// MLIR Open Meeting - May 2023

CATALYST

An AOT/JIT compiler for hybrid quantum programs

// What you need to know

01

Quantum Computing

Circuits

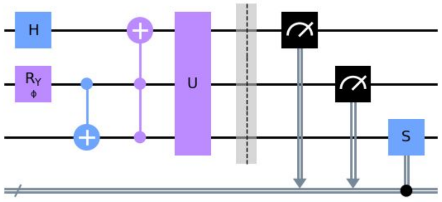

People think this way?

Programming

```
@qml.qnode(dev)
def circuit(phi):
    qml.RX(phi, wires=0)
    qml.RY(2 * phi, wires=1)
    qml.CNOT(wires=[1, 2])
    return qml.expval(qml.PauliZ(0))

def cost(x):
    phi = np.arcsin(x)
```

Yes, and code too!

Applications

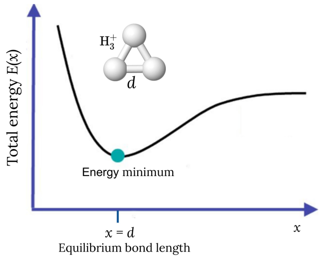

Ok, now what?

Quirks

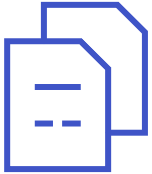

Where is the copy button?

// Quantum Circuits

Memory

Logic Gates

Measurement

Dynamicism?

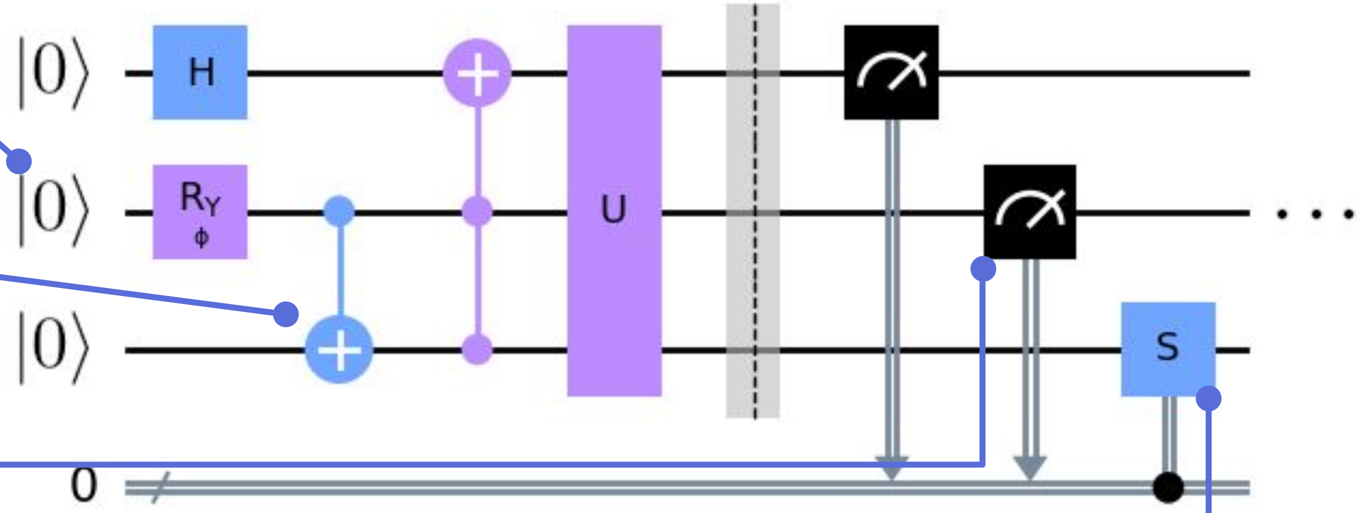

// Quantum Programming

PENNYLANE

Python eDSL

Execution infrastructure

Classical meta-programs

Integrate *Quantum Nodes* into workflows

```python
import numpy as np
import pennylane as qml

dev = qml.device("default.qubit", wires=3)

@qml.qnode(dev)
def circuit(phi):
    qml.RX(phi, wires=0)
    qml.RY(2 * phi, wires=1)
    qml.CNOT(wires=[1, 2])
    return qml.expval(qml.PauliZ(0))

def cost(x):
    phi = np.arcsin(x)
    return circuit(phi)

qml.grad(cost)(0.5)
```

Applications of Quantum Computing

// Factoring

Challenge assumptions about computational complexity

// Search

Last resort in the absence of structure

// Simulation

Nature computing nature

- Ground state estimation
- Unitary evolution

// Hybrid models

Variational algorithms

- Quantum neural networks
- Differentiable programming

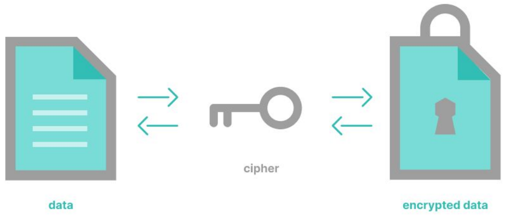

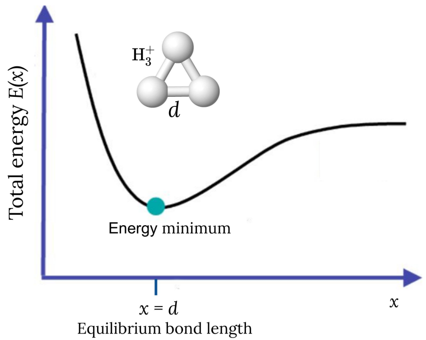

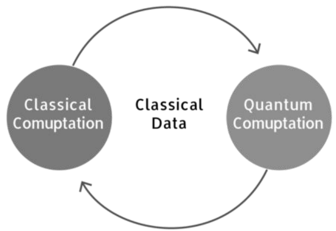

02

Towards a modern Quantum Compilation architecture

// What are early frameworks doing wrong?

PennyLane

Qiskit

Cirq

ProjectQ

```
def qft(n):
    circuit = QuantumCircuit(n)
    for k in range(n-1, -1, -1):
        circuit.h(k)
        for qb in range(k):
            circuit.cp(np.pi/2**(k-qb), qb, k)
    return circuit
```

n = 4

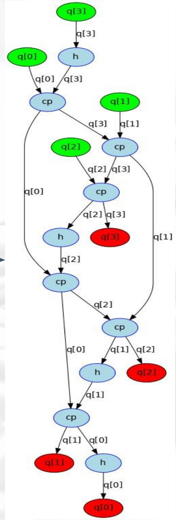

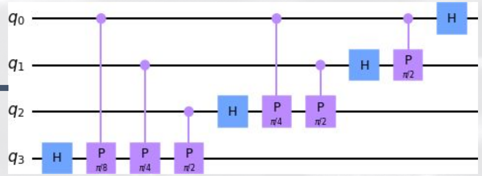

Optimization

Device execution

// Scale it up

Compiling Shor's algorithm for large numbers

Modern RSA ~2000 bit keys

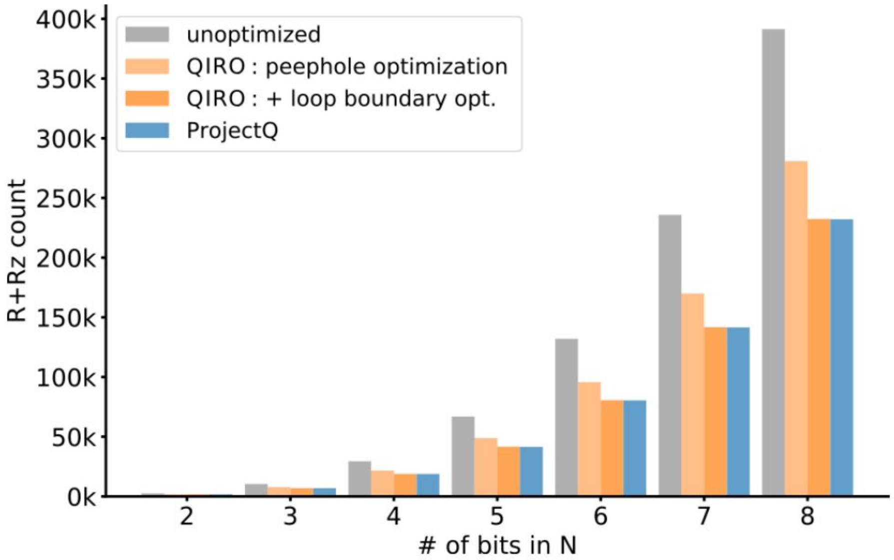

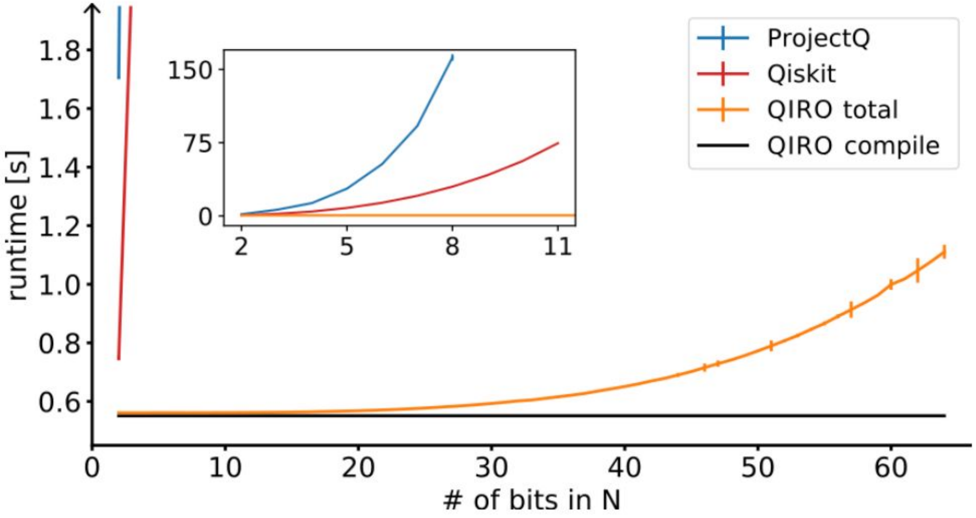

Flat representation of program grows ∝ n^4

→ order of 10^3 years in compile time

Dynamic representation of program

→ constant compile time (order of seconds)

David Ittah, Thomas Häner, Vadym Kliuchnikov, and Torsten Hoefler. 2022. *QIRO: A Static Single Assignment-based Quantum Program Representation for Optimization.* ACM Transactions on Quantum Computing 3, 3, Article 14. DOI

// Emergence of Quantum IRs

Early MLIR dialects

Ittah (21/01) arXiv:2101.11030
McKaskey (21/01) arXiv:2101.11365
McKaskey (21/09) arXiv:2109.00506
Peduri (21/09) arXiv:2109.02409
Guo (22/05) arXiv:2205.03866

Quantum Intermediate Representation (QIR)

Introducing QIR - Q# Blog (microsoft.com)

Industry adoption

QCOR (ornl.gov)

Introducing Catalyst (pennylane.ai)

CUDA Quantum (nvidia.com)

The future

Low-level code generation

...

Quantum Assembly or MLIR?

OpenQASM

```
OPENQASM 2.0;
include "qelib1.inc";
qreg q[4];
h q[3];
cp(pi/8) q[0],q[3];
cp(pi/4) q[1],q[3];
```

2.0: Textual description of static circuit

3.0: Added (limited) classical instructions & dynamicism

No in-memory representation or transformation infrastructure

MLIR

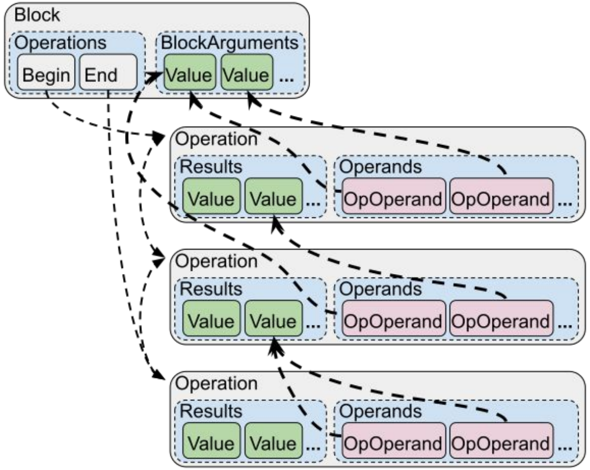

Rich compilation ecosystem

Accommodate multiple domain-specific abstractions side-by-side

LLVM or MLIR?

QIR

```
define void @BellPair(%Qubit* %q1, %Qubit* %q2) {
entry:
  call void @__quantum__qis__h(%Qubit* %q1)
  call void @__quantum__qis__cnot(%Qubit* %q1,
                                  %Qubit* %q2)
  ret void
}
```

Opaque pointer types

Operations as functions

MLIR

```
func @BellPair(%q1: !quantum.bit, %q2: !quantum.bit)
{
    quantum.h %q1 : !quantum.bit
    quantum.cnot %q1, %q2 : !quantum.bit,
                            !quantum.bit
    return
}
```

Extensible types, operations, attributes, assembly, verifiers, structured ops, ...

03

Catalyst

Reimagining the quantum computing stack

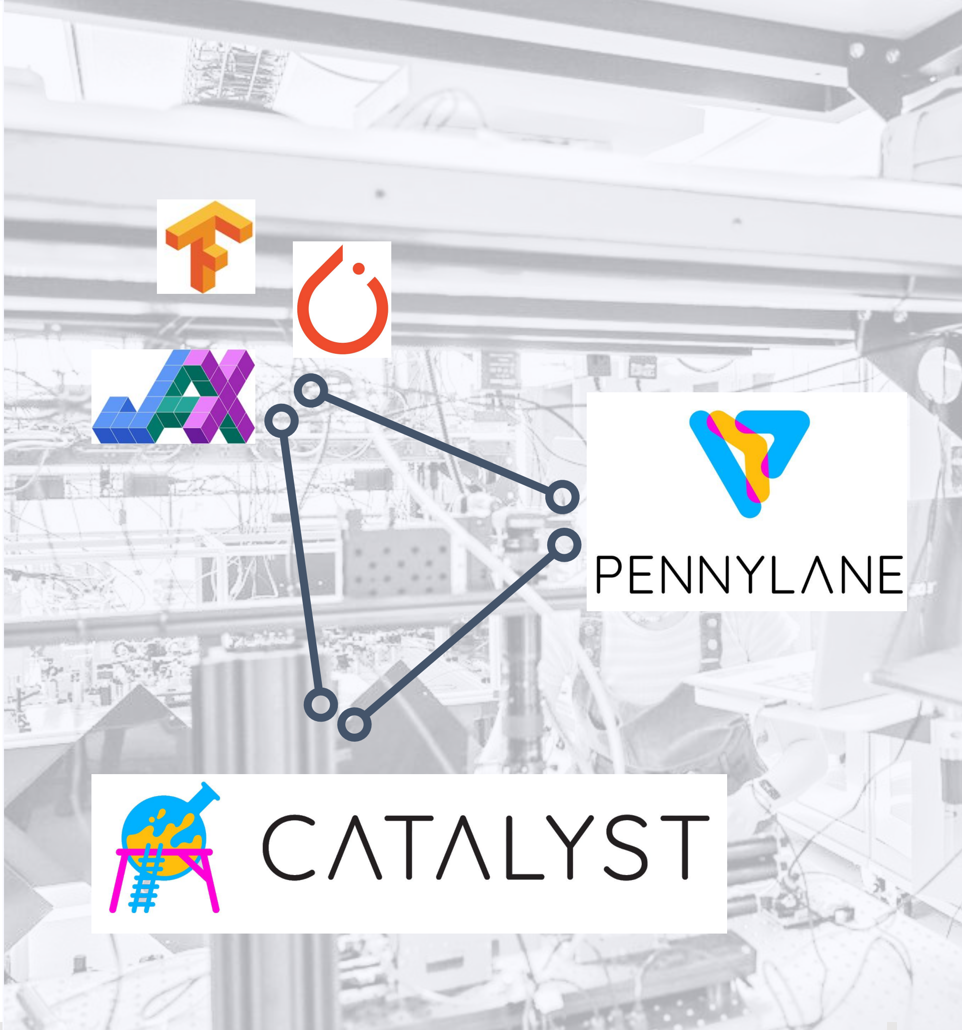

Current Issues

- Reducing latency between classical and quantum components
- Compilation bottlenecks - scaling to very large circuits
- Dynamic quantum programs - quantum execution adapts to intermediate results, unbounded programs
- Heterogenous execution - distribute computation across host machine, accelerators, and quantum computers
- Unified architecture for compiling quantum & classical

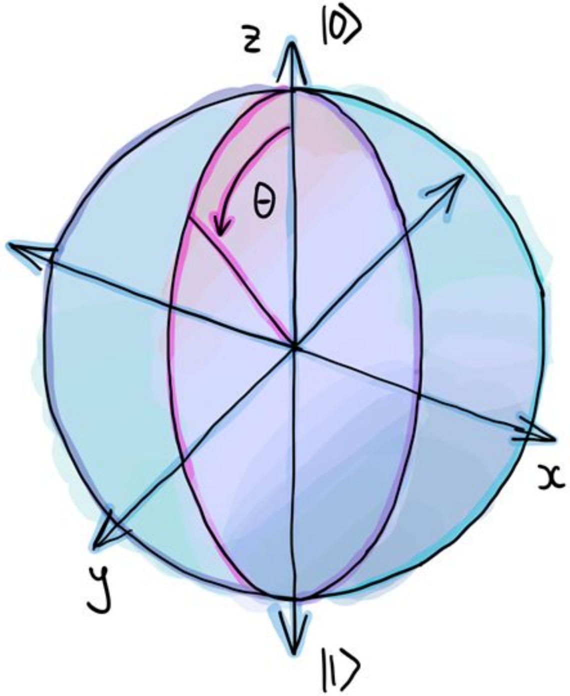

Current Issues

Hardware/simulator device

JAX just-in-time compilation

Python control flow

```
import pennylane as qml
import jax
from jax import numpy as jnp

dev = qml.device('braket.aws.qubit', device_arn=...)

@jax.jit
@qml.qnode(dev, interface="jax")
def circuit(params):
    qml.RX(params[0], wires=0)

    for i in range(0, 3):
        qml.CRX(params[i + 1], wires=[i, i + 1])

    qml.Hadamard(wires=3)

    return qml.expval(qml.PauliZ(1) @ qml.PauliZ(2))

>>> params = jnp.array([[1.6, 1.2, -0.3, 1.0], [0.8, -0.543, 0.1, 0.6]]).T
>>> circuit(params)
DeviceArray([0.6791972, 0.9782419], dtype=float32)

>>> cost = lambda params: jnp.sum(jnp.sin(params))
>>> jax.grad(cost)(params)
DeviceArray([[-0.02919955,  0.6967067 ],
             [ 0.3623577 ,  0.8561624 ],
             [ 0.9553365 ,  0.9950042 ],
             [ 0.5403023 ,  0.8253356 ]], dtype=float32)
```

Easy interface

```python
import pennylane as qml
from pennylane import numpy as np

dev = qml.device("lightning.qubit", wires=1)

@qml.qnode(dev)
def circuit(phi1, phi2):
    qml.RX(phi1, wires=0)
    qml.RY(phi2, wires=0)
    return qml.expval(qml.PauliZ(0))

def cost(x, y):
    return np.sin(np.abs(circuit(x, y))) - 1

dcost = qml.grad(cost, argnum=[0, 1])
```

```python
import pennylane as qml
from catalyst import qjit, grad
from jax import numpy as jnp

dev = qml.device("lightning.qubit", wires=1)

@qml.qnode(dev)
def circuit(phi1, phi2):
    qml.RX(phi1, wires=0)
    qml.RY(phi2, wires=0)
    return qml.expval(qml.PauliZ(0))

@qjit
def cost(x, y):
    return jnp.sin(jnp.abs(circuit(x, y)[0])) - 1

>>> cost(0.53, 0.12)
-0.24437858702920867

>>> qjit(grad(cost, argnum=[0, 1]))(0.53, 0.12)
(Array(-0.32874746), Array(-0.06765491))
```

// The Catalyst Stack

Frontend:

- Program capture (tracing)
- Extend PL with dynamic elements

MLIR:

- Hybrid autodiff
- Circuit optimizations
- Classical optimization

CodeGen:

- Single source
- Leverage LLVM infrastructure

Execution:

- Support Device-Host interactions
- Dynamic instruction dispatch
- Heterogeneous environment

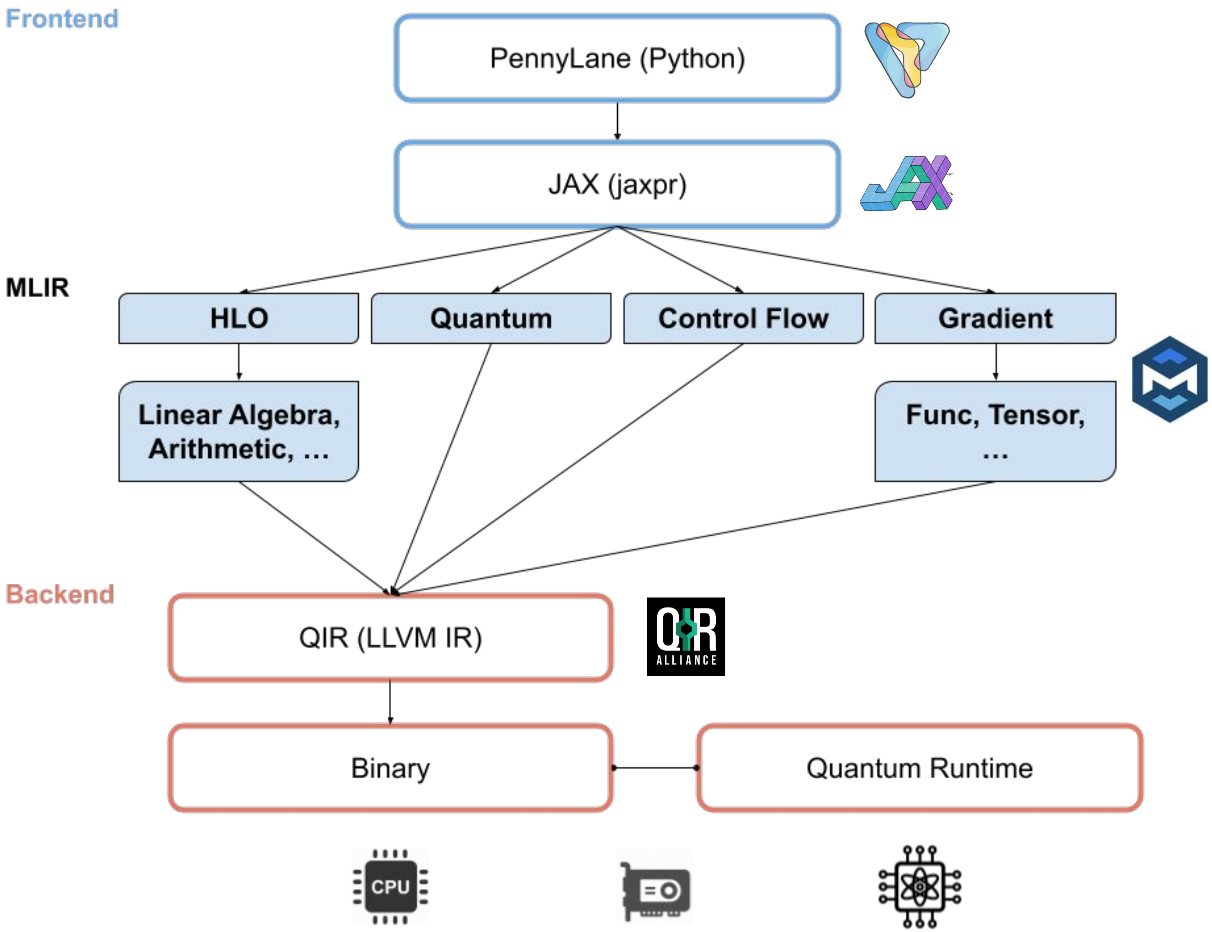

// PL + JAX = <3

Converting Python to MLIR

- ✓ Extending the JAX program representation

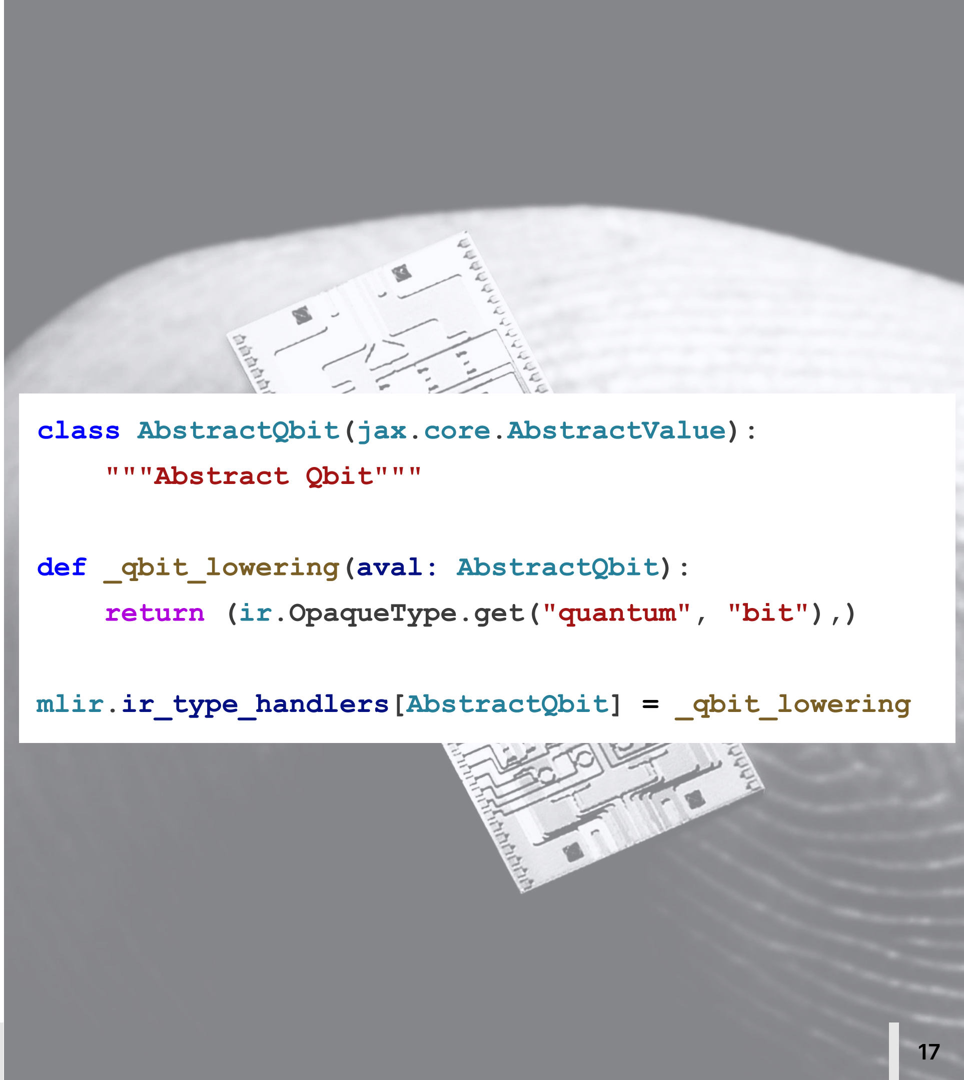

```python
class AbstractQbit(jax.core.AbstractValue):
    """Abstract Qbit"""

def _qbit_lowering(aval: AbstractQbit):
    return (ir.OpaqueType.get("quantum", "bit"),)

mlir.ir_type_handlers[AbstractQbit] = _qbit_lowering
```

// PL + JAX = <3

Converting Python to MLIR

- Extending the JAX program representation
- Conversion to value semantics

```python
unitary_p = jax.core.Primitive("unitary")
unitary_p.multiple_results = True

def unitary(matrix, *qubits):
    """Instantiate JAX primitive."""
    return unitary_p.bind(matrix, *qubits)

def _unitary_abstract_eval(matrix, *qubits):
    return (AbstractQbit(),) * len(qubits)
```

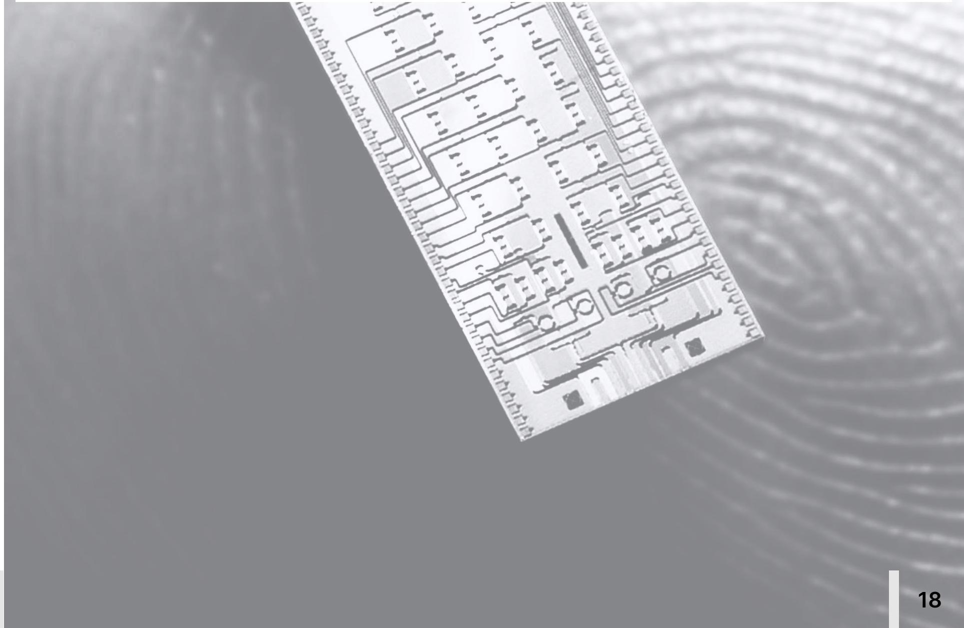

// PL + JAX = <3

Converting Python to MLIR

- ✓ Extending the JAX program representation
- ✓ Conversion to value semantics
- ✓ Custom MLIR lowerings

```
unitary_p = jax.core.Primitive("unitary")
unitary_p.multiple_results = True

def unitary(matrix, *qubits):
    """Instantiate JAX primitive."""
    return unitary_p.bind(matrix, *qubits)

def _unitary_abstract_eval(matrix, *qubits):
    return (AbstractQbit(),) * len(qubits)
```

```
def _unitary_lowering(ctx, matrix: ir.Value,
        *qubits: Tuple[ir.Value]):
    ctx.allow_unregistered_dialects = True

    mlir_op = QubitUnitaryOp([q.type for q in
qubits], matrix, qubits)

    return mlir_op.results

unitary_p.def_abstract_eval(_unitary_abstract_eval)
mlir.register_lowering(unitary_p, _unitary_lowering)
```

// MLIR

The Quantum IR

- ✓ Quantum dialect

```
def QubitUnitaryOp : Gate_Op<"unitary", [NoMemoryEffect]> {
    let summary = "Apply an arbitrary unitary matrix";

    let arguments = (ins
        AnyTypeOf<[
            2DTensorOf<[Complex<F64>]>, MemRefRankOf<[Complex<F64>], [2]>
        ]>:$matrix,
        Variadic<QubitType>:$in_qubits
    );

    let results = (outs
        Variadic<QubitType>:$out_qubits
    );

    let assemblyFormat = [{
        `(` $matrix `:` type($matrix) `)` $in_qubits attr-dict `:`
        type($out_qubits)
    }];
}
```

// MLIR

The Quantum IR

- Quantum dialect
- Optimizations

```
LogicalResult Fusion::match(UnitaryOp op)
{
    ValueRange qbs = op.getInQubits();
    Operation *parent = qbs[0].getDefiningOp();

    if (!isa<UnitaryOp>(parent))
        return failure();

    for (auto qb : qbs)
        if (qb.getDefiningOp() != parent)
            return failure();

    return success();
}
```

// MLIR

The Quantum IR

- ✓ Quantum dialect
- ✓ Optimizations

```
void Fusion::rewrite(UnitaryOp op, PatternRewriter &rewriter)
{
    ValueRange qbs = op.getInQubits();
    UnitaryOp parent = cast<UnitaryOp>(qbs[0].getDefiningOp());

    Value m1 = op.getMatrix();
    Value m2 = parent.getMatrix();

    Value res = rewriter.create<linalg::MatmulOp>(op.getLoc(),
        {m1, m2}).getResult();

    rewriter.updateRootInPlace(op, [&] { op->setOperand(0, res); });
    rewriter.replaceOp(parent, parent.getResults());
}
```

// MLIR

The Quantum IR

- Quantum dialect
- Optimizations
- Gradient dialect

```
def GradOp : Gradient_Op<"grad", [
        DeclareOpInterfaceMethods<CallOpInterface>,
        DeclareOpInterfaceMethods<SymbolUserOpInterface>]> {
    let summary = "Compute partial derivative tensors of a function.";

    let arguments = (ins
        StrAttr:$method,
        FlatSymbolRefAttr:$callee,
        Variadic<AnyType>:$operands,
        AnyIntElementsAttr:$diffArgIndices
    );

    let results = (outs
        Variadic<AnyTypeOf<[AnyFloat, RankedTensorOf<[AnyFloat]>]>>
    );

    let assemblyFormat = [{
        $method $callee `(` $operands `)` attr-dict `:`
        functional-type($operands, results)
    }];
}
```

// Quantum Autodiff

hybrid quantum function
$R^n \rightarrow R$

Real function computed via quantum execution

```
@qml.qnode(dev)
def circuit(phi):
    qml.RX(phi, wires=0)
    qml.RY(2 * phi, wires=1)
    qml.CNOT(wires=[1, 2])
    return qml.expval(qml.PauliZ(0))
```

Classical function for argument processing

```
def gate_args(phi):
    return phi, 2 * phi
```

```
def circuit(arg1, arg2):
    qml.RX(arg1, wires=0)
    qml.RY(arg2, wires=1)
    qml.CNOT(wires=[1, 2])
    return qml.expval(qml.PauliZ(0))
```

Quantum execution, typically producing expectation values

// Quantum Autodiff

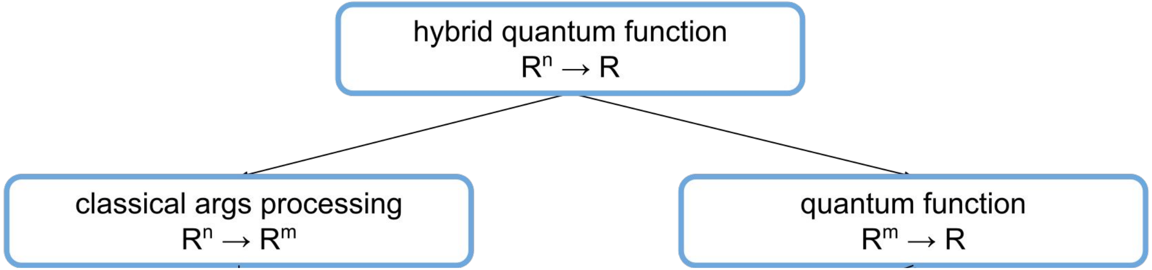

Classical function for argument processing

```
def gate_args(phi):
    return phi, 2 * phi
```

```
def circuit(arg1, arg2):
    qml.RX(arg1, wires=0)
    qml.RY(arg2, wires=1)
    qml.CNOT(wires=[1, 2])
    return qml.expval(qml.PauliZ(0))
```

Quantum execution, typically producing expectation values

// Quantum Autodiff

Perform shifting dynamically at runtime

WIP:

Enzyme integration for hybrid gradient architecture

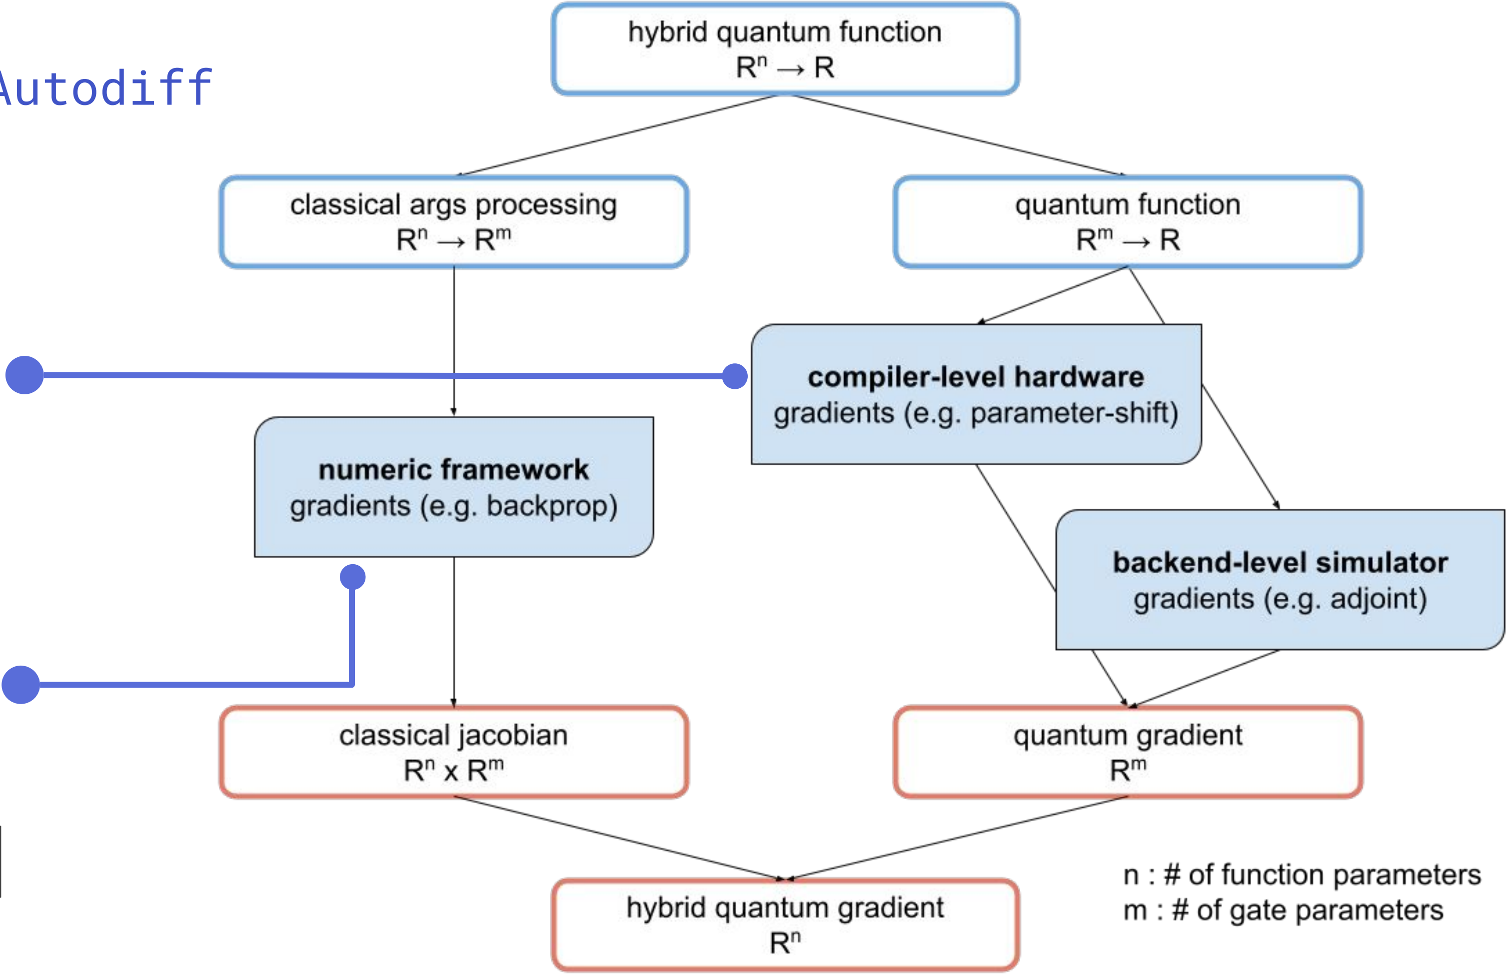

n : # of function parameters
m : # of gate parameters

What next?

04

Device compilation & execution

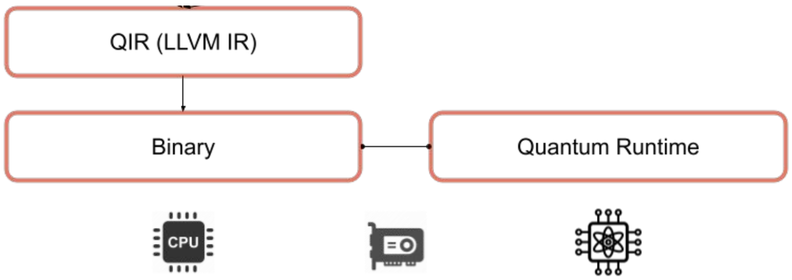

Device-specific compilation, hardware execution, compilation for QPU - co-processor systems

Optimizations

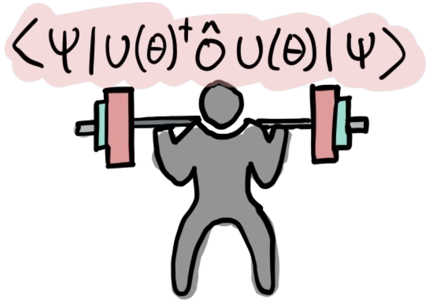

Moving out of beta → optimizing for speed

Quantum compilation algorithms in MLIR

Thank you

XANADU

pennylane.ai
Twitter → @PennyLane.ai

David Ittah
davidi@xanadu.ai

GitHub
https://github.com/PennyLaneAI/catalyst