QPyTorch: A Low-Precision Arithmetic Simulation Framework

Tianyi Zhang, Zhiqiu Lin, Guandao Yang, Christopher De Sa

Cornell University

Our Framework

- Faciliates low-precision training research
- Supports various number formats and rounding options
- Implements state-of-the-art low-precision training techniques

Low-Precision Computation

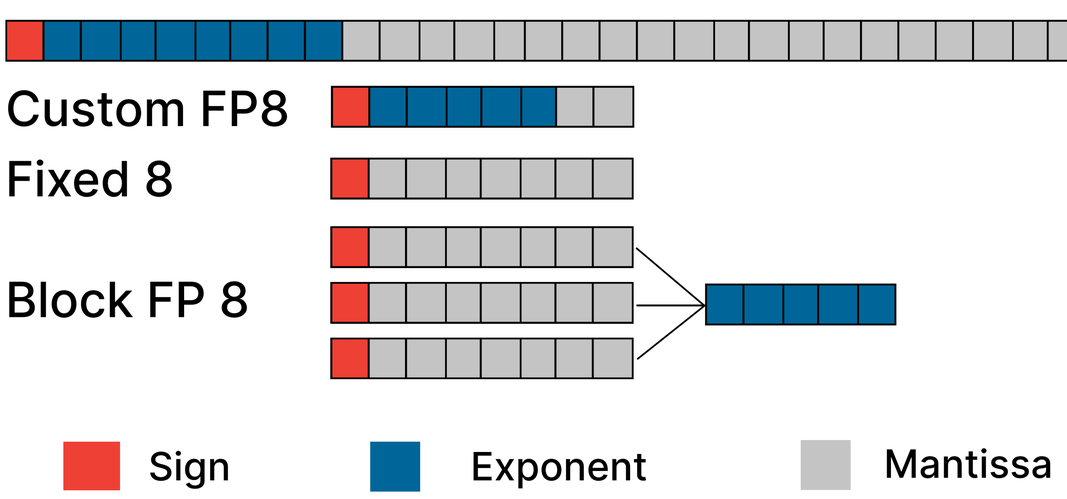

QPyTorch can simulate various number formats:

- floating point: lower precision than single precision
- fixed point: arbitrary precision
- block floating point: lower precision than single precision

Quantization and Rounding

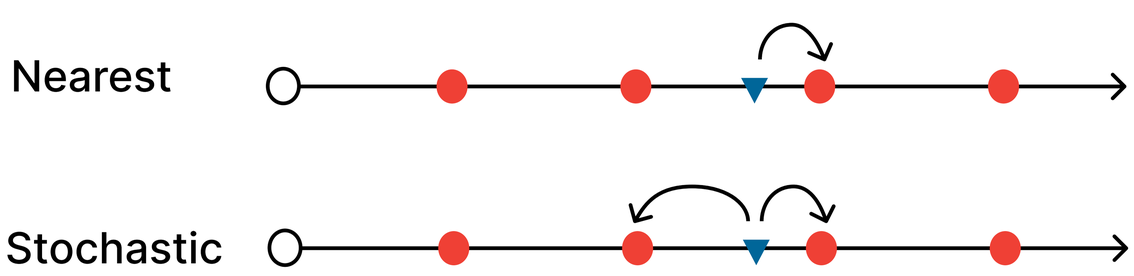

QPyTorch can simulate various rounding strategies

Efficient Simulation

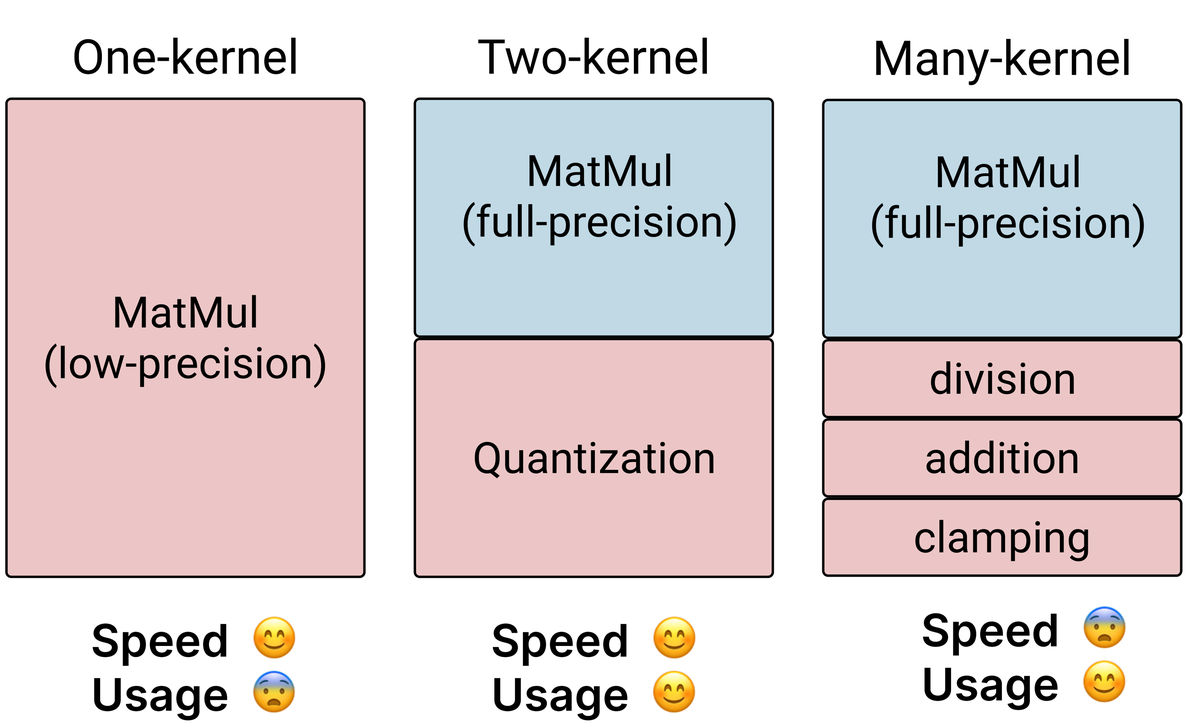

Experiment

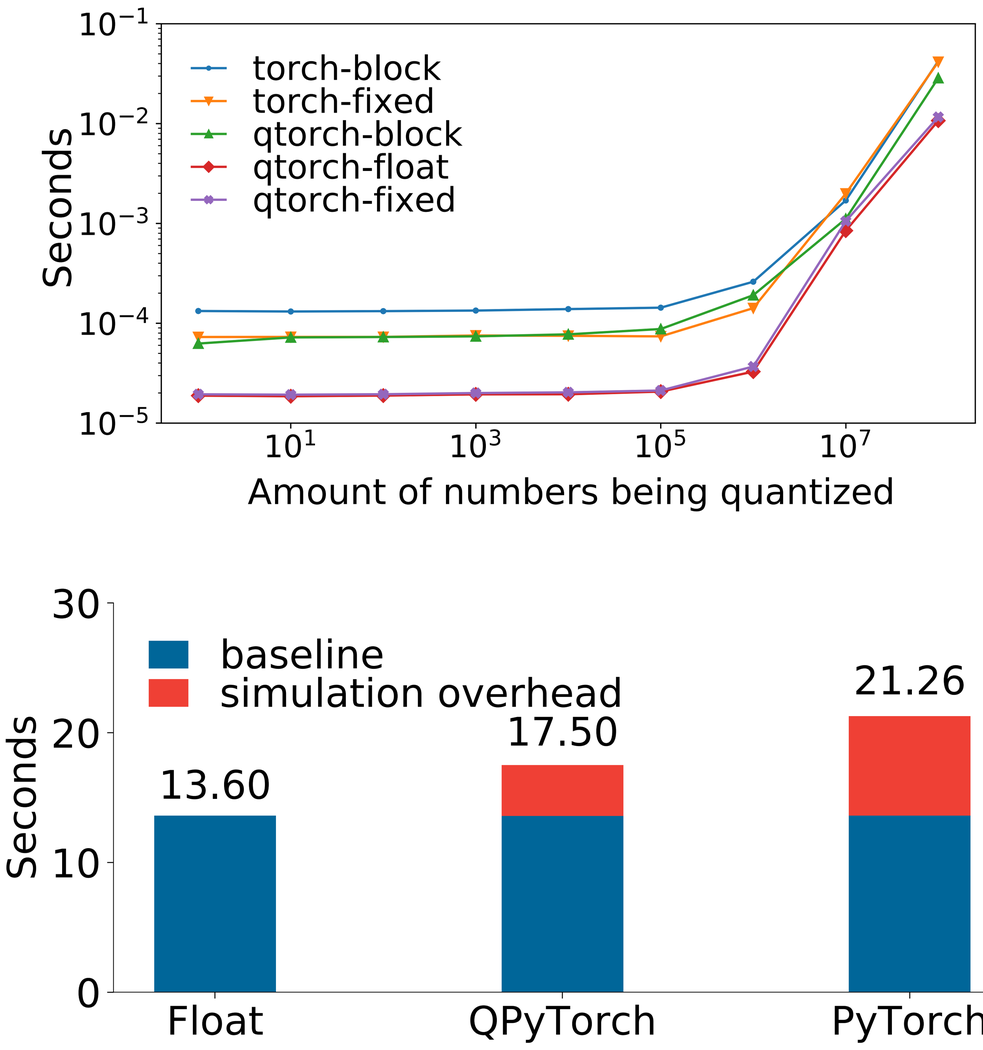

Front-end Interface

Code Example

```
prec = FloatingPoint(bit=8)
acc_prec = FloatingPoint(bit=16)

model = net()
model = lower(model, layer_type=[“conv”],
              forward_number=prec,
              backward_number=prec)

optim = SGD(model.parameters(), lr=0.01)
optim = OptimLP(optim, weight_quant=prec,
                grad_quant=prec,
                acc_quant=acc_prec)
```

Installation

```
pip install qtorch
```

The 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing @ NeurIPS 2019