
Overcoming Oscillations in Quantization-Aware Training

Markus Nagel^{*1} Marios Fournarakis^{*1} Yelysei Bondarenko¹ Tijmen Blankevoort¹

Abstract

When training neural networks with simulated quantization, we observe that quantized weights can, rather unexpectedly, oscillate between two grid-points. The importance of this effect and its impact on quantization-aware training (QAT) are not well-understood or investigated in literature. In this paper, we delve deeper into the phenomenon of weight oscillations and show that it can lead to a significant accuracy degradation due to wrongly estimated batch-normalization statistics during inference and increased noise during training. These effects are particularly pronounced in low-bit (≤ 4 -bits) quantization of efficient networks with depth-wise separable layers, such as MobileNets and EfficientNets. In our analysis we investigate several previously proposed QAT algorithms and show that most of these are unable to overcome oscillations. Finally, we propose two novel QAT algorithms to overcome oscillations during training: *oscillation dampening* and *iterative weight freezing*. We demonstrate that our algorithms achieve state-of-the-art accuracy for low-bit (3 & 4 bits) weight and activation quantization of efficient architectures, such as MobileNetV2, MobileNetV3, and EfficientNet-lite on ImageNet. Our source code is available at <https://github.com/qualcomm-ai-research/oscillations-qat>.

1. Introduction

Quantization is one of the most successful methods for optimizing neural networks for efficient inference and on-device execution while maintaining a high accuracy. By compressing the weights and activations from the regular

^{*}Equal contribution ¹Qualcomm AI Research, an initiative of Qualcomm Technologies, Inc.. Correspondence to: Markus Nagel <markusn@qti.qualcomm.com>, Marios Fournarakis <mfournar@qti.qualcomm.com>.

32-bit floating-point format to more efficient low bit fixed-point representations, such as INT8, we can reduce power consumption and accelerate inference when deploying neural networks on edge devices (Horowitz, 2014).

Despite its clear power and latency benefits, quantization comes at the cost of added noise due to the reduced precision. However, researchers in recent years have shown that neural networks are robust to this noise and can be quantized to 8-bits with a minimal drop in accuracy using *post-training quantization* techniques (PTQ) (Nagel et al., 2019; Banner et al., 2019; Nahshan et al., 2020). PTQ can be very efficient and, generally, only requires access to a small calibration dataset, but suffers when applied to low-bit quantization (≤ 4 -bits) of neural networks. Meanwhile, *quantization-aware training* (QAT) has become the de-facto standard method for achieving low-bit quantization while maintaining near full-precision accuracy (Krishnamoorthi, 2018; Esser et al., 2020; Nagel et al., 2021). By simulating the quantization operation during training or fine-tuning, the network can adapt to the quantization noise and reach better solutions than with PTQ.

In this paper, we focus on the *oscillations* of quantized weights that occur during quantization-aware training. This is a little-known and under-investigated phenomenon in the optimization of quantized neural networks, with significant consequences for the network during and after training. When using the popular *straight-through estimator* (STE) (Bengio et al., 2013) for QAT, weights seemingly randomly oscillate between adjacent quantization levels leading to detrimental noise during the optimization process. Equipped with this insight, we investigate recent advances in QAT that claim improved performance and assess their effectiveness in addressing this oscillatory behavior.

An adverse symptom of weight oscillations is that they can corrupt the estimated inference statistics of the batch-normalization layer collected during training, leading to poor validation accuracy. We find that this effect is particularly pronounced in low-bit quantization of efficient networks with depth-wise separable layers, such as MobileNets or EfficientNets, but can be addressed effectively by re-estimating the batch-normalization statistics after training.

While batch-normalization re-estimation overcomes one significant symptom of the oscillations, it does not tackle its

root cause. To this end, we propose two novel algorithms that are effective at reducing oscillations: *oscillation dampening* and *iterative weight freezing*. By addressing oscillations at their source, our methods improve accuracy beyond the level of batch-normalization re-estimation. We show that both methods achieve state-of-the-art results for 4 and 3-bit quantization of efficient networks, such as MobileNetV2, MobileNetV3, and EfficientNet-lite on ImageNet.

2. Oscillations in QAT

We first investigate why weights oscillate in quantization-aware training and how this phenomenon affects neural network training in practice.

2.1. Quantization-aware training

One of the most effective ways to quantize a neural network is by training the network with simulated quantization. During the forward pass, floating-point weights and activation are quantized using the quantization function $q(\cdot)$. It takes input vector \mathbf{w} and returns quantized output $\widehat{\mathbf{w}}$ given by:

$$\widehat{\mathbf{w}} = q(\mathbf{w}; s, n, p) = s \cdot \text{clip} \left(\left\lfloor \frac{\mathbf{w}}{s} \right\rfloor, n, p \right), \quad (1)$$

where $\lfloor \cdot \rfloor$ is the *round-to-nearest* operator, $\text{clip}(\cdot, \alpha, \beta)$ is a clipping function with lower and upper bounds α and β , respectively, s is a scaling factor, and n and p the lower and upper quantization thresholds. In this formulation, the quantized weights $\widehat{\mathbf{w}}$ are the ones used during inference, whereas the original floating-point weights \mathbf{w} only serve as proxies for the optimization and are commonly referred to as *latent weights* or *shadow-weights*.

A fundamental challenge in the QAT formulation is that the rounding function in equation (1) does not have a meaningful gradient, which makes gradient-based training impossible. One of the most popular techniques for alleviating this issue involves using the straight-through estimator (STE) (Bengio et al., 2013; Hinton, 2012) to approximate the true gradient during training. In practice, this means that we approximate the gradient of the rounding operator as 1 within the quantization limits. We can thus define the gradient of the loss \mathcal{L} with respect to \mathbf{w} as:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}} = \frac{\partial \mathcal{L}}{\partial \widehat{\mathbf{w}}} \cdot \mathbf{1}_{n \leq \mathbf{w}/s \leq p}, \quad (2)$$

where $\mathbf{1}$ is the indicator function that is 1 if \mathbf{w} falls within the quantization grid, and 0 otherwise, such that there is no gradient outside of the representable quantization region. The STE gradient approximation has been widely adopted in recent literature closing the gap between quantized and full-precision accuracy for a wide range of tasks and networks (Gupta et al., 2015; Zhou et al., 2016; Krishnamoorthi, 2018; Jain et al., 2019; Esser et al., 2020; Bhalgat et al., 2020; Nagel et al., 2021).

2.2. Oscillations

Despite its wide adaptation and enormous success in QAT, the STE has a counter-intuitive and very interesting side-effect. The STE causes implicit stochasticity during the optimization process due to the latent weights oscillating around the decision boundary between adjacent quantized states. This phenomenon was recently also observed by Défossez et al. (2021).

To illustrate this, we propose a simple toy regression example that we will refer to frequently throughout this paper. We start with an optimal floating-point weight, w_* , as target for our 1D toy regression problem and sample a data vector \mathbf{x} from a distribution X with bounded variance such that, $\mathbb{E}[\mathbf{x}^2] = \sigma^2$, where $0 < \sigma^2 < \infty$. Then we optimize the least-squares problem

$$\min_w \mathcal{L}(w) = \mathbb{E}_{\mathbf{x} \sim X} \left[\frac{1}{2} (\mathbf{x}w_* - \mathbf{x}q(w))^2 \right], \quad (3)$$

where $q(\cdot)$ is the quantizer from equation (1) and w is the latent weight. We optimize the objective using the STE formulation for the gradients. From figure 1 (left), we can see that as the latent weight w approaches the optimal value w_* , it starts oscillating around the decision threshold between the quantization levels above w_\uparrow and below w_\downarrow the optimal value, as opposed to converging to the region closer to the optimal value $q(w_*)$.

The weight oscillates around the decision threshold due to the fact that the gradient in equation (2) is constant above the threshold, pushing the latent weight down towards w_\downarrow , and constant below the threshold pushing the latent weight up towards w_\uparrow . Please refer to appendix A.1 for an analytical expression of this gradient. The induced oscillations happen irrespective of the learning rate. In appendix A.3 we show that decreasing the learning rate reduces the amplitude of the oscillations but do not affect their frequency.

The frequency of the oscillation is depends on the distance of the optimal value from its closest quantization level, $d = |w_* - q(w_*)|$. Let's assume that $q(w_*) = w_\uparrow$, then the gradient below the threshold is $k = s/d - 1$ times larger than the gradient above, where s is scaling factor or quantization step-size and $s > d$. Interpreting the gradient as the velocity of the latent weight, w needs to travel for k iterations before it crosses the threshold and its velocity/gradient is reversed. In appendix A.2 we show experimentally that the oscillation frequency is indeed directly proportional to closeness of w_* to $q(w_*)$.

It is interesting to note that this behavior bears similarities to stochastic rounding (Gupta et al., 2015), where the closeness of the latent weight to the quantization level is related to the probability of rounding to that level. However, in STE the source of the stochasticity stems from the discrete nature




Figure 1. Oscillation of a single weight in toy-regression problem for various gradient estimator methods: STE (left), EWGS (J. Lee, 2021) (middle), and DSQ (Gong et al., 2019) (right).




Figure 2. Progression of weights in the first channel of the first depth-wise separable layer (conv.1.0) during the last 1000 iterations of training MobileNetV2 on ImageNet. The top plot shows the integer weights and the bottom plot zooms in on the decision boundary between 0 and 1 for the corresponding latent weights.

of the gradients and not from sampling. We note that oscillations are not unique to the vanilla STE, but are present in several variations of the STE proposed in literature, a small subset of which we present in figure 1.

2.3. Oscillations in practice

These oscillations are not just an unfortunate side-effect of this fictitious toy example. They are indeed present in larger neural networks with significant implications for their optimization. Figure 2 shows the progression of 3-bit quantized weights in a depth-wise separable layer of MobileNetV2 close to convergence, trained using LSQ (Esser et al., 2020) on ImageNet. We observe that many of the weights appear to randomly oscillate between two adjacent quantization levels. In figure 3, we also see that after the supposed convergence of the network, a large fraction of the latent weights lie right at the decision boundary between grid points. This further reinforces the observation that a significant proportion of weights oscillate and do not converge.

We identify two major issues associated with the oscillations in neural network training: Wrong estimation of the batch-




Figure 3. Weight distribution in depth-wise separable layer of third inverted residual block of MobileNetV2 (conv.3.1). Histogram of latent weights (left). Histogram of the distance of latent weight from their closest quantization grid-point ($w_{\text{int}} - w/s$) (right).

normalization (Ioffe & Szegedy, 2015) inference statistics, and an adverse effect on the optimization of the network.

2.3.1. THE EFFECT ON BATCH-NORMALIZATION

During training batch-normalization layers track the exponential moving average (EMA) of the mean and variance for each layer’s output, so that it can be used during inference as an approximation of the real sample statistics. When training full-precision neural networks we can expect that weights will change very slowly close to convergence. As a result, the output statistics of each layer are expected to be quite stable across iterations and therefore the EMA is a good estimate of the statistics. However, oscillations in QAT can lead to rapid changes of the integer weights (see in figure 2) leading to significant distribution shift between iterations, even close to convergence. The sudden and large changes in the output distribution induced by oscillations can corrupt the EMA statistics leading to a significant degradation in accuracy. In fact, there are two factors that amplify this effect: the weight bit-width and the number of weights per output channel.

The lower the bit-width b , the larger the distance becomes between quantization levels, as it is proportional to $1/2^b$. As the oscillating weights move from one quantization level to the other, they cause a proportionally larger shift in the output distribution. The second important factor is the number

Overcoming Oscillations in QAT

Network	Layer	$\max D_{\text{KL}}$	$\mathbb{E}[D_{\text{KL}}]$
ResNet18	layer1.0.conv1	0.0059	0.0002
ResNet18	layer1.0.conv2	0.0130	0.0014
ResNet18	layer3.0.conv1	0.0006	0.0001
MobileNetV2	conv3.0 (PW)	0.7858	0.0292
MobileNetV2	conv3.1 (DW)	55.3782	1.25464
MobileNetV2	conv3.2 (PW)	0.0065	0.0012
MobileNetV2	conv10.0 (PW)	0.0037	0.0004
MobileNetV2	conv10.1 (DW)	27.2618	0.2900
MobileNetV2	conv10.2 (PW)	0.0267	0.0034

Table 1. Discrepancy between estimated and actual population statistics of batch-normalization following convolutional layers in ResNet18 and MobileNetV2 with weights quantized to 3 bits. We report the maximum and mean KL-divergence across the output channels of each layer. DS: depth-wise separable, PW: point-wise.

of weights per output channel. The smaller the number of weights, the larger the contribution of individual weights to the final accumulation. When the number of accumulations increases, the effects of oscillations average out due to the law of large numbers.

In table 1, we use the KL-divergence¹ to quantify the discrepancy between population and estimated statistics. We indeed observe that KL divergence is much larger for depth-wise separable layers than in point-wise convolutions in MobileNetV2 and the full convolutions in ResNet18.

A cheap and straightforward solution to this problem is to re-estimate the batch-normalization statistics with a small subset of data after training. This method, which we call *batch-normalization (BN) re-estimation*, is sometimes used in stochastic quantization formulations (Peters & Welling, 2018; Louizos et al., 2019). However, we argue that it is also essential in deterministic QAT formulations due to the oscillating weights.

In table 2, we see that BN re-estimation not only improves the final quantized accuracy for MobileNetV2 but also reduces the variance among different seeds. We further observe that the gap in accuracy widens as the bit-width goes down for MobileNetV2, which is not the case for ResNet18.

2.3.2. THE EFFECT ON TRAINING

Next to harming the BN statistics, oscillations can also have a negative effect on the training process itself. To illustrate this we first show that a converged MobileNetV2 with 3-bit weights can achieve a lower training loss (and higher valida-

¹We assume that population p and estimated q outputs are normally distributed, such that $D_{\text{KL}}(p, q) = \log \frac{\sigma_2^2}{\sigma_1^2} + \frac{\sigma_1^2 + (\mu_1 - \mu_2)^2}{2\sigma_2^2} - \frac{1}{2}$, where $p(x) = \mathcal{N}(\mu_1, \sigma_1)$ and $q(x) = \mathcal{N}(\mu_2, \sigma_2)$.

Network	Bits	pre-BN	post-BN
ResNet18	4	70.15 ^{0.03}	70.20 ^{0.02}
ResNet18	3	69.63 ^{0.01}	69.70 ^{0.05}
MobileNetV2	8	71.79 ^{0.07}	71.89 ^{0.05}
MobileNetV2	4	68.99 ^{0.44}	71.01 ^{0.05}
MobileNetV2	3	64.97 ^{1.23}	69.50 ^{0.04}

Table 2. Validation accuracy (%) on ImageNet before (pre-BN) and after BN re-estimation (post-BN) for networks trained using low bit weight quantization (Esser et al., 2020) for 20 epochs. We report the average of 3 seeds and the STD in the superscript.

Method	Train Loss	Val. Acc. (%)
Baseline	1.3566	69.50
SR (mean + std)	1.3547 ^{0.0053}	69.58 ^{0.09}
SR (best)	1.3391	69.85
AdaRound	1.3070	70.12
Freezing	-	70.33

Table 3. Ablation study on the effect of oscillations on QAT. Training loss and validation accuracy (%) for MobileNetV2 with 3-bit weight quantization. SR is stochastic rounding of oscillating weights and AdaRound the binary optimization of oscillating weights.

tion accuracy) if we randomly sample oscillating weights between their two oscillating states. To this end, we sample all weights that were oscillating at the end of training with a probability proportional to the time spent at each quantized state, i.e. $p(w_{\uparrow}) = \mathbb{E}_t [w^t = w_{\uparrow}]$. We calculate the expectation using an exponential moving average over the integer weights as done in line 15 of algorithm 1.

In table 3, we present the results of this experiment. We observe that the mean training loss over the sampled networks is similar to that of the final converged model. However, many samples achieve a lower training loss and the best randomly sampled network shows a significantly lower training loss. We also perform a binary optimization of the oscillating weights using an adaptation of AdaRound (Nagel et al., 2020). We optimize the rounding of all layers simultaneously on the final task loss, akin to what is done in the literature with simulated annealing to solve binary optimization problems (Kirkpatrick et al., 1983). We see that this binary optimization significantly improves over the best stochastic sample and the original converged network. This suggests that weight oscillations prevent the network from converging to the best local minimum during training, and can be detrimental to the optimization process.

Finally, we show that preventing oscillations earlier in training using our oscillation freezing technique (cf. section 4.3) leads to an even higher validation accuracy than the binary

optimization of the oscillating weights. This indicates that the oscillations not only prevent QAT from converging to the best local minimum at the end of training, but can also lead the optimizer towards sub-optimal directions earlier in training.

3. Related work

Earlier work in QAT (Cai et al., 2017; Yin et al., 2019) investigated different variations of the STE and demonstrate experimentally that the derivative of the clipped ReLU is a better approximation than the vanilla STE (derivative of identity). In practice, the former is the same as setting the gradient outside the quantization grid to zero, as per equation (2), and is the most widely used formulation in QAT literature (Choi et al., 2018; Esser et al., 2020; Uhlich et al., 2020).

Despite the success of the STE trick in practice, a lot of prior work has focused on the non-optimality of the STE as a gradient approximator. Défossez et al. (2021) claim that the STE causes instability during training due to weights oscillating leading to *biased* gradients and weights. For this reason, they propose replacing the rounding operation with additive Gaussian noise in the forward pass to simulate quantization noise. Whereas weight oscillations are indeed an important side-effect of the STE, the idea that the STE yields *biased* gradients in a deterministic training setting is opaque, as we discuss in more detail in appendix A.4. Nguyen et al. (2020) use a toy example to claim that solutions under STE can lead to a sub-optimal quantized solution but do not recognize that the quantized weight will oscillate around the optimum rather than converge to the wrong value. Based on their observation, they propose an absolute-cosine regularization for the full-precision weights to make them more quantization-friendly but do not compare to other competing QAT methods.

Other work focuses on alternative *gradient estimators* for back-propagation. Despite seeming differences in their motivation and formulations, all these methods in one way or another propose adaptations of STE that boil down into two main categories: *multiplicative* or *additive*. Multiplicative methods apply a scaling to the data-gradient, whereas additive methods add an input-independent term to the gradient through regularization. We provide more details and derivations in appendix A.1.

Multiplicative J. Lee (2021) (EWGS) acknowledge that the STE generates *coarse gradients* that treat all latent weights within a quantization bin equally regardless of their distance from its centre and propose scaling the gradient by this distance. However, they do not provide any evidence on why and how these coarse gradient harm optimization. Kim et al. (2020) arrive to a similar formulation but are motivated

by robust quantization. They claim that clustering latent weights around bin centres makes them more robust when quantized to different bit widths. In both these methods, the gradient scales linearly with the distance of the latent weight from its quantized value.

Gong et al. (2019) (DSQ) try to alleviate the *gradient error* that the STE causes during low-bit quantization but do not provide any evidence about the source or magnitude of this error. They propose using a series of hyperbolic tangent functions to simulate the rounding function in the backward pass. This leads to a non-linear scaling of the STE gradient, with the gradient increasing with the distance from the bin centre. Yang et al. (2019) also approximate the rounding function with a stack of sigmoid functions to address the *gradient mismatch* problem. They parameterize the sigmoid with a temperature, which increases during training to better approximate the step function.

As we show in the appendix, these multiplicative methods can affect the magnitude of the oscillations but cannot stop them from happening. In figure 1, we visualize the effect of EWGS and DSQ gradients on the 1D regression problem, described in section 2.2. It is evident, that neither of these methods overcome oscillations.

Additive Various regularizations have been proposed in literature to make the weight distribution more quantization-friendly. Chai (2021) apply an L^2 regularization on the difference between quantized and floating-point weights, while Han et al. (2021) apply bin regularization to enforce the weights to cluster at the quantization grid points. Their bin regularization method is similar to our proposed *oscillation dampening* method described later in section 4.2. However, their motivation is to form sharper weight distribution around the bin centre that reduce the quantization error rather than preventing oscillations. In contrast to the multiplicative methods, these additive methods add a term to the gradient and can thus dampen oscillations by pulling the latent weights towards the grid points.

Park & Yoo (2020) recognize that the STE can cause activation instability during training of low-bit quantized MobileNets and propose progressively freezing whole layers during training. Chen et al. (2019) approximate the gradient using a small neural network, whose parameters are learnt using meta-learning but is limited to binary weights.

4. Overcoming oscillations in QAT

Now that we have established that oscillations can have a negative effect during optimization, especially for low-bit quantization, we focus on how to overcome them. First we introduce a metric for quantifying oscillations and then we propose two novel techniques aimed at preventing oscillations at their source during quantization-aware training.

4.1. Quantifying Oscillations

Before we can address the oscillations, we need a way of detecting and measuring them during training. We propose calculating the frequency of oscillations over time using an exponential moving average (EMA). We can then define a minimum frequency as a threshold for oscillating weights. For an oscillation to occur in iteration t it needs two satisfy two conditions:

1. The integer value of the weight needs to change, thus $\mathbf{w}_{\text{int}}^t \neq \mathbf{w}_{\text{int}}^{t-1}$ where $\mathbf{w}_{\text{int}}^t = \text{clip}\left(\lfloor \frac{\mathbf{w}}{s} \rfloor, n, p\right)$ is the integer value of the quantized weight.
2. The direction of the change in the integer domain needs to be the opposite than that of its previous change, thus $o^t = \text{sign}(\Delta_{\text{int}}^t) \neq \text{sign}(\Delta_{\text{int}}^{\tau})$ where τ is the iteration of the last change in integer domain and $\Delta_{\text{int}}^t = \mathbf{w}_{\text{int}}^t - \mathbf{w}_{\text{int}}^{t-1}$, the direction of the change.

We then track the frequency of oscillations over time using an exponential moving average (EMA):

$$f^t = m \cdot o^t + (1 - m) \cdot f^{t-1}. \quad (4)$$

4.2. Oscillation dampening

When weights oscillate, they always move around the decision threshold between two quantization bins. This means that oscillating weights are always close to the edge of the quantization bin. In order to dampen the oscillatory behavior, we employ a regularization term that encourages latent weights to be close to the center of the bin rather than its edge. We define our dampening loss similar to weight decay as:

$$\mathcal{L}_{\text{dampen}} = \|\widehat{\mathbf{w}} - \text{clip}(\mathbf{w}, sn, sp)\|_F^2, \quad (5)$$

where s , n and p are the quantization parameters as defined in equation (1) and $\widehat{\mathbf{w}}$ are the bin centers. Note, as the bin centers are our optimization target, no gradients propagate back through this term. Our final training objective is now: $\mathcal{L} = \mathcal{L}_{\text{task}} + \lambda \mathcal{L}_{\text{dampen}}$. We choose to apply our bin regularization in the latent weight domain such that the resulting gradient

$$\frac{\partial \mathcal{L}_{\text{dampen}}}{\partial \mathbf{w}} = 2(\mathbf{w} - \widehat{\mathbf{w}}) \cdot \mathbf{1}_{sn \leq \mathbf{w} \leq sp} \quad (6)$$

is independent of the scale s and, therefore, indirectly independent of the bit-width². We further clip our latent weights to the range of the quantization grid, such that only weights that do not get clipped during quantization receive a regularization effect. This is important to avoid any harmful

²A similar loss defined in integer domain ($\|\mathbf{w}_{\text{int}} - \mathbf{w}/s\|_F^2$) would lead to a scaling of gradient with $1/s^2$ and make the regularization on \mathbf{w} dependent on the scale and bit-width.

Algorithm 1 QAT with iterative weight freezing

```

1: Init:  $f^0 \leftarrow \mathbf{0}, b \leftarrow \mathbf{0}, \Delta^\tau \leftarrow \mathbf{0}, \mathbf{w}_{\text{EMA}(\text{int})}^0 \leftarrow \mathbf{w}_{\text{int}}^0$ 
2: for  $t = 1, \dots, T$  do
3:   Calculate gradient  $g^t = \frac{\partial \mathcal{L}}{\partial \mathbf{w}}$ 
4:   Optimizer update for weights  $\mathbf{w}^t[-b]$  using  $g^t$ 
5:    $\mathbf{w}_{\text{int}}^t \leftarrow \text{clip}\left(\lfloor \frac{\mathbf{w}^t}{s} \rfloor, n, p\right)$ 
6:    $\Delta_{\text{int}}^t \leftarrow \mathbf{w}_{\text{int}}^t - \mathbf{w}_{\text{int}}^{t-1}$ 
7:    $o^t \leftarrow (\text{sign}(\Delta_{\text{int}}^t) \neq \text{sign}(\Delta_{\text{int}}^\tau)) \odot (\Delta_{\text{int}}^t \neq \mathbf{0})$ 
8:    $f^t \leftarrow m \cdot o^t + (1 - m) \cdot f^{t-1}$ 
9:   for  $i = 1, \dots, N$  do
10:    if  $f_i^t > f_{\text{th}}$  then
11:       $b_i \leftarrow \text{True}$ 
12:       $\mathbf{w}_i^t \leftarrow s \cdot \lfloor \mathbf{w}_{\text{EMA}(\text{int})_i}^{t-1} \rfloor$ 
13:    end if
14:  end for
15:   $\mathbf{w}_{\text{EMA}(\text{int})}^t \leftarrow m \cdot \mathbf{w}^{t-1} + (1 - m) \cdot \mathbf{w}_{\text{EMA}(\text{int})}^{t-1}$ 
16:   $\Delta_{\text{int}}^\tau[o^t] \leftarrow \Delta_{\text{int}}^t[o^t]$ 
17: end for

```

interactions with quantization scale gradient in LSQ-based range learning (Esser et al., 2020). A drawback of such a regularization is that it does not only affect weights that oscillate but can also hinder the movement of weights that are not in an oscillating state.

4.3. Iterative freezing of oscillating weights

We propose another, more targeted, approach to prevent weights from oscillating by freezing them during training. In this method, we track the oscillations frequency per weight during training, as described in equation (4). If the oscillation frequency of any weight exceeds a threshold f_{th} , that weight gets frozen until the end of training. We apply the freezing in the integer domain, such that potential change in the scale s during optimization does not lead to a different rounding.

When a weight oscillates, it does not necessarily spend an equal amount of time at both oscillating states. As we show in the toy example in section 2.2, the likelihood of a weight being in each state linearly depends on distance of that quantized state from the optimal value. As a result, the expectation of all quantized values over time will correspond to the optimal value. Once the frequency of a weight exceeds the threshold, it could be in either of the two quantized states. To freeze the weight to its more frequent state we keep a record of the previous integer values using an exponential moving average (EMA). We then assign the most frequent integer state to the frozen weight by rounding the EMA.

We summarize our proposed iterative weight freezing in algorithm 1. Note, this algorithm can be used in combination with any gradient-based optimizer and is not limited to a

particular quantization formulation or gradient estimator. The idea of freezing weights on an iteration level is closely related to iterative pruning (Zhu & Gupta, 2018), where small weights are iteratively pruned (frozen to zero).

5. Experiments

In this section, we evaluate the effectiveness of our proposed methods for overcoming oscillations and compare them against other QAT methods on ImageNet (Russakovsky et al., 2015). We focus on low-bit quantization, 4 and 3 bits, of efficient networks with depth-wise separable convolutions, where the effect of weight oscillation is most pronounced. We start with our ablation studies, where we only quantize the weights in a network and, finally, present results for low-bit weight and activation quantization.

5.1. Experimental setup

Quantization We follow the example of existing QAT literature and apply LSQ-type (Esser et al., 2020) weight and activation quantization: we quantize all weights to a low bit-width while keeping the weights of the first and last layer to 8-bits, and quantize the input to all layers except for the normalizing layers. We use per-tensor quantization (Krishnamoorthi, 2018) and learn the quantization scaling factor.

Optimization In all cases, we start from a pre-trained full-precision network and instantiate the weight and activation quantization parameters using MSE range estimation (Nagel et al., 2021). We use SGD with a momentum of 0.9 and train using a cosine annealing learning-rate decay. We train for 20 epochs with only weight quantization for the ablation studies. For weight and activation quantization in section 5.3, we train all models for 90 epochs. Depending on the network and quantization bit-width we train with a learning rate of either 0.01 or 0.0033.

5.2. Ablation studies

Oscillation dampening In table 4 we study how the strength of the dampening loss affects the network’s final accuracy and the proportion of oscillating weights at the end of training. In the first three rows, we observe that as we increase the coefficient λ , the proportion of oscillating weights decreases and the gap between pre and post-BN re-estimation accuracy closes. However, too much dampening harms the final accuracy, suggesting that excessive regularization inhibits beneficial movement of weights between quantization levels.

Our solution to this problem is to gradually increase the regularization strength during training. This allows the latent weights to move more freely in the first stages of training,

Regularization	pre-BN	post-BN	Osc.(%)
Baseline	64.97 ^{1.23}	69.50 ^{0.04}	4.93
$\lambda = 10^{-4}$	65.97 ^{1.52}	69.65 ^{0.08}	2.18
$\lambda = 10^{-3}$	66.99 ^{1.41}	69.96 ^{0.12}	0.21
$\lambda = 10^{-2}$	68.04 ^{1.04}	68.57 ^{0.07}	0.01
$\lambda = \cos(0, 10^{-4})$	64.47 ^{1.59}	69.61 ^{0.07}	2.64
$\lambda = \cos(0, 10^{-3})$	68.79 ^{1.31}	70.37 ^{0.06}	1.63
$\lambda = \cos(0, 10^{-2})$	70.18 ^{0.18}	70.26 ^{0.08}	1.11

Table 4. Effect of oscillation dampening strength and schedule in the final accuracy of MobileNetV2 with 3-bit weight quantization. *post-BN*: accuracy after BN re-estimation, *pre-BN*: accuracy before BN re-estimation. *Osc.*: percentage of oscillating weights at the end of training ($f > 0.005$). We report the average of 3 seeds and the STD in the superscript.

Method	pre-BN	post-BN	Osc.(%)
Baseline	64.97 ^{1.23}	69.50 ^{0.04}	4.93
$f_{th} = 0.02$	68.13 ^{2.14}	69.96 ^{0.04}	2.93
$f_{th} = 0.015$	69.79 ^{0.07}	70.13 ^{0.05}	1.23
$f_{th} = 0.01$	69.12 ^{0.53}	69.18 ^{0.47}	0.06
$f_{th} = \cos(0.04, 0.015)$	69.51 ^{0.15}	69.96 ^{0.03}	2.33
$f_{th} = \cos(0.04, 0.01)$	69.97 ^{0.06}	70.33 ^{0.07}	0.04

Table 5. Effect of oscillation frequency threshold (f_{th}) on training of MobileNetV2 with 3-bit weight quantization. *post-BN*: accuracy after BN re-estimation, *pre-BN*: accuracy before BN re-estimation, *Osc.*: percentage of oscillating weights at the end of training ($f > 0.005$). We report the average of 3 seeds and the STD in the superscript.

while reducing harmful oscillations close to convergence by applying a stronger regularization. We find that a cosine annealing schedule for λ works well in practice. Han et al. (2021) also notice that such a regularization is harmful at the early stages of training but instead adopt a two-stage optimization process. We see that such a strategy can significantly dampen the oscillations while not harming the accuracy. The best damping configuration improves almost 1% over the post-BN re-estimation baseline and more than 5% over the pre-BN re-estimation baseline.

In figure 4 (left), we also visualize the effect of dampening on the latent weight distribution of the same depth-wise separable layer as in figure 3. As intended, the latent weights are now clustered around the quantization bin centres with hardly any weights at the decision boundary.

Iterative weight freezing In table 5, we demonstrate the effectiveness of our iterative weight freezing algorithm for various freezing thresholds. Using a constant threshold throughout training, we see that the amount of remaining




Figure 4. Histogram of the distance of latent weights of layer conv.3.1 of MobileNetV2 from their closest quantization grid-point after training with oscillation dampening (left) or iterative weight freezing (right).

Method	W/A	Val. Acc. (%)
Full-precision	32/32	71.7
LSQ* (Essex et al., 2020)	4/4	69.5 (-2.3)
PACT (Choi et al., 2018)	4/4	61.4 (-10.3)
DSQ (Gong et al., 2019)	4/4	64.8 (-6.9)
EWGS (J. Lee, 2021)	4/4	70.3 (-1.6)
LSQ + BR (Han et al., 2021)	4/4	70.4 (-1.4)
LSQ + Dampen (ours)	4/4	70.5 (-1.2)
LSQ + Freeze (ours)	4/4	70.6 (-1.1)
LSQ* (Essex et al., 2020)	3/3	65.3 (-6.5)
LSQ + BR (Han et al., 2021)	3/3	67.4 (-4.4)
LSQ + Dampen (ours)	3/3	67.8 (-3.9)
LSQ + Freeze (ours)	3/3	67.6 (-4.1)

Table 6. Comparison of oscillation dampening and freezing with other QAT methods for MobileNetV2 on ImageNet classification. We report validation accuracy (%) and the difference to the corresponding full precision baseline. Note, both Han et al. (2021) and J. Lee (2021) use a slightly higher baseline model. *results from Han et al. (2021).

oscillations significantly reduces with a decreasing threshold and only some low frequency oscillations remain in the network. We also see that the pre-BN re-estimation accuracy is closer to the the post-BN re-estimation accuracy, as one would expect when there are fewer oscillations at the end of training.

However, we note that if the oscillation threshold becomes too low, too many weights get frozen in the early stages of training reducing the final accuracy. To address this, we apply an annealing schedule to the freezing threshold similar to the one used in dampening. This allows us to use a stronger freezing threshold and freeze almost all oscillations towards the end of training, when they are most disruptive. The best freezing threshold improves almost 1% over the post BN re-estimation baseline and more than 5% over the pre BN re-estimation baseline. Its accuracy is on par with the oscillation dampening while having significantly less remaining oscillations (0.04% vs 1.11%).

Method	W/A	Val. Acc. (%)
Full-precision	32/32	65.1
LSQ* (Essex et al., 2020)	4/4	61.0
LSQ + BR (Han et al., 2021)	4/4	61.5
LSQ + Dampen (ours)	4/4	63.7
LSQ + Freeze (ours)	4/4	63.6
LSQ* (Essex et al., 2020)	3/3	52.0
LSQ + BR (Han et al., 2021)	3/3	56.0
LSQ + Dampen (ours)	3/3	59.0
LSQ + Freeze (ours)	3/3	58.9

Table 7. Comparison of oscillation dampening and freezing with other QAT methods for MobileNetV3-Small on ImageNet classification. We report validation accuracy (%). *results from Han et al. (2021).

Method	W/A	Val. Acc. (%)
Full-precision	32/32	75.4
LSQ* (Essex et al., 2020)	4/4	72.3
LSQ + Dampen (ours)	4/4	73.5
LSQ + Freeze (ours)	4/4	73.5
LSQ* (Essex et al., 2020)	3/3	69.7
LSQ + Dampen (ours)	3/3	71.1
LSQ + Freeze (ours)	3/3	71.0

Table 8. Comparison of oscillation dampening and freezing for EfficientNet-lite on ImageNet classification. We report validation accuracy (%). *our re-implementation.

In figure 4 (right), we see how iterative weight freezing changes the latent weight distribution of layer conv.3.1 of MobileNetV2. The majority of the latent weights are now frozen at the bin centres removing the peaks observed at the decision boundary in figure 3.

5.3. Comparison to other QAT methods

We compare our methods for overcoming oscillation to other QAT alternatives and demonstrate their effectiveness in low-bit quantization of popular efficient neural networks. To compare to existing methods in the literature we quantize both weights and activations. In table 6, we show the results for MobileNetV2 and demonstrate that both our algorithms outperform all competing QAT techniques in literature for 3 and 4-bit quantization. We also achieve state-of-the-art results for MobileNetV3-Small and EfficientNet-lite in tables 7 and 8, respectively. In all cases, our oscillation prevention methods significantly improve (> 1%) over the commonly used LSQ baseline, showing the general applicability of our methods to other efficient networks.

We note that oscillation dampening leads to an increase of

approximately 33% in training time compared to the LSQ baseline. A similar trend was also observed by (Han et al., 2021). On the other hand, iterative weight freezing has a negligible computational overhead while achieving similar performance.

6. Conclusion

In this work, we showed that weight oscillations induced by the straight-through estimator and its variations proposed in literature can adversely affect the performance of quantized networks when trained with quantization-aware training. The effect of weight oscillations is more pronounced in low-bit quantization of light-weight networks, such as MobileNetV2 that include depth-wise layers with only a few weights per output channel. We demonstrate that these oscillations not only corrupt the running batch-normalization statistics used during inference but also harm the optimization process itself, leading to sub-optimal solutions. We propose two methods to tackle this problem: *oscillations dampening* and *iterative weight freezing*, and show that both methods lead to state-of-the-art accuracy on ImageNet classification for a range of efficient models quantized to low-bits.

References

- Banner, R., Nahshan, Y., and Soudry, D. Post training 4-bit quantization of convolutional networks for rapid-deployment. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Bhalgat, Y., Lee, J., Nagel, M., Blankevoort, T., and Kwak, N. Lsq+: Improving low-bit quantization through learnable offsets and better initialization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020.
- Cai, Z., He, X., Sun, J., and Vasconcelos, N. Deep learning with low precision by half-wave gaussian quantization. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Chai, S. M. Quantization-guided training for compact tiny{ml} models. In *Research Symposium on Tiny Machine Learning*, 2021.
- Chen, S., Wang, W., and Pan, S. J. Metaquant: Learning to quantize by learning to penetrate non-differentiable quantization. In *Neural Information Processing Systems (NeurIPS)*, 2019.
- Chmiel, B., Banner, R., Hoffer, E., Yaacov, H. B., and Soudry, D. Logarithmic unbiased quantization: Practical 4-bit training in deep learning. *arXiv preprint arXiv:2112.10769*, 2022.
- Choi, J., Wang, Z., Venkataramani, S., Chuang, P. I., Srinivasan, V., and Gopalakrishnan, K. PACT: parameterized clipping activation for quantized neural networks. *arXiv preprint arxiv:805.06085*, 2018.
- Défossez, A., Adi, Y., and Synnaeve, G. Differentiable model compression via pseudo quantization noise. *arXiv preprint arXiv:2104.09987*, 2021.
- Esser, S. K., McKinstry, J. L., Bablani, D., Appuswamy, R., and Modha, D. S. Learned step size quantization. In *International Conference on Learning Representations (ICLR)*, 2020.
- Fan, A., Stock, P., Graham, B., Grave, E., Gribonval, R., Jegou, H., and Joulin, A. Training with quantization noise for extreme model compression. In *International Conference on Learning Representations (ICLR)*, 2021.
- Gong, R., Liu, X., Jiang, S., Li, T., Hu, P., Lin, J., Yu, F., and Yan, J. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. *International Conference on Computer Vision (ICCV)*, 2019.
- Gupta, S., Agrawal, A., Gopalakrishnan, K., and Narayanan, P. Deep learning with limited numerical precision. In *International Conference on Machine Learning (ICML)*, 2015.
- Han, T., Li, D., Liu, J., Tian, L., and Shan, Y. Improving low-precision network quantization via bin regularization. In *International Conference on Computer Vision (ICCV)*, 2021.
- Helwegen, K., Widdicombe, J., Geiger, L., Liu, Z., Cheng, K.-T., and Nusselder, R. Latent weights do not exist: Rethinking binarized neural network optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Hinton, G. Neural networks for machine learning, lectures 15b. 2012.
- Horowitz, M. 1.1 computing’s energy problem (and what we can do about it). In *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pp. 10–14, 2014.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, 2015.

- J. Lee, D. Kim, B. H. Network quantization with element-wise gradient scaling. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Jain, S. R., Gural, A., Wu, M., and Dick, C. Trained uniform quantization for accurate and efficient neural network inference on fixed-point hardware. *arXiv preprint arXiv:1903.08066*, 2019.
- Kim, J., Yoo, K., and Kwak, N. Position-based scaled gradient for model quantization and pruning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Kirkpatrick, S., Gelatt, C., and Vecchi, M. Optimization by simulated annealing. *Science (New York, N.Y.)*, 220: 671–80, 06 1983.
- Krishnamoorthi, R. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018.
- Louizos, C., Reisser, M., Blankevoort, T., Gavves, E., and Welling, M. Relaxed quantization for discretized neural networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- Nagel, M., van Baalen, M., Blankevoort, T., and Welling, M. Data-free quantization through weight equalization and bias correction. In *International Conference on Computer Vision (ICCV)*, 2019.
- Nagel, M., Amjad, R. A., Van Baalen, M., Louizos, C., and Blankevoort, T. Up or down? Adaptive rounding for post-training quantization. In *International Conference on Machine Learning (ICML)*, 2020.
- Nagel, M., Fournarakis, M., Amjad, R. A., Bondarenko, Y., van Baalen, M., and Blankevoort, T. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021.
- Nahshan, Y., Chmiel, B., Baskin, C., Zheltonozhskii, E., Banner, R., Bronstein, A. M., and Mendelson, A. Loss aware post-training quantization. *arXiv preprint arXiv:1911.07190*, 2020.
- Nguyen, H. D., Alexandridis, A., and Mouchtaris, A. Quantization aware training with absolute-cosine regularization for automatic speech recognition. In *Interspeech*, 2020.
- Park, E. and Yoo, S. PROFIT: A novel training method for sub-4-bit mobilenet models. In *European Conference on Computer Vision (ECCV)*, 2020.
- Pervez, A., Cohen, T., and Gavves, E. Low bias low variance gradient estimates for hierarchical boolean stochastic networks. In *International Conference on Machine Learning (ICML)*, 2020.
- Peters, J. W. T. and Welling, M. Probabilistic binary neural networks. *arXiv preprint arXiv:1809.03368*, 2018.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 2015.
- Uhlich, S., Mauch, L., Cardinaux, F., Yoshiyama, K., Garcia, J. A., Tiedemann, S., Kemp, T., and Nakamura, A. Mixed precision dnns: All you need is a good parametrization. In *International Conference on Learning Representations (ICLR)*, 2020.
- Yang, J., Shen, X., Xing, J., Tian, X., Li, H., Deng, B., Huang, J., and Hua, X.-s. Quantization networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Yin, P., Lyu, J., Zhang, S., Osher, S. J., Qi, Y., and Xin, J. Understanding straight-through estimator in training activation quantized neural nets. In *International Conference on Learning Representations (ICLR)*, 2019.
- Zhou, S., Ni, Z., Zhou, X., Wen, H., Wu, Y., and Zou, Y. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016.
- Zhu, M. and Gupta, S. To prune, or not to prune: exploring the efficacy of pruning for model compression. In *International Conference on Learning Representations, ICLR, Workshop Track Proceedings*, 2018.

A. Appendix

A.1. Gradients for 1D regression problem

In this section, we analytically derive the gradient descent (GD) updates for the latent weight the 1D toy regression example of section 2.2. We compare the gradient from the vanilla STE with other methods from the literature (cf. section 3) and our proposed oscillation dampening solution (cf. section 4.2). We define \bar{w} as the boundary between the two quantization levels closest to the optimal value w_* , such that $\bar{w} = (w_\uparrow + w_\downarrow)/2$, where $w_\uparrow = q(w_* + \frac{s}{2})$ and $w_\downarrow = q(w_* - \frac{s}{2})$.

The gradient of the mean-squared loss from equation (3) w.r.t. to the weight is given by:

$$\frac{\partial \mathcal{L}(w)}{\partial w} = \sigma^2 (q(w) - w_*) \frac{\partial q(w)}{\partial w}, \quad (7)$$

$$= \begin{cases} \sigma^2 (w_\uparrow - w_*) \frac{\partial q(w)}{\partial w}, & \text{if } w \geq \bar{w} \\ \sigma^2 (w_\downarrow - w_*) \frac{\partial q(w)}{\partial w}, & \text{if } w < \bar{w} \end{cases} \quad (8)$$

Assuming $\sigma = 1$, a learning rate of η , and that no clipping takes place we write the GD update rules for subset of methods discussed in section 3. For the PSG (Kim et al., 2020), the update rule includes an arbitrarily small constant, $\epsilon > 0$, while EWGS (J. Lee, 2021) introduce a $\delta \geq 0$ scaling factor.

STE	$w_{t+1} = \begin{cases} w_t - \eta (w_\uparrow - w_*), & \text{if } w \geq \bar{w} \\ w_t - \eta (w_\downarrow - w_*), & \text{if } w < \bar{w} \end{cases}$
PSG (Kim et al., 2020)	$w_{t+1} = \begin{cases} w_t - \eta (w_\uparrow - w_*) (w_\uparrow - w + \epsilon), & \text{if } w \geq \bar{w} \\ w_t - \eta (w_\downarrow - w_*) (w - w_\downarrow + \epsilon), & \text{if } w < \bar{w} \end{cases}$
EWGS (J. Lee, 2021)	$w_{t+1} = \begin{cases} w_t - \eta (w_\uparrow - w_*) (1 + \delta (w - w_\uparrow)), & \text{if } w \geq \bar{w} \\ w_t - \eta (w_\downarrow - w_*) (1 - \delta (w - w_\downarrow)), & \text{if } w < \bar{w} \end{cases}$
Dampening (sec. 4.2)	$w_{t+1} = \begin{cases} w_t - \eta (w_\uparrow - w_* + 2\lambda (w - w_\uparrow)), & \text{if } w \geq \bar{w} \\ w_t - \eta (w_\downarrow - w_* + 2\lambda (w - w_\downarrow)), & \text{if } w < \bar{w} \end{cases}$

If we look at the equations for PSG and EWGS more carefully, we see that they introduce an *always positive* scaling term to the STE gradient (assuming sufficiently small δ). Therefore, these methods can only change the magnitude of the update and not its direction, meaning they cannot stop oscillations from happening. Because the methods apply a scaling to the STE gradient, we call them *multiplicative* variations in section 3.

On the other hand, our dampening method adds an extra term to the gradient, which is of opposite sign to the STE gradient. Thus, this *additive* variation of the STE can indeed prevent oscillations by changing the direction of the gradient.

A.2. Relation of the distance and frequency

In toy example of section 2.2, we explained how the frequency of the oscillations depends linearly on distance of the optimal weight value w_* from its closest quantization level $q(w_*)$. We can, thus, define the oscillation frequency as:

$$f = \frac{|q(w_*) - w_*|}{s}, \quad (9)$$

where s is the quantization scaling factor or size of the quantization bin. The frequency depends linearly on the distance $d = |q(w_*) - w_*|$ due to the squared error loss. In figure 5, we demonstrate this dependency experimentally by varying the distance d and observing the effect on the oscillation frequency.

A.3. Different learning rates and multiplicative regularization methods

In figure 6, we plot the latent weight in our toy regression example for different learning rates assuming STE gradients. We can see that the change in learning rate only influences the amplitude of the oscillations and not the oscillation frequency.




Figure 5. Oscillation of single weight in toy-regression problem for optimal values. The distance between the closest quantization point and the optimal value directly corresponds to the oscillation frequency.




Figure 6. Oscillation of single weight in toy-regression problem for various learning rates (with STE).

This is also apparent from equation (9), which shows that the frequency is independent of the learning rate. Theoretically, this suggests that *multiplicative* variations of the STE, such as EWGS and PSG, cannot prevent oscillations from occurring. Assuming a sufficiently small learning rate, weights will in the limit still converge to the decision threshold and start oscillating. While oscillating, the distance of the oscillating weight from its two closest quantization levels is roughly constant and equal to $s/2$. This means that close to the decision threshold the weight update rules for PSG and EWGS become:

$$\begin{array}{l} \text{PSG (Kim et al., 2020)} \\ \text{EWGS (J. Lee, 2021)} \end{array} \left| \begin{array}{l} w_{t+1} = \begin{cases} w_t - \eta \left(\frac{s}{2} + \epsilon \right) (w_{\uparrow} - w_*) , & \text{if } w \geq \bar{w} \\ w_t - \eta \left(\frac{s}{2} + \epsilon \right) (w_{\downarrow} - w_*) , & \text{if } w < \bar{w} \end{cases} \\ w_{t+1} = \begin{cases} w_t - \eta \left(1 - \frac{s\delta}{2} \right) (w_{\uparrow} - w_*) , & \text{if } w \geq \bar{w} \\ w_t - \eta \left(1 - \frac{s\delta}{2} \right) (w_{\downarrow} - w_*) , & \text{if } w < \bar{w} \end{cases} \end{array} \right.$$

Both methods now apply a constant and positive scaling to the gradient. Given that typically $s < 1$, and that for EWGS $\delta < 1$, the combined gradient scaling is < 1 effectively reducing the learning rate and consequently the oscillation amplitude. This analysis is only valid if the oscillation amplitude is small, which is typical towards the end of training. However, in the earlier stages of training, where the weights move more freely, the dynamics of the oscillations might be more complicated, as the gradient or velocity of the weight depends on the distance from the closest quantization bin.

A.4. The bias of the STE

In the original paper on the straight-through estimator, Bengio et al. (2013) note that “...(the STE) is clearly a biased estimator”. Several authors, such as Chmiel et al. (2022); Fan et al. (2021), have repeated this statement and have argued that this *bias* is a fundamental drawback of STE when applied to quantization, without defining what this actually means in the context of quantization-aware training.

The bias of the STE in the original formulation occurs in a probabilistic setting. Let’s assume we want to optimize a loss $\mathcal{L} = \mathbb{E}_{p(z|\theta)} [f(z)]$ where z is latent parameter coming from a discrete distribution $p(z|\theta)$. To learn the parameter θ we need

to calculate the derivative $\nabla_{\theta} \mathbb{E}_{p(z|\theta)} [f(z)]$. However, due to the discrete nature of p_z this derivative does not exist. Instead the STE provides us with the convenient approximation $\mathbb{E}_{p(z|\theta)} [\nabla_z f(z) \nabla_{\theta} \tilde{z}]$, where \tilde{z} is the *unquantized* version of z . In this case, indeed the STE is in most cases a biased estimator (Pervez et al., 2020), as the expectation of the gradient does not match with the gradient of the expectation.

However, it is unclear what the bias of the estimator means in the QAT setting. Quantized networks do not have an inherent probabilistic property, unless explicitly defined so, e.g. in Bayesian neural networks or when training with stochastic rounding (Gupta et al., 2015; Louizos et al., 2019). The gradients for the floating-point shadow weights are indeed *wrong* or *biased*, but these weights are by themselves meaningless because they are not used in the forward pass. As described in (Helwegen et al., 2019) the underlying weights of a neural network when trained with STE are not representative for the network’s performance, but rather can be interpreted as a reservoir that slowly accumulates small gradient updates. Therefore, we believe that comparing to floating-point gradient in the conventional QAT setting can be misleading and nonconstructive.