Sustainable Quantum Computing: Analyzing Costs and Carbon Emissions

Teodor CERVINSKI[1], Cristian-Valeriu TOMA[1], Claudiu BRANDAȘ[2], Marius POPA[1]
[1]Bucharest University of Economic Studies
[2]West University of Timisoara
teodor.cervinski@csie.ase.ro, cristian.toma@ie.ase.ro, claudiu.brandas@e-uvt.ro, marius.popa@ie.ase.ro

As quantum computing progresses from theory to practical systems, its environmental and economic sustainability remains underexplored. This paper analyzes both the carbon footprint and cost implications of quantum computing technologies, with attention to their alignment with global sustainability goals. Key contributors to energy consumption, emissions, and operational costs are identified across the quantum computing lifecycle: hardware manufacturing, cryogenic cooling, runtime power demands, and quantum circuit simulation. The study employs a qualitative methodology, integrating a comprehensive review of scientific literature, environmental assessments, cost analysis reports, and benchmarking frameworks. It further presents a comparative analysis of quantum and classical high-performance computing (HPC), evaluating energy efficiency, environmental impact, and cost-effectiveness across realistic scenarios. By addressing both sustainability and economic dimensions, this research provides new insights for developers and policymakers, supporting the advancement of greener and more cost-effective quantum technologies.

Keywords: *Quantum computing, Sustainability, Carbon footprint, Operational costs, Energy consumption, Environmental impact*
DOI: 10.24818/issn14531305/29.3.2025.03

1 Introduction

Quantum computing is rapidly emerging as one of the most transformative technologies of the 21st century. It holds the potential to revolutionize fields such as cryptography, optimization, and climate modeling. In contrast to classical computers, which operate on bits governed by binary logic, quantum computers use qubits—quantum bits—that exploit fundamental principles of quantum physics, including superposition and entanglement. These properties enable quantum systems to process information in ways that can significantly surpass classical computing capabilities for specific problem domains [1].

As investment in quantum technologies accelerates, and prototypes scale toward commercially relevant systems, new challenges emerge—not only in technical performance, but also in economic viability and environmental sustainability. Understanding the costs and carbon impacts associated with quantum computing is essential to ensure its responsible and scalable adoption.

While most research in quantum computing has focused on computational capabilities and algorithmic performance, the environmental and economic implications of this technology remain relatively underexplored. Growing concerns about the sustainability of emerging digital infrastructures have prompted both academic and industry communities to examine these impacts, particularly in quantum systems based on superconducting qubits. Significant energy consumption arises from the need for cryogenic cooling, error correction protocols, and specialized control electronics [2].

As interest in quantum computing expands among governments, private companies, and research institutions, critical questions are emerging regarding its alignment with climate goals, cost-effectiveness, and the principles of responsible innovation [3].

This paper addresses the central research question: What are the environmental and economic impacts of quantum computing—particularly its carbon footprint and cost structures—and how can the development of quantum technologies be aligned with sustainability principles? To explore this, we adopt a qualitative research methodology, drawing on insights from peer-reviewed scientific literature, computational models for carbon emissions, cost analysis studies, and life cycle assessment frameworks.

2 Energy and Cost Drivers Across Quantum Computing Architectures

Quantum computers leverage principles of quantum mechanics, such as superposition and entanglement, to enable qubits to process information in ways that can surpass classical computers for specific problem domains. Several quantum hardware architectures are currently under active development, each with distinct implications for energy consumption and operational costs. The most prominent approaches include superconducting qubits, trapped ions, topological qubits, and photonic systems. These architectures vary significantly in their energy requirements, cooling demands, and supporting infrastructure, all of which contribute to both their carbon footprint and economic viability.

Among current quantum hardware platforms, superconducting qubit architectures are the most widely implemented. Pioneered by companies such as D-Wave, IBM, and Google—with early contributions from academic institutions including MIT—these systems have advanced rapidly in recent years. However, superconducting architectures are also among the most energy-intensive, primarily due to stringent cooling requirements. They operate at temperatures near absolute zero (~15 millikelvin), maintained by complex and energy-demanding dilution refrigerators. The cost of these refrigerators varies widely depending on size, cooling power, and customization, with standard research-grade systems typically ranging from $300,000 to over $1 million USD. The energy required for cryogenic cooling represents a substantial portion of the total operational energy consumption of these systems. Annual electricity costs for refrigeration alone are estimated between $10,000 and $50,000, depending on the type of qubits used and the duration of usage. In addition, significant energy and cost contributions arise from the necessary control electronics, classical interface systems, and qubit initialization processes, all of which impact both the carbon footprint and the overall economic efficiency of superconducting quantum computers.

Quantum error correction is another critical factor influencing both the energy efficiency and economic viability of quantum computing systems. Due to the inherent instability and noise in quantum hardware, error correction mechanisms are essential for maintaining computational reliability over time. These mechanisms detect and correct errors such as bit-flip, phase-flip, or combined errors. Implementing error correction requires substantial overhead: additional physical qubits must be dedicated to protecting logical qubits, and repeated gate operations significantly increase computational workload. This added complexity, known in the literature as *error correction overhead*, leads to higher energy consumption and greater demands on control infrastructure, which in turn escalate both operational costs and carbon footprint [5].

Moreover, recent research emphasizes that the entire life cycle of a quantum computing system must be considered when assessing its total carbon footprint. Focusing solely on runtime energy consumption provides an incomplete picture of the technology's environmental impact. A comprehensive evaluation should include the emissions and costs associated with hardware manufacturing, infrastructure development, cooling systems, maintenance, and end-of-life disposal [6]. This broader perspective is essential for accurately determining the true sustainability profile and long-term economic implications of quantum computing technologies.

In addition to energy demands, the operational costs of quantum computing systems represent an important consideration for both sustainability and economic viability. These costs include electricity consumption for cryogenic cooling, control electronics, and supporting infrastructure, which can accumulate significantly over time. For superconducting quantum computers, the annual electricity cost for operating a dilution refrigerator alone is estimated to range between **$10,000 and $50,000**, depending on system utilization and local energy prices [2, 4, 8]. Control electronics add further energy overhead, with total system power consumption for mid-scale platforms such as IBM Osprey or Google Sycamore reported at **20–35 kW**, translating into thousands of dollars in additional yearly electricity expenses. Moreover, maintenance and service contracts for cryogenic systems and control hardware further contribute to ongoing operational expenditures. As quantum computers scale toward fault-tolerant architectures, these costs are expected to rise substantially due to increased qubit counts, more complex control systems, and higher error correction overhead [10], [11]. Understanding these economic factors is essential for assessing the future competitiveness and sustainability of quantum technologies.

Table 1. Power Consumption of IBM Osprey, Google Sycamore, and D-Wave Advantage Quantum Computers.

Subsystem	**IBM Osprey 433 qubits (gate model)**	**Google Sycamore 53 qubits (gate model)**	**D-Wave Advantage 5000+ qubits (annealer)**	**Estimated Cost Impact**
Dilution Refrigerator	≈25 kW (one large dilution refrigerator)	≈20–23 kW (cooling system): • ~10 kW for the cryo-compressor • ~10-13 kW for cooled water support	≈15–25 kW (dilution refrigerator)	CapEx $300K–$1M + OpEx $10K–50K/year
Qubit Control Electronics	~5–10 kW (estimated) ~10 kW for older models, a few kW for newer models	≈3 kW – Microwave control racks and electronics for support	<1 kW (approx.) – The annealer uses DC/low frequency control currents while the chip itself uses negligible power	~$100K–$500K (hardware) + OpEx varies
Qubit Initialization/Readout	Negligible –Initialization is passive, readout uses minimal power	Negligible – Part of the supporting electronics mentioned above	Negligible – Minimal power overhead	Negligible
Error Corretion / Overhead	~0 kW (N/A) - No error correction mechanism is implemented	~0 kW (N/A) - Sycamore' 2019 run did not use error correction codes [6]	~0 kW (N/A) The annealers do not employ error correction	N/A, but projected to add substantial CapEx and OpEx in future fault-tolerant systems

From the data presented in Table 1, it is evident that cooling systems represent the primary source of energy consumption in superconducting qubit quantum computers. Although error correction overhead is not yet a significant factor in current experimental systems, this is likely to change as quantum hardware scales. In future large-scale, fault-tolerant quantum computers, Quantum Error Correction (QEC) is expected to introduce substantial additional demands. Studies suggest that QEC could dominate both the energy consumption and hardware complexity of such systems, potentially increasing resource requirements by an order of magnitude compared to uncorrected qubit operations [4]. This growth in hardware and operational demands will also translate into significant cost increases, both in terms of additional capital expenditures (CapEx) for expanded cryogenic and control infrastructure, and higher operational expenditures (OpEx) due to greater power consumption and system maintenance needs. Addressing these scaling challenges is therefore essential for ensuring both the sustainability and economic viability of future quantum computing architectures [4], [10], [11].

As research advances, alternative quantum processor architectures are emerging as strong contenders. Each of these platforms offers distinct advantages and limitations, addressing some of the challenges associated with superconducting qubit systems—such as extreme cooling requirements—while introducing new technical and economic considerations of their own. Factors such as hardware complexity, energy efficiency, scalability, and cost of implementation vary widely across architectures, making it essential to evaluate both their environmental and financial impacts when assessing their long-term potential.

Regarding energy consumption, alternative quantum computing architectures exhibit distinct demands and requirements. In addition to technical factors, their respective capital and operational costs must also be considered when evaluating their overall sustainability:

- **Trapped Ion Qubits.** These systems confine ions in an ultra-high vacuum chamber at near-room temperature, eliminating the need for dilution refrigeration. However, they rely on an extensive laser infrastructure to manipulate the qubits. Scaling such systems will require thousands of individually controlled laser channels [7], significantly increasing both energy consumption and hardware complexity. The capital costs of high-precision laser systems and optical components are substantial, while operational costs include continuous power for laser cooling and control.
- **Photonic Qubits.** As the name implies, this approach uses photons as qubits. Energy consumption is primarily driven by the devices used to generate, manipulate, and detect photons. Among current architectures, photonic systems may offer one of the most favorable sustainability profiles, with lower cooling requirements and potentially lower operational energy costs compared to superconducting platforms. However, specialized optical components and single-photon sources can introduce high capital costs, particularly at scale.
- **Topological Qubits.** This architecture seeks to encode quantum information in exotic quasiparticles (e.g., Majorana zero modes). Like superconducting qubit systems, topological qubits require cryogenic environments with extreme cooling demands, implying comparable levels of energy consumption. While their intrinsic error resistance could eventually reduce overhead costs in error correction, the technology remains in early experimental stages, and current capital and operational costs are high due to complex materials and fabrication requirements.

3 Measured Power Use and Real Cost of Running Quantum Workloads

As a practical example, this section considers the execution of an XOR (exclusive OR) operation for simple encryption and decryption of data. While such basic logical operations are not typically recommended for execution on a quantum computer—being more efficiently handled by classical processors such as CPUs or GPUs—this test case serves to illustrate the current performance characteristics, power consumption, and actual cost of running quantum circuits on available quantum computing platforms. In quantum computing circuits, loop constructs such as *for* or *while* statements are not implemented within the quantum hardware itself. Instead, any iterative logic is executed on the accompanying classical processor, while the quantum computer processes only the explicitly defined quantum circuit. In the example presented in Figure 1, a total of 8 qubits are used to represent one input byte from a file (*input2*), and another 8 qubits represent one byte of a symmetric key (*input1*). The XOR operation alters the state of the *input2* quantum register based on the key. The corresponding quantum assembly sequence for each byte-level XOR operation is written in QASM 2.0, as shown below.

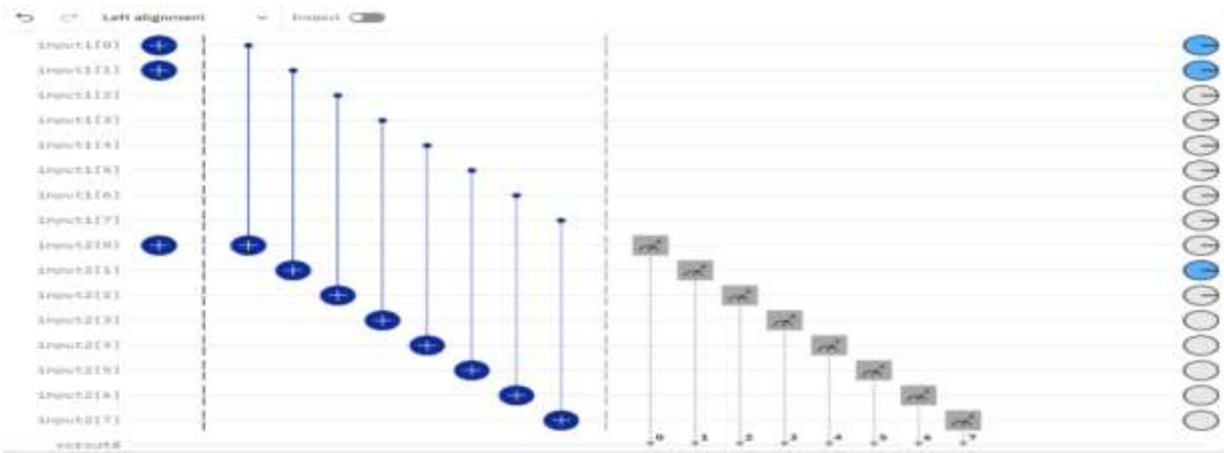

Fig. 1. Quantum circuit in QASM 2.0 for applying XOR function on 8 qubits (00000011 XOR 00000001 -> 00000010) using CNOT quantum gates

```
OPENQASM 2.0;
include "qelib1.inc";
qreg input1[8]; // Define quantum register for first input array
qreg input2[8]; // Define quantum register for second input array
creg xorout[8];
x input1[0];
x input1[1];
x input2[0];
barrier input1, input2;
cx input1[0], input2[0];
cx input1[1], input2[1];
cx input1[2], input2[2];
cx input1[3], input2[3];
cx input1[4], input2[4];
cx input1[5], input2[5];
cx input1[6], input2[6];
cx input1[7], input2[7];
barrier input1, input2;
measure input2[0] -> xorout[0];
measure input2[1] -> xorout[1];
measure input2[2] -> xorout[2];
measure input2[3] -> xorout[3];
measure input2[4] -> xorout[4];
measure input2[5] -> xorout[5];
measure input2[6] -> xorout[6];
measure input2[7] -> xorout[7]
```

Additional source code for this example is available in the authors' GitHub repository, implemented in QASM 3.0 (via Qiskit), as well as in Java Strange (FX) and Q# [9]. Figure 2 presents the result of executing the quantum circuit from Figure 1 on a real quantum processor: IBM's 127-qubit *Kyiv* quantum computer. This quantum computing circuit was tested on several real quantum computing platforms, including IBM's 127-qubit

Kyiv quantum computer, Microsoft's *Quantinuum* system, and Rigetti's quantum hardware. The results of these tests, including execution times and associated operational costs, are summarized in Table 2. The comparison from Table 2 is relevant for a big amount of bytes and several running sessions, because in quantum the impact could be higher due to measurement operations on real qubits and error corrections.

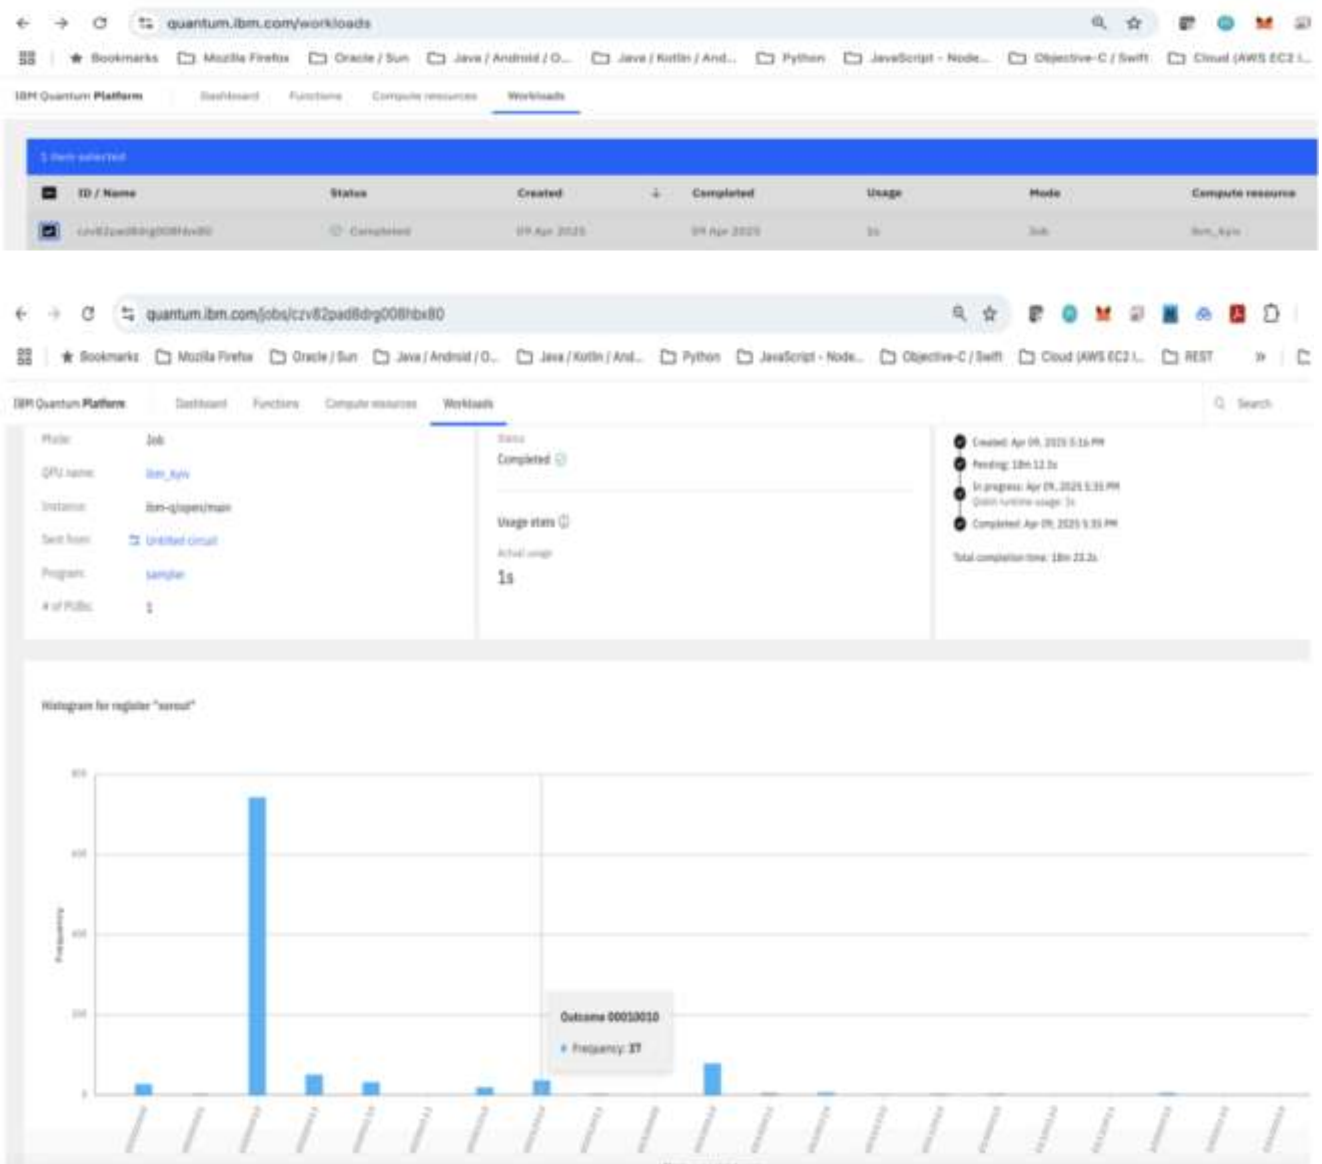

Fig. 2. Time of execution for the quantum circuit in QASM 2.0 for applying XOR function on 8 qubits (00000011 XOR 00000001 -> 00000010) using CNOT

Table 2. Content details

Subsystem	IBM Kyiv	MS Quantinuum	Rigetti
Byte-level XOR for one byte of key and file content (without optimization)	≈1 second(s)	≈1 second(s)	≈1 second(s)
Char level (2 bytes) XOR for 2 bytes of key and file content (without optimization)	≈1.2 seconds	≈1.4 seconds	≈1.1 seconds

Two Chars level (2*2 bytes) XOR for 4 bytes of key and file content (without optimization)	≈2 seconds	≈2.1 seconds	≈1.7 seconds

At the time of testing, IBM's 127-qubit Kyiv platform was accessed under the free academic tier provided through IBM Quantum Experience. Similarly, the tests on Microsoft Quantinuum and Rigetti were conducted under academic/research access programs. While no direct costs were incurred for these runs, commercial access to these platforms is priced on a per-shot or per-execution basis, with rates varying by provider and platform scale. As quantum computing moves toward wider commercialization, these cost structures will become an increasingly important factor in assessing economic viability.

Unfortunately, detailed information regarding the actual power consumption of these cloud-based quantum platforms is not publicly disclosed and cannot be easily measured by end users. However, it is possible to infer resource and time usage based on the execution characteristics of specific algorithms and circuits. From an empirical standpoint, Rigetti's platform demonstrates notable optimization when using the Quil programming language, which—although more complex than QASM—is specifically designed to enhance performance for algorithms such as the Quantum Approximate Optimization Algorithm (QAOA). Such optimizations can influence both execution efficiency and, in commercial settings, the cost of running quantum workloads.

4 Conclusions and Sustainability Assessment of Energy and Costs in Quantum Computing

The field of quantum computing systems is still in an early stage of development, and there is not yet consensus on standard metrics for assessing energy utilization, operational costs, or overall sustainability. Nevertheless, researchers have begun to systematically evaluate various quantum hardware platforms in terms of energy consumption, carbon footprint, and full life cycle impacts. Increasingly, these assessments are also considering economic factors such as capital expenditures (CapEx) and operational expenditures (OpEx), which will be critical to understanding the long-term viability and scalability of quantum computing technologies [10], [11], [12].

Superconducting quantum computers, being among the most advanced architectures to date, are also the most extensively studied in sustainability evaluations. These systems exhibit a notable energy overhead even when idle, primarily due to the continuous operation of cryogenic cooling systems. An in-depth life cycle assessment (LCA) by Cordier et al. (2023) [8] established an environmental profile for a superconducting quantum computer and compared it to a functionally equivalent classical supercomputer. In this scenario, the quantum system demonstrated a lower overall impact on climate change, environmental damage, and human health—largely because of the substantial energy consumption associated with the classical supercomputer's thousands of processing cores.

However, this advantage is only realized under specific conditions and scaling regimes. The benefits emerge when the superconducting quantum computer (SCQ) achieves problem-solving efficiency that offsets its baseline power consumption. The study also highlights that as systems scale to fault-tolerant architectures; the introduction of Quantum Error Correction (QEC) hardware imposes significant additional overhead. For example, supporting ~100 logical qubits require extensive auxiliary hardware, which substantially increases both the environmental footprint and the associated capital and operational costs. Additional racks and infrastructure required for QEC can partially negate the sustainability gains of quantum approaches. This underscores the critical importance of addressing error correction efficiency and hardware optimization in the future development of scalable and economically viable quantum computing architectures.

References

[1] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.

[2] N. Arora and P. Kumar, "Sustainable quantum computing: Opportunities and challenges of benchmarking carbon in the quantum computing lifecycle," *arXiv preprint*, arXiv:2408.05679, 2024.

[3] S. Cordier, K. Thibault, M. L. Arpin, and M. Ben Amor, "Scaling up to problem sizes: An environmental life cycle assessment of quantum computing," *arXiv preprint*, arXiv:2411.00118, 2024.

[4] S. Krinner, S. Lazar, A. Remm, et al., "Engineering cryogenic setups for 100-qubit scale superconducting circuit systems," *EPJ Quantum Technology*, vol. 6, no. 1, p. 2, 2019.

[5] C. Gidney and M. Ekerå, "How to factor 2048-bit RSA integers in 8 hours using 20 million noisy qubits," *Quantum*, vol. 5, p. 433, 2021.

[6] F. Arute, K. Arya, R. Babbush, et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, pp. 505–510, 2019.

[7] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, "Trapped-ion quantum computing: Progress and challenges," *Applied Physics Reviews*, vol. 6, no. 2, p. 021314, 2019.

[8] C. Cordier, F. Heymans, and P. Vallet, "Life-cycle assessment of superconducting quantum computers," *Journal of Industrial Ecology*, vol. 27, no. 2, pp. 251–265, 2023.

[9] C. Toma, "Quantum computing and post quantum cryptography lecture notes," [Online]. Available: https://github.com/critoma/quantum/.

[10] C. Blank and S. Kordy, "Economics of quantum computing," *arXiv preprint*, arXiv:2106.11191, 2021.

[11] F. Gyongyosi and S. Imre, "A survey on quantum computing technology and its future perspectives," *ACM Computing Surveys*, vol. 51, no. 3, pp. 38:1–38:34, 2018.

[12] European Commission, "Quantum technologies flagship strategic research and innovation agenda (SRIA) 2024," European Commission, 2024.

Teodor CERVINSKI has a bachelor's degree in Economic Informatics and a MSc. degree in Cyber Security from the Department of Economic Informatics and Cybernetics at Bucharest University of Economic Studies, Romania. His research interests are Computer Network, Computational Cryptography, Application Security, and Quantum Computing.

Cristian TOMA has graduated from the Faculty of Cybernetics, Statistics and Economic Informatics, Economic Informatics bachelor program, within University of Economic Studies Bucharest in 2003. He has graduated from the BRIE master program in 2005, with practical stage in Germany and graduated from the PhD stage in 2008. He is cofounder of the IT&C | Cyber Security Master Program (www.ism.ase.ro) and cofounder of the SECITC – The International Conference on Security for Information Technology and Communications (www.secitc.eu). Now, he is professor in Department of Economic Informatics and Cybernetics and his work focuses on the IoT - Internet of Things, Crypto Blockchain, eMbedded/Secure Elements/Mobile app, Cloud/High Performance Computing, Computer anti-viruses and viruses, Robotics, Artificial Intelligence, Quantum Computing and Post Quantum Cryptography, and computational cryptography - cybersecurity.

Claudiu BRANDAS is Professor, PhD at the West University of Timisoara, Faculty of Economics and Business Administration, Department of Business Information Systems and Statistics. He earned his PhD from "Babes-Bolyai" University of Cluj-Napoca, the Faculty of Economics in Decision Support Systems conception and design. Currently, his research interests include DSS (Decision Support System), Business Intelligence, Collaborative Systems, Business Information Systems Analysis and Design, Business Process Modeling, Information Systems Control and Audit, Artificial Intelligence, Cybersecurity and Software Project Management.

Marius POPA has graduated the Faculty of Cybernetics, Statistics and Economic Informatics in 2002. He holds a PhD diploma in Economic Cybernetics and Statistics. He joined the staff of Academy of Economic Studies, teaching assistant in 2002. Currently, he is Associate Professor in Economic Informatics field and branches within Department of Economic Informatics and Cybernetics at Faculty of Cybernetics, Statistics and Economic Informatics from Bucharest University of Economic Studies. He is the author and co-author of 9 books and over 140 articles in journals and proceedings of national and international conferences, symposiums, workshops in the fields of data quality, software quality, informatics security, collaborative information systems, IT project management, software engineering.