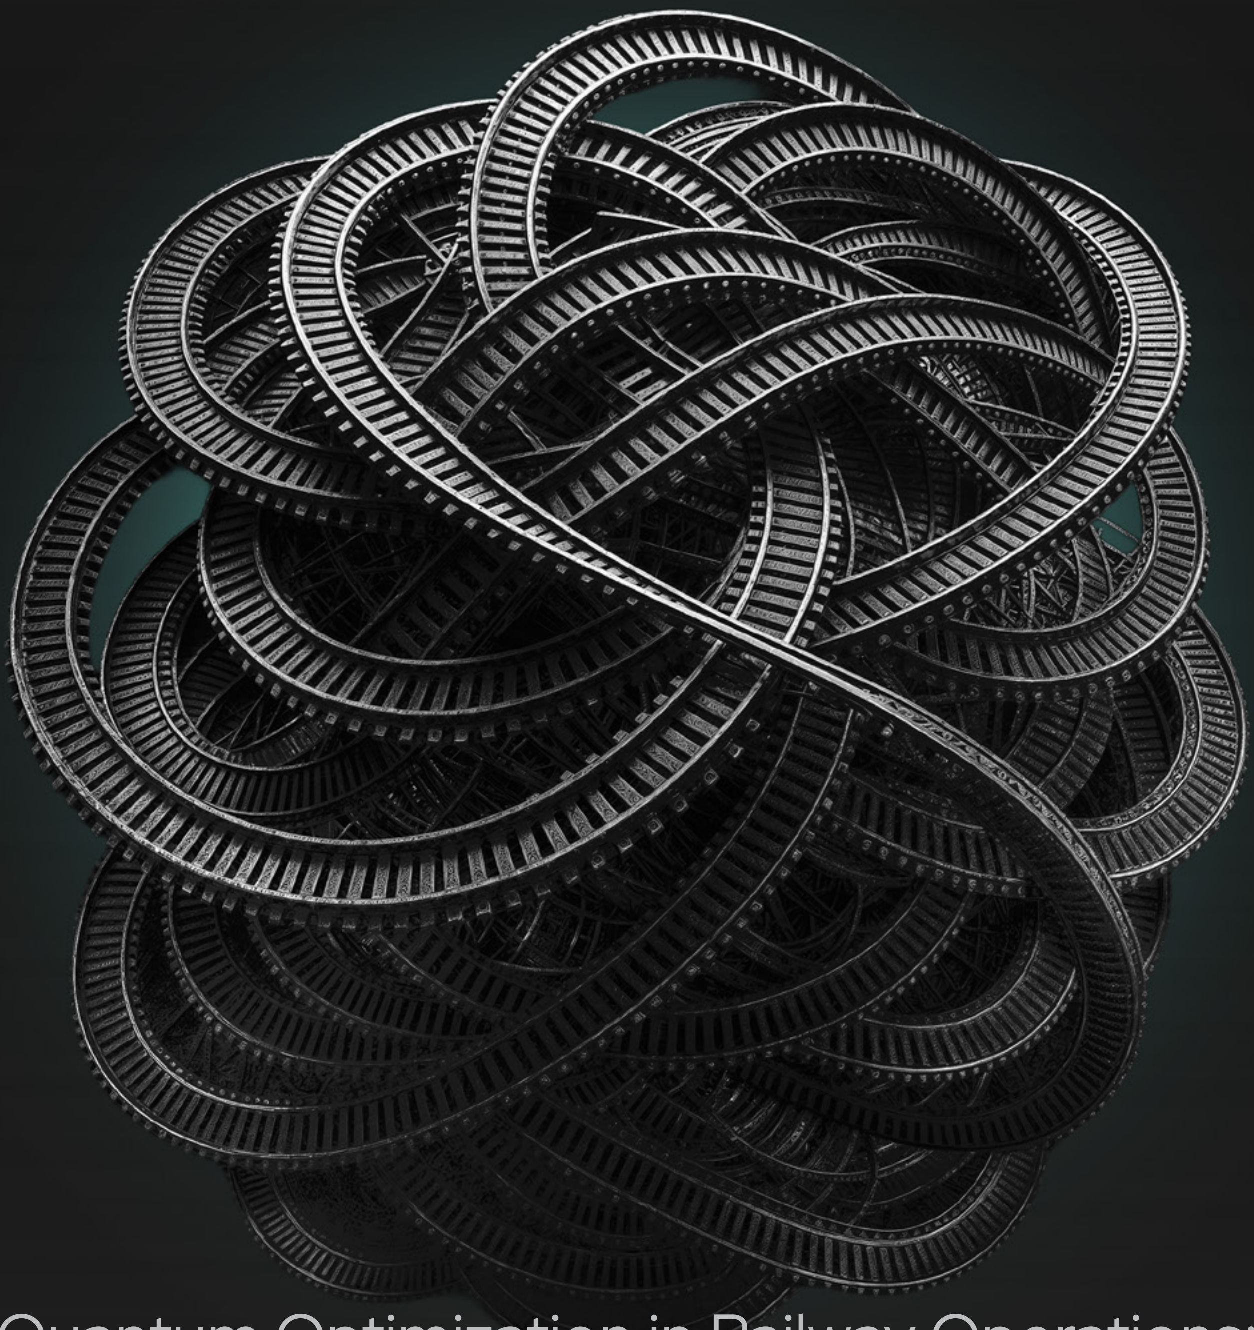

Quantum Optimization in Railway Operations:

A Hybrid Approach to Rolling Stock Planning

Technical Whitepaper

iqm.tech

Summary

Rolling stock planning is the problem of assigning trains to scheduled services in a way that minimises operational cost while satisfying a range of hard constraints. It is one of the more computationally demanding problems in railway management, and as networks grow, classical solvers struggle to find good solutions within practical time limits. This whitepaper describes a collaboration between IQM Quantum Computers and Deutsche Bahn to test a hybrid quantumclassical approach to this problem. Working from a real operational dataset, the two organisations developed and benchmarked an algorithm that embeds the Quantum Approximate Optimization Algorithm (QAOA) within a classical divide-andconquer framework. This paper reports what the approach achieved, where the current limitations lie, and what the results indicate about the path to quantum advantage in industrial scheduling.

The operational complexity of large-scale rail networks

Running a national railway requires coordinating an enormous number of moving parts simultaneously. Every day, operators must assign hundreds or thousands of trains to scheduled services, ensure no train exceeds its service interval before reaching a maintenance facility, keep the fleet cycling efficiently across the network, and do all of this in a way that keeps costs down and passengers moving. Small inefficiencies compound quickly at scale.

Rolling stock planning sits at the core of this challenge. Unlike timetabling, which determines when and where services run, rolling stock planning determines which physical train delivers each service and in what sequence, while satisfying a set of hard operational constraints around maintenance, distance, and network topology. The quality of this assignment is measured in empty kilometres: the distance trains travel without carrying passengers. Empty kilometres represent energy consumed, crew deployed, and capacity consumed with no revenue to show for it. Minimising them, within all the constraints the network imposes, is the central objective.

It is a hard problem in the formal sense. The number of feasible train cycles grows rapidly with network size, the constraints interact in complex ways, and finding a provably optimal solution requires exhaustive search that quickly becomes computationally prohibitive. In practice, large networks rely on heuristic solvers that find good solutions without guaranteeing the best one.

Empty kilometers

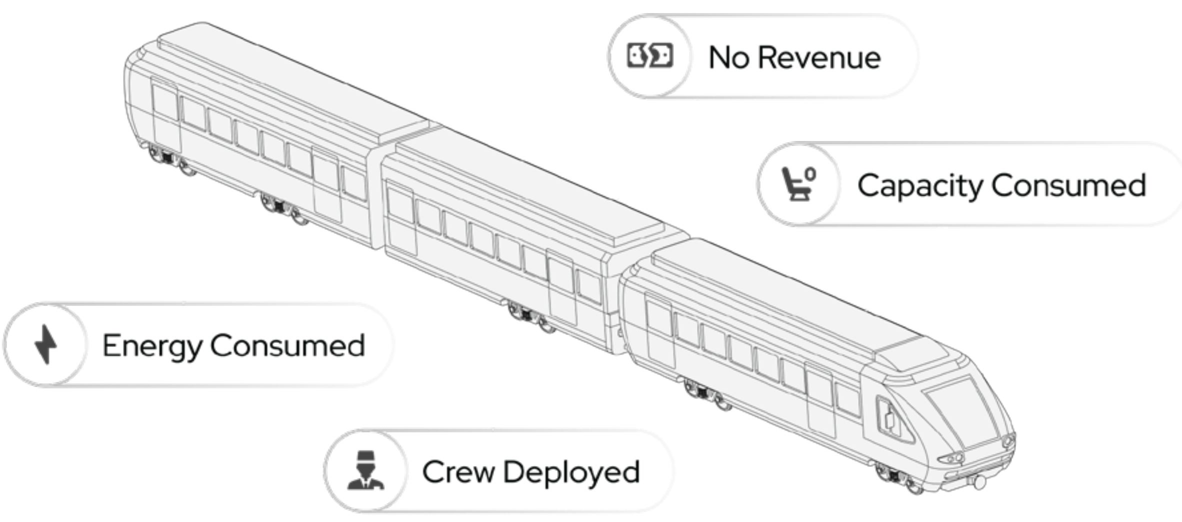

The rolling stock planning problem

For this collaboration, Deutsche Bahn provided a scheduling instance drawn from real operational data: 190 trips across five stations (Cologne, Munich, Berlin, Frankfurt, and Hamburg) over two days, with Hamburg serving as the mandatory maintenance location and a maximum cycle length of 4,000 km before each train's service interval. Each train had to complete a closed cycle of trips, return to its starting point, and pass through Hamburg for maintenance, all while the set of selected cycles collectively covered every scheduled trip in the timetable.

IQM reformulated the problem as a Maximum-Weight Independent Set (MWIS) problem on a conflict graph. Each feasible train cycle is represented as a node; two nodes are connected by an edge if the corresponding cycles cannot be selected together because they cover the same scheduled trip. Node weights reflect cycle quality, rewarding passenger kilometres and penalising empty ones. Finding the best rolling stock assignment is then equivalent to finding the highestweight set of nodes with no edges between them.

The 190-trip instance produced approximately 98,500 feasible cycles, giving a conflict graph too large for any generic exact solver to handle directly. This is what motivated the hybrid approach: rather than attempting to solve the full graph, the algorithm works through it in stages, each time selecting a manageable subgraph, solving it, and incorporating the result before moving on.

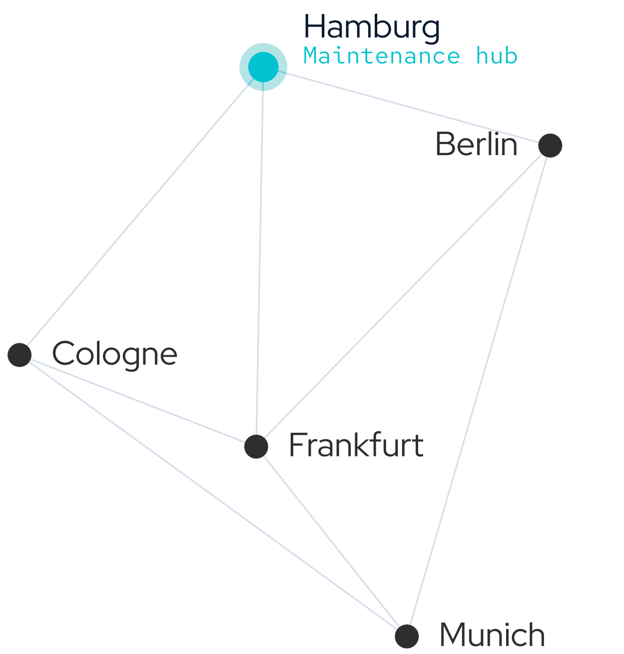

Problem instance	
Scheduled trips	190
Stations	5
Planning horizon	2days
Maintenance	Hamburg
Max systel length	4,000 km
Feasible cycles	≈ 98,500

A hybrid quantum-classical approach

The algorithm works as a divide-and-conquer loop. At each iteration, a subgraph of a fixed size is selected from the remaining conflict graph and solved using a designated solver, either classical or quantum. The cycles selected by the solver are added to the overall solution, those cycles and all their conflicting neighbours are removed from the graph, and the process repeats until the graph is empty. The quantum component, QAOA, plugs into the subgraph-solving step.

The choice of which subgraph to select at each step matters considerably. Selecting subgraphs at random tends to include many low-quality cycles, producing poor partial solutions that constrain the rest of the algorithm. Selecting purely by node weight collapses the approach into something close to a greedy heuristic. The best-performing strategy, identified empirically, was to prioritise cycles by the number of passenger-carrying trips they contain. This draws the solver toward high-quality regions of the solution space without becoming as myopic as pure greedy selection.

An important post-processing step converts raw QAOA output into valid solutions. Because QAOA samples may include conflicting cycles, a pruning procedure removes the lowest-quality conflicting node at each step until the result is a valid independent set. At small subgraph sizes, pruning contributes substantially to solution quality; its relative contribution is expected to decrease as subgraph sizes grow and the conflict structure becomes denser.

Results and key findings

The algorithm was tested across multiple solver types and subgraph sizes. Solution quality was measured in empty kilometres and fleet size. Three findings stand out.

1. The hybrid framework works at real problem scale

The divide-and-conquer structure successfully applied QAOA to a problem instance well beyond what any quantum processor could handle in a single pass. By operating on subgraphs of 20 nodes while the surrounding classical loop managed a graph of nearly 100,000 nodes, the algorithm produced feasible, good-quality solutions without requiring hardware that does not yet exist. This is the practical value of the hybrid approach: it makes quantum computation relevant to industrial problems today, not only once fault-tolerant processors arrive.

2. Larger subgraphs produce better solutions

Using a classical exact solver to isolate the effect of subgraph size, the experiments found a clear and statistically significant relationship: larger subgraphs consistently produced fewer empty kilometres. Because this effect was measured with a classical solver, it is properly understood as a property of the algorithmic architecture rather than of quantum hardware specifically. The relevance for quantum is that as QPUs scale to handle larger subgraphs, the hybrid framework is expected to translate that hardware progress directly into better scheduling outcomes.

3. The full pipeline was validated on quantum hardware

QAOA circuits at circuit depth $p = 1$ were executed on IQM quantum hardware on subgraphs of up to 20 nodes, confirming that the complete pipeline works correctly from problem formulation through hardware execution to post-processing. At this depth and subgraph size, the quantum results were comparable to classical baselines, which is consistent with what theory predicts for shallow circuits on small instances. The hardware runs were not intended to demonstrate advantage at this stage; they confirmed that the integration is sound and established the baseline from which performance will be measured as circuit depth and qubit count increase.

Scalability and disruption management

This collaboration addressed the deterministic form of rolling stock planning: assigning trains to a fixed timetable under known constraints with time to find a good solution. In daily operations, Deutsche Bahn also faces the stochastic form of the problem, where unexpected events such as train faults, track closures, or cascading delays require rapid reassignment of rolling stock on timescales of minutes. This places much stricter demands on solver speed.

A similar hybrid architecture could in principle be applied to disruption management problems of this kind. Its iterative structure would allow partial solutions to be updated incrementally as the situation evolves, rather than requiring a full restart from scratch, though this has not been tested and would require further research. The quantum subroutine can also be swapped out or upgraded independently as hardware improves, without redesigning the surrounding framework. As QPU latency falls and circuit fidelity rises, the same algorithm becomes applicable to progressively more time-sensitive scheduling problems without structural changes.

“The hybrid-classical approach allows us to use present-day quantum computers to solve an industrial-size problem, as opposed to miniature toy problems. The classical part of the algorithm preprocesses the problem and divides it into small ‘chunks’ that the quantum part can solve.”

– Dr. Jiri Guth Jarkovsky
Quantum Application Engineer,
IQM Quantum Computers

Outlook

The collaboration between IQM and Deutsche Bahn produced a technically grounded starting point for quantumassisted railway optimisation. The hybrid framework performed correctly, solution quality scaled predictably with subgraph size, and the full pipeline was validated on real quantum hardware. The remaining gap is hardware scale: the subgraph sizes at which QAOA is expected to outperform classical heuristics are not yet within reach of current processors.

The most direct path to closing that gap runs through improvements in gate fidelity, qubit count, and qubit connectivity. As those improve, the same algorithmic architecture will support larger subgraphs, and the scaling results from this work give a concrete estimate of what those improvements are worth in operational terms. Alongside hardware progress, further algorithm work on dynamic cycle generation, adaptive subgraph selection, and extensions to larger networks will continue to raise the ceiling.

Railway operators that begin engaging with hybrid quantum optimisation now are better placed to benefit from those advances when they arrive. The problem structure is well understood, the algorithmic framework is in place, and the data requirements are defined. What changes as hardware improves is how much of that framework the quantum solver can carry.

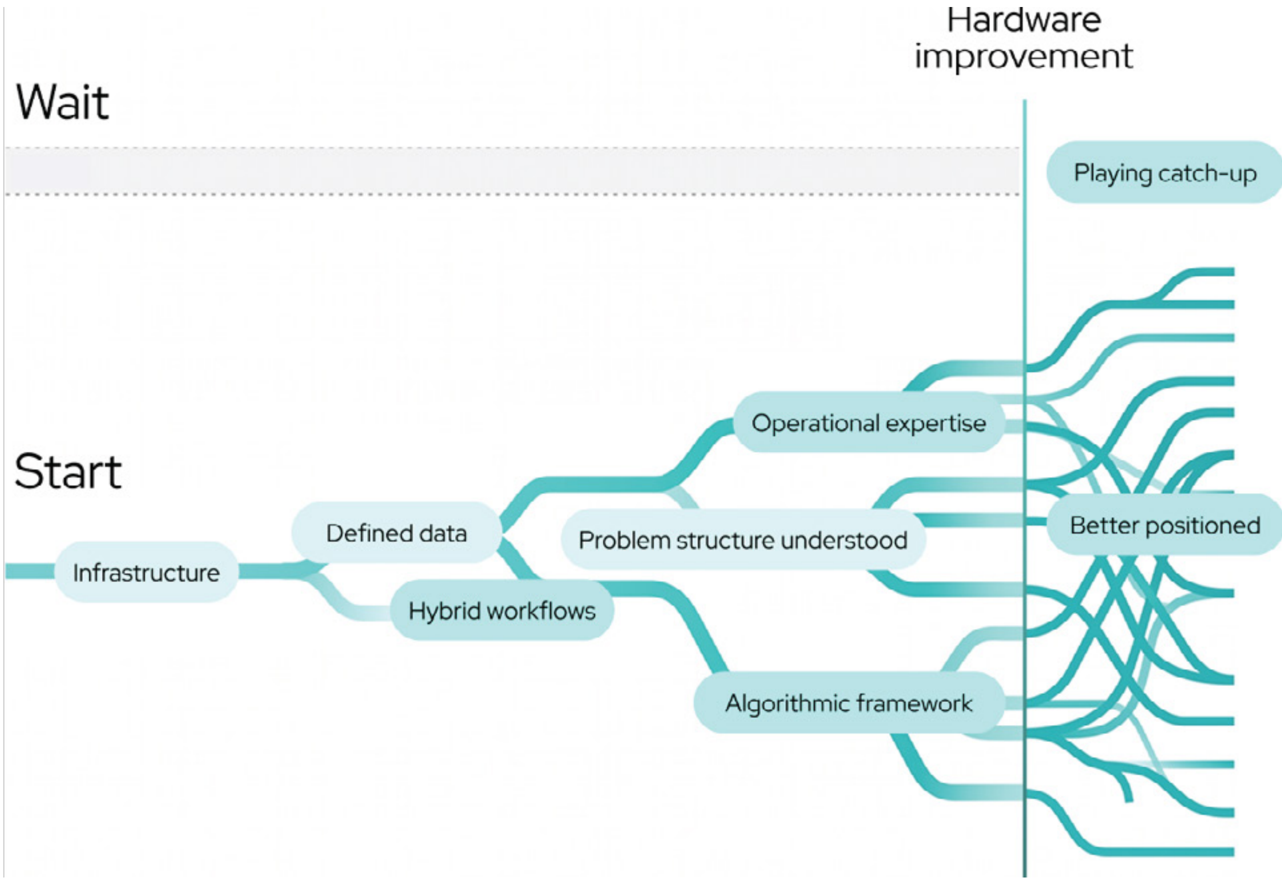

"What we need at the end from the quantum vendors is a 24/7 system to use it in production. When you see the cloud vendors, they offer a system with an availability of 99.98 percent. That is what we need from the quantum vendors as well. Otherwise, you can use it only for short experiments, but you cannot use it in production."

– Manfred Rieck
Head of Quantum Tech,
Deutsche Bahn

About Deutsche Bahn

Organisation

Deutsche Bahn AG is one of the world's largest transport and logistics companies, headquartered in Berlin. With around 230.000 employees and operations in more than 130 countries, DB is Germany's primary rail operator for both passengers and freight. DB Systel, its technology subsidiary, is responsible for the group's digital and IT infrastructure, including the operational decisionsupport systems used across the network.

Operations

Deutsche Bahn runs long-distance, regional, and urban rail services carrying over four million passengers daily in Germany. Freight operations are handled through DB Cargo, one of Europe's largest rail freight carriers. DB Engineering & Consulting provides rail infrastructure expertise to operators and governments internationally. Across this scale of operations, rolling stock planning is a continuous and computationally intensive activity, making it a natural focus for research into advanced optimisation methods.

`deutschebahn.com`

About IQM

Quantum Computers.
Delivered.

IQM Quantum Computers (Nasdaq: IQMX) is a global leader in superconducting quantum computers, delivering full-stack quantum systems and cloud platform access to enterprises, research institutions, universities, high-performance computing centers, and national laboratories worldwide. IQM's on-premises deployment model gives customers direct ownership and control of their quantum infrastructure. Founded in 2018 and headquartered in Finland, with major operations in Munich, IQM employs over 400 people and operates across Europe, Asia, and North America. IQM is the first publicly listed European quantum company on Nasdaq Stock Market.

iqm.tech

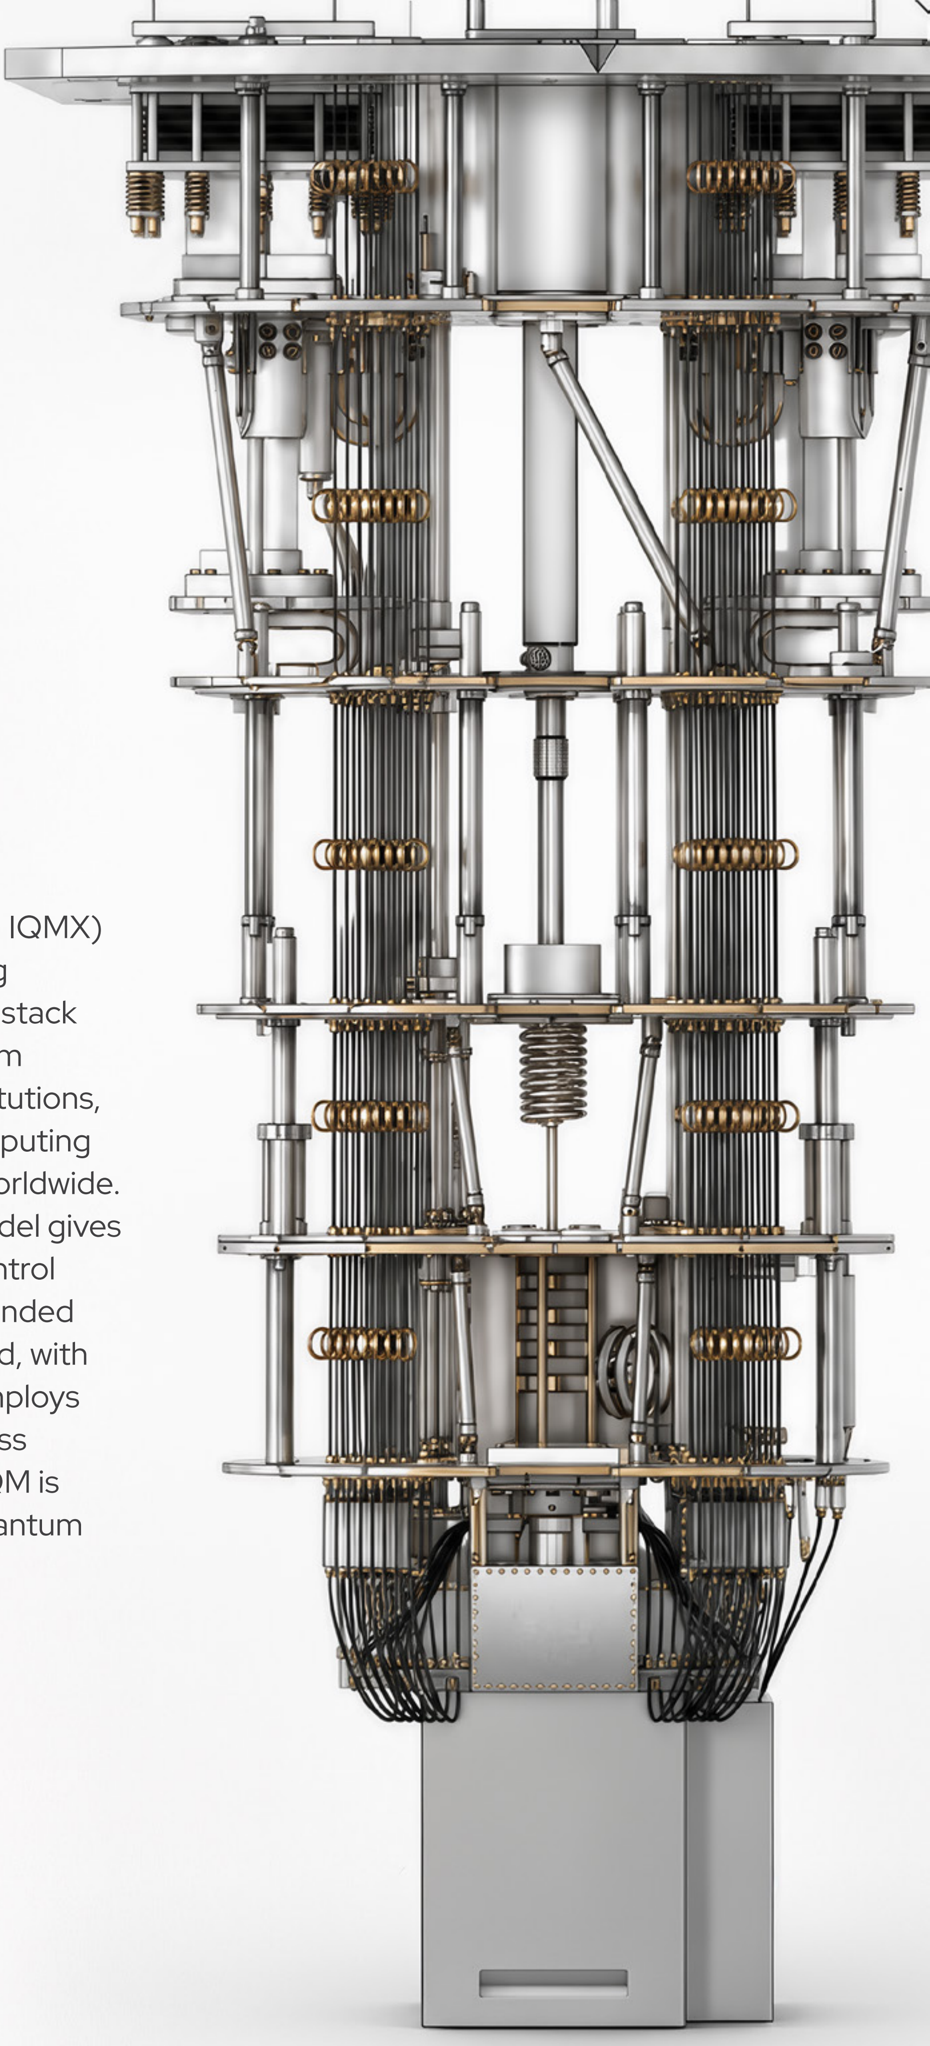

Contributors

Jiri Guth Jarkovsky
IQM Quantum Computers

Patricia Bickert
Deutsche Bahn AG / DB Systel

Elisabeth Wybo
IQM Quantum Computers

Martin Leib
IQM Quantum Computers