High Performance Computing, Artificial Intelligence, and Quantum Computing Convergence

A Technical Whitepaper on Photonic Quantum Computing Integration in HPC centres

Quandela's team

October 2025

Table of Contents

Introduction

High Performance Computing (HPC), and Artificial Intelligence (AI) form the backbone of today's scientific discovery and technological innovation. They enable large-scale simulation of physical systems, from climate dynamics to nuclear fusion, and with AI excel at extracting patterns and predictions from vast datasets. Together, they define the current frontier of computational capabilities. Yet both face fundamental limitations: HPC struggles with the exponential growth of complexity in many-body systems, while AI models are increasingly bottlenecked by data quality, energy consumption, and interpretability.

Quantum computing now enters as a third, complementary paradigm. By exploiting superposition, entanglement, and interference, quantum processors can naturally tackle problems that overwhelm even the largest supercomputers and the most advanced AI models. Crucially, their role is not to replace but to augment HPC and AI workflows. Quantum hardware integrated into with HPC and AI establishes a three-cornered computational infrastructure. Each paradigm addresses the weaknesses of the others — HPC provide scale, AI provides adaptivity, and quantum provides fundamentally new pathways through intractable problems.

This convergence opens the door to capabilities that today remain out of reach: accelerating the discovery of new drugs and advanced materials, enabling climate models that capture the full complexity of planetary systems, powering autonomous technologies that operate safely in any environment, and securing global communications with provably unbreakable encryption. By uniting HPC, AI, and quantum computing, we are setting the stage for a new era where the limits of simulation, prediction, and security are redefined.

In this paper, we take six practical use-cases as a common thread to explore how near-term photonic quantum computers can already be applied today. We show how these algorithms map onto current and upcoming hardware, and how HPC centres can derive value by installing quantum processing units now rather than waiting for fault-tolerant machines. The discussion is structured around three pillars: hardware, where we outline the roadmap of scalable photonic processors; software, where we introduce the tools that enable simulation and seamless integration; and algorithms, where we detail the methods that underpin the use-cases. Together, these sections provide a comprehensive account of the state-of-the-art capabilities of Quandela's photonic quantum computers and how they can be brought into operation in existing computational infrastructures.

Use-Cases

To understand the value of photonic quantum computing in practice, it is essential to move beyond abstract benchmarks and examine concrete problems where existing methods struggle. In this section we introduce six representative use-cases, chosen to reflect both the diversity of industrial and scientific challenges and the range of algorithmic families that address them. They span critical infrastructure resilience, autonomous systems, cybersecurity, financial stability, and secure communications. Each is first presented from a

classical perspective, highlighting why current classical computing approaches face limitations and later from a quantum perspective, showing how quantum algorithms and hardware can address the classical limitations.

Use-Case 1: Structural integrity of critical infrastructure

Cracks in structures such as dams, pipelines, and pressure vessels evolve under stress in ways that are difficult to predict with classical simulation tools[1]. Finite-element models can in principle describe crack propagation, but resolving the fine-scale stress fields near a fracture tip demands extremely high mesh densities and thus vast computational resources. Each refinement multiplies the number of degrees of freedom, quickly exceeding what even large systems can handle. Moreover, when the initial crack geometry is uncertain, many simulations must be repeated with slightly different configurations, further compounding the cost. As a result, operators must often rely on simplified approximations or empirical safety margins rather than detailed fracture evolution models, leaving uncertainty in the true structural integrity of critical infrastructure.

Current HPC systems + AI methods therefore fail to provide timely and sufficiently precise fracture modelling at realistic scales, leaving room for alternative approaches.

Use-Case 2: Autonomous driving and perception under rare conditions

For self-driving vehicles, perception systems must work flawlessly not just in ideal daylight but in rare conditions like heavy rain, snow, or night-time glare. These edge cases are both safety-critical and underrepresented in training datasets. Classical generative models attempt to fill this gap by synthesizing new training images, but they often "hallucinate" — introducing unrealistic artefacts that reduce the quality of the generated data and bias downstream models[2]. This results in systems that may misinterpret a shadow as an obstacle, or worse, fail to recognize a pedestrian in the dark. Such shortcomings are a major barrier to deploying autonomous systems at scale.

Existing classical generators therefore cannot yet deliver the realism and reliability needed to rebalance training data for autonomy under rare conditions.

Use-Case 3: Cybersecurity and malware detection

Malware that slips past detection can lead to devastating real-world consequences: service outages, data theft, or breaches of critical infrastructure. Classical classifiers trained on past malware families can identify known threats, but adversaries constantly generate new, slightly altered variants designed to evade them. The imbalance between abundant "benign" data and scarce examples of novel malware strains further limits classical models[3]. When an attack is missed, organisations may suffer cascading effects — from financial losses and reputational damage to systemic risks in sectors like healthcare or energy where digital resilience is vital.

Despite heavy use of AI, the imbalance and adaptability of malware datasets continue to make robust detection a challenge beyond current computing systems capabilities.

Use-Case 4: Financial stability and fallen angel prediction

In financial markets, companies that are downgraded from investment grade to "junk" status — so-called fallen angels — trigger sudden volatility. Institutional investors such as pension funds or insurers are often obliged by regulation to sell downgraded assets, leading to fire sales and systemic ripples[4]. Classical predictive models struggle with this task because fallen angels are rare events, creating highly imbalanced datasets. As a result, early warning systems are prone to miss subtle signals of impending downgrades. For financial institutions, this means exposure to unexpected losses and the inability to reallocate capital proactively.

Because of the rarity and imbalance, classical forecasting models remain unreliable in predicting these critical events.

Use-Case 5: Secret key generation for secure communications

The security of encrypted communication, from online banking to diplomatic exchanges, depends on secret keys that must be truly unpredictable. If attackers can even slightly anticipate key values, they can weaken or break the encryption. Classical approaches rely on pseudo-random number generators seeded by deterministic inputs, or on hardware noise sources that are difficult to certify as private and tamper-proof. This leaves a residual risk that the randomness is not fully secure. For governments, banks, and critical industries, the ability to generate verifiably unpredictable keys is therefore not just a technical detail but a cornerstone of long-term digital security[5].

Classical sources of randomness remain vulnerable to bias and certification issues, preventing them from guaranteeing uncompromised key security.

Use-Case 6: Sustainable blockchain consensus

Modern blockchain networks rely on Proof-of-Work mechanisms that require solving large cryptographic puzzles to validate transactions and add new blocks. This process, while ensuring security and decentralisation, is extremely energy-intensive: vast computing power is devoted to generating hashes that have no purpose beyond consensus[6]. As networks grow, their electricity consumption scales accordingly, already matching that of some countries. Efforts to replace Proof-of-Work with less demanding schemes such as Proof-of-Stake often introduce other issues, including wealth concentration and reduced robustness against certain types of attacks. Moreover, classical mining approaches remain locked in an arms race of hardware acceleration, where small efficiency gains translate into significant capital and energy costs. This makes it increasingly difficult to maintain both scalability and sustainability in blockchain systems without compromising on decentralisation or security also known as the "blockchain trilemma".

With these six practical use-cases in mind this paper will explore the photonic near-term technology of Quandela and how the machines of today can already be used in a meaningful way.

Road to Utility

To understand how photonic quantum computing will deliver value, it is necessary to connect the hardware roadmap with the algorithms and use-cases that motivate it. This section is structured in three parts: first, we describe the near-term capabilities, outlining how successive generations of machines expand in qubit number and new modules; second, we introduce native photonic algorithms, which exploit the unique properties of light and extend beyond qubit encodings; and finally, we look towards fault-tolerant quantum computing (FTQC), the long-term destination of the hardware roadmap.

Near-term roadmap

First, we outline how the operational workflow of Quandela's photonic quantum computer to highlight near-term capabilities and advances with respect to other platforms[7]. More detailed information can be found in Appendix A.

At its core, the photonic quantum processing unit (QPU) consists of three main modules:

1. **Photon Generation** — single photons are produced on demand by semiconductor integrated devices (quantum dots), acting as high-speed, on-demand qubit generator emitting high-quality photons.
2. **Quantum Interconnect** —Photons are routed to different spatial outputs and modules via fast switches and optical fibres (DMX in the figure below).
3. **Manipulation** & **Detection** — photons interfere inside a reconfigurable photonic integrated circuit (PIC), composed of optical elements that implement programmable transformations. The photons are detected at the outputs, and their distribution encodes the result of the computation.

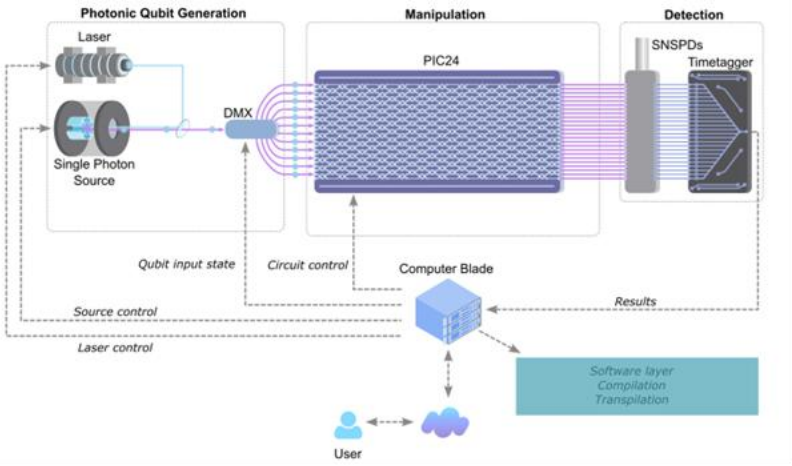

Fig. 1: In the photonic qubit generation, the laser excites the single photon source to generate a stream of single photons. The quantum interconnect (DMX) routes the photons to the different input ports of the photonic integrated circuit (PIC). The photons enter the PIC simultaneously and interfere with each other based on the programmed circuit. At the outputs of the PIC superconducting nanowire single photon detectors (SNSPDs) detect the photon distribution. Everything is controlled by a computer blade and can be remotely accessed by a user via our cloud platform.

The number of photons of a QPU is the number of single photons that can be injected into the photonic integrated circuit (PIC) simultaneously. The number of spatial in- and outputs, called optical modes depend on the size of the PIC.
A photonic qubit can be created from one photon and two optical modes; this is called dual rail encoding.

QPUs with 6 and 12 qubits have been installed in external HPC centres and are available via the Quandela cloud platform. These devices allow researchers to begin developing, testing, and validating algorithms under real experimental conditions. Future generations — 24 and 48 qubits— will unlock new capabilities by increasing the number of qubits and by introducing new modules to the architecture of Fig 1.

The new modules are:

- **Photon-number-resolving (PNR) detectors:** allows to distinguish between one or several photons in the same output mode in contrast to the existing threshold detectors.
- **Entangled input states:** the photons entering the chip can be already entangled and thus some parts of the computation can be done outside of the chip.
- **Feed-forward:** allows adaptive measurement schemes where results from one part of the circuit determine the later parts of the circuit.

The six uses-case described in the previous section have been picked to illustrate different aspects of the near-term capabilities of the photonic QPUs.

Use-Case	Algorithm	New QPU modules
Structural integrity of critical infrastructure	*Quantum Variational Differential Equation Solver (QVDES)*	Can be executed on today's QPUs
Secret key generation for secure communications	*Quantum Random Number Generation (QRNG)*	
Financial stability and fallen angel prediction	*Quantum Reservoir Computing (QRC)*	
Autonomous driving and perception under rare conditions	*Quantum-enhanced Generative Algorithms (QEGA)*	Entangled Input States and Photon Number Resolving Detectors
Cybersecurity and malware detection	*Quantum Convolutional Neural Networks (QCNNs)*	Entangled Input States, Photon Number Resolving Detectors and Feed-Forward
Sustainable blockchain consensus	*Quantum Proof-of-Work* (QPoW)	Photon Number Resolving Detectors

Table 1: Six use-case mapped to an algorithm which can solve it and to the QPU capabilities to test

As we can see in table 1, the use-cases or more precisely their correspondent algorithms can partly be already tested on the QPUs available today. For the other algorithms, we will see in the next section, that we can simulate the missing modules already today and develop the code to be run as soon as the new modules are available.

An important consideration is also that with each new generation of QPU the size of the problem related to each use-case which can be run on the QPU is also increasing. This means that for the three algorithms (QVDES, QRNG and QRC), which can be executed on the current 6 and 12 qubit machines, the 24 and 48 qubit machines will be able to solve larger problem instances coming closer to quantum utility. As we will see in later section the 48 qubit machine will be the first which can no longer be simulated classically and thus will demonstrate a quantum advantage. When we move to the 100 qubit machine and then to the logical qubit machines the practical advantage will become fully visible. Each machine starting from today's 6 and 12 qubit machines are helping to test algorithms and to train the workforce of tomorrow.

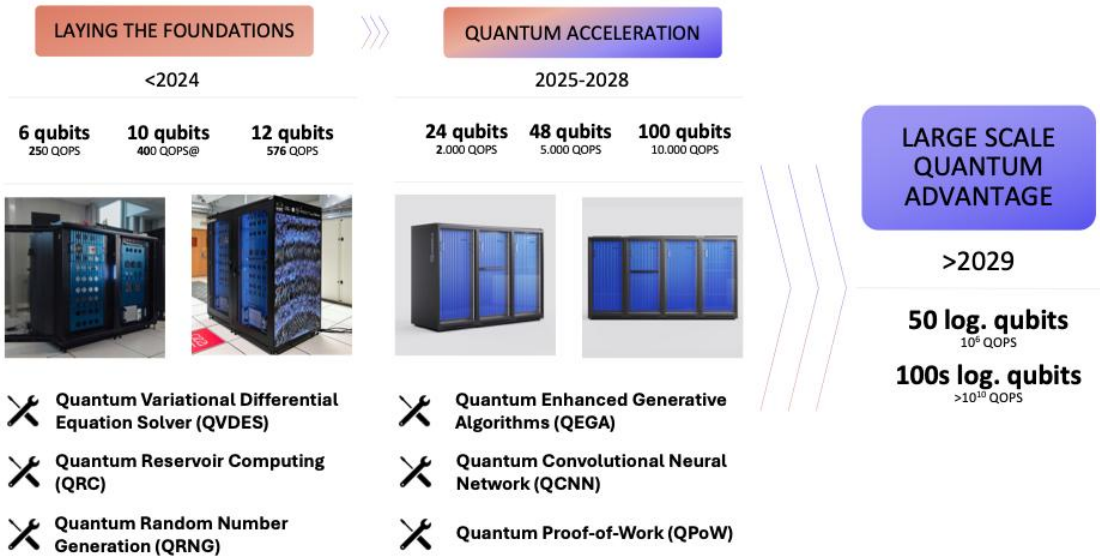

Fig. 2: What the next generation of photonic quantum computers can do.

Native Photonic algorithms

When describing quantum computers, one often counts resources in terms of qubits — the two-level quantum systems that form the building blocks of most platforms. In photonics there is a second possibility to make quantum computations and thus build photonic native algorithms[8]. These algorithms do not rely on qubits but make use of the quantum properties of the photons directly. Specifically, the concept of boson sampling[9] (see Appendix B) and extend it with the new modules of feed forward and entangled input states as described in the previous section.

Several of the use-cases and their corresponding algorithms mentioned above rely precisely on such photon-native algorithms:

- **Quantum Reservoir Computing (QRC):** exploits the natural interference of photons across many modes to process time-series data.
- **Quantum Convolutional Neural Networks (QCNNs):** use feed-forward measurements and hierarchical interference patterns to extract features from high-dimensional data.
- **Quantum Enhanced Generative Algorithms (QEGA):** integrate quantum correlations as structured "annotations" that can improve the realism and diversity of synthetic data.
- **Quantum Proof-of-Work (QPoW):** uses boson sampling directly as a hard task for nodes in the blockchain to solve.

They represent the most promising path to useful applications on near-term devices as they do not require fault tolerant quantum computing and error correction.

The photon native algorithms are more than pure boson sampling; with the modules feed-forward and entangled input states they are even harder to simulate classical and open the door for new applications (see Appendix B).

Towards fault-tolerant quantum computing

While this paper focuses on near-term photonic quantum processors, it is important to emphasise that Quandela's ultimate goal is fault-tolerant quantum computing (FTQC)[10]. Fault tolerance will allow large-scale, error-corrected computations capable of addressing problems far beyond the reach of classical systems.

The QPUs driving Quantum Acceleration in figure 2 — from 24 to 100 qubits — are part of this near-term vision. They represent the building blocks of future fault-tolerant architectures: the same modules (photon sources, interferometers, detectors, and feed-forward) building up near term algorithms will underpin the error-corrected photonic platforms. Each generation adds capabilities that are directly relevant to FTQC, such as entanglement generation, feed forward, and error correction. Scaling is ensured by modularity and optical networking, guaranteeing an ever-evolving system within the HPC centre infrastructure. This contrasts with monolithic approaches, such as superconducting based systems, ions and cold atoms, that are neither modular, nor do they possess an optical networking ability.

For HPC centres, this means that installing and working with near-term QPUs today is not a detour into a dead-end technology. It is an investment in the skills, infrastructure, and algorithmic expertise that will be directly transferrable to full-scale fault-tolerant quantum computers. Early deployments provide immediate educational and research value, while at the same time preparing HPC centres for the transition to FTQC as hardware matures. And by utilising the photonic native algorithms one will be able to reach utility even before the fault tolerant era.

Access and Software Stack

Use-Case Discovery Toolkit

The algorithms described in the previous section are not abstract proposals. They can already be explored today through simulation environments that reproduce the behaviour of photonic quantum processors. This means that HPC centres do not need to wait for large-scale hardware to begin engaging with quantum computing. The same workflows that will later be deployed on machines with 24 or 48 photons can already be tested, benchmarked, and refined using software tools that run efficiently on classical HPC clusters and HPC centres

The main value of todays machines comes from the capability to test and benchmark new algorithms. The QPUs enable developers to validate their algorithms in a realistic environment and thus speed up the discovery process. Quandela's software stack is built with this in mind and acts as a **use-case discovery toolkit**. At its core is **MerLin**, an open-source python library bringing together quantum computing and artificial intelligence.

Discovering new algorithms requires two things: **prototyping many algorithms** and **testing the developed prototypes**. To increase the amount of prototyping MerLin is embedded with well-known machine learning libraries like PyTorch, TensorFlow and Jax, making it possible for non-quantum experts to take part in the discovery of use-cases for quantum machine learning (QML) algorithms. To address direct testing of the new use-cases MerLin connects with Quandela's QPUs and automates the compilation for the end-user. In the prototyping phase running large simulations on GPUs is made possible by photonic simulators from the **Perceval** open-source python library[11], which underpins MerLin. Perceval makes it possible to design photonic circuits and facilitates the connection to Quandela's QPUs via the **Quandela Cloud**.

The three pillars **MerLin**, **Perceval** and **Quandela Cloud** define the use-case discovery toolkit. With it, quantum computing can create value in the near-term and is accessible to everyone.

Remote QPU Access

The **Quandela Cloud** provides a straightforward interface to real photonic hardware. Through a web portal or API, users can submit jobs and retrieve the results. The workflow is aligned with Perceval and MerLin: a circuit or algorithm build in python can be executed either on a simulator or on a real QPU simply by switching the backend. This continuity ensures that once an algorithm is validated in simulation, it can be deployed to hardware without modification. The same applies to noise simulators which emulate the exact noise of the available QPUs.

Today, the Cloud offers access to a 6 and a 12 qubit QPU one installed in Canada, and one installed in France. The algorithm from the 5th use-case, certified quantum random number generation (QRNG), is inbuild in the Cloud and enables the user to retrieve certified random numbers on demand.

The Quandela Cloud is part of the Quandela Hub which also includes a training centre and a community section. The learning paths in the training centre provide a fun and engaging introduction on how to submit ones first job to the QPU and give in depth lessons of the algorithms and use-case Quandela has built. If one still has some questions the community forum lets the user engage with the photonic community to share knowledge and get inspired.

Simulation

Simulating the developed prototypes is the first step in testing their capabilities. For this it is interesting to simulate prototypes which go beyond the current hardware capabilities so that one can already prepare for the next QPU once available. The software stack gives the user already today the possibility to simulate the extra modules mentioned in the road to utility section (feed-forward, entangled input states and PNR detectors). Additionally, it is already possible today to run simulations with as many photons and optical modes as one like, only constraint by classical compute capabilities. The classical limits however are already reached around 35-40 photons on pure boson sampling even when using HPC resource and GPUs. Simulating more than 10 photons on a personal computer becomes already challenging[12].

Perceval and the Quandela Cloud provide multiple simulation backends relevant for different scenarios (see Appendix C).

Installation and Integration in HPC centres

A key strength of photonic quantum processors is their ease of deployment in external HPC centres. Quandela's QPUs have been designed from the outset to operate with minimal infrastructure requirements, ensuring that they can be installed and integrated without disrupting existing workflows.

Low energy demand and simple cooling

Unlike superconducting, photonic processors operate mostly at room temperature. The only components that require cryogenics — the photon source and detectors — are contained within a compact, self-sufficient cryostat that is integrated into the rack. No external cooling systems are needed, and the entire QPU consumes only around 3 kW for a 6-qubit machine, comparable to a household electric hob. This makes it possible to connect the system directly to a standard wall socket, without requiring specialised power or cooling infrastructure.

High availability with automated calibration

Quandela's systems have demonstrated 92% availability in continuous cloud service and in third-party HPC centres such as OVHcloud, Exaion, and TGCC[13]. This reliability is achieved through automated calibration routines that maintain stable photon quality and circuit performance without manual intervention. As a result, the QPU operates consistently over long periods, with performance metrics remaining within specification.

Compact footprint

The full system occupies only a few square metres, significantly less than many other quantum computing platforms that require large dilution refrigerators or vacuum chambers. This small footprint means that a QPU can be installed in a standard HPC centre room alongside classical racks without modification of the facility.

Seamless software integration

The QPU comes with a built-in control layer that connects directly to Perceval, Quandela's open-source photonic circuit framework. From the cloud user's perspective, the QPU behaves like any other accelerator: circuits can be submitted through Perceval, executed locally or on the Quandela Cloud, and results are returned via a standard API. The same mechanism also allows integration with a HPC centre's existing access system.

Benefits for HPC centres

The value of deploying a photonic quantum processor in a HPC centre is not limited to preparing for future fault-tolerant machines. Even today, QPUs provide tangible benefits across two complementary dimensions: educational and integration value, and strategic research value.

Educational and integration value

Installing a QPU today gives researchers and engineers the opportunity to gain hands-on experience with quantum computing under real conditions. Instead of relying solely on

simulations, they can test new algorithms directly on hardware, explore realistic noise behaviour, and learn how to make efficient use of quantum resources. This accelerates the development of quantum solutions for the six use-cases highlighted in this paper and trains the next generation of users to think in hybrid quantum–classical terms.

For HPC centre staff, early access is equally valuable. Operating a QPU today means learning how to integrate it into schedulers, APIs, and AI workflows alongside CPUs and GPUs. This knowledge does not expire: the integration work carried out now will be reused for the next generation of machines, such as the 48-photon processor, and later for fully fault-tolerant QPUs. By starting today, HPC centres ensure that once larger devices are available, they will be usable immediately rather than requiring years of infrastructure catch-up.

Strategic research value

Beyond training and integration, installing a QPU today positions cloud providers strategically for the future. Early adopters can begin prototyping algorithms, validating workflows, and developing workloads that will later scale up to practical problem sizes. This preparation shortens the time to utility: by the time fault-tolerant devices become available, centres that already have tested applications and mature software stacks will be able to deploy them from day one.

The trajectory mirrors the early history of GPUs in HPC centre. Initially, GPUs were installed for niche acceleration tasks. But when AI workloads exploded, centres with GPU expertise and infrastructure were ready to lead, while others were left behind. QPUs are at a similar stage: not yet dominating workloads, but already indispensable for building the foundations of future breakthroughs. By working today on domains such as finance (rare-event prediction), cybersecurity (malware detection), energy (infrastructure resilience), and mobility (autonomous driving), centres ensure that when larger machines arrive, they have concrete, validated workloads that make quantum utility real.

In short: waiting means falling behind. Installing QPUs today gives HPC centres both the skills and the research portfolio to be leaders when quantum computing matures — just as those who invested early in GPUs defined the trajectory of AI today.

Algorithms

Financial stability (fallen angel prediction) → Quantum Reservoir Computing (QRC)
QRC exploits the natural interference of photons across many modes to create a high-dimensional dynamical reservoir[14]. What gives QRC its potential advantage is that the quantum correlations generated in the reservoir can reveal subtle patterns even when data is scarce and imbalanced, conditions under which classical models often fail. Independent research has shown that with only a handful of photons and modes, QRC can already match or exceed strong classical baselines. In collaboration with Crédit Agricole CIB, Quandela has benchmarked hybrid quantum–classical classifiers that improve precision on rare-event prediction. Current 6–12 photon devices can already run meaningful prototypes, while larger 24-photon processors will expand the model expressivity.

Critical infrastructure resilience → Quantum Variational Differential Equation Solver (QVDES)
QVDES reformulates fracture mechanics as an energy-minimisation problem, where a parametrised quantum circuit encodes nodal displacements and minimises the elastic energy of the structure. This approach, developed jointly with EDF, has been experimentally demonstrated on Quandela's photonic quantum processor Ascella[15]. By encoding spatial information directly into quantum amplitudes and using a QLOQ (qubit logic on qudits) representation[16], the solver achieves compact yet expressive models that scale efficiently with mesh refinement. Coarser quantum meshes can be used as warm starts for finer ones, mitigating optimisation difficulties and enabling progressive convergence of crack profiles. Current 6–12 photon processors already reproduce benchmark cases, while upcoming 24-photon systems will support higher-resolution meshes and more accurate stress-intensity estimates.

Cybersecurity (malware detection) → Quantum Convolutional Neural Networks (QCNNs)
QCNNs use layered photonic circuits and feed-forward operations to extract hierarchical features from imbalanced malware datasets, capturing correlations classical methods miss[17]. In collaboration with Orange, simulations of QCNNs with up to 15 photons have demonstrated feasibility. Execution on hardware requires feed-forward, which will become available in the 24–48 photon generations. Thanks to MerLin, these models can already be trained and validated in simulation, ready for deployment as hardware matures.

Autonomous driving (rare-condition perception) → Quantum-Enhanced Generative Algorithms (QEGAs)
Quantum annotations — compact, structured correlations derived from entangled photons — can stabilise classical image-to-image models and reduce artefacts/hallucinations in synthetic night-time scenes. This approach was validated in the Airbus–BMW global challenge, showing improved realism[18]. While current machines cannot yet provide entangled multi-photon inputs, the 48-photon generation will unlock this capability. In the meantime, MerLin allows these pipelines to be fully simulated and prepared.

Secure communications (cryptographic keys) → Quantum Random Number Generation (QRNG)
Certified QRNG protocols use the fundamental randomness of quantum measurements to generate verifiably unpredictable numbers. This does not require scaling photon numbers and is already implemented in today's 6- and 12-photon devices, with certified QRNG services available via the Quandela Cloud. This makes QRNG the earliest use-case to reach direct utility[19].

Sustainable blockchain consensus → Quantum Proof-of-Work (QPoW)
QPoW replaces energy-intensive classical mining with a photonic sampling process that proves computational effort through genuine quantum statistics. The algorithm encodes block information into the inputs of a linear optical circuit and performs a boson-sampling operation, whose output distribution cannot be efficiently mimicked by classical machines[20]. Each sample thus serves as a verifiable "proof" of work, validated by coarse-grained statistical checks rather than redundant computation. Because photons interact only via interference, the energy cost scales with photon generation rather than with sustained computation, offering a secure and environmentally sustainable consensus mechanism. Future multi-photon processors will enable higher validation throughput and stronger resistance to classical spoofing.

Conclusion

Quantum computing is often spoken of as a distant prospect, tied to the eventual arrival of fault-tolerant machines. This white paper shows a different picture: useful steps can already be taken today. With photonic processors, HPC centres can install systems that run at room temperature, consume only a few kilowatts, and integrate seamlessly into existing workflows. These machines are modular and upgradeable: a 6 or 12 qubits processor installed now can grow step by step into the 24, 48 and 100 qubits systems that unlock advanced algorithms such as quantum convolutional neural networks and quantum-enhanced generative models.

Installing such systems now has two immediate consequences. First, it shortens the time from laboratory research to practical utility, as integration work and algorithm development can proceed in parallel with hardware advances. Second, it ensures that HPC centres build the expertise and infrastructure needed for fault-tolerant quantum computing, when it becomes available.

In this sense, near-term quantum deployment is less about demonstration and more about preparation: establishing reliable access, developing applications, and training users in hybrid workflows. The trajectory is similar to the early adoption of GPUs in scientific computing, where HPC centres that experimented early were best positioned when the technology matured.

The opportunity, therefore, is not speculative but practical. By incorporating quantum processors today, HPC centres contribute to building the groundwork for a computing paradigm that will, in time, complement both classical simulation and AI. The HPC centres that engage early will be the ones ready to use this capability when it becomes indispensable.

References

1. Alshoaibi, A. M. & Fageehi, Y. A. Advances in Finite Element Modeling of Fatigue Crack Propagation. *Appl. Sci.* **14**, 9297 (2024).
2. Winter, K. *et al.* Generative AI for Autonomous Driving: A Review. Preprint at https://doi.org/10.48550/arXiv.2505.15863 (2025).
3. Almajed, H., Alsaqer, A. & Frikha, M. Imbalance Datasets in Malware Detection: A Review of Current Solutions and Future Directions. *Int. J. Adv. Comput. Sci. Appl. Ijacsa* **16**, (2025).
4. Belloni, M., Helmersson, T., Jarmuzek, M., Mosk, B. & Nikolic, F. Understanding what happens when "angels fall". https://www.ecb.europa.eu/press/financial-stability-publications/fsr/focus/2020/html/ecb.fsrbox202011_03~578f4f74dc.en.html (2020).
5. OVHcloud, the first global player to improve website access security with a quantum computer. https://corporate.ovhcloud.com/en-gb/newsroom/news/ovhcloud-improve-website-access-security-quantum-computer/ (2025).
6. Zimba, A., Phiri, K. O., Mulenga, M. & Mukupa, G. A systematic literature review of blockchain technology and energy efficiency based on consensus mechanisms, architectural innovations, and sustainable solutions. *Discov. Anal.* **3**, 14 (2025).
7. Maring, N. *et al.* A versatile single-photon-based quantum computing platform. *Nat. Photonics* **18**, 603–609 (2024).
8. Salavrakos, A., Maring, N., Emeriau, P.-E. & Mansfield, S. Photon-native quantum algorithms. *Mater. Quantum Technol.* **5**, 023001 (2025).
9. The computational complexity of linear optics | Proceedings of the forty-third annual ACM symposium on Theory of computing. https://dl.acm.org/doi/10.1145/1993636.1993682.

10. Gliniasty, G. de *et al.* A Spin-Optical Quantum Computing Architecture. *Quantum* **8**, 1423 (2024).
11. Heurtel, N. *et al.* Perceval: A Software Platform for Discrete Variable Photonic Quantum Computing. *Quantum* **7**, 931 (2023).
12. Heurtel, N., Mansfield, S., Senellart, J. & Valiron, B. Strong simulation of linear optical processes. *Comput. Phys. Commun.* **291**, 108848 (2023).
13. Maring, N. *et al.* One nine availability of a Photonic Quantum Computer on the Cloud toward HPC integration. Preprint at https://doi.org/10.48550/arXiv.2308.14582 (2023).
14. Sakurai, A., Hayashi, A., Munro, W. J. & Nemoto, K. Quantum optical reservoir computing powered by boson sampling. *Opt. Quantum* **3**, 238–245 (2025).
15. Remond, U. *et al.* Quantum remeshing and efficient encoding for fracture mechanics. Preprint at https://doi.org/10.48550/arXiv.2510.14746 (2025).
16. Lysaght, L., Goubault, T., Sinnott, P., Mansfield, S. & Emeriau, P.-E. Quantum circuit compression using qubit logic on qudits. Preprint at https://doi.org/10.48550/arXiv.2411.03878 (2024).
17. Monbroussou, L. *et al.* Photonic Quantum Convolutional Neural Networks with Adaptive State Injection. Preprint at https://doi.org/10.48550/arXiv.2504.20989 (2025).
18. BMW Group and Airbus reveal winners of Quantum Computing Challenge. https://www.press.bmwgroup.com/global/article/detail/T0446782EN/bmw-group-and-airbus-reveal-winners-of-quantum-computing-challenge?language=en.
19. Certified Randomness in Tight Space | PRX Quantum. https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.5.020348.
20. Singh, D. *et al.* Proof-of-work consensus by quantum sampling. *Quantum Sci. Technol.* **10**, 025020 (2025).

21. Somaschi, N. *et al.* Near optimal single photon sources in the solid state. *Nat. Photonics* **10**, 340–345 (2016).
22. Salavrakos, A. Error-mitigated photonic quantum circuit Born machine. *Phys. Rev. A* **111**, (2025).

Appendix A – Hardware Architecture

A.1 Photon sources

At the heart of Ascella lies the quantum dot — a nanometre-scale semiconductor structure sometimes described as an "artificial atom"[21]. When excited by a laser pulse, it emits exactly one photon, with very high purity and indistinguishability. This deterministic behaviour contrasts with older spontaneous parametric down-conversion (SPDC) methods, which produce photons only probabilistically.

Quandela's expertise in semiconductor nanostructures allows the production of bright and stable quantum dots with reproducible performance. To ensure purity and stability, the quantum dots are cooled to 4 K using compact rack-compatible cryo-modules. This level of cooling is modest compared to the millikelvin temperatures required for superconducting qubits.

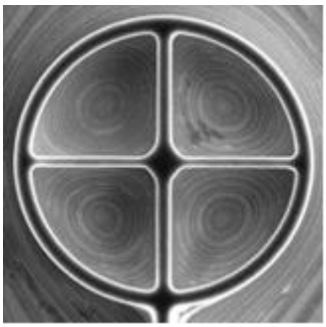

Fig. 3: Close up view of one quantum dot surrounded by the nanostructure controlling the photon emission. The diameter is 20 nanometres, and the quantum dot is right in the centre of the cross.

A.2 Photonic Integrated Circuit (PIC)

The PIC is the core computational stage. It consists of a reconfigurable network of beam splitters and phase shifters, integrated in a stable optical chip.

- **Beam splitters** divide and recombine optical paths, creating superpositions.
- **Phase shifters** tune the optical path length and act as programmable parameters.

Together, these components implement arbitrary linear-optical transformations. By programming the phase shifters, one can define the quantum circuit executed on the photons.

The PIC is fully reconfigurable and stable over long timescales, with automatic calibration ensuring reproducibility.

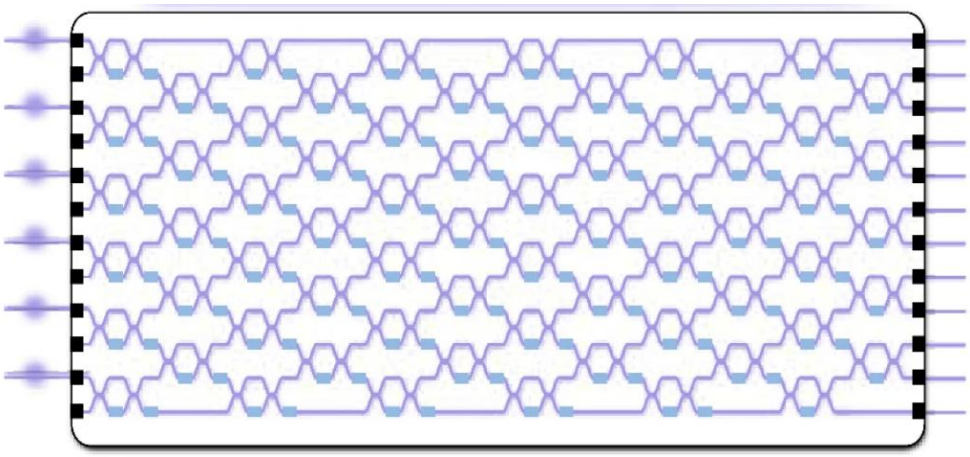

Fig. 4: Graphic of the photonic integrated chip with 12 inputs and 12 outputs. The blue lines represent optical waveguides, the blue crossings represent phase shifters, and the orange squares represent programmable phase shifters. The 6 orange circles on the left indicate 6 photons entering the chip.

A.3 Detectors

At the output, photons are registered by an array of single-photon detectors. Two types are used:

- **Threshold detectors** — indicate the presence or absence of a photon in a mode.
- **Photon-number-resolving detectors (PNRDs)** — can distinguish whether one, two, or more photons arrive simultaneously in a mode.
- **PNRDs** are particularly important for advanced algorithms (e.g. QCNNs, generative models), as they provide richer measurement outcomes than binary threshold detection.

Like the photon source, superconducting nanowire single photon detectors (SNSPDs) are cooled to 4 K to reduce noise and increase efficiency.

A.4 Performance and benchmarking

Ascella performance is measured both at the gate level (when using dual-rail qubit encoding) and at the hardware level (intrinsic photon metrics).

- **Gate-based fidelities:**

  1-qubit gates: 99.6% ± 0.1%

  2-qubit gates: 99.2% corrected (99.5% mitigated)

  3-qubit gates: 86% ± 1.2%

- **Photon quality metrics:**

  Photon purity: > 99%

  Indistinguishability: 94% - 98%

  Clock rate – shots/sec: up to MHz

- **Operational benchmarks:**

  Availability in external HPC centres: 92% (OVHcloud, Exaion, GENCI).

  Power consumption: ~3 kW for a 6-photon system (comparable to a domestic electric hob).

  Calibration: every 6 to 9 months. (the system includes an auto-optimization system to keep constant fidelity metrics)

A.5 Photon loss and mitigation

Photon loss remains the main challenge for scaling. Every stage — source, coupling, propagation, detection — reduces the probability that all (N) photons survive to the output. The n-photon coincidence rate can be approximated as:

$$C = \frac{RR}{N}\,\eta^N$$

where (RR) is the laser repetition rate, (N) is the photon number, and (η) is the transmission efficiency.

Mitigation strategies include:

- **Source improvements:** brighter quantum dots and better collection efficiency.
- **Integrated optics:** reducing coupling and propagation losses.
- **Detector efficiency:** advances in superconducting nanowire detectors.
- **Recycling and error-mitigation protocols**: re-using partial measurement outcomes to compensate for lost photons[22].

Thanks to these advances, robust multi-photon experiments and stable long-term operation are already possible, and each roadmap generation builds on these improvements.

Appendix B - Boson sampling

B1. Fock states and Fock space

A Fock state describes exactly how many photons occupy each optical mode. For example, the state

$|2,0,1\rangle$ represents two photons in the first mode, none in the second, and one in the third.

The collection of all such states forms the Fock space. For (N) photons distributed across (M) modes, the dimension of the Fock space is given by:

$$\dim = \binom{N + M - 1}{N}$$

This grows rapidly with both photon number and mode number, and it is this combinatorial explosion that makes photonic systems powerful for computation and difficult to simulate classically.

B2. Linear optics

In a photonic processor, photons evolve through beam splitters and phase shifters. These elements implement linear transformations of the optical modes. The mathematics is simple to describe but costly to compute at scale, because the output probabilities depend on calculating matrix permanents, a task known to be #P-hard.

B3. Boson sampling

The simplest and most fundamental native photonic algorithm is boson sampling: injecting multiple indistinguishable photons into an interferometer and sampling the output distribution. While boson sampling itself has no direct practical use, it demonstrates how even modest numbers of photons and modes generate distributions that are exponentially hard for classical computers to simulate.

More advanced algorithms, such as Quantum Reservoir Computing (QRC), Quantum Convolutional Neural Networks (QCNNs), and Qauntum Enhanced Generative Models (QEGMs), build on the same photonic-native principles but extend them to structured, application-oriented tasks. These are the algorithms central to the use-cases discussed in this paper.