NetQIR: An Extension of QIR for Distributed Quantum Computing

Jorge Vázquez-Pérez[1], F. Javier Cardama[1], César Piñeiro[2], Tomás F. Pena[1,2], Juan C. Pichel[1,2], and Andrés Gómez[3]

[1]Centro Singular de Investigación en Tecnoloxías Intelixentes (CiTIUS)
[2]Departamento de Electrónica e Computación, Universidade de Santiago de Compostela
[3]Galicia Supercomputing Center (CESGA)

arXiv:2408.03712v1 [quant-ph] 7 Aug 2024

***Abstract*—The rapid advance of quantum computing has highlighted the need for scalable and efficient software infrastructures to fully exploit its potential. While current quantum processors have significant scalability problems due to the limited number of qubits on each chip, distributed quantum computing offers a promising solution by networking multiple quantum processing units (QPUs). To support this paradigm, robust intermediate representations (IRs) are crucial for translating high-level quantum algorithms into executable instructions across distributed systems. This paper introduces NetQIR, an extension of Microsoft's Quantum Intermediate Representation (QIR), specifically designed for distributed quantum computing. NetQIR is designed to meet the specific needs of distributed quantum systems by incorporating functions to manage quantum and classical communications between QPUs. The main objective is to facilitate the development of new distributed compilers by improving the integration and execution of quantum programmes in a distributed infrastructure, taking advantage of modular architectures to improve scalability. By extending QIR to support distributed quantum computing, NetQIR aims to complement and add capabilities to an already supported quantum IR and, at the same time, take advantage of the tools previously created for QIR. Throughout this paper the specification of the intermediate representation is introduced, including the basic instructions necessary to enable distributed quantum computing in an abstract form independent of the target machine.**

I. Introduction

The evolution of computing has progressed from simple mechanical calculators to modern-day classical computers, which have significantly transformed numerous fields including science, engineering, and everyday life. Despite these advances, classical computers face limitations in solving certain complex problems efficiently, such as factoring large numbers, simulating quantum systems, and optimising large-scale systems [12], [22]. This has led to the emergence of quantum computing, which leverages the principles of quantum mechanics to process information in fundamentally new ways, offering the potential to solve these intractable problems more efficiently than classical computers can achieve [8], [27].

Over the last few years, the development of a comprehensive software stack for quantum computing has gained importance in order to be able to program quantum devices in a scalable and easy way. This software stack includes quantum high-level languages, compilers, and runtime environments, designed to enable the programming and execution of quantum algorithms on quantum devices [17], [20]. High-level quantum programming languages such as Q# [32], Quipper [16] or Qiskit [2] facilitate the development of quantum algorithms by abstracting the complexities of quantum hardware [30].

For the efficient development of these software tools, quantum code compilers will play a crucial role. A compiler is a software that translates high-level languages into low-level instructions that quantum processors can execute [1]. In classical computing, the concept of IR is introduced as an abstract-machine code to facilitate the development of new compilers [31]. This concept is extended in the world of quantum computing to allow a common IR for target high-level languages and a starting point for low-level instructions. The main objectives of using an IR is to optimise the quantum code and ensure compatibility with various hardware backends [18], [24].

Despite these advances, one of the critical challenges in quantum computing remains the scalability and noise of the qubits. Current quantum hardware is limited by the number of qubits that can be reliably maintained and manipulated on a single chip, thus complicating the development of more complex algorithms [3]. These limitations have led to the development of new computing approaches, one of which is distributed quantum computing (DQC) in modular architectures where multiple QPUs are networked together to work on a problem collaboratively [6], [21], [36]. DQC uses both quantum and classical communications to distribute and synchronise computations across QPUs, thereby potentially overcoming the scalability constraints of individual quantum chips [5], [23], [28].

DQC introduces the need for an IR that can efficiently handle the complexities of distributing quantum operations across multiple QPUs in addition to represent quantum communications as classical. Currently there is no distributed quantum IR that allows interoperability between other languages or systems, therefore this requires the extension of an IR for quantum computing such as Microsoft's QIR, with the goal of extending existing features with new communications and distributed computing directives.

NetQIR, our proposed extension to QIR, aims to address these challenges by providing an IR specifically designed for DQC. The main objective is to adapt QIR, an IR commonly

used because it is based on LLVM IR and allows taking advantage of the optimisations already developed for these systems, thus obtaining an IR with greater interoperability and machine-abstract independence. In order to properly explain it, first, in Section II the related work to this one is presented. Here both DQC IRs and other distributed quantum works are presented. Then, in Section III the protocols in quantum communication are presented: *teledata* and *telegate*. After all the concepts required are introduced, in Section IV the NetQIR data types and functions are introduced and thoroughly explained. After this, and looking to exemplify the difference between the functions presented in the previous section, in Section V a representative example is presented with two different implementations shown and compared. Finally, in Section VI, a set of conclusions in stated, along with the future work that this article leaves ahead.

II. Related work

Regarding the related work, not a lot of IRs have been considered for DQC in contrast with the notable amount of them developed for monolithic quantum computing. In fact, not a lot of high level languages—or even libraries—have also been developed. A more exhaustive analysis on this can be found in a review of the DQC state of the art by Barral et al. [5].

From monolithic quantum computing we will highlight only one IR and that has already been mentioned in the previous section. It represents the core and the starting point of our work: Microsoft's QIR [26]—-from now on, it will be referred to as QIR—. This IR is based on the LLVM IR [14] in an attempt to integrate quantum computation into the LLVM infrastructure. LLVM is a versatile framework for building compilers and code transformation tools. It lets developers write high-level language code that can be efficiently compiled to machine code for various architectures, with extensive support for code optimisation and analysis. Among the multiple tools and softwares included in LLVM there is the aforementioned LLVM IR which bridges the gap among the multiple front-ends to the architectures. Back to QIR, this approach of extending the LLVM IR to include quantum directives aims to achieve an integration with the classical compilation stack in order to benefit from its advanced tools to facilitate the path of producing high efficient quantum instructions. This work will extend QIR—and, therefore, it will further extend the LLVM IR—by adding the necessary directives to perform quantum communications. Throughout this manuscript this extension will be explained and exemplified.

Now, regarding the quantum communications, there are some works that are considered the basis of this one. First of all, there is the Quantum Message Passing Interface (QMPI) [17]. As its own name implies, it represents an adaptation of the classical Message Passing Interface (MPI) [13] protocol for quantum communications. In this work, they observed a clear analogy between the classical send and receive directives with the quantum send and receive directives and they took advantage of this analogy to define them. Moreover they differ between copy semantics and move semantics. These terms are usually referred to, in the literature, as telegate or remote gate—for the copy semantics—and teledata or teleport—for the move semantics—. We will understand the difference when talking about the communication directives in the following section [34].

They also defined collective operations, which do not present the same analogy to classical ones. This is mainly due to a reason: the non-cloning theorem. The fact that arbitrary quantum data cannot be copied erases the possibility of having collective operations in quantum communication that operate on data as collectives operations do in classical communications. For instance, a classical broadcast sends a copy of a specified data to every node on the communicator. A quantum broadcast cannot be implemented following this scheme because it would not be allowed to send a copy of the quantum data to each node. QMPI tries to solve these problems regarding collective operations, as it also tries this work. In future sections our collective operations are presented and the pertinent comparison with the ones introduced by QMPI is made.

In terms of IRs for DQC two are found: InQuIR [25] and NetQASM [9]. The first one, InQuIR, is developed from the starting point solely as an IR for DQC. Their primary motivation stemmed from the absence of a dedicated IR for distributed quantum systems. InQuIR stands out for formally defining the grammar of the IR. Using this formal syntax it defines the operational semantics of the IR, which allows to define and predict how the InQuIR programs will behave under several circumstances. They even propose some important examples as deadlocks and qubit exhaustion, and they propose a roadmap solving these type of inconveniences. But InQuIR has several drawbacks: the syntax definition is quite inconsistent, mixing quantum and classical data without differentiating, and it provides a too low level approach with explicit generation of the Einstein-Podolsky-Rosen (EPR) pairs and instructions that acknowledge the architecture of the machine, which should be left to the backend, not to the IR. The InQuIR it also contributes with a toy compiler which, by giving it a QASM code and an architecture returns the corresponding InQuIR code. This exemplifies the abstraction problem that InQuIR presents: an IR should be platform independent.

As an alternative, as we already mentioned, there is NetQASM [9]. NetQASM is not so clearly defined as an IR from the get-go as InQuIR is. This is because it is much more than that. NetQASM presents an abstract model of the architecture composed by an application layer—responsible for the classical communications between nodes—and a so called quantum network processing unit (QNPU)—responsible for the quantum computations and communications—. This hints the scope of NetQASM: the Quantum Internet. This means, it is thought for quantum networks and sets aside the *between-cores* communications which do not need these extra layers that NetQASM presents. Moreover, NetQASM presents a basic language, called *vanilla*, and a set of variations specifically designed for the different quantum architectures,

called *flavours*. As they themselves state, the vanilla version acts as an IR and the different flavours act as the assemblies. NetQASM, as an IR for DQC, has the main disadvantage that requires the aforementioned abstract model so much network-oriented. It also does not consider conditional gates, which are constantly employed in quantum communication protocols, as part of the IR. What they actually do is performing a measurement, sending the result to the application layer and waiting until the application layer returns a subroutine with the gate—if the measurement was 1—or without the gate—in the opposite case—.

As it can be noticed, every work here mentioned has some aspects that, from our point of view, suppose or could suppose a drawback for an IR. We will try to resolve all of these inconveniences presented in the different works here exposed. Along the work there will be references to these problems in the precise points where our proposal presents a response to one, or more, of them.

III. Communication protocols

Quantum communication is an exploitation of one of the main features of quantum computing: entanglement. Because of this, there are a large number of variants for communicating quantum information using specific protocols, each with its advantages and disadvantages [10], [11].

This section explains the two communication protocols considered to perform quantum information sharing for DQC: *teledata* and *telegate*. In this case, both techniques use an entangled EPR pair, with one qubit of the pair residing on a QPU and the other qubit located on a physically separated QPU. Those EPR pairs are employed to create a link between the two QPUs, allowing, through sending and receiving classical information from measurements of specific qubits, the quantum data to travel from one QPU to the other.

Figure 1 shows the basic structure of the *teledata* [7] (in Figure 1a) and the *telegate* [15] (in Figure 1b). The objective is the same in both techniques, starting from a state $|a\rangle = \alpha|0\rangle + \beta|1\rangle$ it is necessary that the QPU_2 (remote) can compute using this information through the EPR pair $|\Phi^+\rangle$. Each protocol is explained below:

The effect of the first one, the *teledata* protocol, is to transmit the state of the qubit $|a\rangle$ in the QPU_1 to an empty qubit in QPU_2. This transmission involves a teleportation of the quantum state, so that at the origin the qubit collapses when measured and is transmitted to the destination qubit.

On the contrary, the effect of the latter, the *telegate* protocol, is to generate a pair in the state $\alpha|00\rangle + \beta|11\rangle$ where the first qubit of the pair is in the QPU_1 and the second qubit is in QPU_2. This second qubit is used as a control qubit for a controlled operation. Taking into account that the qubit intended for use as control is of the form $|a\rangle = \alpha|0\rangle + \beta|1\rangle$, it can be observed that using the second qubit of the pair will have the same effect, performing a controlled operation in QPU_2 with the state of the qubit in QPU_1.

The main difference that can be observed between teledata and telegate is that in the former the status is transferred and the computation is performed locally. On the other hand, in the case of telegate, the state is not transferred since quantum gates are performed remotely.

Table I compares both techniques by evaluating four characteristics of interest for this type of protocol:

1) *Collapsed qubit:* indicates whether the source qubit collapses once the protocol is executed, thus being necessary a reset of the qubit.
2) *Entanglement result:* refers to the scope affected by the entanglement generated between the remote and local qubit. This entanglement can be local to the computation node or global to the distributed system.
3) *Measures:* how are the measures to be taken to implement the protocol.
4) *Number of synchronizations:* the number of synchronizations between the QPUs needed just to execute the communication protocol.

Characteristics	Collapsed qubit	Entanglement result	Measures	Number of synchronisations
Teledata	✓	Local	Local - Local	1
Telegate	✗	Global	Local - Remote	2

TABLE I: Comparative features between teledata and telegate techniques.

As can be seen, in the case of *teledata* the qubit collapses when sending the information, requiring a future reset. This occurs because the quantum state is completely transferred to the target node, therefore the operations will be performed locally and, as a consequence, the generated entanglement will also be local. On the other hand, the measurements are performed at the same time in the two qubits local to the QPU, therefore it will only be necessary to synchronise the two QPUs once.

In the case of *telegate* it is completely different, since the quantum information is shared as a reference, without destroying the qubit, so it will not be necessary to reset it at any time. Since a reference is shared, it also implies that the entanglement generated will be global to the distributed system. Finally, an initial measurement is performed locally and the last one in the remote QPU, therefore two isolated synchronisations between the QPUs will be necessary. This second synchronisation, also known as Cat-DisEnt, is up to the compiler to decide when it should be applied, especially when the qubit is no longer used.

As can be seen after observing both protocols, both have advantages and disadvantages and there is no clear predominant option, being very variable the decision taken depending on the problem to be solved, therefore the proposed specification will take into account both options.

IV. Functions and data types added by NetQIR

In this section we are going to cover the different functions and data types added by NetQIR. In a lot of ways these present a resemblance with MPI. This is intentional and happens for two main reasons: the years of research behind

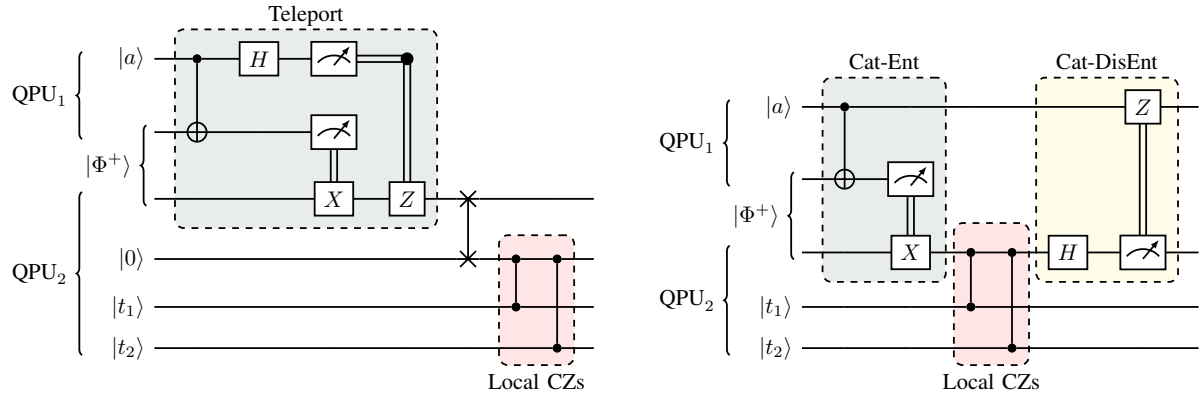

(a) Teledata.

(b) Telegate.

Fig. 1: Examples of teledata and telegate circuits for the application of CZs.

classical communications and the intention of facilitating the understanding of NetQIR directives in a high-performance computing (HPC) environment (that is, precisely, where they are going to be employed).

First of all, classical communications research come a long way. This is something which cannot be simply ignored and, in fact, the mistakes made in this area have to be taken into account to avoid committing them again. MPI was thought to be a robust, long-term standard for enabling communication in parallel computing environments, but the aim of most researchers at the time was to find an automatic parallelisation technique that avoided the need of explicitly formulate the communication [4], [19], [33]. In DQC something similar is happening with the automatic partitioning of circuits, which is equivalent to an automatic graph partitioning problem and this is considered a NP-hard problem [29]. And this is the main reason why NetQIR adopts a similar approach to MPI in classical communications. Nevertheless, a lot of the literature [5] regarding DQC revolves around finding automatic circuit partitioning methods employing a graph representation. In this work, the automatic partitioning methods are not completely discarded. In fact, the developer can implement or use them provided that they transform the code to the NetQIR syntax.

And, secondly, employing a MPI-like syntax in NetQIR can also facilitate the understanding of the directives. This is due to the fact that HPC and quantum computing are two very different fields that need each other. In this case, NetQIR is an HPC software that aims to target quantum computing devices. If this software follows a syntax similar to a widespread software as MPI is, the learning curve will be much more affordable than if NetQIR is completely different. Being an extension of QIR also reiterates on this idea, because a lot of the data types and functions employed are going to be the same (for instance, quantum gates and measurements).

Once this similarity with MPI has been addressed and justified, is time to talk about the implemented functions and data types. A specification of the functions and their arguments has been developed in order to formally display the NetQIR extension [35]. Here we will further explain the concepts specified there.

First, we will talk about the data types added. The data types added on top of the QIR data types are groups and communicators, referred to as `%Group` and `%Comm`. These two types encapsulate the same concepts as in classical computing, meaning that they are responsible of defining the scope of the functions and grouping subsets of nodes. The only difference is that the nodes, in this case, are quantum nodes, i.e., QPUs.

Regarding the functions, they can be grouped in three categories: *state functions*, *data type functions* and *communication functions*.

- *State functions*. There are only two functions for this type which are called `__netqir__init` and `__netqir__finalize`. As their names indicate, they are responsible for initialising and finalising, respectively, the distributed environment.
- *Data type functions*. These type of functions are responsible of interacting with the new data types: groups and communicators. This is, basically, create and obtain information about them and the nodes contained by them—as, for instance, information about the topology[1] of a communicator—.
- *Communication functions*. These are the quantum communication directives and, so, they they represent the most important group of functions. As we have already mentioned, they include point-to-point and collective communication directives. But this will be further explained in the following section.

In order to simplify the manuscript, the data type functions are not going to be further explained. The reason for doing this is because the concept of these functions is very similar to those in the MPI standard. Groups and communicators are widely employed in classical communications and so are their associated directives. For further reading about these functions

[1] In this context, when talking about *topology* the concept referred to is not the quantum topology, but the node topology inside a communicator.

and their use we refer the reader to the aforementioned specification [35].

A. Communication functions

Now the bread and butter of this work is arrived at: quantum communication directives. As already hinted, the communication directives will be of two types: point-to-point and collective. The first one, the point-to-point directives, are responsible for communication solely between two nodes, while collective directives represent one-to-many communication. Point-to-point directives will be explained first and then collective ones, in order to progress from a simpler to a more complicated perspective.

1) Point-to-point communication: Point-to-point communication in quantum computing is exactly the same as in classical computing in the sense that one node sends/receives information to/from another node. The only thing that changes is the fact that in the classical case the information is purely classical and in the quantum one could be classical or quantum information.

In Table II we can see these functions. They are divided in two subgroups: the sending and the receiving functions. The names of the subgroups are self-explanatory: the first one is responsible of sending the information and the second one of receiving it. Even more, we can notice that for every sending function there is a receiving function as a counterpart. This fact is not arbitrary, this is due to the fact that for every send function in a node, there has to be the specific counterpart receive function in the destination node, and vice-versa. If not, the result might be arbitrary or, depending on the compiler, a compilation error could be thrown.

Talking about the sending functions, first we will talk about the most basic function: `__netqir__qsend`. It represents the part of the circuit in the sending QPU (that in Figure 1 corresponds to QPU_1) of both the telegate and the teledata. We say both because this function leaves to the compiler the choice between teledata and telegate. If the user wants to use one of the two protocols specifically it needs to use `__netqir__qsend_teledata` or `__netqir__qsend_telegate`, depending on whether the desired protocol is teledata or telegate, respectively. Lastly, we specified the function `__netqir__measure_send`. This corresponds to the sending of one classical bit from a measurement. This function is implemented to allow the user in the frontend to develop their own quantum communication protocol. This way, if any other protocol arises, it can be implemented and compared with the telegate and teledata without the need of using another IR. Moreover, each of these recently explained functions have a version with the "array" tag in their name. This tag just indicates that, instead of just sending one qubit, an array of qubits is being sent. This means that it just changes the type of data of the function. Now, regarding the receive functions, each sending function in Table II has at the same level its receiving counterpart. So the explanation that we just made for every sending function is analogous for its receiving counterpart.

It must be specified that, for instance, if in one node the function employed is `__netqir__qsend_teledata`, then the receiving node has to employ `__netqir__qrecv_teledata`. Using `__netqir__qrecv_telegate` would end up in a bad result because two parts from distinct protocols would not correctly match and, therefore, give an unexpected result. The only functions that could be employed with more flexibility are the `__netqir__qsend` and the `__netqir__qrecv`. This is because they do not specify which protocol they are using and, therefore, if the other node does, it simply changes to the specific protocol. Again, with the array functions happens exactly the same, but with a change in the data type.

Something that can be noticeable in this variations on the sending and receiving functions is the fact that both the protocols and the data sent could be specify inside the `__netqir__qsend` and the `__netqir__qrecv` functions, without the need of specifying so many different directives. This is done in order to facilitate possible compiler optimisation. If these variations were flaged with an argument of a function, they would be unknown for the compiler. This could present a problem for the compilers, because they may not be able to perform some optimisations due to the impossibility of knowing the type of data sent or the protocol employed. This way everything is transparent for the compiler in case it is of interest for a given optimisation.

Also, we mentioned in the related work that InQuIR had some problems: problems in the syntax definition and in the low level approach that they took by specifying the creation and the entanglement swapping between cores. It can be noted that with the solution presented in this work the syntax is not a concern because is inherited from the LLVM IR and, more important, the EPR generation and the entanglement swapping directives are completely separated from the IR. They are left as a responsibility of the compiler to the backend, that is the one that actually knows the architecture of the machine below (and, therefore, its connectivity). We also said that NetQASM presented some inconveniences: the tax of the abstraction model and the non-inclusion of the conditional gates as a possible operation. Here, we never imposed an abstract model for the programs and, regarding conditional gates, they are implemented as a classical `if` instruction, but inserting inside the `if` the conditioned quantum gates. We even allow conditional gates between nodes by adding the function `__netqir__measure_send` and `__netqir__measure_recv` (and its "array" counterparts) to send the results of a measurement to another node, as has been already explained.

2) Collective communication: As previously mentioned, `qsend` and `qrecv` are the essential primitives to perform distributed computing between QPUs, but they are not always going to be the most efficient or convenient use, which is why collective communication directives are introduced.

Collective communication in NetQIR is designed to support sophisticated interaction patterns essential for distributed quantum computing. Unlike point-to-point communications,

Point-to-point communication functions			
Sending functions		**Receiving functions**	
__netqir__qsend_array	(Array*, i32, i32, Comm*)	__netqir__qrecv_array	(Array**, i32, i32, Comm*)
__netqir__qsend_array_teledata	(Array*, i32, i32, Comm*)	__netqir__qrecv_array_teledata	(Array**, i32, i32, Comm*)
__netqir__qsend_array_telegate	(Array*, i32, i32, Comm*)	__netqir__qrecv_array_telegate	(Array**, i32, i32, Comm*)
__netqir__qsend	(Qubit*, i32, Comm*)	__netqir__qrecv	(Qubit**, i32, Comm*)
__netqir__qsend_teledata	(Qubit*, i32, Comm*)	__netqir__qrecv_teledata	(Qubit**, i32, Comm*)
__netqir__qsend_telegate	(Qubit*, i32, Comm*)	__netqir__qrecv_telegate	(Qubit**, i32, Comm*)
__netqir__measure_send_array	(Array*, i32, i32, Comm*)	__netqir__measure_recv_array	(i1*, i32, i32, Comm*)
__netqir__measure_send	(Qubit*, i32, Comm*)	__netqir__measure_recv	(i1*, i32, i32, Comm*)

Collective communication functions			
__netqir__scatter	(Array*, i32, Array*, i32, i32, Comm*)	__netqir__expose	(Qubit*, i32, Comm*)
__netqir__scatter_teledata	(Array*, i32, Array*, i32, i32, Comm*)	__netqir__expose_array	(Array*, i32, i32, Comm*)
__netqir__scatter_telegate	(Array*, i32, Array*, i32, i32, Comm*)		
__netqir__gather	(Array*, i32, Array*, i32, i32, Comm*)	__netqir__reduce	(Array*, i32, Array*, i32, i32, Comm*)
__netqir__gather_teledata	(Array*, i32, Array*, i32, i32, Comm*)	__netqir__reduce_teledata	(Array*, i32, Array*, i32, i32, Comm*)
__netqir__gather_telegate	(Array*, i32, Array*, i32, i32, Comm*)	__netqir__reduce_telegate	(Array*, i32, Array*, i32, i32, Comm*)

TABLE II: NetQIR functions: point-to-point and collective.

collective communications involve operations where multiple quantum processing units (QPUs) participate in a coordinated manner.

The functions specified are intended to be similar to those used in classical distributed computing, with the objective of facilitating the understanding of the high-performance computing user. These functions are related to the `scatter`, `gather`, `reduce` and, finally, `expose` operations. It should be noted that, due to the non-cloning theorem of quantum computing, the function specifications do not exhibit the same behaviour as their classical alternatives. For example, it is not possible to send a copy of a quantum value to one or more other QPUs, so a `broadcast` function is meaningless, but it can be replaced by the `expose` function: instead of sending a copy, a reference to the qubit is exposed. This allows the communicator QPUs to modify the exposed qubit using quantum theory.

First of all, the directives that facilitate the distribution and collection of quantum information collectively are `__netqir__scatter` and `__netqir__gather` with similar operation to their respective classical functions. The `scatter` function distributes an array of qubits from one QPU to several others, enabling parallel processing of quantum data. On the contrary, the `gather` function collects qubits from multiple QPUs and consolidates them into a single QPU.

It is important to note that, like the point-to-point functions, the collective directives also have a *teledata* or *telegate* option that the user or the compiler can choose depending on the characteristics previously explained. Figure 2 shows the evolution of a quantum system when applying a scatter and a gather using the teledata technique, observing how the computation node loses the data when sending it due to the non-cloning theorem.

Moreover, the `__netqir__reduction` directive has been added. This allows collecting information from a set of remote qubits to which an operation is applied to obtain a final result.

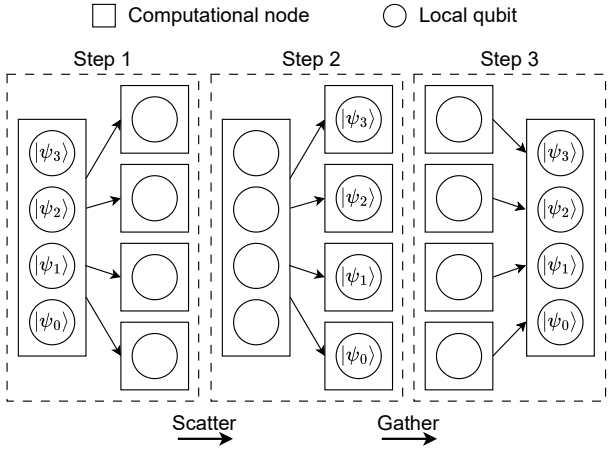

Fig. 2: Example on the use of a scatter teledata on a qubit array (step 1 to 2) and the use of the inverse gather teledata operation (step 2 to 3).

This type of directive could be replaced by others that have already been defined, such as a `scatter-gather` or `qsend-qrecv`, but the specific use of `reduce` implies two main advantages: simplifying code complexity and enhancing computational efficiency.

Firstly, it reduces code complexity. It is easy to see if a multi-controlled CNOT is performed, in which the target node would have to do as many `qrecv` as there are remote nodes. With the `reduce` directive, only this function would be called, making the code more readable.

Secondly, it enhances computational efficiency. In terms of efficiency, there are purely collective and reduction operations, such as the multi-controlled gates or the qubit parity check, which can be optimised in a more specific way if the compiler

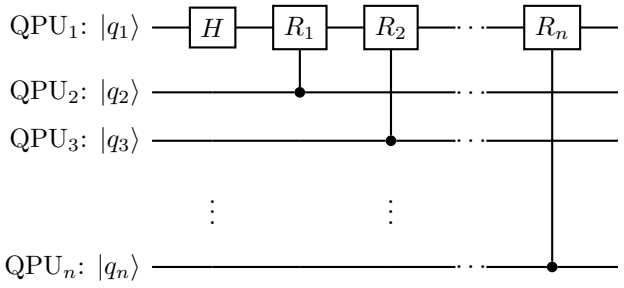

Fig. 3: Use case for using the `__netqir__expose` directive on the $|q_1\rangle$ qubit because it is the target of the rest of the remote qubits.

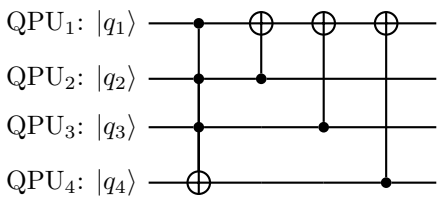

Fig. 4: Use case for using the `__netqir__expose` directive on the $|q_1\rangle$ qubit because it is the target of the rest of the remote qubits.

is told so, instead of with a succession of `qsend` and `qrecv` without semantic meaning.

Finally, in classical computing there is the `broadcast` function with the main objective of send a copy of a value to all nodes. In quantum computing it is not possible to make copies of quantum information due to the non-cloning theorem, as explained above, for this reason the `__netqir__expose` directive has been defined, with a similar operation to the classical `broadcast`. This function is based on the use of the *telegate* communication technique to share a reference of the qubit(s) to another set of QPUs. Therefore, this means that if a QPU modifies the value of the qubit, this modification will be visible to the entire distributed system, differing from the classical `broadcast` in which copies of values are shared.

The contribution of the directive is to be able to facilitate the sharing of a set of qubits local to the QPU to the full complement of communicator compute nodes. This will be very useful to improve the efficiency of different operations, especially in those where all nodes need to use a qubit as a target or control (as shown in Figure 3), as for example in the case of the distributed QFT algorithm.

V. Comparative methods

Now, in order to compare point-to-point directives and collective directives, an example code of NetQIR will be displayed. The example in question is based on Figure 4. This circuit is a clear example of how the reduce and the expose directives, both collectives directives, can significantly decrease the size of the code. First, a multi-CNOT, as it was hinted before, is the perfect operation for a reduce directive. This operation is followed by several CNOTs, each of them with a control qubit in a different node, which represent the perfect case for applying an expose operation as it was already shown in 3.

Regarding the code, in Figure 5 it can be seeing the circuit under consideration implemented by using only point-to-point directives. In order to simplify and reduce the size of the example code several LLVM elements have been eliminated: the "declare" statements at the end of the file, the definition of the data types as opaque types (for `%Qubit`, `%Array`...) and the implementation of the loops, which require lots of lines of code in LLVM IR. It also needs to be mentioned that, because the development in QIR is at an early stage, there is not defined a multi-CNOT gate. Therefore, it is left as indicated in the example code. Back to the code, a lot of lines of code and even the use of loops is employed for performing the required operations to construct the circuit. On the contrary, in Figure 6, a much simpler code is shown. It eliminates the need for loops and, therefore, improves the amount of code displayed for this operation. The specific way of performing the expose and the reduce are left to the compiler. If the compiler desires to, it can even translate the collective operations to the point-to-point analogous. But the fact that the compiler can decide how to perform this operations provides versatility to the code and, even better, possibilities for optimisations and efficient implementations at low level.

VI. Conclusions

As a conclusion to this work, a novel IR for distributed quantum systems is proposed, called NetQIR, as an extension of the already employed QIR. As every IR, it is thought to be employed as an intermediate point for the frontend and the backend. This was considered the whole time while developing the specification in order to achieve the suitable level of abstraction, as shown in comparisons with InQuIR and the vanilla version of NetQASM—the other IRs for distributed quantum systems that are available nowadays, at the best of our knowledge—. This was facilitated by being an extension of an already specified one as QIR is.

In this search for the suited abstraction, the definition of MPI-like functions as *send* and *recv* resulted in a desirable path for creating simple primitives to achieve an exchange of communication in point-to-point communications. Along with these directives came the ones associated with one-to-many communication, also referred to in this work as collective operations. These allowed for a higher abstraction layer from the backend for those specific cases of communication.

As future work, implementing a compilation toolchain with NetQIR as the intermediate code would be perfect for testing and improving the IR. For instance, this could allow a whole group of optimisations to be implemented and tested. These optimisations could go from a simpler approach as the elimination of redundant communication operations to a more

```
define void @main(i32 noundef %0, ptr noundef %1) #0 {
   entry:
      ; Variable allocation
      %2 = alloca i32, align 4
      ; Init the NetQIR communication and get the rank of the process
      %3 = call i32 @__netqir__init(i32 noundef %0, ptr noundef %1)
      %4 = call i32 @__netqir__comm_rank(%Comm* @netqir_comm_world, ptr %2)

   ; ---------------------- REDUCTION DIRECTIVE ----------------------
   ; Choose if it is the process receiving or sending the qubit
      %5 = load i32, ptr %2, align 4
      %6 = icmp eq i32 %5, 4
      br i1 %6, label %7, label %9
   ; Process sending
   7:
      %8 = call i32 @__netqir__qsend(%Qubit* %a, i32 noundef 1,
                              %Comm* @netqir_comm_world)
      br label %exit1
   ; Process receiving
   9:
      ; <<<<<<<<<<<<<<<< This code is inserted in a loop or repeated three times
      %10 = call i32 @__netqir__qrecv(%Qubit** %b, i32 noundef 0,
                               %Comm* @netqir_comm_world)
      %B = load %Qubit*, %Qubit** %b, align 8
      ; >>>>>>>>>>>>>>>>
      ; Here it would go the QIR instruction for a multicontrolled CNOT
   exit1:

   ; ---------------------- EXPOSE DIRECTIVE ----------------------
   ; Choose if it is the process receiving or sending the qubit
      %11 = load i32, ptr %2, align 4
      %12 = icmp eq i32 %11, 0
      br i1 %12, label %13, label %15
   ; Process sending
   13:
      %14 = call i32 @__netqir__qsend(%Qubit* %d, i32 noundef 1,
                              %Comm* @netqir_comm_world)
      br label %exit2
   ; Process receiving
   15:
      ; <<<<<<<<<<<<<<<< This code is inserted in a loop
      %10 = call i32 @__netqir__qrecv(%Qubit** %e, i32 noundef 0,
                               %Comm* @netqir_comm_world)
      %E = load %Qubit*, %Qubit** %e, align 8
      call void @__quantum__qis__cnot__body(%Qubit* %E, %Qubit* %f)
      ; >>>>>>>>>>>>>>>>
   ; Leave the if-else
   exit2:
      ; End of the program
      %11 = call i32 @__netqir__finalize()
}
```

Fig. 5: Example employing point-to-point directives.

```
define void @main(i32 noundef %0, ptr noundef %1) #0 {
   entry:
      ; Variable allocation
      %2 = alloca i32, align 4
      ; Init the NetQIR communication and get the rank of the process
      %3 = call i32 @__netqir__init(i32 noundef %0, ptr noundef %1)
      %4 = call i32 @__netqir__comm_rank(%Comm* @netqir_comm_world, ptr %2)

   ; ----------------------- REDUCTION DIRECTIVE -----------------------
   ; Choose if it is the process receiving or sending the qubit
      %5 = load i32, ptr %2, align 4
      %6 = icmp eq i32 %5, 4
      br i1 %6, label %7, label %9

   ; Process sending
   7:
      %8 = call i32 @__netqir__reduce(%Array* %a, i32 noundef 1,
                              %Array* null, i32 noundef NETQIR_MULTICNOT,
                              i32 noundef 4, %Comm* @netqir_comm_world)
      br label %exit
   ; Process receiving
   9:
      %9 = call i32 @__netqir__reduce(%Array* %b, i32 noundef 1,
                              %Array* %result, i32 noundef NETQIR_MULTICNOT,
                              i32 noundef 4, %Comm* @netqir_comm_world)

   ; ----------------------- EXPOSE DIRECTIVE -----------------------
   ; Choose if it is the process receiving or sending the qubit
      %11 = load i32, ptr %2, align 4
      %12 = icmp eq i32 %11, 4
      br i1 %12, label %13, label %15
   ; Process sending
   13:
       %8 = call i32 @__netqir__expose(%Qubit* null, i32 noundef 0,
                             %Comm* @netqir_comm_world)
      br label %exit
   ; Process receiving
   15:
      %8 = call i32 @__netqir__expose(%Qubit** null, i32 noundef 0,
                             %Comm* @netqir_comm_world)
      %11 = call i32 @__quantum__qis__cnot__body(%Qubit* null, %Qubit* %a)
   ; Leave the if-else
   exit2:
      ; End of the program
      %11 = call i32 @__netqir__finalize()
}
```

Fig. 6: Example employing collective directives.

complicated one as optimising the mapping of the qubit to a specific architecture given the IR code.

REFERENCES

[1] A. V. Aho, M. S. Lam, R. Sethi, and J. D. Ullman, *Compilers: Principles, Techniques, and Tools (2nd Edition)*. USA: Addison-Wesley Longman Publishing Co., Inc., 2006.

[2] G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher, F. J. Cabrera-Hernández, J. Carballo-Franquis, A. Chen, C.-F. Chen, J. M. Chow, A. D. Córcoles-Gonzales, A. J. Cross, A. Cross, J. Cruz-Benito, C. Culver, S. D. L. P. González, E. D. L. Torre, D. Ding, E. Dumitrescu, I. Duran, P. Eendebak, M. Everitt, I. F. Sertage, A. Frisch, A. Fuhrer, J. Gambetta, B. G. Gago, J. Gomez-Mosquera, D. Greenberg, I. Hamamura, V. Havlicek, J. Hellmers, Łukasz Herok, H. Horii, S. Hu, T. Imamichi, T. Itoko, A. Javadi-Abhari, N. Kanazawa, A. Karazeev, K. Krsulich, P. Liu, Y. Luh, Y. Maeng, M. Marques, F. J. Martín-Fernández, D. T. McClure, D. McKay, S. Meesala, A. Mezzacapo, N. Moll, D. M. Rodríguez, G. Nannicini, P. Nation, P. Ollitrault, L. J. O'Riordan, H. Paik, J. Pérez, A. Phan, M. Pistoia, V. Prutyanov, M. Reuter, J. Rice, A. R. Davila, R. H. P. Rudy, M. Ryu, N. Sathaye, C. Schnabel, E. Schoute, K. Setia, Y. Shi, A. Silva, Y. Siraichi, S. Sivarajah, J. A. Smolin, M. Soeken, H. Takahashi, I. Tavernelli, C. Taylor, P. Taylour, K. Trabing, M. Treinish, W. Turner, D. Vogt-Lee, C. Vuillot, J. A. Wildstrom, J. Wilson, E. Winston, C. Wood, S. Wood, S. Wörner, I. Y. Akhalwaya, and C. Zoufal, "Qiskit: An open-source framework for quantum computing," 2 2019. [Online]. Available: https://doi.org/10.5281/zenodo.2562111

[3] C. G. Almudever, L. Lao, X. Fu, N. Khammassi, I. Ashraf, D. Iorga, S. Varsamopoulos, C. Eichler, A. Wallraff, L. Geck, A. Kruth, J. Knoch, H. Bluhm, and K. Bertels, "The engineering challenges in quantum computing," *Proceedings of the 2017 Design, Automation and Test in Europe, DATE 2017*, pp. 836–845, 5 2017.

[4] H. E. Bal, J. G. Steiner, and A. S. Tanenbaum, "Programming languages for distributed computing systems," *ACM Computing Surveys (CSUR)*, vol. 21, pp. 261–322, 9 1989.

[5] D. Barral, F. J. Cardama, G. Díaz, D. Faílde, I. F. Llovo, M. M. Juane, J. Vázquez-Pérez, J. Villasuso, C. Piñeiro, N. Costas, J. C. Pichel, T. F. Pena, and A. Gómez, "Review of distributed quantum computing. from single qpu to high performance quantum computing," 4 2024. [Online]. Available: https://arxiv.org/abs/2404.01265v1

[6] R. Beals, S. Brierley, O. Gray, A. W. Harrow, S. Kutin, N. Linden, D. Shepherd, and M. Stather, "Efficient distributed quantum computing," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 469, 5 2013. [Online]. Available: http://dx.doi.org/10.1098/rspa.2012.0686orviahttp://rspa.royalsocietypublishing.org.

[7] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, "Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels," *Phys. Rev. Lett.*, vol. 70, pp. 1895–1899, Mar 1993. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.70.1895

[8] H. Buhrman, R. Cleve, and A. Wigderson, "Quantum vs. classical communication and computation," in *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, 1998, pp. 63–68.

[9] A. Dahlberg, B. V. D. Vecht, C. D. Donne, M. Skrzypczyk, I. T. Raa, W. Kozlowski, and S. Wehner, "Netqasm—a low-level instruction set architecture for hybrid quantum–classical programs in a quantum internet," *Quantum Science and Technology*, vol. 7, p. 035023, 6 2022. [Online]. Available: https://iopscience.iop.org/article/10.1088/2058-9565/ac753fhttps://iopscience.iop.org/article/10.1088/2058-9565/ac753f/meta

[10] D. Ferrari, A. S. Cacciapuoti, M. Amoretti, and M. Caleffi, "Compiler design for distributed quantum computing," *IEEE Transactions on Quantum Engineering*, vol. 2, 2021.

[11] D. Ferrari, S. Carretta, and M. Amoretti, "A modular quantum compilation framework for distributed quantum computing," *IEEE Transactions on Quantum Engineering*, vol. 4, 2023.

[12] R. P. Feynman, "Simulating physics with computers," *International Journal of Theoretical Physics*, vol. 21, pp. 467–488, 6 1982. [Online]. Available: https://link.springer.com/article/10.1007/BF02650179

[13] M. P. Forum, "Mpi: A message-passing interface standard," 1994.

[14] L. Foundation, "Llvm assembly language reference manual." [Online]. Available: https://releases.llvm.org/2.6/docs/LangRef.html

[15] D. Gottesman and I. L. Chuang, "Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations," *Nature*, vol. 402, no. 6760, p. 390–393, Nov. 1999. [Online]. Available: http://dx.doi.org/10.1038/46503

[16] A. S. Green, P. L. Lumsdaine, N. J. Ross, P. Selinger, and B. Valiron, "Quipper: A scalable quantum programming language," *ACM SIGPLAN Notices*, vol. 48, pp. 333–342, 6 2013. [Online]. Available: https://dl.acm.org/doi/10.1145/2499370.2462177

[17] T. Haner, D. S. Steiger, T. Hoefler, and M. Troyer, "Distributed quantum computing with qmpi," *International Conference for High Performance Computing, Networking, Storage and Analysis, SC*, 11 2021. [Online]. Available: https://dl.acm.org/doi/10.1145/3458817.3476172

[18] K. Hietala, R. Rand, S.-H. Hung, X. Wu, and M. Hicks, "Verified optimization in a quantum intermediate representation," 4 2019. [Online]. Available: https://arxiv.org/abs/1904.06319v4

[19] C. A. Hoare, "An axiomatic basis for computer programming," *Communications of the ACM*, vol. 12, pp. 576–580, 10 1969. [Online]. Available: https://dl.acm.org/doi/10.1145/363235.363259

[20] B. Koen, A. Sarkar, T. Hubregtsen, M. Serrao, A. A. Mouedenne, A. Yadav, A. Krol, I. Ashraf, and C. G. Almudever, "Quantum computer architecture toward full-stack quantum accelerators," *IEEE Transactions on Quantum Engineering*, vol. 1, 2020.

[21] S. W. Loke, "From distributed quantum computing to quantum internet computing: an overview," 8 2022. [Online]. Available: https://arxiv.org/abs/2208.10127v2

[22] I. L. Markov, "Limits on fundamental limits to computation," *Nature 2014 512:7513*, vol. 512, pp. 147–154, 8 2014. [Online]. Available: https://www.nature.com/articles/nature13570

[23] R. V. Meter, T. D. Ladd, A. G. Fowler, and Y. Yamamoto, "Distributed quantum computation architecture using semiconductor nanophotonics," *International Journal of Quantum Information*, vol. 8, pp. 295–323, 6 2009. [Online]. Available: http://arxiv.org/abs/0906.2686http://dx.doi.org/10.1142/S0219749910006435

[24] T. S. Metodi and S. D. Gasster, "Design and implementation of a quantum compiler," *https://doi.org/10.1117/12.852548*, vol. 7702, pp. 251–262, 4 2010. [Online]. Available: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7702/77020S/Design-and-implementation-of-a-quantum-compiler/10.1117/12.852548.fullhttps://www.spiedigitallibrary.org/conference-proceedings-of-spie/7702/77020S/Design-and-implementation-of-a-quantum-compiler/10.1117/12.852548.short

[25] S. Nishio and R. Wakizaka, "Inquir: Intermediate representation for interconnected quantum computers," 2 2023. [Online]. Available: https://arxiv.org/abs/2302.00267v1

[26] *QIR Specification*, QIR Alliance: https://qir-alliance.org. [Online]. Available: https://github.com/qir-alliance/qir-spec

[27] H. Riel, "Quantum computing technology," *Technical Digest - International Electron Devices Meeting, IEDM*, vol. 2021-December, pp. 1.3.1–1.3.7, 2021.

[28] S. Rodrigo, S. Abadal, E. Alarcon, and C. G. Almudever, "Will quantum computers scale without inter-chip comms? a structured design exploration to the monolithic vs distributed architectures quest," *2020 35th Conference on Design of Circuits and Integrated Systems, DCIS 2020*, 11 2020.

[29] S. Schwartz, "An overview of graph covering and partitioning," *Discrete Mathematics*, vol. 345, p. 112884, 8 2022.

[30] M. A. Serrano, J. A. Cruz-Lemus, R. Perez-Castillo, and M. Piattini, "Quantum software components and platforms: Overview and quality assessment," *ACM Computing Surveys*, vol. 55, p. 164, 12 2022. [Online]. Available: https://dl.acm.org/doi/10.1145/3548679

[31] J. Stanier and D. Watson, "Intermediate representations in imperative compilers," *ACM Computing Surveys (CSUR)*, vol. 45, 7 2013. [Online]. Available: https://dl.acm.org/doi/10.1145/2480741.2480743

[32] K. Svore, A. Geller, M. Troyer, J. Azariah, C. Granade, B. Heim, V. Kliuchnikov, M. Mykhailova, A. Paz, and M. Roetteler, "Q#: Enabling scalable quantum computing and development with a high-level dsl," in *Proceedings of the Real World Domain Specific Languages Workshop 2018*, ser. RWDSL2018. New York, NY, USA: Association for Computing Machinery, 2018. [Online]. Available: https://doi.org/10.1145/3183895.3183901

[33] A. S. Tanenbaum and M. v. Steen, *Distributed Systems: Principles and Paradigms (2nd Edition)*. USA: Prentice-Hall, Inc., 2006.

[34] R. Van Meter, K. Nemoto, and W. Munro, "Communication links for distributed quantum computation," *IEEE Transactions on Computers*,

vol. 56, no. 12, p. 1643–1653, Dec. 2007. [Online]. Available: http://dx.doi.org/10.1109/TC.2007.70775

[35] J. Vázquez-Pérez and F. J. Cardama, "Netqir/netqir-spec: v0.0.1," Jul. 2024. [Online]. Available: https://doi.org/10.5281/zenodo.13142521

[36] R. Wakizaka, "Towards reliable distributed quantum computing on quantum interconnects," *ACM International Conference Proceeding Series*, pp. 114–116, 3 2023. [Online]. Available: https://dl.acm.org/doi/10.1145/3594671.3594691