INTERNATIONAL CONFERENCE ON QUANTUM COMPUTING

Institut Henri Poincaré, 11 rue Pierre et Marie Curie, 75005 PARIS

12 - 16 May 2025

FRANCE 2030 PROGRAMME DE RECHERCHE QUANTIQUE

QuanTiP Domaine de Recherche et d'Innovation Majeur

QUANDELA

Région îledeFrance

anr Inria

ALICE & BOB

cea cnrs

ihp Institut Henri Poincaré

C12

Detailed Program & Abstracts

Monday 12/05/2025

Welcome and Introduction — *Mo 13:30-14:00*

TBA (invited)

Natalie Brown (Quantinuum) — Mo 14:00-15:00

On the role of coherence for quantum computational advantage

Hugo Thomas et al. (Quandela, LIP6, ENS, inria, Edinburgh) — Mo 15:00-15:30

Quantifying the resources available to a quantum computer appears to be necessary to separate quantum from classical computation. Among them, entanglement, magic and coherence are arguably of great significance. We introduce path coherence as a measure of the coherent paths interferences arising in a quantum computation. Leveraging the sum-over-paths formalism, we obtain a classical algorithm for estimating quantum transition amplitudes, the complexity of which scales with path coherence. As path coherence relates to the hardness of classical simulation, it provides a new perspective on the role of coherence in quantum computational advantage. Beyond their fundamental significance, our results have practical applications for simulating large classes of quantum computations with classical computers.

Pauli Propagation: a computational framework for simulating noisy and noiseless quantum circuits

Armando Angrisani (EPFL) & Victor Martinez (ENS Lyon, IBM France) — Mo 16:00-16:45

Understanding the capabilities of classical simulation methods is essential for identifying scenarios where quantum computers offer a genuine advantage. This not only ensures that quantum resources are deployed efficiently — only when truly necessary — but also opens the door to offloading certain subroutines onto classical hardware.
Motivated by these considerations, our recent works develop classical algorithms for simulating and surrogating quantum circuits. Our approach is grounded in *Pauli Propagation*, a computational framework that leverages Pauli-path methods within the Heisenberg picture. While previous works on Pauli-path simulation have primarily targeted circuits subject to depolarizing noise, we significantly broaden the scope by extending classical simulability to both noiseless circuits and those affected by arbitrary local noise, including non-unital channels.
Taken together, our results demonstrate that Pauli Propagation is a powerful and scalable framework for the classical simulation and surrogation of quantum circuits. In contrast to other techniques — such as stabilizer simulation or tensor network approaches — Pauli Propagation can maintain tractable computational costs even in the presence of high-dimensional or chaotic quantum dynamics.

Universal algorithm for transforming Hamiltonian eigenvalues

Mio Murao (Tokyo University) — Mo 16:45-17:15

In this work, we provide a new way of manipulating Hamiltonians, by transforming their eigenvalues while keeping their eigenstates unchanged. We develop a universal algorithm that deterministically implements any desired (suitably differentiable) function on the eigenvalues of any unknown Hamiltonian, whose positive-time and negative-time dynamics are given as a black box. Our algorithm uses correlated randomness to efficiently combine two subroutines---namely controlization and Fourier series simulation---exemplifying a general compilation procedure that we develop. The time complexity of our algorithm is significantly reduced via said compilation technique compared to a naïve concatenation of the subroutines and outperforms similar methods based on the quantum singular value transformation.

Roundtable featuring session speakers — *Mo 17:15-17:45*

Tuesday 13/05/2025

Harnessing Schrödinger cat states for quantum computing and metrology in bosonic cQED (invited)

Yvonne Gao (National University of Singapore) — Tue 9:00-10:00

Cat states, with their distinctive phase-space interference features, are promising candidates for a wide range of quantum information processing tasks. In this talk, I will present our recent work on manipulating the non-Gaussian characteristics of mesoscopic cat states residing in superconducting cavities, aimed at achieving protection against photon loss and enabling quantum-enhanced parameter estimation. These demonstrations are carried out using small yet highly versatile bosonic cQED devices, consisting of only a single cavity coupled to a nonlinear ancillary transmon. Our results highlight the rich dynamics of bosonic cat states and underscore their significant potential in quantum information science.

Individual solid-state nuclear spin qubits with coherence exceeding seconds

James O'Sullivan (CEA) — Tue 10:00-10:30

I will present a new platform for quantum information processing, consisting of ^{183}W nuclear spin qubits adjacent to an Er^{3+} impurity in a CaWO4 crystal, interfaced via a superconducting resonator and detected using a microwave photon counter at 10 mK. We study two nuclear spin qubits with T_2* of 0.8(2) s and 1.2(3) s, T_2 of 3.4(4) s and 4.4(6) s, respectively. We demonstrate single-shot quantum non-demolition readout of each nuclear spin qubit using the Er^{3+} spin as an ancilla. We introduce a new scheme for all-microwave single- and two-qubit gates, based on stimulated Raman driving of the coupled electron-nuclear spin system. We realize single- and two-qubit gates on a timescale of a few milliseconds, and prepare a decoherence-protected Bell state with 88% fidelity and T_2* of 1.7(2) s. I will also present recent results investigating a spin-9/2 niobium nucleus in the same $CaWO_4$ crystal, controlled and read out via an Er ancilla.

Quantum cellular automata: structure & quantum simulation of QED

Pablo Arrighi (INRIA Saclay) — Tue 11:00-11:30

Quantum cellular automata (QCA) consist in arrays of identical finite-dimensional quantum systems, evolving in discrete-time steps by iterating a finite-depth, translation-invariant quantum circuit. I will give an overview of their mathematical structure, showing that all non-signalling unitary operators in discrete space are of that form. I will then showcase what quantum cellular automata can achieve in terms of reformulation and quantum simulation of quantum field theories. In particular, I will explain the core ideas behind QCA-based quantum algorithms for the quantum simulation of the Dirac Equation, the Dirac Equations in Curved spacetime, Fermions as qubits even in 2D and 3D, and quantum electrodynamics.

Auxiliary-assisted cooling of many-body systems via a stochastic mechanism

George Mouloudakis et al. (FU Berlin) — Tue 11:30-12:00

Many important optimization tasks can be mapped to the search of ground states of effective Hamiltonians, allowing for the solution of computationally hard problems in physics, finance, logistics and other fields. In quantum annealing, an initial easy-to-prepare Hamiltonian, cooled down to its ground state, gets gradually changed towards a more complex one, whose ground state represents the solution to a specific optimization problem. Despite its evident success, quantum annealing faces the challenge of diabatic transitions that tend to excite the systems towards energetically higher states, ruining the performance of adiabatic protocols. In this work we propose a new mechanism to cool down diabatic transitions in many-body chains by coupling one of their ends to an auxiliary qubit that is frequently being reset to its ground state. In contrast to previously existing protocols that rely on knowledge of the time-dependent spectrum, our protocol relies on stochastic choices of the auxiliary qubit's energy at each reset. Using our protocol, we demonstrate the ability to effectively cool static chains down to their ground states, independent of their initial state or size. Our results are also generalized to time-dependent Hamiltonians, showing the efficiency of our protocol in improving the performance of several quantum annealing tasks.

Roundtable featuring session speakers, Paris quantum company C12 — *Tue 12:00-12:30*

Composing Quantum Programs (invited)

Stacey Jeffery (CWI Amsterdam) — Tue 14:00-15:00

Program composition is something we take for granted in designing classical algorithms, but when quantum subroutines are called in superposition, analyzing the complexity of the resulting program is not straightforward. It turns out that by using the right quantum program paradigm, a superposition of different length programs can be implemented in their average cost, just as you would expect when you call programs at random classically. For quantum programs, we can even do something that is not possible with classical programs: composed bounded-error programs without incurring log factors.

Efficiency bounds of the quantum simulation of the nonlinear Burgers' equation

Henri Pinsolle, Eric Cancés, Daniel Huerga, Florent Renac (ENPC, ONERA) — Tue 15:00-15:30

Numerical solutions of partial differential equations (PDEs) are generically very resource-intensive and quantum algorithmic strategies may provide an alternative. A recently proposed approach consists in mapping a linear PDE to a Schrödinger-type equation amenable to quantum simulation *[S. Jin, N. Liu, Y. Yu, Phys. Rev. A 108, 032603 (2023)]*. In this talk, we will consider the quantum simulation of the one dimensional nonlinear Burgers' equation, a toy model of non-compressible fluid dynamics hosting shock waves, which can be linearised through the Cole-Hopf transformation. We propose a numerical analysis allowing us to derive efficiency bounds preserving the correctness of the solution, even after the nonlinear mapping. This analysis can be applied to extract efficiency bounds of other nonlinear PDEs related to the Burgers' equation.

Quantum amplitude estimation from classical signal processing

Farrokh Labib et al. (Unitary Foundation, BQP, Phasecraft, Quantonation) — Tue 16:00-16:30

We demonstrate that the problem of amplitude estimation can be mapped directly to a problem in signal processing called direction of arrival (DOA) estimation. The DOA task is to determine the direction of arrival of an incoming wave with the fewest possible measurements. The connection between amplitude estimation and DOA allows us to make use of the vast amount of signal processing algorithms to post-process the measurements of the Grover iterator at predefined depths. Using an off-the-shelf DOA algorithm called ESPRIT together with a compressed-sensing based sampling approach, we create a phase-estimation free, parallel quantum amplitude estimation algorithm with a worst-case sequential query complexity of $\sim 4.2/\varepsilon$ and a parallel query complexity of $\sim 0.26/\varepsilon$ at 95% confidence. This performance is statistically equivalent and a 15x improvement over previous state of the art, for sequential and parallel query complexity respectively, which to our knowledge is the best published result for amplitude estimation.
The approach presented here provides a simple, robust, parallel method to performing QAE, with many possible avenues for improvement borrowing ideas from the wealth of literature in classical signal processing.

Roundtable featuring session speakers — *Tue 16:30-17:00*

Poster Session 1.A — *Tue 17:00-17:45*

Benchmarking Quantum Computers: Towards a Standard Performance Evaluation Approach

Rubén Peña Guzmán et al. (U.Bilbao - Basque Country, Spain)

The technological development of increasingly larger quantum processors on different quantum platforms raises the problem of how to fairly compare their performance, known as quantum benchmarking of quantum processors. In this work, we briefly review the most important aspects of both classical processor benchmarks and the metrics comprising them, providing precise definitions and analyzing the quality attributes that they should exhibit. Additionally, we review some of the most important metrics and benchmarks for quantum processors proposed in the literature, assessing what quality attributes they fulfill. Finally, we propose general guidelines for quantum benchmarking and its associated test suite.

Quantum walk search and Schwinger model implementation via QCA over NISQ cQED processor

Andrea Mammola et al. (C12 ; U.Aix-Marseille)

Demonstrating quantum advantage in the NISQ era requires resource-efficient, hardware-tailored implementations of meaningful algorithms. We present two such implementations—Quantum Walk (search) and the Schwinger model—via Quantum Cellular Automata (QCA) on C12 Quantum Electronics' circuit QED (cQED) hardware. QCA's close-to-physics definition aligns naturally with NISQ constraints, and C12's processor architecture natively supports required unitaries and allows for all-to-all connectivity. Using C12's noisy emulator Callisto, we simulate QCA-based Quantum Walk (search) on cycle and torus graphs, benchmarking their performance against standard approaches. These results highlight the potential of the QCA framework and cQED hardware for practical NISQ-era quantum dynamics and field theory simulations.

Combinatorial optimization with Rydberg atoms: the barrier of interpretability

Christian de Correc, Thomas Ayral, Corentin Bertrand (Eviden)

Analog quantum computing with Rydberg atoms can solve hard combinatorial problems like the Maximum Independent Set (MIS) on Unit Disk (UD) graphs. Generalizing beyond UD-MIS requires embeddings with ancilla qubits, which are only guaranteed to work if the embedded problem is solved exactly. We show, both numerically and analytically, that approximate solutions—common under imperfect annealing—rarely map back to valid solutions of the original problem. This challenges the reliability of current embedding schemes, such as the crossing lattice, under realistic conditions. Our findings stress the need for embeddings that remain meaningful even when annealing is not ideal.

Quantum error correction beyond Pauli noise : Coding bases sparse representation

Thomas Tuloup (Eviden)

We present a method of simulating quantum error correcting codes under non-Pauli noises based on a code-basis sparse representation of the state vector. Our approach represents quantum states in a basis defined by the common eigenvectors of the stabilizers of the code, along with the logical operator's eigenvector the state is prepared in. This sparse representation significantly reduces computational overhead while maintaining the ability to model arbitrary noise channels. Our simulator enables Monte Carlo stochastic simulations to compute error correction thresholds for the rotated surface code under different non-Pauli idling noise models up to distance d=11, allowing us to analyze the performance of large-scale quantum codes under complex noise conditions that more closely reflect experimental reality.

Digital controllability and resource scaling in frustrated Ising models: from QAOA to continuous-time quantum annealing

Ruiyi Wang, Vincenzo Roberto Arezzo, Kiran Thengil, Giovanni Pecci, Giuseppe Santoro
(SISSA Trieste)

Exponentially small spectral gaps are commonly regarded as a major limitation for quantum optimization algorithms, especially in adiabatic and variational settings. In this work, we investigate a frustrated Ising ring—first introduced by Roberts et al. (PRA 2020)—which features such a gap due to a single antiferromagnetic bond. Surprisingly, we demonstrate that the system remains digitally controllable via the Quantum Approximate Optimization Algorithm (QAOA), with the number of variational layers scaling only quadratically with the system size. We trace this efficient performance to the free-fermionic structure of the model and provide evidence that similar behavior holds for generic free-fermion Hamiltonians. This suggests that small spectral gaps do not necessarily entail exponential overhead in gate-based variational protocols. Finally, we explore continuous-time evolution for the same model, establishing connections between digital controllability and diabatic annealing strategies.

Error correctable microwave dual rail qubits

Vahid Shaghaghi, Iivari Pietikainen, Radim Filip, Steven Girvin (1), Ondrej Cernotik
(Palacky University Olomouc ; (1) Yale University)

A dual-rail qubit, where information is encoded in the single-photon subspace of two superconducting microwave cavities, has been shown to combine simple gate design and better scaling performance for surface code than conventional transmon qubits. However, having a common error state in this protocol introduces erasure, which cannot be corrected on a single physical qubit. We analyse a four-photon variant of the dual-rail qubit to overcome this obstacle and benefit from the linear nature of cavity modes. This encoding uses the same hardware as the single-photon variant, with one of the cavities coupled dispersively to a transmon ancilla and a beam-splitter coupler between the cavities. It enables photon-loss errors to be detected and corrected. We provide proposals and detailed numerical simulations for state preparation and measurement, single-qubit and two-qubit gates, and error correction. With these tools in hand, the prospects for fault-tolerant operation and reaching the break-even point will be discussed.

Dequantization and Expressivity in Photonic Quantum Fourier Models

Hugo Thomas et al. (Quandela, LIP6, ENS, Edinburgh, EPFL)

In this work, we study the models emerging from linear optical circuits and their augmented version with non-linearity, such as feedforward adaptivity or state injection. These architectures are keen to be used for near term application of quantum computing on learning tasks. We show that they are described by Fourier-type sums, and this allows us to investigate the classical simulation of these models with random Fourier methods.

Poster Session 1.B — *Tue 17:45-18:30*

Towards the Entanglement of a Superconducting qubit to Erbium defects

Kritika Mundeja et al. (National University of Singapore)

Superconducting circuits are promising candidates for building quantum processors due to their scalability, design flexibility, fast gate operations and high fidelity. However, their operation at microwave frequencies presents challenges for integration with existing telecommunication infrastructure which relies on optical fibres for long-distance signal transmission at room temperature. To bridge this frequency gap, we aim to achieve fast and reliable optical-microwave entanglement using rare-earth ions as an intermediary.
Our initial focus is on establishing entanglement between an Erbium ion ensemble and a Transmon qubit through a cavity. We leverage the Zeeman splitting of Erbium ions to align with the microwave regime of our superconducting cavity. Additionally, we employ chirped pulses to efficiently retrieve the excitation stored within the ensemble. Through this poster, we present our entanglement protocol and results, demonstrating spin-resonator cooperativity, long spin coherence times and high Transmon qubit coherence in a magnetic field. Our work represents a significant step toward the seamless integration of superconducting circuits with optical quantum networks, paving the way for scalable quantum communication and hybrid quantum computing architectures.

SFFT-based Homogenization: Using Tensor Trains to Enhance FFT-based Homogenization

Sascha Hauck, Matthias Kabel, Ali Mazen, Nicolas Gauger
(Fraunhofer Institute ; Kaiserslautern University ; Multiverse Computing)

Homogenization is a key technique for approximating the macroscopic properties of materials with microscale heterogeneity. The FFT-based Homogenization method has gained widespread usage due to its computational efficiency and accuracy in handling complex microstructures. However, despite its advantages, the method is limited by speed and memory constraints, particularly when applied to high-resolution discretizations. These limitations affect its scalability and efficiency, especially in large-scale simulations or when dealing with highly detailed microstructures. These challenges arise from the fundamental reliance on the Fast Fourier Transform, which imposes inherent restrictions on further advancements. In this paper, we propose a novel SFFT-based Homogenization algorithm that utilizes a Quantized Tensor Train variant of the Quantum Fourier Transform. This method is tailored to the geometry under consideration and offers significant improvements in time complexity and memory efficiency compared to the traditional FFT-based approach while remaining executable on classical hardware. The method is applicable only if a suitable Quantized Tensor Train representation exists for the stiffness operator associated with the underlying geometry.

Nonreciprocal Quantum Batteries

Borhan Ahmadi, Pawel Mazurek, Pawel Horodecki, Shabir Barznajeh

(Gdansk University; Institute of Science and Technology of Austria)

Nonreciprocity, arising from the breaking of time-reversal symmetry, has become a fundamental tool in diverse quantum technology applications. It enables directional flow of signals and efficient noise suppression, constituting a key element in the architecture of current quantum information and computing systems. Here we explore its potential in optimizing the charging dynamics of a quantum battery. By introducing nonreciprocity through reservoir engineering during the charging process, we induce a directed energy flow from the quantum charger to the battery, resulting in a substantial increase in energy accumulation. Despite local dissipation, the nonreciprocal approach demonstrates a fourfold increase in battery energy compared to conventional charger-battery systems. This effect is observed in the stationary limit and remains applicable even in overdamped coupling regimes, eliminating the need for precise temporal control over evolution parameters. Our result can be extended to a chiral network of quantum nodes, serving as a multi-cell quantum battery system to enhance storage capacity. The proposed approach is straightforward to implement using current state-of-the-art quantum circuits, both in photonics and superconducting quantum systems. In a broader context, the concept of nonreciprocal charging has significant implications for sensing, energy capture, and storage technologies or studying quantum thermodynamics.

Generative-Based Algorithm for Data Clustering on Hybrid Classical-Quantum NISQ Architecture

Julien Rauch (Université Paris Saclay)

Clustering is a well-established unsupervised machine-learning approach to classify data automatically. In large datasets, the classical version of such algorithms performs well only if significant computing resources are available (e.g., GPU). An alternative approach relies on integrating a quantum processing unit (QPU) to alleviate the computing cost. This is achieved through the QPU's ability to exploit quantum effects, such as superposition and entanglement, to natively parallelize computation or approximate multidimensional distributions for probabilistic computing (Born rule). In this paper, we propose first a clustering algorithm adapted to a hybrid CPU-QPU architecture while considering the current limitations of noisy intermediate-scale quantum (NISQ) technology. Secondly, we propose a quantum algorithm that exploits the probabilistic nature of quantum physics to make the most of our QPU's potential. Our approach leverage on ideas from generative machine-learning algorithm and variational quantum algorithms (VQA) to design an hybrid QPU-CPU algorithm based on a mixture of so-called quantum circuits Born machines (QCBM). We hope to achieve accurate data clustering and acceleration on the NISQ architectures scheduled to be available in the next few years. Finally, we analyse our results and summarize the lessons learned from exploiting a CPU-QPU architecture for data clustering.

Optimizing unitary coupled cluster wave functions on quantum hardware : error bound and resource-efficient optimizer

Martin Plazanet and Thomas Ayral (Eviden)

Simulating quantum many-body physics is the prime motivation behind the inception of quantum computing. To this end, many methods for quantum chemistry calculations have been developed, some of which with noisy intermediate scale quantum (NISQ) devices in mind. In this work, we study the projective quantum eigensolver (PQE) method, recently introduced by [1] in order to overcome the shortcomings of the most widespread NISQ method, the variational quantum eigensolver (VQE). We first show that one can derive a rather tight upper bound on the energy error. We then introduce a mathematical study of the classical optimization itself, and derive a novel residue-based optimizer. Using molecules LiH and BeH2 as benchmarks, we present numerical evidence of the superiority of our scheme over both VQE and the original optimization introduced in [1], thus establishing the viability of PQE as a very promising alternative to VQE.

[1] PRX Quantum 2 (2021), 10.1103/prxquantum.2.030301.

Silicon Spin Qubit Circuits: Modeling and Evaluation of Energetic Efficiency

Konstantina Koteva (Institut Néel ; University of Singapore)

As quantum devices scale, noise and rising cooling costs threaten performance. We present a full-stack model for silicon spin qubits (EDSR and ESR), connecting gate operations, measurement, and cryogenic power to computational fidelity. Using experimental data, we benchmark energy consumption for a 4-qubit VQE and a 20-qubit random circuit. Through the Metric-Noise-Resource (MNR) framework, we identify optimal qubit temperatures and driving frequencies to minimise power use without compromising success. Our results highlight practical strategies for energy-efficient, scalable quantum computing.

Efficient cross-device verification via Pauli sampling for highly-entangled, highly-doped states

Jose Carrasco (FU Berlin)

Cross-device verification (a.k.a. distributed inner product estimation) allows two remote parties to estimate inner products $tr(\rho\sigma)$, with each having black-box access to copies of ρ and σ, respectively. When the states ρ and σ exhibit low entanglement or can be prepared with few non-Clifford gates, this task can be reduced to independently learning efficient classical descriptions of each state using established techniques, and sharing this description in order to compute the overlap. In our work, we argue that efficient cross-device verification is also possible in more complex scenarios where tensor network and stabilizer-based methods are insufficient. Specifically, we introduce a class of highly entangled real states that cannot be approximated by circuits with *log*-many non-Clifford gates, and prove that Bell sampling enables efficient inner product estimation for these states. Notably, these findings are robust against preparation errors. We present possible applications of the results in quantum cryptography and verification.

Refs: arXiv:2405.06544, arXiv:2501.11688

Wednesday 14/05/2025

Exploiting quantum photonics for quantum computing and quantum machine learning (invited)

Philip Walther (Vienna University) — We 9:00-10:00

After providing a brief overview of recent advancements in the generation and processing of multi-photon states, I will show the potential of photonic quantum machine learning. After presenting a quantum-enhanced reinforcement learning using a tunable integrated processor, I will discuss our development of a so-called quantum memristor for single photons. These devices, which can mimic the behavior of neurons and synapses, hold great promise for the realization of quantum neural networks. I will also present how photonic processors can implementing quantum-enhanced kernels for machine learning tasks. At the end I will change topic by briefly discussing the flexibility of photonic systems for tasks that require non-standard quantum computer architectures.

Exponentially Enhanced Scheme for the Heralded Qudit GHZ State in Linear Optics

Seungbeom Chin et al. (Okinawa Inst. Sc. and Tech.; Korea Inst. Sc. and Tech.) — We 10:00-10:30

High-dimensional multipartite entanglement plays a crucial role in quantum information science. However, existing schemes for generating such entanglement become complex and costly as the dimension of quantum units increases. In this work, we overcome the limitation by proposing a significantly enhanced linear optical heralded scheme that generates the d-level N-partite GHZ state with single-photon sources and linear operations. We design our scheme based on the graph picture of linear quantum networks, introduced in *Quantum 5,611 (2021)* and *npj Quantum Information 10(1), 67 (2024)*. Our scheme requires dN photons, which is the minimal required photon number, with substantially improved success probability from previous schemes. It employs linear optical logic gates compatible with any qudit encoding system and can generate generalized GHZ states with installments of beamsplitters. With efficient generations of high-dimensional resource states, our work opens avenues for further exploration in high-dimensional quantum information processing. This talk is based on *Physical Review Letters 133 (25), 253601 (2024)*.

General approach to quantum information with quantum optics from a superselection rule perspective

Perola Milman et al. (U. Paris Diderot) — We 11:00-11:30

We show that explicitly incorporating a phase reference into the quantum state of the field permits its unambiguous definition and the evaluation of its resourcefulness in quantum information. This approach, in particular, allows for the identification of non-classical features of the field through a unified physical and computational criterion based on particle entanglement in first quantization. It also provides a general framework for describing arbitrary bosonic quantum information encodings, along with the necessary and sufficient Gaussian and non-Gaussian resources they require.

Minimizing resource overhead in fusion-based quantum computation using hybrid spin-photon devices

Shane Mansfield (Quandela) — We 11:30-12:00

We present three schemes for constructing a (2,2)-Shor-encoded 6-ring photonic resource state for fusion-based quantum computing, each relying on a different type of photon source. We benchmark these architectures by analyzing their ability to achieve the best-known loss tolerance threshold for fusion-based quantum computation using the target resource state. More precisely, we estimate their minimum hardware requirements for fault-tolerant quantum computation in terms of the number of photon sources to achieve on-demand generation of resource states with a desired generation period. Notably, we find that a group of 12 deterministic single-photon sources containing a single matter qubit degree of freedom can produce the target resource state near-deterministically by exploiting entangling gates that are repeated until success. The approach is fully modular, eliminates the need for lossy large-scale multiplexing, and reduces the overhead for resource-state generation by several orders of magnitude compared to architectures using heralded single-photon sources and probabilistic linear-optical entangling gates. Our work shows that the use of deterministic single-photon sources embedding a qubit substantially shortens the path toward fault-tolerant photonic quantum computation.

Roundtable featuring session speakers, Paris quantum company Quandela — *We 12:00-12:30*

Free afternoon for
- scientific discussions: the amphitheater remains available
- company lab visits: upon registration & screening

Cocktail dinner for all conference participants — *We 18:30-20:30*

Thursday 15/05/2025

Quantum Machine Learning: Hype vs Reality (invited)

Marco Cerezo de la Roca (Los Alamos) — Th 9:00-10:00

Quantum Machine Learning (QML) broadly refers to integrating the learning methodology og classical machine learning methods with quantum computational capabilities for coherent data analysis. Despite significant initial excitement surrounding this technology, it has become increasingly clear that QML is plagued with issues that prevent its deployment for realistic, large-scale problems. In this talk we will go over recent advancements in QML and address the central question: *Is there evidence that QML is genuinely useful?*

A unifying account of warm start guarantees for patches of quantum landscapes

Hela Mhiri et al. (LIP6 - EPFL) — Th 10:00-10:30

Barren plateaus are fundamentally a statement about quantum loss landscapes on average but there can, and generally will, exist patches of barren plateau landscapes with substantial gradients. Previous work has studied certain classes of parameterized quantum circuits and found example regions where gradients vanish at worst polynomially in system size. Here we present a general bound that unifies all these previous cases and that can tackle physically-motivated ansätze that could not be analyzed previously. Concretely, we analytically prove a lower-bound on the variance of the loss that can be used to show that in a non-exponentially narrow region around a point with curvature the loss variance cannot decay exponentially fast. This result is complemented by numerics and an upper-bound that suggest that any loss function with a barren plateau will have exponentially vanishing gradients in any constant radius subregion. Our work thus suggests that while there are hopes to be able to warm-start variational quantum algorithms, any initialization strategy that cannot get increasingly close to the region of attraction with increasing problem size is likely inadequate.

Quantum Circuit Optimization with Differentiable Projected Entangled Pair States for Ground State Preparation

Baptiste Anselme Martin (Eviden) — Th 11:00-11:30

The interplay between quantum computers and tensor networks have been increasingly popular, and can provide pathways to overcome difficult problems inherent to quantum algorithms, such as preparing relevant initial states for further computations. In this work, we utilize tensor networks to optimize quantum circuits for ground state preparation. Specifically, we employ differentiable Projected Entangled Pair States (PEPS) across various topologies to simulate and optimize parameterized quantum circuits for model Hamiltonians. Our approach enables the preparation of ground states with high energy accuracy, even for large qubit systems and connectivities that are beyond one dimension. Furthermore, we demonstrate that PEPS-based optimization may help mitigate the barren plateau phenomenon by providing a warm-start initialization with enhanced gradient magnitudes. We believe this work pushes the potential of quantum computing by leveraging classical pre-processing for both NISQ experiments and FT algorithms and helps to identify which tasks are best suited for classical or quantum resources.

Optimal permutation generation with local quantum gates for quadratic assignment problems

Dylan Laplace Mermoud (ENSTA) — Th 11:30-12:00

In this talk, we present a quantum algorithm that generates all permutations that can be spanned by one- and two-qubits permutation gates, which can be directly read as an output of the circuit. The construction of the circuits follows from group-theoretical results, most importantly the Bruhat decomposition of the group generated by CNOT gates. We next use the longest Coxeter word of the Bruhat order of the symmetric group, and assemble everything be leveraging the properties of the presentations of semidirect products of nicely presented groups. The circuits obtained are of linear depth and quadratic size. We use these circuits to build ansatze to tackle optimization problems such as quadratic assignment problems, and two of their refinements, graph isomorphism and heaviest k-subgraph problems. We finally present some results about the performance of the algorithm on the instances of the QAPLib obtained via simulation.

Roundtable featuring session speakers — *Th 12:00-12:30*

Reducing the overhead of quantum error correction (invited)

Aleksander Kubica (Yale University) — Th 14:00-15:00

Fault tolerance (FT) and quantum error correction (QEC) are essential techniques to building reliable quantum computers from imperfect components. Optimizing the resource and time overheads needed to implement QEC is one of the most pressing challenges that will facilitate a transition from the NISQ era to the FT era. In this talk, I will present a few intriguing ideas that can significantly reduce these overheads, such as erasure qubits, single-shot QEC and algorithmic FT. I will also discuss novel methods of adapting QEC protocols in the presence of defective qubits and gates.

BBGKY hierarchy for quantum error mitigation

Theo Saporiti et al. (CEA) — Th 15:00-15:30

The confinement/deconfinement phase transition of QCD at finite densities is still numerically inaccessible by classical computations. Quantum computers, with their potential for exponential speedup, could overcome this challenge. However, their current physical implementations are affected by quantum noise. In my talk, I will introduce a novel quantum error mitigation technique based on a BBGKY-like hierarchy, which is applicable to any arbitrary digital quantum simulation. The core idea is to improve zero-noise extrapolations by incorporating additional constraints from the hierarchy equations associated to the digital spin system. Our results indicate that the mitigation scheme systematically improves the quality of the (1+1)-Schwinger model measurements.

Teleporting quantum errors: Knill error correction in the era of modern quantum processors

Michael Vasmer et al. (INRIA Paris, U.Waterloo, Perimeter Institute) — Th 16:00-16:30

Quantum computing hardware has continued its advance in recent years, with modern processors capable of executing larger circuits with reduced noise. However, these processors remain too limited to run large-scale algorithms such as Shor's algorithm. To achieve scalable fault-tolerant quantum computation, it is widely believed that quantum error correction will be essential. In this talk, we focus on an approach to quantum error correction based on quantum teleportation, due to Knill.This approach offers many advantages, such as reducing the impact of leakage and coherent errors, and simplifying the decoding problem. Notably, we show that in Knill's approach the decoding problem for circuit-level noise is essentially the same as the decoding problem for channel noise, which means that code and decoder pairs optimised for channel noise also perform well under circuit noise. We illustrate this property by decoding a family of (high-rate) lifted product codes using belief propagation, without additional decoding layers such as OSD.

Roundtable featuring session speakers, Paris quantum company Alice&Bob — *Th 16:30-17:00*

Poster Session 2.A — *Th 17:00-17:45*

Noisy gates for simulating quantum computers

Michele Vischi (U Trieste)

We present a novel method for simulating the noisy behavior of quantum computers, which allows to efficiently incorporate environmental effects in the driven evolution implementing the gates acting on the qubits. We show how to modify the noiseless gate executed by the computer to include any Markovian noise, hence resulting in what we will call a noisy gate. We compare our method with the IBM qiskit simulator, and show that it follows more closely both the analytical solution of the Lindblad equation as well as the behavior of a real quantum computer, where we ran algorithms involving up to 18 qubits; as such, our protocol offers a more accurate simulator for NISQ devices. The method is flexible enough to potentially describe any noise, including non-Markovian ones. The noise simulator based on this work is available as a python package at the link, https://pypi.org/project/quantum-gates.

Orbit dimensions of Fock states in quantum optics

Eliott Mamon (LIP6)

Both linear and Gaussian quantum optics are bases of platforms explored for quantum computing. However, by default they are sub-univeral, and thus both require additional costly experimental ingredients to be able to universally drive quantum states in Fock space. In the meantime, preparable states are only able to attain a strict subspace of state space under arbitrary linear or Gaussian unitaries, called the orbit of the state. As preparing specific bosonic states is also experimentally challenging, understanding the structure of orbits of preparable states under these sub-universal unitaries is crucial to gain insight on the state space that current optical platforms can explore.
We propose using the dimension of these orbits as a simple yet powerful measure of their "richness", i.e. how many states an initial state can reach under linear or Gaussian unitaries. This dimension is computable for a wide class of states from the unitary group's Lie algebra, and is invariant under the action of the group. We demonstrate this approach for Fock basis states and show how orbit dimension offers both a diagnostic for state convertibility and an experimental witness of non-Gaussianity.

Efficient quantum circuits for quantum state preparation, diagonal operators and multi-controlled gates

Julien Zylberman (Observatoire de Paris)

Many quantum algorithms rely on the assumption that efficient quantum routines exist for tasks such as quantum state preparation (the process of encoding classical data into qubit states), unitary and non-unitary diagonal operators, and multi-controlled operations. Implementing these routines on a quantum computer necessitates the synthesis of quantum circuits, where efficiency is gauged by circuit size, depth, and the number of ancilla qubits required. However, existing methods for exact quantum circuit synthesis, particularly for quantum state preparation and diagonal operators, often scale exponentially with the number of qubits. In this poster, we present novel routines with efficient complexity scalings leveraging parallelization and approximation methods [1][2]. Additionally, we present the first quantum circuit for n-controlled gates with sublinear-in-n circuit depth without ancilla [3].

Ref: [1] PRA 109(4), 042401 [2] ACM 10.1145/3718348 [3] Nature Communications 15:5886

Advantage of Quantum Machine Learning from General Computational Advantages

Natsuto Isogai et al. (Tokyo University)

An overarching milestone of quantum machine learning (QML) is to demonstrate the advantage of QML over all possible classical learning methods in accelerating a common type of learning task as represented by supervised learning with classical data. In supervised learning, in terms of the PAC learning model, learning complexity is defined for a class of functions, called a concept class, reduced to the computational complexity when the class has only one function. Therefore, advantages of computational complexity does not straight forwardly imply advantages of learning complexity. The provable advantages of QML in supervised learning have been known so far only for the learning tasks designed for using the advantage of specific quantum algorithms, i.e., Shor's algorithms. Here we explicitly construct an unprecedentedly broader family of supervised learning tasks with classical data to offer the provable advantage of QML based on general quantum computational advantages in the PAC learning model, progressing beyond Shor's algorithms.

QHC: Ultra Minimal Quantum Hardware with Simplified Circuitry

Omid Faizy (LCMCP, Sorbonne University)

Quantum computing relies on preparing a quantum core in a specific state and measuring the resulting outcome after quantum evolution. In Quantum Hamiltonian Computing (QHC), these essential processes—state preparation (input) and measurement (output)—are fully integrated within the quantum core, embedded directly in its graph topology. This approach sets QHC apart from classical computers, which feature spatially distinct hardware for input, processing, and output, and from qubit-based quantum computers, which, while unifying processing within a quantum core, still employ separate input and output registers. By embedding all computational functions into the quantum graph, QHC achieves an ultra-minimal hardware design with simplified circuitry, eliminating the need for distinct registers and interconnects. Furthermore, the performance of QHC, including its response time to inputs and the strategies for extracting results, is intricately tied to the topology of the quantum graph. This dependence highlights the pivotal role of graph structure in defining the capabilities of this innovative computing paradigm.

One-to-one Correspondence between Deterministic Port-Based Teleportation and Unitary Estimation

Satoshi Yoshida (Tokyo University)

Port-based teleportation is a variant of quantum teleportation, where the receiver can choose one of the ports in his part of the entangled state shared with the sender, but cannot apply other recovery operations. We show that the optimal fidelity of deterministic port-based teleportation (dPBT) using $N=n+1$ ports to teleport a d-dimensional state is equivalent to the optimal fidelity of d-dimensional unitary estimation using n calls of the input unitary operation. From any given dPBT, we can explicitly construct the corresponding unitary estimation protocol achieving the same optimal fidelity, and vice versa. Using the obtained one-to-one correspondence between dPBT and unitary estimation, we derive the asymptotic optimal fidelity of port-based teleportation given by $1 - O(d^4N^{-2}) \leq F \leq 1 - \Omega(d^4N^{-2})$, which improves the previously known result given by $1 - O(d^5N^{-2}) \leq F \leq 1 - \Omega(d^2N^{-2})$. We also show that the optimal fidelity of unitary estimation for the case $n \leq d - 1$ is $F = (n+1)/d^2$, and this fidelity is equal to the optimal fidelity of unitary inversion with $n \leq d - 1$ calls of the input unitary operation even if we allow indefinite causal order among the calls.

Towards an error-protected qubit based on pinhole Josephson junctions

Tien Nguyen Dinh, Joël Griesmar (Ecole Polytechnique)

Transmon-based devices have demonstrated logical qubits that surpass the performance of their physical components. However, scaling quantum error correction (QEC) requires large hardware overhead. Alternatively, intrinsically noise-resilient qubits could reduce such overhead. In this work, we focus on the realisation of a cos 2φ qubit based on high-transmission junctions termed pinhole junctions—a design in which the two ground states reside in distinct parity subspaces, naturally protecting the qubit from local noise and errors. Our immediate objective is to demonstrate quantum coherence by observing Rabi and Ramsey oscillations in the protected regime. Further work will focus on implementing high-fidelity single- and two-qubit gates and extending the architecture to one- and two-dimensional arrays to explore exotic quantum states, such as Gottesman-Kitaev-Preskill states.

Analysis, Synthesis and Measurement of Quantum Linear Time-Invariant Systems

Jacques Ding (IN2P3, CNRS)

From first principles, we develop a general framework to quantize, synthesize and measure any multimode Linear Time-Invariant system from its classical transfer function, thus revealing its fundamental quantum noise, without Markovian assumption or state-space representation. We determine the inherent Lie group structure of such systems. We present a tomography scheme that generalizes homodyne detection for frequency-dependent quantum states.

Poster Session 2.B — *Th 17:45-18:30*

Benign Overfitting with Quantum Kernels

Joachim Tomasi, Sandrine Anthoine, Hachem Kadri (LIS, U. Aix-Marseille)

Quantum kernels quantify similarity between data points by measuring the inner product between quantum states. By embedding data into quantum systems, quantum kernel feature maps, which may be classically intractable to compute, can capture complex patterns. However, designing effective quantum feature maps remains a major challenge. Many quantum kernels suffer from exponential concentration, resulting in near-identity kernel matrices that fail to capture meaningful data correlations, leading to overfitting and poor generalization. Inspired by benign overfitting in classical machine learning, we introduce local–global quantum kernels, derived from both subsystem and full-system measurements. Through numerical experiments, we show that our quantum kernels exhibit benign overfitting, that is, they generalize well despite interpolating the training data.

Finding problems with potential quantum supremacy using graphs

Sebastian Grillo et al. (U.A. Asunciòn, Paraguay)

A high L1 spectral norm is essential for quantum decision trees to resist classical simulation. Therefore, problems efficiently solved by quantum trees but not classically must involve outputs with high spectral norm. This paper models single-query quantum decision trees as weighted dynamical graphs (WDGs), where the norm corresponds to the sum of edge weights. Based on their equivalence to degree-2 polynomials, WDGs map variables to vertices and monomials to edges, with an alternative matrix form also provided. Two optimization problems are introduced to analyze the norm: one maximizes it under a bounded range, the other minimizes the range with fixed norm. Uniformly distributing edge weights per vertex is shown to boost the norm. Finally, combining trees can produce sequences with growing spectral norms, hinting at exponential quantum advantage.

Tweaking Data Encoding Complexity in Quantum Machine Learning

Hillol Biswas (WAPCOS Limited, India)

Quantum machine learning, whether in Quantum Classical (QC) or Quantum Quantum (QQ) computing regime, requires large amounts of data to be stored correctly in quantum circuits to obtain the desired performance and computation time advantage. A hybrid encoding performed well on a synthetic dataset after delving deeper into the present data encoding technique. Compared to standard encoding, a novel hybrid encoding run by traditional optimisers such as COBYLA and L_BFGS_B in VQC achieved accuracy of 0.95 and 0.90, respectively. A 10,000-sample VQC training and testing yielded a candidate compared to Estimators and Sampler QNN; nevertheless, computation time is significantly longer than that of the conventional technique. The work intends to highlight the complexity of data encoding for VQC in QML so that QML can emerge as a candidate in the current era of classical machine learning.

From Spreading to Localization: Real-Time Dynamics in a 2D Spin Model

Alessandro Santini (Ecole Polytechnique)

We explore the real-time dynamics of a two-dimensional quantum spin system, focusing on how its behavior changes with varying interactions. Starting from a simple ordered configuration, we use advanced variational techniques based on neural-network quantum states to simulate the system's evolution. Our study reveals a dynamical transition between regimes characterized by distinct spreading behaviors of correlations. For weak anisotropy, excitations propagate freely, while for strong anisotropy, they remain localized—signaling the emergence of a confinement-like phenomenon. These results offer new insights into non-equilibrium dynamics in two dimensions and highlight the potential of machine learning tools for tackling complex quantum systems.

A Gradient-Free Approach to Variational Quantum Algorithms with HOPSO

Ijaz Mohammad (Slovak Academy of Sciences)

Variational Quantum Algorithms (VQAs) offer a promising approach for near-term quantum computing but are hindered by hardware noise and the issue of vanishing gradients, commonly referred to as barren plateaus. To address these challenges, we introduce a physics-inspired, gradient-free optimizer designed for Variational Quantum Eigensolver (VQE) applications called Harmonic Oscillator-Based Particle Swarm Optimization (HOPSO). HOPSO shows robust performance on molecular systems such as H_2 and LiH, excelling particularly in noisy environments and high-dimensional problem spaces. In comparison to other optimizers, it offers greater stability, improved tunability, and more reliable convergence. These results highlight HOPSO's potential to enhance the VQA performance on NISQ-era devices, contributing to the development of more efficient hybrid quantum-classical algorithms.

Turning qubit noise into a blessing: automatic state preparation and long-time dynamics for impurity models on quantum computers

Corentin Bertrand et al. (Eviden)

Noise is often regarded as a limitation of quantum computers. In this work, we show that in the dynamical mean field theory (DMFT) approach to strongly-correlated systems, it can actually be harnessed to our advantage. Indeed, DMFT maps a lattice model onto an impurity model, namely a finite system coupled to a dissipative bath. While standard approaches require a large number of high-quality qubits in a unitary context, we propose a circuit that harvests amplitude damping to reproduce the dynamics of this model with a blend of noisy and noiseless qubits. We find compelling advantages with this approach: a substantial reduction in the number of qubits, the ability to reach longer time dynamics, and no need for ground state search and preparation. This method would naturally fit in a partial quantum error correction framework.

Collision-assisted Information Scrambling on a Configurable Photonic Chip

Shuyi Liang (Shanghai Jiao Tong University)

Quantum interference and entanglement are in the core of quantum computations. The spread of information in the quantum circuit helps to mitigate the circuit depth. Although information scrambling in the closed systems has been proposed and tested in the digital circuits, how to measure the evolution of quantum correlations between systems and environments remains an open question. Here, we propose a codesigned photonic circuit to investigate the information scrambling in an open quantum system by implementing the collision model with cascaded Mach-Zehnder interferometers. We numerically simulate the photon propagation and find that the tripartite mutual information strongly depends on the system-environment and environment-environment interactions. We further use compressed sensing to reduce the number of observables and the number of shots required to reconstruct the density matrix. Our results provide a reconfigurable photonic platform for simulating open quantum systems and pave the way for exploring controllable dissipation and non-Markovianity in discrete-variable photonic computing.

Public event (in French, for non-expert audience):

Interactive video visit of the Quantum Integrated Circuits lab at ENS, Paris
animated by Diego Ruiz (@stream_theory)

Friday 16/05/2025

A mixed-species atom array for quantum computing (invited)

Giulia Semeghini (Harvard University) — Fr 9:00-10:00

In this talk, we will explore recent advancements in quantum computing using Rydberg atom arrays and present new opportunities enabled by the use of a dual-species array based on a mixture of alkali and alkaline-earth atoms. Trapped arrays of interacting Rydberg atoms have become a leading platform for quantum information processing and quantum simulation due to their large system size and programmability. The addition of a second atomic species opens new possibilities for implementing selective qubit control, engineering asymmetric inter- and intra-species Rydberg interactions, and exploring novel architectures for quantum computing. Additionally, continuous replenishment of both atomic species is central to improving scalability and circuit depths, while decreasing cycle times. We will present our ongoing efforts toward creating a new continuously reloaded, programmable atom array based on a mixture of Rb and Yb, and discuss its applications—from combining coherent and dissipative dynamics for quantum information processing to simulating lattice gauge theories.

Fast and Error-Correctable Quantum RAM

Francesco Cesa, Hannes Bernien and Hannes Pichler (IQOQI, U. Innsbruck; U. Chicago) — Fr 10:00-10:30

Quantum devices can process data in a fundamentally different way than classical computers. To leverage this potential, many algorithms require the aid of a quantum Random Access Memory (QRAM), i.e. a module capable of efficiently loading large datasets onto the quantum processor. However, a realization of this fundamental building block is still outstanding due to many crucial challenges, including incompatibility with current quantum hardware and quantum error-correction.

In this talk, I will present a novel QRAM design, that enables fast and robust QRAM calls, naturally allows for fault-tolerant and error-corrected operation, and can be integrated on present hardware. This places a long missing, fundamental component of quantum computers within reach of currently available technology. Our proposal employs a special quantum resource state that is consumed during the QRAM call, after being assembled efficiently in a dedicated module. I will explain the key points of our work, and show how the long standing challenges that prevented the deployment of QRAM so far are overcome within our scheme. Concretely, I will discuss detailed blueprints and quantitative estimations for modern neutral-atom processors, where our proposed QRAM module finds a particularly efficient implementation. Preprint at arXiv:2503.19172

When Quantum and Classical Models Disagree: Learning Beyond Minimum Norm Least Square & Subspace Preserving Quantum Convolutional Neural Network Architectures

Léo Monbroussou et al. (LIP6, Sorbonne Université) — Fr 11:00-11:30

Variational Quantum Circuits (VQCs) are central to many Quantum Machine Learning algorithms and are among the leading candidates for demonstrating useful quantum advantage. However, these methods face key theoretical challenges, including barren plateaus (vanishing gradients) and difficulties in rigorously proving model expressibility. In the first part of this talk, I will address a critical bottleneck in VQC design: how can we construct quantum models that resist classical approximation via Fourier sampling methods? I will begin by reviewing recent results on quantum Fourier models and introduce our recent work, which proposes a general theoretical framework for quantum advantage in regression problems.

In the second part of the talk, I will present an alternative approach to designing practical and expressive VQCs, illustrated through the Subspace-Preserving Quantum Convolutional Neural Network. By enforcing symmetry-preserving constraints in the computation, this architecture provides strong polynomial advantages and is supported by theoretical arguments regarding its expressibility, trainability, and scalability, while avoiding existing surrogate methods.

Towards Practical Quantum Neural Network Diagnostics with Neural Tangent Kernels

Francesco Scala et al. (U.Basel, IBM Zurich, U.Pavia) — Fr 11:30-12:00

In this work, we propose a practical framework allowing to employ the Quantum Neural Tangent Kernel (QNTK) for approximating Quantum Neural Networks (QNNs) performance before training. We show how a critical learning rate and a characteristic decay time for the average training error can be estimated from the spectrum of the QNTK evaluated at the initialization stage. We then show how a QNTK-based kernel formula can be used to analyze, up to a first-order approximation, the expected inference capabilities of the quantum model under study. We validate our proposed approach with extensive numerical simulations, using different QNN architectures and datasets. Our results demonstrate that QNTK diagnostics yields accurate approximations of QNN behavior for sufficiently deep circuits, can provide insights for shallow QNNs. This approach enables detecting – hence also addressing – potential shortcomings in model design.

Roundtable featuring session speakers — *Th 12:00-12:30*

End of Conference