

DoubleSqueeze: Parallel Stochastic Gradient Descent with Double-pass Error-Compensated Compression

Hanlin Tang, *Xiangru Lian*, Chen Yu, Tong Zhang, Ji Liu

Presenter: Xiangru Lian



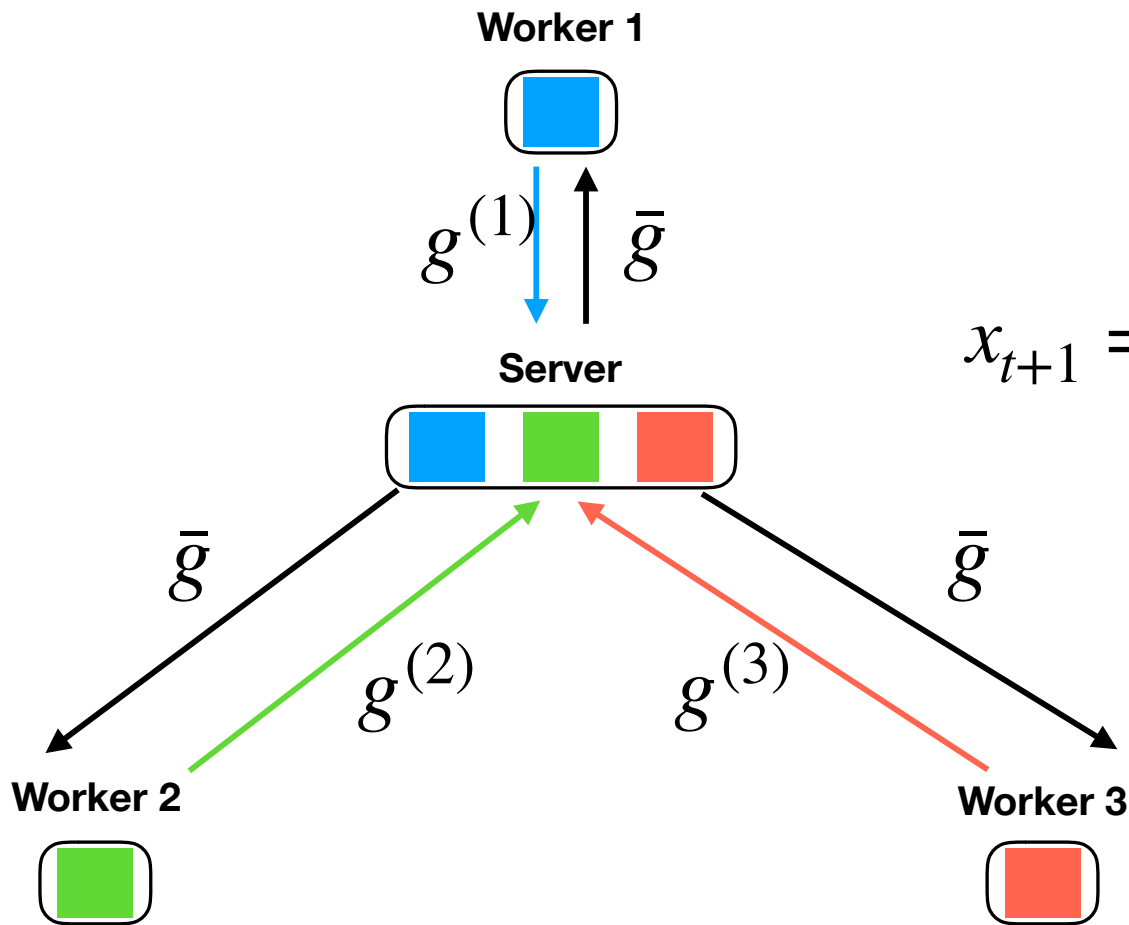
UNIVERSITY of
ROCHESTER



香港科技大學
THE HONG KONG
UNIVERSITY OF SCIENCE
AND TECHNOLOGY



Compressed SGD (existing algorithms)



$$x_{t+1} = x_t - \frac{\gamma}{n} \sum_{i=1}^n C_{\omega} [g^{(i)}]$$

Compression Operator :


- 1bit Quantization
- Clipping
- Top-k Sparsification

Compressed SGD introduces error:

$$1.2 \rightarrow 1; \quad \textit{error} = -0.2$$

We can do better by compensating this error:

$$1.2 \rightarrow 1; \quad \textit{error} = -0.2$$



Next Step

$$\text{Next_Grad} \leftarrow \text{Next_Grad} - \textit{error}$$

DoubleSqueeze

High Level: Compensating Error for **Both Server and Workers**

Worker i :

$$g^{(i)} \leftarrow \nabla F(x; \xi^{(i)}), \quad v^{(i)} \leftarrow C_\omega [g^{(i)} + \delta^{(i)}], \quad \delta^{(i)} \leftarrow g^{(i)} + \delta^{(i)} - v^{(i)}$$

Server :

$$\bar{g} \leftarrow \frac{1}{n} \sum_{i=1}^n v^{(i)}, \quad \bar{v} \leftarrow C_\omega [\bar{g} + \bar{\delta}], \quad \bar{\delta} \leftarrow \bar{g} + \bar{\delta} - \bar{v}$$

On All Workers (Model Update):

$$x \leftarrow x - \gamma \bar{v}$$

Convergence Rate

Assumptions:

$f(x)$ **Non Convex, with L-Lipschitz Gradient;**

$$\mathbb{E}_{\xi \sim \mathcal{D}_i} \|\nabla F(x; \xi) - \nabla f_i(x)\|^2 \leq \sigma^2, \forall i, \forall x;$$

$$\|C_\omega[x] - x\|^2 \leq \zeta^2$$

T: Total Iterations



$$\mathbb{E} \|\nabla f(\bar{x}_T)\|^2 \lesssim \frac{1 + \sigma}{\sqrt{nT}} + \frac{\zeta^{\frac{2}{3}}}{T^{\frac{2}{3}}} \quad \text{(DoubleSqueeze)}$$

$$\mathbb{E} \|\nabla f(\bar{x}_T)\|^2 \lesssim \frac{1 + \sigma}{\sqrt{nT}} + \frac{\zeta}{\sqrt{T}} \quad \text{(Compressed SGD)}$$


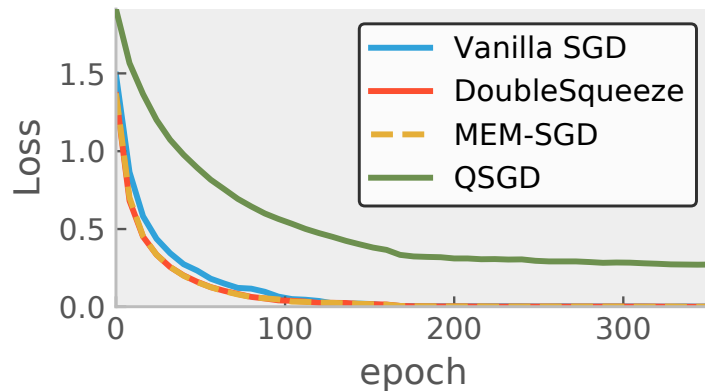
Experiments

ResNet-18 on CIFAR-10. 8 Nvidia 1080Ti GPUs. 1 GPU per worker.


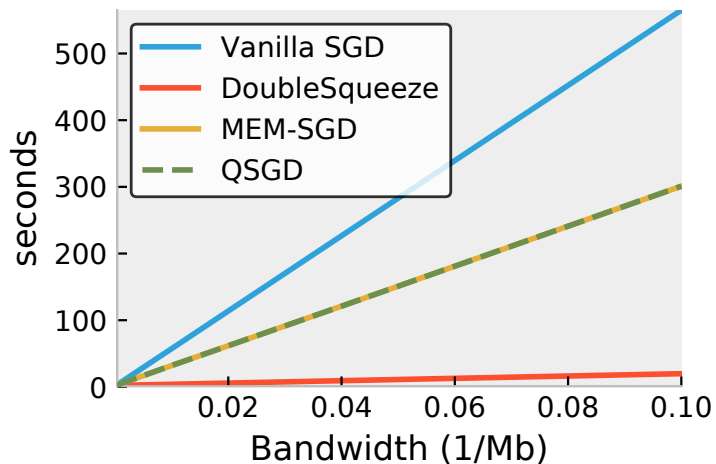
1Bit Quantization:

Top-k Sparsification:

Convergence Rate



Per-Epoch Time



Thanks

Welcome to Pacific Ballroom #99 to see the poster
for more detail