.github.io/2025/>

References:

- [1] L. Burgholzer, et al., “Forward and Backward Constrained Bisimulations for Quantum Circuits Using Decision Diagrams”, *ACM Transactions on Quantum Computing* 6(2), 2025. DOI: 10.1145/3712711
- [2] R. Bødker Christensen, et al., “Analysis and Verification of Quantum Communication Protocols in UPPAAL”, *CAV 2026*, 2026. DOI: 10.1007/978-3-032-32537-2_18
- [3] F. Bauer-Marquart, S. Leue, and C. Schilling, “symQV: Automated Symbolic Verification of Quantum Programs”, *FM 2023*, 2023. DOI: 10.1007/978-3-031-27481-7_12

Please contact:

Kim Guldstrand Larsen
Aalborg University, Denmark
kgl@cs.aau.dk

Max Tschaikowski
Sapienza University of Rome, Italy
max.tschaikowski@uniroma1.it

Cold Light in Orbit: A Photonic Path to Quantum Computing in Space

by Andreas M. Hein (SnT, University of Luxembourg) and
Carl Shneider (SnT, University of Luxembourg)

Orbital data centres are moving from vision to hardware, yet almost every concept still assumes conventional CMOS hardware. We argue that a cryogenically cooled photonic satellite constellation is a promising near-term architecture for computing in orbit and the natural on-ramp to photonic quantum processing and quantum networking in space.

From vision to hardware

In little more than a year, orbital data centres have gone from a niche subject to a crowded field. A start-up flew a commercial graphics processor to orbit in 2025; Kepler Communications launched orbital cloud infrastructure in 2026; Google published its Project Suncatcher concept, a formation of satellites carrying conventional tensor processors linked by dense coherent optical inter-satellite links [1], [L2]. The Horizon Europe ASCEND study and a 2025 report by the European Space Policy Institute examined far larger installations as a lower-carbon alternative to terrestrial facilities.

Nearly all of these studies share one assumption: the computing payload is based on an off-the-shelf terrestrial data-centre rack. Our own exploration of more than 20,000 candidate orbital computing architectures suggests this is exactly where the hidden leverage is located [3]. It is the computing paradigm, not the solar array and not the radiator, that decides whether an orbital data centre will have a future in space. We evaluated the potential of in-memory computing, neuromorphic processors, temporal computing, photonic and optical computing, quantum processors, and longer-horizon reversible computing schemes. Among them, hybrid photonic-electronic hardware showed order-of-magnitude potential, because it avoids much of the radiation shielding mass that electronics demand and because vacuum is an unusually good medium for moving photons.

A cold photonic constellation

We selected one architecture out of that trade space that combines photonic with quantum computing: a cryogenically cooled photonic mesh satellite constellation (Figure 1). Each node carries a photonic matrix-multiplication accelerator, a Mach-Zehnder interferometer mesh or a microring weight bank in silicon nitride or thin-film lithium niobate. It is held at 70–80 K by combined passive and active cooling. Neighbouring nodes exchange activations over coherent free-space optical links, with no conversion back to electronics at the chip boundary. Indicative targets are 10 to 100 nodes, links of one to ten terabits per second across separations of 100 to 1000 km, and a constellation power budget of 10 to 100 kW.

The low temperatures in space, combined with the cryogenic cooling on board, suppress the thermal drift in phase shifters




Figure 1: Architecture of the cryo-cooled photonic mesh satellite constellation. Each satellite hosts a photonic matrix-multiplication node held at 70–80 K; nodes are interconnected by coherent optical inter-satellite links.

that limits photonic accelerators on the ground, while also lowering optical loss. The choice of 70–80 K rather than something colder is deliberate: it is reachable with single-stage Stirling or pulse-tube coolers and avoids the roughly tenfold power penalty of the 4 K regime.

Falsifiable, not aspirational

Cryogenics in orbit is expensive, so the project is built around claims that can be refuted. Our goal is to validate whether the constellation can deliver energy per useful inference at least twice lower than a room-temperature orbital baseline of equal throughput, counted after cooler efficiency, parasitic loads, vibration isolation and radiator area; whether phase-shifter drift stays below the accuracy threshold for 90% of an orbit without active recalibration; and whether coherent links sustain a terabit per second at representative inter-satellite distances. Project Suncatcher is treated as the primary benchmark rather than as a foil: cold photonics is only worth flying if its advantage survives the cooling bill.

Where quantum enters

A cryogenic, low-loss, optically interconnected node provides close to optimal conditions for photonic quantum information processing. Superconducting nanowire single-photon detectors, deterministic photon sources and quantum memories still need to be demonstrated in space. The cold stage, the optical packaging, the coherent inter-satellite links and the acquisition and pointing technologies have already been operated in space. A constellation built for classical optical inference in the early 2030s seems a credible pathway towards photonic quantum processors and distributed quantum networking as those components mature towards 2045.

Method, partners and outputs

We intend to characterise and size the computing payload, including the full cooling chain; build an integrated spacecraft and constellation model; and explore the trade space against

terrestrial photonic accelerators, room-temperature orbital alternatives and conventional data centres.

The work builds on existing collaborations with TU Berlin, which leads the OMLO work on optical computing for machine learning in orbit [2], and the Silesian University of Technology, working on quantum computing for space applications.

Links:

- [L1] Space Systems Engineering (SpaSys) group, SnT, University of Luxembourg:
<https://www.uni.lu/snt-en/research-groups/spasys/>
- [L2] Google Research, “Exploring a space-based, scalable AI infrastructure system design”:
<https://research.google/blog/exploring-a-space-based-scalable-ai-infrastructure-system-design/>

References:

- [1] B. Agüera y Arcas et al., “Towards a future space-based, highly scalable AI infrastructure system design”, arXiv:2511.19468, 2025.
- [2] F. Kübler et al., “Concept of an in-orbit AI-system based on optical computing”, CEAS Space J., vol. 18, 2026.
- [3] A. M. Hein et al., “Orbital Data Centers: A Techno-Economic and Sustainability Analysis”, forthcoming, 2026.

Please contact:

Andreas M. Hein
SnT, University of Luxembourg, Luxembourg
andreas.hein@uni.lu

European Research and Innovation

Crete Looks to the Stars: SPACE-Crete Secures European Excellence Funding

by Panagiotis Tsakalides (FORTH-ICS) and Vassilis Charmandaris (FORTH)

The Foundation for Research and Technology – Hellas (FORTH) will establish the Cretan Center of Excellence in Emerging Space Technologies, structured around intelligent space systems, optical and quantum communications, and quantum sensors. The €35 million Horizon Europe Teaming project ranked 2nd among 61 proposals from across Europe.

The SPACE-Crete proposal, coordinated by Panagiotis Tsakalides, Professor at the Department of Computer Science of the University of Crete and Director of the Institute of Computer Science at FORTH, has been selected under the European call HORIZON-WIDERA-2025-ACCESS-01-01: Teaming for Excellence. With a total budget of €35 million over six years, it ranked second among 61 proposals from across Europe and first among the Greek submissions. The project starts on 1 January 2027.

The SPACE-Crete project will establish the Cretan Center of Excellence in Emerging Space Technologies (CCEST), which will operate as an independent entity with the long-term goal of financial sustainability. Leveraging the combined expertise of FORTH's Institutes of Computer Science (ICS), Astrophysics (IA), and Electronic Structure and Laser (IESL), CCEST will be structured around three closely interconnected research pillars: intelligent space systems, optical and quantum communications, and quantum sensors.

Focusing on space communications, space surveillance and monitoring, and Earth observation, the Center will, for the first time in Southeastern Europe, create a fully integrated space ecosystem. This ecosystem will span the entire innovation spectrum, from fundamental research and new ESA-certified infrastructures, including the upgraded Skinakas Observatory, to technology maturation and industrial uptake.

Strategic Partnerships

This vision will be realized through strategic collaboration with three leading European space research organizations. CEA, France's largest public research organization, participates through the Space Architecture & Systems and CosmoStat laboratories. DLR, Germany's leading aerospace research center, contributes expertise in quantum technologies and satellite navigation. In addition, OHB-Hellas, a member of the German OHB SE group, the third-largest space technology company in Europe, participates as a key industrial partner. At the national level, the University of Crete plays a central role, leading the educational mission of CCEST. Through its School of Sciences and its close collaboration with FORTH, it will serve as the core for the Center's graduate programs, helping



Skinakas Observatory in Crete, which will be upgraded as part of SPACE-Crete to support optical and quantum communications and ground-to-satellite links. © Skinakas Observatory (University of Crete / FORTH)

train the next generation of space scientists and engineers on the island.

Impact for Greece and Crete

Greece has not yet secured a position in the European space economy that matches its research base. CCEST aims to change this reality by addressing a structural gap that has held back Southeastern Europe for decades. At the national level, CCEST will act as a central coordination hub for emerging space technologies, directly supporting the country's space strategy and complementing national investments in microsatellites and ground stations. As a trusted partner of ESA, it will strengthen Greece's participation in the European space program, expanding opportunities for involvement in large-scale European initiatives.

For Crete, the impact will be direct. The Center is expected to create more than 120 highly skilled jobs, attracting both Greek researchers from the diaspora and international scientific talent. The upgraded Skinakas Observatory will evolve into a leading European infrastructure for developing optical and quantum communication technologies, with strong prospects of becoming the country's main operational hub for ground-to-satellite telecommunications. Our vision is to transform Crete into the 'SPACE Island of Southeastern Europe'. At the same time, a dedicated innovation hub will foster the growth of deep-tech startups, integrating space technologies into Crete's economic development strategy.

Education and human capital development are at the core of CCEST's mission. The Center will offer a comprehensive range of educational opportunities, from professional training and lifelong learning programs for students and professionals to international Master's and PhD programs. Together, these initiatives will cultivate a dynamic ecosystem of researchers and entrepreneurs, supporting the growth of the Greek space sector and helping reverse the long-standing brain drain from

Crete and Greece. SPACE-Crete builds on our scientific legacy and expands it.

SPACE-Crete and FORTH acknowledge the stakeholders whose co-investment, alongside European funding, makes the realization of CCEST possible: the Ministry of Development, the Ministry of Digital Governance, the Region of Crete, and Motor Oil Hellas. Combined with European funding, the total investment in CCEST amounts to €35 million, making it one of the largest research and innovation investments undertaken in Crete. This co-investment confirms that space technologies are now a strategic priority for Greece and Crete.

For the ERCIM community, CCEST will open access to ESA-grade testing facilities and to an optical ground station testbed at Skinakas. FORTH welcomes collaboration with ERCIM member institutes in computational sensing, optical and quantum communications, and intelligent space systems.

Links:

[L1] <https://www.space-crete.com/>

[L2] <https://skinakas.physics.uoc.gr/en/home/>

Please contact:

Panagiotis Tsakalides

Institute of Computer Science, FORTH, Greece

tsakalid@ics.forth.gr

SYNAPSE: An Integrated Platform for Cyber Risk and Resilience Management

by Massimo Cossentino, Giuseppe Papuzzo, Vishal Sonth and Fabio Martinelli (ICAR-CNR)

Cyberattacks against Europe's digital infrastructures are growing in scale and sophistication, yet organisations still manage technical risk, situational awareness, incident response and cyber-insurance in separate silos. The Horizon Europe project SYNAPSE integrates all of these into a single platform, combining AI-driven threat intelligence, standardised incident-response orchestration and blockchain-based cyber-insurance.

Operators of essential services (OES) and their supply chains face a rapidly expanding threat landscape, while the NIS2 Directive requires continuous, evidence-based cybersecurity and resilience management. In practice, most organisations – especially SMEs – still assess technical vulnerabilities, business impact and insurance exposure through disconnected tools: threat intelligence is rarely correlated with an organisation's own assets, incident-response procedures are often static rather than executable, and cyber-insurance is still largely underwritten through questionnaires rather than through continuous monitoring of security posture.

SYNAPSE is a three-year Horizon Europe project (Grant Agreement No. 101120853) coordinated by CNR, bringing to-

gether 15 partners from 8 countries. Its goal is to integrate technical and economic risk assessment, AI-enhanced situational awareness, coordinated incident response, preparedness training and cyber-insurance enablement into a single platform, validated at Technology Readiness Level 7 across four real-world pilots – energy (Germany), a healthcare supply chain (Cyprus), a cybersecurity insurer, and a fourth pilot selected through the project's open call.

The SYNAPSE architecture (Figure 1) integrates its capabilities as interoperable building blocks. A Multi-layer Security & Privacy Monitoring toolset supplies raw evidence to Integrated Risk Management, which combines technical risk assessment with a model-driven economic risk assessment – translating technical findings into asset valuation, impact cost and insurance-relevant figures, as detailed in Figure 2 – and to SYNAPSE Smart Contracts, which use the Hyperledger Fabric blockchain to support dynamic, evidence-based cyber-insurance policies and claims. The same evidence feeds AI-enhanced Situational Awareness, whose outputs trigger Incident Response & Business Continuity, centred on a Cybersecurity Orchestration, Automation & Response (COAR) engine that executes machine-readable “SYNAPSE Playbooks”, complemented by a Cyber Range that lets teams rehearse the same playbooks used in real incidents as part of hands-on Preparedness training. An Information Exchange, Alerting & Reporting layer lets SYNAPSE instances at different organisations and national cybersecurity authorities share playbooks, models and indicators, enabling shared situational awareness and coordinated response rather than single-organisation defence alone. A role-based dashboard presents the relevant subset of this information to operators, analysts and insurance experts.




Figure 1: SYNAPSE conceptual architecture.




Figure 2: The SYNAPSE approach to integrated cybersecurity risk & resilience management.

Several components push beyond the state of the art. An AI-based CTI Extractor turns unstructured threat reports into structured STIX 2.1 intelligence while preserving a traceable link back to the source passage, so analysts can verify or correct AI-generated content rather than trusting it blindly. A companion Threat Hunting Engine detects malicious infrastructure ahead of reactive feeds — in longitudinal validation, roughly a third of the infrastructure it flagged remained unseen by mainstream sources three days later, with a false-positive rate of 0.31%. SYNAPSE researchers have also published peer-reviewed results on LLMs for security triage: one study achieved up to 97% accuracy inferring phishing intent from email content alone [2]; another shows a constrained, multi-role “agentic” LLM workflow raises malicious-alert classification accuracy from near 0% to over 90% [3]. On the response side, SYNAPSE’s COAR engine executes playbooks built on the OASIS CACAO standard [1]. To the consortium’s knowledge, it is the first fully open orchestration engine of its kind, allowing response procedures to be automated, shared and rehearsed via the Cyber Range across organisations.

By combining technical and economic risk assessment, AI-driven threat intelligence, standards-based response orchestration and cyber-insurance enablement in one federatable platform, SYNAPSE aims to help operators of essential services meet the resilience expectations of NIS2 and strengthen the collective resilience of Europe’s digital infrastructures, contributing to a more autonomous European cybersecurity capacity. Updates and public deliverables are available on the project website [L1].

Links:

[L1] <https://www.synapse-project.eu>

References:

- [1] B. Jordan, A. Thomson, “CACAO Security Playbooks Version 1.0”, OASIS Committee Specification 02, <https://docs.oasis-open.org/cacao/security-playbooks/v1.0/security-playbooks-v1.0.html>
- [2] E. Eilertsen, V. Mavroeidis, G. Grov, “LLM-Powered Intent-Based Categorization of Phishing Emails”, IEEE CSR 2025.
- [3] E. Eilertsen, V. Mavroeidis, G. Grov, “Towards Agentic Investigation of Security Alerts”, IEEE BigData 2025.

Please contact:

Massimo Cossentino
CNR-ICAR, Italy
massimo.cossentino@icar.cnr.it

OpenSwarm: Coordinated Smart Devices for a Sustainable Future

by Giulia Petrarulo, Filip Maksimovic and Thomas Watteyne (Inria)

OpenSwarm, an EU-funded project coordinated by Inria that concluded in April 2026, has developed and validated a new open-source approach for enabling large groups of small, low-power wireless devices to work together intelligently, efficiently and autonomously.

Running from January 2023 to April 2026, OpenSwarm addressed a fundamental challenge in large-scale swarm systems: how to enable highly constrained devices to communicate, collaborate and make use of artificial intelligence (AI), while minimising energy consumption and simplifying the way such systems are programmed and managed. The project's results now provide a set of reusable technologies for researchers and developers working on energy-aware swarms of collaborative smart devices.

OpenSwarm brought together advances across three closely connected areas: efficient networking and management of smart nodes; collaborative, energy-aware AI; and energy-aware programming of swarms. Together, these contributions move beyond conventional approaches in which connected devices largely operate as isolated units or depend on centralised infrastructure.

Scientific advances enabling new capabilities

The project's key scientific contribution is an integrated approach to building collaborative smart nodes: small, resource-constrained devices that can interpret the data they generate, share information with neighbouring devices and collectively adapt their behaviour. Before OpenSwarm, achieving this combination at large scale was difficult because the capabilities required for intelligent collaboration (e.g. wireless communication, localisation, etc.) typically compete for limited power. OpenSwarm demonstrated that these capabilities can be coordinated as part of a common, energy-aware architecture.

A further achievement was demonstrating these capabilities at large scale for constrained swarm devices. The technologies were validated on two large laboratory testbeds: the 900-node KiloBot testbed, which assessed the versatility and portability of the framework, including operation on severely constrained robots, and the 725-node DotBot testbed, which focused on low-power networking, energy management and localisation (Figure 1).

From research results to reusable technologies

The technologies developed were integrated into the OpenSwarm "package", an open-source architecture underpinning all swarm capabilities and enablers, including software and other documents. Through this single reference point, the project's results can be easily re-used by researchers and developers interested in energy-aware swarms of collaborative smart nodes, serving as a basis for further advances in this type of technologies.

The project's results were tested and demonstrated through five real-world Proofs of Concept covering four broad applica-



Figure 1: The OpenSwarm testbed of 725 DotBots.



Figure 2: Demonstration at the Swarm Workshop in Brussels, 26–27 November 2025.

tion areas: cities and communities, environmental monitoring, industrial and occupational safety, and mobility.

Taken together, these demonstrations showed the versatility of the OpenSwarm approach across very different environments and operational requirements. The use cases explored how collaborative, energy-aware swarms could support renewable energy communities, assist human workers during harvesting, monitor ocean noise pollution, improve health and safety in industrial production environments, and operate in moving networks such as trains. Videos presenting the features and outcomes of each Proof of Concept are available on the project website [L1].

Collaboration turning into impact

OpenSwarm also established a coherent framework for coordinating a heterogeneous software ecosystem across its partners, supporting traceability across partners' developments. This collaborative approach helped the consortium meet its initial scientific and outreach targets and define additional indicators to assess the performance and outcomes of its demonstrations more precisely.

The project resulted in 22 journal articles in open-access journals, 62 conference papers at leading venues and four awards. Beyond publications, partners contributed to standardisation, education, open science and outreach through activities including hackathons, workshops and roadshows. These efforts are intended to extend the project's impact and support the continued adoption and development of its results (Figure 2).

A foundation for the next generation of smart nodes

OpenSwarm's results point towards a next generation of distributed systems in which connected devices are equipped with intelligence to interpret the data they generate, collaborate in a decentralized manner, and communicate efficiently.

By demonstrating these capabilities at scale, integrating them into an open-source architecture and validating them across diverse real-world scenarios, OpenSwarm has helped establish a foundation for future energy-efficient and trustworthy swarm systems. Its results contribute to broader European priorities around sustainability, clean energy and digital transformation, while providing researchers and developers with technologies they can continue to reuse and advance.

OpenSwarm brought together ten partners from Austria, Belgium, France, Germany, Ireland, Portugal and the UK, combining expertise in swarm robotics, IoT, low-power wireless communications, nanoelectronics and global Internet standards. The consortium included research leaders such as Inria, IMEC, KU Leuven and the University of Sheffield, market leaders including Analog Devices and Siemens, and SME Ingeniarius, which applies novel technologies to environmental protection.

This project has received funding from the European Union's Horizon Europe Framework Programme under Grant Agreement No. 101093046. Disclaimer: The content of this article reflects the views of the authors and not necessarily those of the European Commission.

Links:

[L1] <https://openswarm.eu/>

[L2] EU Cordis webpage:

<https://cordis.europa.eu/project/id/101093046>

References:

- [1] J. Argote-Gerald, et al., "Design for one, deploy for many: Navigating tree mazes with multiple agents", in Proc. IEEE Int. Symp. Multi-Robot & Multi-Agent Syst. (MRS), Singapore, Dec. 4–5, 2025, doi: 10.48550/arXiv.2510.26900.
- [2] S. Alvarado-Marin, et al., "Automatic lighthouse calibration using conics for indoor robot localization", IEEE Robot. Autom. Lett., vol. 10, 2025, doi: 10.1109/LRA.2025.3575319.
- [3] T. Avé, M. Hutsebaut-Buysse, and K. Mets, "Temporal distillation: Compressing a policy in space and time", Mach. Learn., vol. 114, 2025, doi: 10.1007/s10994-025-06889-9.

Please contact:

Giulia Petrarulo

Inria, France

giulia.petrarulo@inria.fr



SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Call for Proposals

Dagstuhl Seminars and Perspectives Workshops

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is accepting proposals for scientific seminars/workshops in all areas of computer science, in particular also in connection with other fields.

If accepted, the event will be hosted in the seclusion of Dagstuhl's own, well-known, dedicated facilities in Wadern on the western fringe of Germany. Moreover, the Dagstuhl office will assume most of the organisational/administrative work, and the Dagstuhl scientific staff will support the organizers in preparing, running, and documenting the event. Thanks to subsidies the costs are very low for participants.

Dagstuhl events are typically proposed by a group of three to four outstanding researchers of different affiliations. This organizer team should represent a range of research communities and reflect Dagstuhl's international orientation. More information, in particular details about event form and setup, as well as the proposal form and the proposing process, can be found on

<https://www.dagstuhl.de/dsproposal>

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is funded by the German federal and state government. It pursues a mission of furthering world-class research in computer science by facilitating communication and interaction between researchers.

Important Dates

- *Next submission period:*
15 October to 1 November 2026
- *Seminar dates:*
Between October 2027
and September 2028 (tentative).



Call for Participation

ERCIM Forum Beyond Compliance on Digital Ethics in Research

29–30 October 2026, Porto

The fifth edition of the ERCIM Forum Beyond Compliance, dedicated to research ethics in the digital age, will take place on 29–30 October 2026 at INESC TEC in Porto, Portugal.

Rather than focusing only on ethical norms and compliance requirements, the forum explores digital ethics issues as they arise in the practical settings of academic research. It provides an open and inclusive space for exchange with international experts and welcomes anyone concerned with digital ethics in their research activities.

The programme includes three thematic sessions and a keynote:

- **XAI Limitations & Vulnerability**
29 October, morning
What critical gaps remain for uncertainty-aware, meaningful and understandable use of explainable AI in healthcare and industry?
Speakers: Sepideh Pashami (RISE, Sweden) and Ajitha Rajan (Edinburgh University, UK)
- **Keynote on AI in Healthcare**
29 October, morning
Tugce Schmitt (Maastricht University, Netherlands)
- **The Need for National Digital Ethics Committees in Europe**
29 October, afternoon
The session will address the competences national ethics committees need in digital ethics, how committees with differing positions can coordinate, and how competing digital

ethics policies can be addressed without weakening European values. Speakers include Claude Kirchner (National Digital Ethics Committee, France), Vasiliki Mollaki (NCSR “Demokritos”, Greece), Frank Beck Lassen (National Center for Ethics, Denmark) and Yves Moreau (KU Leuven, Belgium).

• AI Companions and Vulnerable Publics

30 October, morning

How can their impacts and safety be researched while ensuring that the research itself is conducted safely? Speakers: Angelica Henestroza (THWS, Germany), Alexandra Prigent (CEA, France) and Jean Gové (Pontifical Gregorian University, Rome).

Participation is free of charge and in person only; no live streaming will be available. The forum will be held at INESC TEC in Porto and will be co-located with the ERCIM Fall Meeting and an AIOLIA training workshop on digital ethics on 28 October. It will follow the ECSS 2026 conference of Informatics Europe, which also includes a workshop on digital ethics. Most talks and summaries will be made available online after the event.

The forum is organised by ERCIM and its Working Group on Digital Ethics, INESC TEC, and the French National Digital Ethics Committee (CCNEN).

Registration:

<https://edu.nl/kvxjj>

Further information:

ERCIM Beyond Compliance:
<https://www.ercim.eu/science/beyond-compliance-digital-ethics-in-research>

ERCIM Fall Meeting:

<https://www.ercim.eu/events/ercim-days/2026-fall-meetings>

AIOLIA training workshop:

<https://www.inesctec.pt/en/events/training-in-digital-ethics>

Training in Digital Ethics

Registrations are now open for a one-day Training in Digital Ethics, organised within the Horizon Europe AIOLIA project (Artificial Intelligence in Human Cognition and Behaviour), in which ERCIM is a project partner. The training is delivered by INESC TEC, an ERCIM member institute, and explores the ethical challenges arising from AI systems that shape human decision-making and interaction.

The course will take place **28 October 2026** at INESC TEC in Porto, Portugal, and is intended for early-career researchers, PhD students, and professors working in computer science, artificial intelligence, and digital technologies. Through lectures, discussions, and practical exercises, participants will gain an understanding of the main ethical principles of trustworthy AI and learn how to identify and address ethical issues throughout the AI system life-cycle.

The afternoon session focuses on applying AI ethics in practice. Using real-world case studies developed within AIOLIA, participants will explore ethical trade-offs, for example, between safety, privacy, explainability, non-discrimination, and human autonomy, and examine practical measures for responsible AI design and governance.

The training will be led by Alexei Grinbaum and Alexandra Prégent (CEA-Saclay), together with Laurynas Adomaitis (RISE Research Institutes of Sweden), internationally recognised experts in AI ethics and governance. Participation is free of charge, but places are limited and registration is required. Researchers interested in integrating ethical reflection into AI research and development are encouraged to apply early.

More information:

<https://www.inesctec.pt/en/events/training-in-digital-ethics>

Van Wijngaarden Awards 2026 Honour Pioneers in Mathematics and Computer Science

CWI has announced the recipients of the 2026 Van Wijngaarden Awards: **Carola-Bibiane Schönlieb** of the University of Cambridge and **Mor Harchol-Balter** of Carnegie Mellon University. The awards recognise scientists whose work has made significant contributions to mathematics and computer science.

Schönlieb is Professor of Applied Mathematics at Cambridge and is known for her research in mathematical imaging and inverse problems. Her work combines partial differential equations, optimisation and machine learning to develop reliable methods for image reconstruction, with applications ranging from medical imaging and molecular biology to materials science and art conservation.

Harchol-Balter holds the Bruce J. Nelson Endowed Chair of Computer Science at Carnegie Mellon University. Her research uses applied mathematics, probability and queueing theory to improve the performance of computing systems, with particular emphasis on reducing response times, computing-resource use and energy consumption.

The Van Wijngaarden Awards are presented every five years by CWI and are named after former CWI director Adriaan van Wijngaarden, one of the pioneers of computer science in the Netherlands. The 2026 awards will be presented on 20 November in Amsterdam during the Van Wijngaarden Soirée, which this year will focus on Europe's AI potential and the role of fundamental research in shaping the future of AI. Former Italian Prime Minister Enrico Letta will give the opening keynote.

More information:

<https://www.cwi.nl/en/news/van-wijngaarden-awards-2026-honour-scientific-pioneers-and-put-dutch-ai-perspective-on-the-european-agenda/>

ERC Starting Grant to Advance Reliable Quantum Error Correction

Jonas Helsen of CWI has received a €1.5 million ERC Starting Grant to develop new methods for analysing and certifying fault-tolerant quantum computers.

Quantum computers could solve certain problems much more efficiently than conventional computers, but their basic information units, known as qubits, are highly susceptible to errors caused by environmental disturbances and imperfect control. As computations become longer and more complex, reliable quantum error correction becomes essential.

Recent progress in quantum hardware has enabled the first steps towards fault tolerant quantum computers, which are designed to detect and correct errors while computations are running. However, methods for assessing how well these error correction mechanisms perform are still developing.

With the five year ERC grant, Helsen, a researcher in CWI's Algorithms & Complexity group, will develop methods for reliably analysing fault tolerant quantum computers. The project will investigate where errors arise, for example in individual qubits, their interactions, or the error correction process, and how these errors can be detected efficiently.

A further goal is to develop algorithms for constructing reliable digital twins of quantum computers. These models could help researchers simulate and test quantum devices despite the difficulty of reproducing quantum behaviour on conventional computers.

The research builds on Helsen's earlier work on testing and characterising physical qubits and aims to provide practical tools for evaluating increasingly sophisticated fault tolerant quantum systems.

Link:

CWI Algorithms and Complexity group: <https://www.cwi.nl/en/research/algorithms-and-complexity/>

ERCIM “Alain Bensoussan” Fellowship Programme

The ERCIM Postdoctoral Fellowship Programme is one of the flagship initiatives of ERCIM. Open to young researchers from around the world, the programme covers a broad range of fields in computer science and applied mathematics.

The fellowship scheme aims to help young scientists deepen their knowledge of European research structures and networks, while gaining valuable experience within leading European research institutions. Fellowships have a duration of 12 months, with the possibility of extension, and are hosted by one of the ERCIM member institutes.

Hosting institutions

Only ERCIM members can host fellows. When an ERCIM member is a consortium, the hosting institute may be any of its member organisations. When an ERCIM member is a funding body, the hosting institute may be one of its affiliated institutions.

Fellowships are offered according to the needs and available funding of the member institutes. Fellows are appointed either through a stipend (a research training agreement) or a work contract, depending on the hosting institute. The type of contract and the

“ The ERCIM Fellowship was my first postdoctoral position and my first opportunity to develop and carry out an independent research proposal. Together with my supervisors, I shaped a project around a research direction I wanted to pursue, deepened my expertise at CWI, and matured as an independent researcher beyond the experience gained during my PhD. The fellowship also provided valuable insight into proposal writing, research funding, and the next steps of an academic career.



Nicolás Benjamín OCAMPO
Former ERCIM Fellow



monthly allowance or salary vary by host institution.

ERCIM encourages applications from researchers both in academia and in industry.

Why apply for an ERCIM Fellowship?

The Fellowship Programme enables talented early-career scientists from all over the world to work on challenging research problems alongside leading European experts. In addition to research excellence, the programme fosters collaboration and knowledge exchange within the European research community.

The programme offers ERCIM Fellows the opportunity to:

- Work with internationally recognised experts;
- Gain a deeper understanding of European research structures and networks;
- Become familiar with the working conditions in leading European research centres;

- Promote cross-fertilisation and cooperation between research groups working in similar areas across Europe.

Equal opportunities

ERCIM is committed to ensuring equal opportunities and promoting diversity. Candidates are not discriminated against on the basis of race, colour, religion, gender, national origin, age, marital status, or disability.

Conditions

Candidates must:

- Have obtained a PhD degree within the last eight years (prior to the application deadline), or be in the final year of their doctoral studies with an outstanding academic record. Proof of the PhD degree must be provided before the start of the fellowship;
- Be fluent in English.

Application deadlines

Applications are accepted twice a year, with deadlines on 30 April and 30 September.

Since its inception in 1991, more than 800 fellows have participated in the programme.

The Fellowship Programme is named in honour of Alain Bensoussan, former president of Inria, one of the three founding institutes of ERCIM.


<https://fellowship.ercim.eu>

“ The ERCIM Fellowship was an excellent opportunity to develop my research independently while collaborating with leading researchers in an international environment. The programme provided strong support for scientific growth, networking, and career development, enabling me to expand my expertise and build valuable collaborations. I highly recommend the ERCIM Fellowship to researchers seeking to broaden their horizons and advance their careers in Europe and beyond.



Adebayo Abayomi-Alli
Former ERCIM Fellow





Showcase your EU Project in **ERCIM News**

Reach up to **10,000 readers**
in Europe's ICT research
and innovation community.

ERCIM News is the quarterly magazine of ERCIM,
connecting researchers, innovators, technology leaders
and policy makers across Europe and beyond.

WHY PUBLISH A SPONSORED CONTRIBUTION?



Visibility in a trusted European
research publication



Reach a specialised audience of
researchers, innovators and decision
makers, and stakeholders in the
European Commission, European
Parliament and national
administrations



Distribution through print and digital
channels



Receive a ready-to-use
dissemination text for Horizon
Europe reporting



Long-term online visibility through
the ERCIM News website

SPONSORED CONTRIBUTION

1000 €

Includes:

- Publication in ERCIM News (print and online)
- Professional editorial presentation
- Permanent online availability
- Social media promotion
- Dissemination-ready reporting

REPORTING-READY TEXT

“ The project was featured in ERCIM News,
reaching researchers, innovators and
technology stakeholders across Europe
through print and online dissemination
channels. ”

CONTACT

Peter Kunz
Editor, ERCIM News
peter.kunz@ercim.eu
ercim-news.ercim.eu



ERCIM – the European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development in information technology and applied mathematics. Its member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.



ERCIM is the European Partner of the World Wide Web Consortium.



Consiglio Nazionale delle Ricerche
Area della Ricerca CNR di Pisa
Via G. Moruzzi 1, 56124 Pisa, Italy
www.iit.cnr.it



I.S.I. – Industrial Systems Institute
Patras Science Park building
Platani, Patras, Greece, GR-26504
www.isi.gr



Centrum Wiskunde & Informatica

Centrum Wiskunde & Informatica
Science Park 123,
NL-1098 XG Amsterdam, The Netherlands
www.cwi.nl



Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and Electrical Engineering, N 7491 Trondheim, Norway
<http://www.ntnu.no/>



Fonds National de la
Recherche Luxembourg

Fonds National de la Recherche
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
www.fnrl.lu



RISE SICS
Box 1263,
SE-164 29 Kista, Sweden
<http://www.ri.se/>



Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
www.ics.forth.gr



SBA Research gGmbH
Floragasse 7, 1040 Wien, Austria
www.sba-research.org/



Fraunhofer ICT Group
Anna-Louisa-Karsch-Str. 2
10178 Berlin, Germany
www.iuk.fraunhofer.de



Eötvös Loránd Research Network
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
www.sztaki.hu/



INESC
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, n° 378,
4200-465 Porto, Portugal
www.inesc.pt



University of Cyprus
P.O. Box 20537
1678 Nicosia, Cyprus
www.cs.ucy.ac.cy/



Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
www.inria.fr



UNIVERSIDAD DE MÁLAGA

Institute for Software Engineering and Software Technology
“José María Troya Linero”, University of Málaga
Calle Arquitecto Francisco Peñalosa, 18, 29010 Málaga
<https://gp.uma.es/itis>