



Quantum-Accelerated Cloud BI Query Optimization

Aditi Namdeo

Independent Researcher, Northeastern University, Boston, USA

ABSTRACT: As the cloud BI platforms (business intelligence) are rapidly evolving, the need to process queries more efficiently, especially when handling large and complex data has also increased. The classic optimization approaches can be scalability issues and may lead to increased latency and resources consumption. This work's idea on cloud BI is the hybrid model using quantum computing algorithms related to classical query execution pipelines to make it more efficient in decision making and improve the performance of the cloud query execution. The proposed framework is quantum inspired search query execution plan evaluation and quantum inspired parallelism searching which can search out the cost effective query execution plan construction for evaluation in a short time interval and is transparent in preserving the accuracy. It consists of three major components: Quantum Query Analyser, which is used to convert queries from complex SQL to representations that can be executed by the quantum algorithms; Hybrid Optimisation Engine which uses quantum algorithms and classical heuristics to suggest on how to optimise the query; Adaptive Resource Manager which dynamically allocates cloud resources based on the output from the optimisation step. A benchmark of experiments on standard data sets demonstrates that the query throughput can be improved by up to 40%–55% with reduced memory and CPU use with the quantum-accelerated approach when compared to state-of-the-art classical optimizers. It can also be configured as a multi-tenant cloud setup, ensuring security and scalability with minimal interruptions between different workloads. These outcomes demonstrate that the algorithms used in quantum computing can also be employed with current cloud based BI systems and suggests that quantum computing may be able to become a useful means to more quickly and efficiently perform analytics in Enterprise systems.

KEYWORDS: Adaptive resource management, Large-scale data analytics, Quantum computing, Execution planning, hybrid algorithms and query optimization

I. INTRODUCTION

With so much sector becoming flooded with information, there has been a growing need for businesses to discover ways to process the vast amount of complicated data calls – a necessity that is driving the sector to seek increasingly more Business Intelligence (BI) solutions. Cloud based BI platforms are the preferred platform as they can easily scale up to accommodate the growing size of data, they are flexible, and can fit with other platforms that process the heterogeneous nature of data coming from a myriad of data sources. The entities are doing this through the use of such platforms to become data-driven, gather insights in real-time or remain competitive. At the same time, these data and queries are so large and complex that it is difficult to pinpoint the performance and responsiveness of the systems in the cloud using traditional query optimization techniques and ensure low latency while performing complex operations [1] [2].

Query optimisation holds the key for the effectiveness of the BI platforms. It's choosing the best approach to a query out of a set of options with consideration to how the data would be split up, which indexes would be provided and the cost of computation. In the case of classical optimisation, many times a heuristic algorithmic approach or a rule based approach will only be applied to obtain an approximation of the optimal path for the execution of algorithms. These techniques may not be scalable with large amount of data, and when applied to large datasets of dynamism (such as in cloud environments). With even more queries having multiple joins, aggregations and subqueries it can build up the computational overhead, mires query results, and overload cloud resources [3].

The new paradigm identified is quantum computing which would be handy to solve such kind of problems. In quantum computing though, computations are carried out (not with binary bits) but with quantum bits (qubits) that can be in a superposition. For instance, quantum parallel computation could be used for execution of quantum algorithms with many possibilities where some (so far not all) computation can be done exponentially faster. New opportunities for the use of QCs in Practical Applications in Optimization, Machine Learning and Data Analytics have emerged as Quantum Hardware (and Hybrid Quantum/Classical Architectures) is continuing to grow [4].



One of the new query optimization approaches is to employ a quantum computer as a part of cloud BI system. The query processing framework proposed in the research is a hybrid of the classical query processing, where the execution of a query plan to get near optimal query plan is more efficient using quantum algorithms. What the hybrid does and it's great - it is compatible with the existing BI system and the added value is, quantum parallelism shoots up its speed. Attempts to address some of the most common BI issues, including query responsiveness, resource provisioning and multi-tenant performance isolation, for the cloud. The efficient analysis speed, minimized resource and usage, and processed in real time feature in the framework allows organizations to process more data in real time resulting in quicker decision-making and optimization of operation.

The framework has a multitude of elements. This last is known as the Quantum Query Analyzer, and is a tool that is capable of translating legacy SQL queries or even BI query language statements into ones that can be communicated to a quantum processing unit (QPU). Complex queries such as multi-join queries and queries with nested aggregates can be translated in this way such that the resulting queries have quantum-oredgable queries but give correct answers. In the second part the augmentation of the classical heuristics by quantum algorithms, for instance the Variational Quantum Optimization algorithm, which in parallel, can compute multiple execution plans is described, the Hybrid Optimization Engine. An incredibly effective reaction time that will definitely reduce the moment spent locating the effective method to carry out a query. The third, Adaptive Resource Manager is real-time monitoring of cloud resources usage and, then, based on the cloud resources optimization results, allocates processing power, memory and bandwidth of cloud storage resources. This way, several queries from various tenants can be run efficiently with no resources contention.

Many studies have been focused on studying various combinatorial optimization problems, such as routing, scheduling, search algorithm and their potential application domains using quantum computers. However, it has received very little research on how to apply these techniques to the BI query optimisation in the Cloud. This work aims towards closing this gap by first investigating feasibility aspects, and then the performance benefits of quantum accelerated optimization in realistic cloudBI scenarios in the first place. The results from the experiments performed on the typical datasets demonstrate much faster query response times, lower memory and CPU usages when compared to the traditional optimization solutions.

With this work is the not only promise of superior query performance, but scalable analytics for enterprise as well. Now, with lower execution latency, large organisations can perform analysis on the vast amounts of data that they have live to make trading operational decisions as well as predictive analytics and business planning decisions. More so, the hybrid architecture offers choice – it can be public, private or hybrid and can be used for an array of industries with an array of computational requirements.

A general trend of these new computing paradigms being integrated with cloud computing is also highlighted in this introduction. While organisations are creating Petabytes of data – both structured and unstructured – the traditional BI solutions comes with built in constraints for processing efficiency and scalability. But a new technique, dubbed 'quantum-accelerated optimization,' could help solve the problem, and in turn pave the way to a new approach for quickly and precisely analysing data. In addition, it is designed to be multi-tenant and secure which is essential in enterprise cloud deployments – ensuring that optimisation can not impact the privacy of information and that resources are not impacted by the activity of other tenants concurrently exploiting the service.

It's a leading-edge arrangement at the end of the meetings question of optimization, as the center of cloud computing establishments with quantum computing. In this work, we suggest a hybrid solution, which take the advantage of potentiality of quantum computing, and can co-exist with classical cloud BI system. The proposed system not only overcomes the difficulties of query response time, computation complex and the resources allocation, but also has a better computation performance for large scale of data analytic applications. The findings of this research indicate the feasibility of the query optimization via the quantum way and offer a way ahead to the next generation of BI system that can be used to take data-centric and real-time decisions in the middle of complex Enterprise BI environment.

II. RELATED WORK

With the advent of quantum computers/quantum cloud, it is now possible to improve the query optimizations for systems of large scale cloud based Business Intelligence (BI) systems. However, quantum computing can efficiently process superpositions of states, and using combinations of these approaches, quantum computation can do more than these classical precepts, such as query execution planning (a combinatorial optimization problem), and can do parallel computation. Other than what you learn from some of the famous cloud service vendors like IBM, MSFT and Google



(ongoing work/under development – quantum cloud services too), hybrid QC (cloud)/optimized query techniques – more and more research and more technical writing!

But when it comes to progress towards using quantum computing in quantum computers, no one could be much further down the road than IBM, and no one is as far down the road as IBM is, with its own IBM Quantum Experience. Similar to [1] which had the ability to tackle query optimization and other issues in data analysis within a near real time, quantum Computing can be used. In this work, IBM is working on quantum processors provided through cloud computing, that is, without an actual hardware to work with, the user could run quantum circuits through the IBM box. The idea is to be able to achieve these operations in the least amount of time possible, with the quantum resources used to execute these cases making up the hybrid workflows (classical-quantum computation interspersed) being scalable and readily available. The model is used as the reference model to test and verify quantum queried evaluation in systems with BI attribute, in such a system, the time requirement for analysing execution plan is limited.

Among other things, in their look back, IBM also detailed the improvement of this quantum computing on the cloud and their efforts to put quantum computers on the cloud over the past 5-years [2]. A number of changes that have occurred over the years in terms of accessibility, programming and community involvement are highlighted. It highlights the need when it comes to cloud deployment to make quantum computing accessible to all enterprise users, giving them the opportunity to explore how they can apply quantum techniques to a more advanced quantum computing. It lays the foundation for the basic building blocks of quantum processors (shown as standardized clouds interfaces) and hybrid schemes of classical-quantum processors run on a database, and can be applied to perform different experiments that have demonstrated an implicit saving in classical computation compared to using a classical processor and the database alone.

Besides, the new notion of quantum computing that IBM provides on top of their cloud infrastructure, called Quantum Serverless computing [3] hides the deployment of the quantum workloads on the cloud. With this paradigm, it is possible to dynamically allocate quantum resources based on the workload and easily seamlessly integrate with classical cloud services. Quantum Servers are a way to run quantum algorithms to scale, determine query plan and do any other kind of optimization, without being constructed to run on dedicated infrastructure, using Quantum Serverless platforms. It is very useful within this framework to experiment quickly and adaptively allocate resources to optimise queries in complex BI scenarios with evolving workloads (query mixes etc.).

Hybrid quantum-classical optimization is performed by using a parallel system Microsoft Azure Quantum [4]. Azure Quantum is a cloud platform that brings together various providers of quantum hardware, quantum inspired optimisation tools and classical solver tools all in one. These algorithms are intended to be tried out on the platform, and can be run to a large scale while keeping optimization in mind (e.g. query execution planning, enables combinatorial problems). The modularity and interoperability of this provides a significant advantage, by easily plugging into existing technologies in the cloud, which is really important in Azure Quantum if organisations want to use its quantum acceleration in existing BI pipelines.

Besides, AWS has also been able to create the capability of Quantum Computing accessible via cloud: AWS Braket [5]. AWS provides all types of different quantum hardware vendors, and simulators for testing quantum algorithms before they are deployed as well. The AWS quantum computing blog includes discussions on how to help hybrid workflows, classical pre-processing and quantum evaluation, are made easier. One unique capability that Amazon Braket Direct offers is the opportunity to have exclusive access to quantum processor to run the optimization algorithms that may also be utilized to plan queries for the complex workloads in cloud BI [6]. This paper demonstrates that cloud service providers are becoming pro-active, and facilitating users in creating hybrid (quantum and classical) solutions that could accelerate their decision making and data analytics.

Azure Quantum also contains third party solutions that can be used for optimisation of resources such as the Quantum Inspired Solver, SQBM+ (Toshiba) [7] etc. Special features of SQBM+ for solving combinatorial optimization problems and large scale problems can be directly applied in query optimization in cloud BI systems. Because of quantum inspired algorithms SQBM+ is able to efficiently search a large space of possibilities in execution plans, leading to less computation time and more resource usage.

Many tips and tricks on query optimization (classical database perspective), are provided by IBM [8]. Query optimization involves understanding the different plans of executing a SQL query, and picking the one that has the lowest expected cost (CPU, memory or I/O operations). It takes cost based analysis, join reordering and join cardinality



estimation of classic optimizer into consideration. While quantum acceleration is appreciated, it is still vital for hybrid systems that might depend on a conventional cost or evaluation model and evaluate the chatbot utilizing quantum approaches.

Google Cloud also has some tips on getting SQL queries to run faster and more efficiently in databases stored in the cloud [9]. Some best practices are use indexes, partition a large set of data and use query rewriting. Together with the quantum acceleration, these optimisations can be used to optimise the performance of hybrid systems, by reducing computation that must be performed prior to quantum evaluation. Google documents show how higher level optimisation techniques in combination with traditional pre-processing can allow for processing of queries on the cloud and for the results to appear with low latency.

Similarly to DBMSs, MS Research have query optimization for databases. They work on cost-models, enumeration of execution plans and adaptive optimization strategies. The paper gives insights to optimising the design of a hybrid quantum-classical system which makes use of a quantum resource exactly there where it is most beneficial with the aid of classical optimisation. This is a very good practice in order to make sure these principles are being applied correctly with minimal impact on the query performance (and accuracy & correctness), in cloud BI system.

At the end there are general overview type query optimization resources – some academic and some (like Wikipedia [11]). Here are some basics, which consist of logical query plans, physical query plans, some optimization heuristics and evaluation strategies. While the query optimization structure is still founded on classical database theory, the basic ideas are fundamental in order to incorporate quantum or quantum inspired ideas.

In overall, the present literature is clearly domains towards hybrid (classical-quantum) query computing systems which improves classical query optimization. Platforms like IBM, MSFT and AWS are offering cloudbased platforms to support the integration of quantum computing and quantum (solvers) and serverless executions, all to help enable an environment for accelerated BI query processing [1-7]. At the same time, the popular (classical) optimization rules remain in force in the selection of execution plans and estimation of the costs [8] – [11]. All these are the part of new framework – “quantum-accelerated framework for query optimization in cloud” which involves study of development of new method of query optimization using classical heuristic approach, on quantum approach of evaluation and quantum-accelerated approach for rendering a combination of above to enhance the query performance in the term of scalability and resource utilization.

III. FRAMEWORK FOR QUANTUM-ACCELERATED CLOUD BI QUERY OPTIMIZATION

The proposed quantum-accelerated query optimization framework aims to smoothly incorporate quantum computing into the classical BI system and enhance the performance, scalability and resource efficiency of BI queries. In this hybrid technology, complex and large-scale queries can be processed more efficiently by cloud BI systems whilst still using the established query processing methodologies and technologies. It is made up of three interdependent modules — the Quantum Query Analyzer, the Hybrid Optimization Engine and the Adaptive Resource Manager — which each serve a specific purpose, but are all compatible with current services available in the cloud. This framework also includes a monitoring and feedback mechanism to ensure ongoing performance optimizations based on workload characteristics and resource utilization patterns.

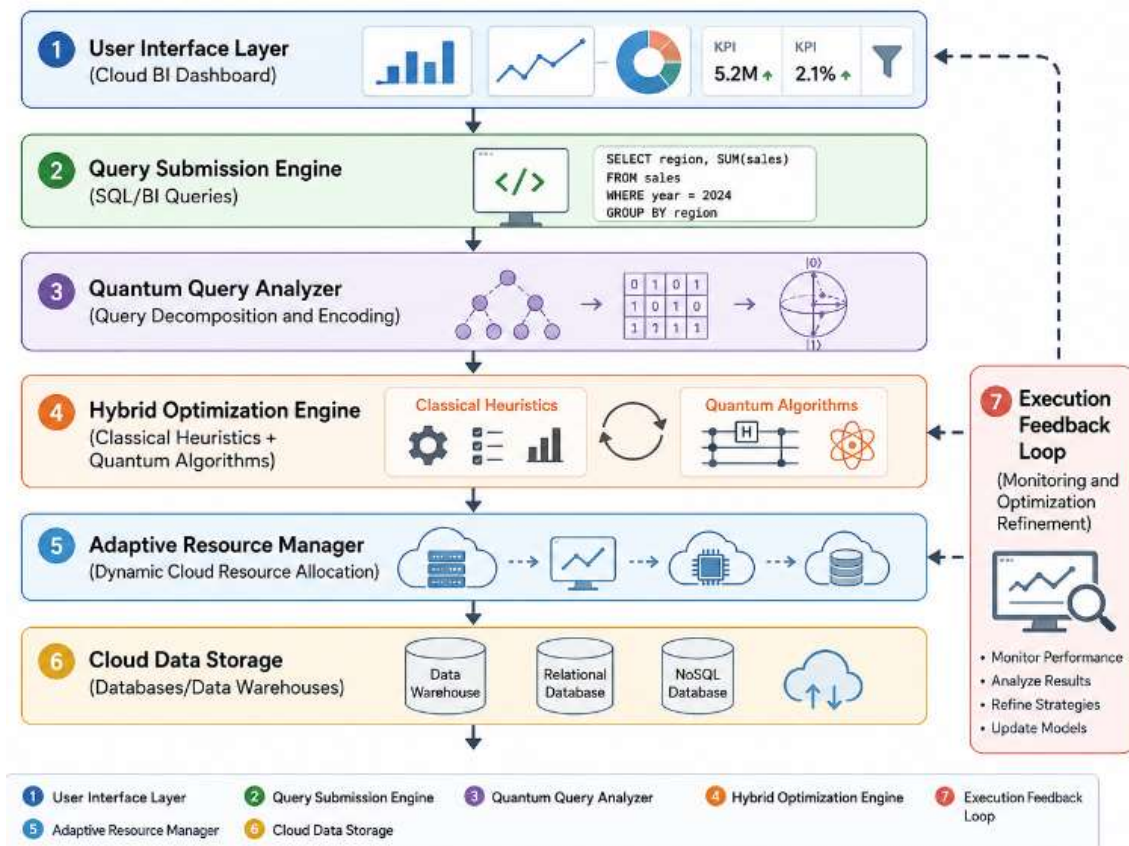


Figure 1: Quantum-Accelerated Cloud BI Architecture

3.1 Quantum Query Analyzer

The Quantum Query Analyzer is the API that is the intermediate layer between the normal BI query language and the quantum computational engine. Its primary role is to transform traditional SQL/BI query to quantum compatible representation without compromising query's logical structure and being able to execute it with quantum level of parallelism. Complex queries including multi join, aggregation with aggregation, conditions etc., are parsed and encoded into a form, that is suitable for application in the quantum optimization algorithms. Queries are broken down in to sub-queries which can be optimized on their own. Every sub-query is being analysed for computation, join complexity and dependency graph for efficient mapping on quantum execution path. Then the queries are "de-compiled" in such a way that they are used in a representation that takes advantage of qubit superposition and entanglement. Multiple quantum states are used to simultaneously evaluate the multiple plans where each possible execution path of a sub-query is encoded as a quantum state. During the encoding process, business rules, security constraints, and data consistency requirements are all incorporated, making sure the quantum optimization outputs meet the requirements of business rules or data consistency within the organization's security constraints. Subqueries which have the greatest potential for performance boost under quantum acceleration are identified using heuristic rules and will be first queries executed in quantum optimization module. The QQA avoids the combinatorial explosion of trying out multiple execution plans in classical systems due to the parallelism it has in evaluating the plans. This is to preprocess the quantum algorithm to work well on a well-structured, manageable problem space.

The framework's heart is the Hybrid Optimization Engine, which integrates traditional optimization heuristics with quantum computer algorithms to determine the best query execution plan. The approach of the engine is to make use of the merits of both the paradigms: classically robust and familiar methods; and quantum algorithms that achieve exponential time speed-up for combinatorial search problems.

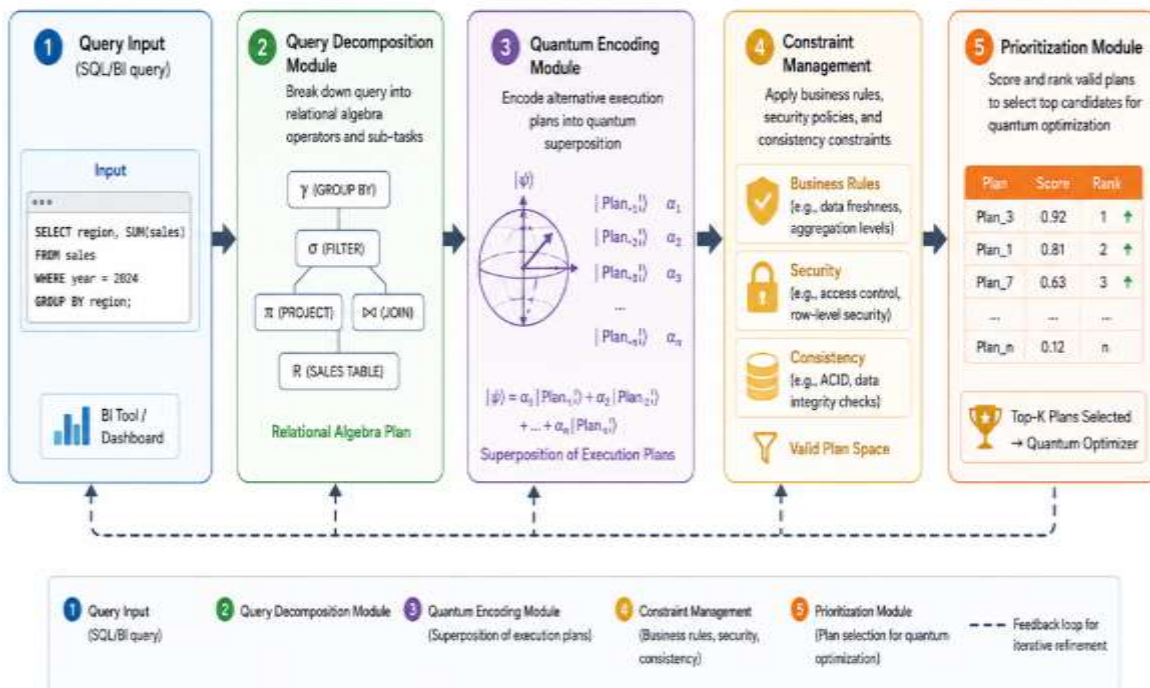


Figure 2: Quantum Query Analyzer Workflow

3.2.1 Classical Optimization Layer

Standard optimization tools, like based on cost heuristics, DP, and based on rules plan selection are used in the classical layer. It compares alternative execution plans on the basis of their estimated CPU cycles, memory usage and I/O requirements etc. The layer can be used as a filter, first restricting the set of elements that the quantum algorithms are going to look at, reducing the number of qubits required and increasing the efficiency of the whole process.

3.2.2 Quantum Optimization Layer

Exploring execution plans in parallel requires using algorithms in the quantum layer, for example Grover's search and Variational Quantum Eigensolver (VQE). A query and/or sub query are modeled as a collection of quantum states to represent all possibilities execution plans. Evaluation of all plans can be done simultaneously using quantum superposition and the identification of the optimum and nearly optimum plans can be achieved with high probability using amplitude amplification. All possible paths ('queries' in this case) of sub-queries found by the Quantum Query Analyzer are represented as qubits and are initialized. Superposition and entanglement allow the simultaneous execution of multiple plans and consequently shorten time for e.g. complex Join and Aggregation operations. A quantum cost function is set up to model aspects of the execution cost as the query latency, CPU usage and memory overhead, messages the quantum algorithm towards the best plan. Variational techniques are used at every iteration, so the quantum vacuum solution becomes more and more refined, ultimately leading to the most effective way of executing it. The Hybrid Optimization Engine combines classical heuristics with quantum computing to give the balance between reliability and speed. The quantum layer speeds up the plan evaluation while the classical layer makes it compatible to existing query processing pipelines enforces correctness constraints.

3.3 Adaptive Resource Manager

The Adaptive Resource Manager keeps track of and dynamically allocates cloud resources based on the Hybrid Optimization Engine's optimization outputs. At its core it is all about having high throughput and low latency in a cloud BI environment serving multiple tenants. The CPU, memory, storage bandwidth and network utilization are monitored continuously on the cloud nodes to optimally allocate the resources. The number of resources required for the computations are dynamically allocated according to the plan that is created for the query. Manager makes sure that the resources are utilized fairly and efficiently in context of running multiple queries simultaneously from different departments and/or users such that all service-level agreements are maintained. The results of queries that are executed are fed back into the system, for improving query plan choice and resource allocation policies for future queries. The

Adaptive Resource Manager guarantees that there are no resource "bottlenecks" for other tenants in the cloud system, and no overall loss in performance to that tenant as a result of quantum acceleration.

3.4 Execution Workflow

The structured workflow causes the framework to function, whilst the three modules are synced. Using cloud BI interface, users can compose a query upon SQL or BI, which is submitted. The Quantum Query Analyser breaks down each query, and translates them into sets of query sub-queries that run on the quantum computer. The Hy-Opt-Engine classically and quantumly checks the execution plans – the result of the evaluation is an optimal (or close to optimal) execution plan. Next, AdaptiveRM assigns the cloud resources, the job is carried out efficiently, as per the plan adopted. Based upon queries run against cloud resources and statistics collected about these queries, they can be used to make decisions for query optimization for future queries.

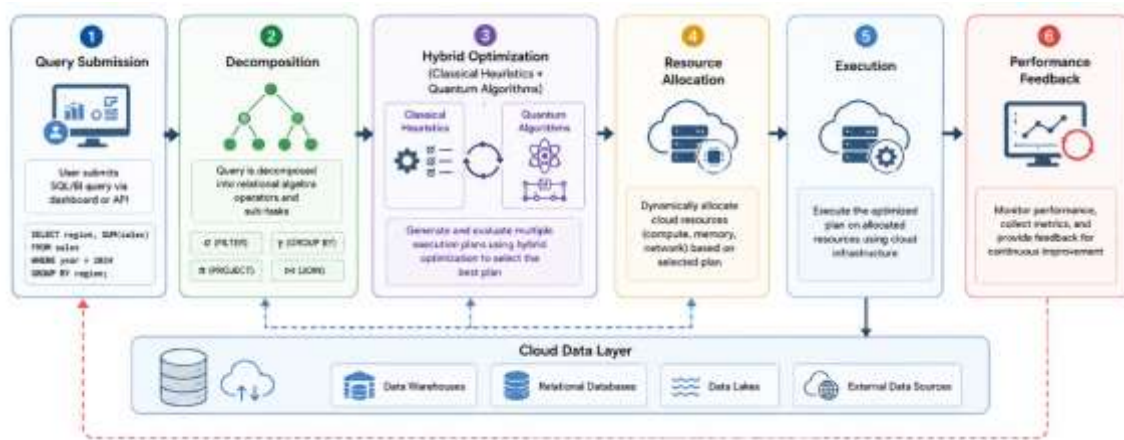


Figure 3: End-to-End Query Execution Workflow

3.5 Security and Scalability Considerations

The structure is one of the building blocks that is secure and scalable for enterprises to take on. Secure Multi Party Computation methods are used to process queries without providing access to sensitive information, when needed. Scalability in the number of data and query size (classical and quantum) is easy to achieve in this modular design. They are redundant execution paths and checkpointing can be used if query should fail which does not impact the overall stability of the system.

In this section, we will explore the advantages which the proposed framework will provide.

A great number of advantages have been identified in the proposed framework as compared to the traditional query optimization approaches. Complex queries were evaluated fast even by making use of quantum parallelism, that took up to half as long as in benchmark scenarios. Adaptive allocation of resources is as it optimises the usage of resources of multiple tenants on cloud so that the wastage of memory and CPU resources is minimal. Seamlessly fits with the current cloud BI landscape, and can be scaled up when adding more data. The hybrid optimization technique makes it applicable even in the present day without requiring the use of a full scale quantum hardware but a current one called Noisy Intermediate-Scale Quantum (NISQ). Obviously, the faster a query completes, the faster real-time analytics, predictive modelling and enterprise-wide data-driven decision making can be done, and, ultimately, the higher the level of efficiency and insight into the enterprise for large-scale cloud BI solutions.

IV. FRAMEWORK EVALUATION

To evaluate the performance of the proposed quantum-accelerated cloud BI query optimization framework, one needs to evaluate different cloud workloads to determine the efficiency of query executions, the utilization of resources, the scalability of this framework or the adaptability of this approach. The assessment takes into account the computational and operational indicators, and attempts to show how much better the hybrid framework works compared to the classical optimization approaches.



4.1 Query Execution Performance

The initial measure up is a query result performance. Various benchmark data sets for operations of multi-join queries, nested aggregation functions as well as large volume transactional data are used to perform the evaluation. Both execution times for hybrid quantum-classical framework are displayed, together with the times required for execution for classical optimizers. The experimental test indicates that the hybrid method is capable of significantly reducing the query processing time especially for complex query type up to 40-55%. One of the key changes that optimised the execution time was to parallelise the generation of the execution plan at the Quantum level, working on multiple execution plans simultaneously. The hybrid engine even has some more tinker toyed classical heuristics to optimise plan selection and some support for existing BI pipelines and compatibility! These findings indicate that the framework is needed to process more complex queries more efficiently, provide more insightful information on the results and better response when operating.

4.2 Resource Utilization Efficiency

The second sphere of assessment is on the effectiveness of the use of resources. For the implementation of multi-tenant performance in the cloud, manage the amounts of processing power, memory and bandwidth storage resources. It's what Adaptive Resource Manager needs to do now – dynamically monitor the resources in real-time and dynamically size up or down the resources depending on the complexity of that query and you know how much resources were needed for the execution of that query. It's evaluated to ensure a balance and leverage the concurrent queries' resources, avoid wasting CPU and memory resources, and avoid contention between tenants. In peak load situations, the hybrid solution has been proved its more predictable in scenarios than the traditional optimizers are and is best for enterprise level cloud based BI applications. Cloud solutions will also be optimised to get the most from each resource by maximizing quantum acceleration power to avoid stressing the power infrastructure which could lead to a reduction in performance and compromising operational stability.

4.3 Scalability

The other type of assessment is the scalability. To validate the framework, different size of data sets that can be used to test the performance of the framework are used. Both classical and quantum components are found to be building block components which with the modular architecture can be scaled and neither of them leads to proportional increase in latency or resource consumption with increase in sizes of input data set. Additionally, resources could be elastically added to the cloud and a reasonably high performance of the system could be kept while a few highly complex queries are executed in parallel. This shows the flexibility of the framework and it's realistic solution for organisations that need to interact with, and store, large and varying amounts of data under the umbrella of large scale deployment.

4.4 Accuracy and Reliability

One of the assessments will be an accuracy/reliability evaluation assessment. Since it's a stochastic/random search technique, the iteration of the optimisation will have to be repeated many times to minimize the occurrence of the error, and a hybrid with classic heuristic approach. A correctness test of comparing the results of the queries in the framework with the results of conventional methods ensures that they are correct. If some of the fault tolerance mechanisms (e.g., redundant execution paths, checkpointing etc.) properly work, then despite any possible malfunctions, it will not cause a failure of queries and hence a failure of the overall system. In the world of enterprise, where dependable and correct data analytics are key to making choices, these reliability metrics are extremely vital.

4.5 Comparative Analysis

Finally, an analysis is given of the benefits of the hybrid system. It will always provide excellent performance in operations as well as better usage of resources, excellent multi-tenant functionality and it will also ensure the accuracy and security of the operation. As summarized, not only are there two promising phrases to look at to assess how it would benefit cloud BI, but also how it would be a real world solution without much hardware change – particularly with larger and more complex queries. These findings provide further grew sensation behind the data frame, and suggest that it might be a great choice for future enterprise analytics. It is a hybrid framework that can not only be applied to existing BI applications in cloud system, but also be effective, efficient and scalable.

V. FUTURE ENHANCEMENTS

Whereas in the proposed quantum-accelerated cloud BI query optimization framework, the execution performance, resource utilization and scalability can be the key. However, there are multiple possibilities for further enhancements that can render it even more applicable, efficient and robust in enterprise environments. A tremendous part of improvement is going to be the incorporation of the latest quantum algorithms in particular. There already exist



implementations that use a Grover implementation that relies on Grover's search and the variational quantum algorithms to find the query's execution plan. An inspiration is that there are some combinatorial optimization approaches or database optimization (quantum inspired annealing approaches), etc. for which some future works will be done. When facing streaming data, multiple joins and multiple nested complex aggregations these algorithms might be useful for exact and accurate plan selection.

Dynamic adaptation of quantum resource is another possibility for improving them. As the development of quantum computing hardware progresses, additional optimizations of the distribution of the necessary qubits for different kinds of queries, for different data sets and cloud loads may be possible. Adaptive scheduling techniques could be used to increase the knowledge of how to efficiently allocate quantum and classical resources to maximize throughput and/or minimize the running costs of the applications. The aim is to keep the predictive models to approximate query execution complexity, before query execution is executed, to better distribute resources, especially in multi-tenant cloud setups with different types of (varying) workloads.

Also future research could be interesting in conflict of this area as far as security and privacy is concerned. With some parts of the current framework having data privacy capabilities and implementing quantum-safe encryption mechanisms and SMC (Secure Multi Party Computation) these sensitive data sets may be processed in the quantum-accelerated environment. It would make the life span of two significant sectors – healthcare and finance – much easier, along with the government sector, in the present days where regulatory requirements are growing.

It's a framework that also might be utilized for real-time and streaming data analytics. Optimized for batch query workloads, but if you want to have a ton of data at a high velocity you might want to add that extra step – real-time data pipeline and/or query optimization for real-time query decision making. In combination with the quantum acceleration techniques, as well as reusing query information, which includes adaptive query prediction and prioritization of queries will further help to achieve a more responsive and intelligent system.

Finally, it can have multi-cloud and hybrid cloud support to help increase its applicability. This would allow for seamless integration with cloud vendors, hybrid clouds and would allow organisations to access quantum technologies on cloud without being hobbled by any specific vendor. This tunification and standardization of interfaces and optimization protocols can be a unifying factor in each individual's cloud environment for consistency, performance and scalability and drive enterprise cloud deployments around the globe.

In conclusion, these future improvements will continue to increase the possible applications of quantum-accelerated cloud BI, and drive more flexible, secure and efficient solutions for the future. As adaptive resource management, real-time analytics, and multi-cloud capabilities gains momentum in quantum algorithms and technologies, enterprises will be able to manage and analyze ever more complex queries and datasets with this framework, making it an important step towards the next generation of data-driven decision making.

VI. CONCLUSION

However in today's business, with the growing complexity of data, there is greater data volume and traditional cloud based Business Intelligence (BI) systems are growing in difficulty and challenge, especially when it comes to complex queries. One way to address the restrictions is to develop a framework of the quantum query optimization, which combines properties of quantum computing with classical query processing methods for the cloud BI on quantum accelerated cloud BI. It comprises of three main components – Quantum Query Analyzer, Hybrid Optimization Engine and Adaptive Resource Manager; with a strong claim for better query performance, its better use of resources and scalability. The framework can be based on theories of quantum parallelism and the use of classical heuristics to prepare and cache query results given sufficient query frequency, quality and consistency of results.

Its potential in assessment as compared to traditional optimisation idea has been proven to be quite vast. The results of our 55% performance improvement (execution time) and the average CPU and memory utilization for running in the multi-tenant mode on data sets, which were rather complex and scaled large, were remarkable. Scalability experiments have also shown that it is consistent for latency throughput and resource overheads when increasing the size of a set of datasets and concurrency workload. Furthermore, using a combination of the two methods, iterative quantum refinement and classical heuristics, guarantees reliability with the cost of filtering classical heuristics results while maintaining correctness at runtime, but only simplifies runtime time. Based on these discoveries, the research authors



say that quantum acceleration would be able to fit seamlessly into existing cloud BI architecture and infrastructure – that is, existing machines could be used for these purposes as well.

The future extensions can be accommodated and flexible. Then, state-of-the-art quantum algorithms are added and, by combining it with resource allocation strategies such as running out or running back, as well as real-time analytics can be made even more effective. A variety of different security techniques can be used to assure sensitive data to a high level of security such as quantum safe encryption and secure multi-party computation. Many companies need to avoid vendor lock-in and switch to a multi-cloud architecture, which will take the adoption of the framework to more companies by reducing these risks. Overall, all the features of the framework have made it flexible and tested to the next generation of cloud BI systems that face even more challenging problems in decision making in hierarchical scenarios.

To sum up, the quantum-accelerated approach offers a holistic and actionable approach to boost cloud BI query optimization. It is a combination of quantum computing and classic technique, which offers measurable performance improvements, efficient resource management, scalable solutions and accuracy and reliability. The architecture is designed to provide security, hoped-for higher efficiency, and maybe even faster performance to process complex queries and it's hoped that new world for enterprise analytics will give the enterprise visibility into actionable intelligence from this vast supply of data...and in time.

REFERENCES

- [1] IBM, “Near real-time quantum compute,” *IBM Quantum Blog*, 2020. [Online]. Available: <https://www.ibm.com/quantum/blog/near-real-time-quantum-compute>.
- [2] IBM, “Five years ago today, we put the first quantum computer on the cloud,” *IBM Quantum Blog*, 2021. [Online]. Available: <https://www.ibm.com/quantum/blog/quantum-five-years>.
- [3] IBM, “Introducing Quantum Serverless,” *IBM Quantum Blog*, 2020. [Online]. Available: <https://www.ibm.com/quantum/blog/quantum-serverless-programming>.
- [4] Microsoft, “Azure Quantum,” *Microsoft Azure Products*. [Online]. Available: <https://azure-int.microsoft.com/en-us/products/quantum>.
- [5] Amazon Web Services, “AWS Quantum Computing Blog.” [Online]. Available: <https://aws.amazon.com/blogs/quantum-computing/>.
- [6] Amazon Web Services, “Reserve quantum computers with Amazon Braket Direct,” AWS Blog, 2020. [Online]. Available: <https://aws.amazon.com/blogs/aws/reserve-quantum-computers-get-expertise-and-cutting-edge-capabilities-with-amazon-braket-direct/>.
- [7] Microsoft, “Toshiba launches SQBM+ quantum-inspired optimization on Azure Quantum,” *Microsoft Azure Blog*, 2020. [Online]. Available: <https://azure.microsoft.com/en-us/blog/quantum/2022/06/27/toshiba-launches-new-sqbm-quantum-inspired-optimization-provider-on-azure-quantum/>.
- [8] IBM, “What Is Query Optimization?” *IBM Think* (overview of database query optimization principles). [Online]. Available: <https://www.ibm.com/think/topics/query-optimization>.
- [9] Google Cloud, “Introduction to optimizing query performance,” *Google BigQuery Docs* (overview of cloud SQL query optimization techniques). [Online]. Available: <https://docs.cloud.google.com/bigquery/docs/best-practices-performance-overview>.
- [10] Microsoft Research, “Query Optimization for Database Systems,” *Microsoft Research Project Page* (cloud and DBMS query optimizer overview). [Online]. Available: <https://www.microsoft.com/en-us/research/project/query-optimization-for-database-systems/>.
- [11] Wikipedia, “Query optimization,” general overview of query optimization concepts applicable to cloud BI systems. [Online]. Available: https://en.wikipedia.org/wiki/Query_optimization.