## Supplementary material for MIME

## Contents of Appendix

A How momentum can help reduce client drift ..... 16
B Proof sketch ..... 17
C Experimental setup ..... 18
C. 1 Description of ablation study ..... 18
C. 2 Description of large scale experiments ..... 18
C. 3 Practicality of experiments ..... 19
C. 4 Hyperparameter search ..... 19
C. 5 Comparison with previous results ..... 20
C. 6 Additional algorithmic details ..... 20
D Additional Adam experiments ..... 21
E Stability of methods to hyper-parameters ..... 21
F Technicalities ..... 22
F. 1 Assumptions and definitions ..... 22
F. 2 Some technical lemmas ..... 23
F. 3 Properties of functions with bounded Hessian dissimilarity ..... 24
G Convergence with a generic base optimizer ..... 25
G. 1 Proof of Theorem I (generic reduction) ..... 25
G. 2 Convergence of MimeSGD and MimeLiteSGD (Corollary II) ..... 28
G. 3 Convergence of MimeAdam and MimeLiteAdam (Corollary III) ..... 31
H Circumventing server-only lower bounds ..... 33
H. 1 Algorithm descriptions ..... 34
H. 2 Bias in updates ..... 34
H. 3 Change in each client update ..... 35
H. 4 Change in each round ..... 37
H. 5 Final convergence rates ..... 41
I Algorithm pseudocodes ..... 42

## A How momentum can help reduce client drift



Figure 2: Client-drift in FedAvg (left) and Mime (right) is illustrated for 2 clients with 3 local steps and momentum parameter $\beta=0.5$. The local SGD updates of FEDAVG (shown using arrows for client 1 and client2) move towards the average of client optima $\frac{x_{1}^{\star}+x_{2}^{\star}}{2}$ which can be quite different from the true global optimum $x^{\star}$. Server momentum $m_{t}$ only speeds up the convergence to the wrong point in this case. In contrast, Mime uses unbiased momentum computed at the server parameter $x_{t}$ and applies it locally at every update. This keeps the updates of MimE closer to the true optimum $x^{\star}$.
In this section we examine the tension between reducing communication by running multiple client updates each round, and degradation in performance due to client drift [32]. To simplify the discussion, we assume a single client is sampled each round and that clients use full-batch gradients.

Server-only approach. A simple way to avoid the issue of client drift is to take no local steps. We sample a client $i \sim \mathcal{C}$ and run SGD with momentum (Mom) with momentum parameter $\beta$ and step size $\eta$ :

$$
\begin{align*}
\boldsymbol{x}_{t} & =\boldsymbol{x}_{t-1}-\eta\left((1-\beta) \nabla f_{i}\left(\boldsymbol{x}_{t-1}\right)+\beta \boldsymbol{m}_{t-1}\right), \\
\boldsymbol{m}_{t} & =(1-\beta) \nabla f_{i}\left(\boldsymbol{x}_{t-1}\right)+\beta \boldsymbol{m}_{t-1} \tag{2}
\end{align*}
$$

Here, the gradient $\nabla f_{i}\left(\boldsymbol{x}_{t}\right)$ is unbiased i.e. $\mathbb{E}\left[\nabla f_{i}\left(\boldsymbol{x}_{t}\right)\right]=\nabla f\left(\boldsymbol{x}_{t}\right)$ and hence we are guaranteed convergence. However, this strategy can be communication-intensive and we are likely to spend all our time waiting for communication with very little time spent on computing the gradients.

FEDAVG approach. To reduce the overall communication rounds required, we need to make more progress in each round of communication. The FedAvg meta algorithm utilizes a base optimizer, a client learning rate and a server learning rate. Each client performs K local update steps of SGD using the client learning rate and communicates the net update (difference between final and initial parameters) to the server. This difference is then treated as a 'pseudo-gradient' and is input into the optimizer (say momentum or Adam) to update the server parameters using the server learning rate. When the base optimizer uses momentum, this momentum is computed at the server level using the pseudo-gradients and is referred to as server momentum.
Starting from $\boldsymbol{y}_{0}=\boldsymbol{x}_{t-1}$, FEDAVG [43] runs multiple SGD steps on the sampled client $i \sim \mathcal{C}$

$$
\begin{equation*}
\boldsymbol{y}_{k}=\boldsymbol{y}_{k-1}-\eta \nabla f_{i}\left(\boldsymbol{y}_{k-1}\right) \text { for } k \in[K] \text {, } \tag{3}
\end{equation*}
$$

and then a pseudo-gradient $\tilde{\boldsymbol{g}}_{t}=-\left(\boldsymbol{y}_{K}-\boldsymbol{x}_{t}\right)$ replaces $\nabla f_{i}\left(\boldsymbol{x}_{t-1}\right)$ in the SGDm algorithm (2). This is referred to as server-momentum since it is computed and applied only at the server level [26]. However, such updates give rise to client-drift resulting in performance worse than the naïve serveronly strategy (2). This is because by using multiple local updates, (3) starts over-fitting to the local client data, optimizing $f_{i}(\boldsymbol{x})$ instead of the actual global objective $f(\boldsymbol{x})$. The net effect is that FEDAVG moves towards an incorrect point (see Fig 2, left). If $K$ is sufficiently large, approximately

$$
\begin{aligned}
\boldsymbol{y}_{K} & \rightsquigarrow \boldsymbol{x}_{i}^{\star}, \text { where } \boldsymbol{x}_{i}^{\star}:=\underset{\boldsymbol{x}}{\arg \min } f_{i}(\boldsymbol{x}) \\
\Rightarrow \mathbb{E}_{i \sim \mathcal{C}}\left[\tilde{\boldsymbol{g}}_{t}\right] & \rightsquigarrow\left(\boldsymbol{x}_{t}-\mathbb{E}_{i \sim \mathcal{C}}\left[\boldsymbol{x}_{i}^{\star}\right]\right) .
\end{aligned}
$$

Further, the server momentum is based on $\tilde{\boldsymbol{g}}_{t}$ and hence is also biased. Thus, it cannot correct for the client drift. We next see how a different way of using momentum can mitigate client drift.

Mime approach. FEDAVG experiences client drift because both the momentum and the client updates are biased. To fix the former, we compute momentum using only global optimizer state as in (2) using the sampled client $i \sim \mathcal{C}$ :

$$
\begin{equation*}
\boldsymbol{m}_{t}=(1-\beta) \nabla f_{i}\left(\boldsymbol{x}_{t-1}\right)+\beta \boldsymbol{m}_{t-1} . \tag{4}
\end{equation*}
$$

To reduce the bias in the local updates, we will apply this unbiased momentum every step $k \in[K]$ :

$$
\begin{equation*}
\boldsymbol{y}_{k}=\boldsymbol{y}_{k-1}-\eta\left((1-\beta) \nabla f_{i}\left(\boldsymbol{y}_{k-1}\right)+\beta \boldsymbol{m}_{t-1}\right) . \tag{5}
\end{equation*}
$$

Note that the momentum term is kept fixed during the local updates i.e. there is no local momentum used, only global momentum is applied locally. Since $\boldsymbol{m}_{t-1}$ is a moving average of unbiased gradients computed over multiple clients, it intuitively is a good approximation of the general direction of the updates. By taking a convex combination of the local gradient with $\boldsymbol{m}_{t-1}$, the update (5) is potentially also less biased. In this way MIME combines the communication benefits of taking multiple local steps and prevents client-drift (see Fig 2, right). Appendix B makes this intuition precise.

## B Proof sketch

In this section, we provide an intuition behind our proof of convergence of MimeMVR. There are three main components: i) how momentum reduces the effect of client drift, ii) how local steps can take advantage of Hessian similarity, and iii) why the SVRG correction improves constants.
Improving the statistical term via momentum. Intuitively, using momentum locally at every client update reduces client drift by incorporating information about other clients from past rounds. Assume that we sample a single client $i_{t}$ in round $t$ and that we use full-batch gradients. Also let the local client update at step $k$ round $t$ be of the form

$$
\begin{equation*}
\boldsymbol{y} \leftarrow \boldsymbol{y}-\eta \boldsymbol{d}_{k} . \tag{6}
\end{equation*}
$$

The ideal choice of update is of course $\boldsymbol{d}_{k}^{\star}=\nabla f(\boldsymbol{y})$ but however this is unattainable. Instead, MIME with momentum $\beta=1-a$ uses $\boldsymbol{d}_{k}^{\text {SGDm }}=\tilde{\boldsymbol{m}}_{k} \leftarrow a \nabla f_{i}(\boldsymbol{y})+(1-a) \boldsymbol{m}_{t-1}$ where $\boldsymbol{m}_{t-1}$ is the momentum computed at the server. The variance of this update can then be bounded as

$$
\begin{aligned}
\mathbb{E}\left\|\tilde{\boldsymbol{m}}_{k}-\nabla f(\boldsymbol{y})\right\|^{2} & \lesssim a^{2} \mathbb{E}\left\|\nabla f_{i_{t}}(\boldsymbol{y})-\nabla f(\boldsymbol{y})\right\|^{2}+(1-a) \mathbb{E}\left\|\boldsymbol{m}_{t-1}-\nabla f(\boldsymbol{y})\right\|^{2} \\
& \approx a^{2} G^{2}+(1-a) \mathbb{E}\left\|\boldsymbol{m}_{t-1}-\nabla f\left(\boldsymbol{x}_{t-2}\right)\right\|^{2} \approx a G^{2} .
\end{aligned}
$$

The last step follows by unrolling the recursion on the variance of $\boldsymbol{m}$. We also assumed that $\eta$ is small enough that $\boldsymbol{y} \approx \boldsymbol{x}_{t-2}$. This way, momentum can reduce the variance of the update from $G^{2}$ to $\left(a G^{2}\right)$ by using past gradients computed on different clients. Of course, this also introduces additional bias into the update. To reduce this bias requires slightly modifying the momentum algorithm similar to [14]. The full analysis is carried out in Appendix H.

Improving the optimization term via local steps. The optimization (second) term in Theorem IV is $\frac{\delta K+L}{\epsilon K}$. In contrast, the optimization term of the server-only methods is $L / \epsilon$. Since in most cases $\delta \ll L L$, the former can be significantly smaller than the latter. This rate also suggests that the best choice of number of local updates is $L / \delta$ i.e. we should perform more client updates when they have more similar Hessians. This generalizes results of [32] from quadratics to all functions.

This improvement is due to a careful analysis of the bias in the gradients computed during the local update steps. Note that for client parameters $\boldsymbol{y}_{k-1}$, the gradient $\mathbb{E}\left[\nabla f_{i_{t}}\left(\boldsymbol{y}_{k-1}\right)\right] \neq \mathbb{E}\left[\nabla f\left(\boldsymbol{y}_{k-1}\right)\right]$ since $\boldsymbol{y}_{k-1}$ was also computed using the same loss function $f_{i_{t}}$. In fact, only the first gradient computed at $\boldsymbol{x}_{t-1}$ is unbiased. Dropping the subscripts $k$ and $t$, we can bound this bias as:

$$
\begin{aligned}
\mathbb{E}\left[\nabla f_{i}(\boldsymbol{y})-\nabla f(\boldsymbol{y})\right] & =\mathbb{E}[\underbrace{\nabla f_{i}(\boldsymbol{y})-\nabla f_{i}(\boldsymbol{x})}_{\approx \nabla^{2} f_{i}(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x})}+\underbrace{\nabla f(\boldsymbol{x})-\nabla f\left(\boldsymbol{y}_{i}\right)}_{\approx \nabla^{2} f(\boldsymbol{x})\left(\boldsymbol{x}-\boldsymbol{y}_{i}\right)}]+\underbrace{\mathbb{E}_{i}\left[\nabla f_{i}(\boldsymbol{x})\right]-\nabla f(\boldsymbol{x})}_{=0 \text { since unbiased }} \\
& \approx \mathbb{E}\left[\left(\nabla^{2} f_{i}(\boldsymbol{x})-\nabla^{2} f(\boldsymbol{x})\right)\left(\boldsymbol{y}_{i}-\boldsymbol{x}\right)\right] \approx \delta \mathbb{E}\left[\left(\boldsymbol{y}_{i}-\boldsymbol{x}\right)\right] .
\end{aligned}
$$

Thus, the Hessian dissimilarity (A2) control the bias, and hence the usefulness of local updates. This intuition can be made formal using Lemma 3. Note that this improved analysis is potentially applicable to any local update methods and is not specific to Mime.

Mini-batches via SVRG correction. In our previous discussion about momentum and local steps, we assumed that the clients compute full batch gradients and that only one client is sampled per round. However, in practice a large number $(S)$ of clients are sampled and further the clients use mini-batch gradients. The SVRG correction reduces this within-client variance since

$$
\operatorname{Var}\left(\nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)-\nabla f_{i}(\boldsymbol{x} ; \zeta)+\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \nabla f_{i}(\boldsymbol{x})\right) \lesssim L^{2}\left\|\boldsymbol{y}_{i}-\boldsymbol{x}\right\|^{2}+\frac{G^{2}}{S} \approx \frac{G^{2}}{S}
$$

Here, we used the smoothness of $f_{i}(\cdot ; \zeta)$ and assumed that $\boldsymbol{y}_{i} \approx \boldsymbol{x}$ since we don't move too far within a single round. Thus, the SVRG correction allows us to use minibatch gradients in the local updates while still ensuring that the variance is of the order $G^{2} / S$. In practical deep learning, this SVRG correction may not very effective [15] and so can be dropped, though it is useful to derive the optimal theoretical rates.

## C Experimental setup

## C. 1 Description of ablation study

We train a 2 hidden layer MLP with 300u-100 neurons on the EMNIST62 (extended MNIST) dataset [12]. The clients' data is separated according to the original authors of the characters [10]. All methods are augmented with momentum-Mime and MimeLite use momentum in the client updates, and the others use server momentum. The momentum parameter is searched over $\beta \in[0,0.9,0.99]$. For Adam, we fix $\beta_{1}=0.9, \beta_{2}=0.99$, and $\epsilon=10^{-3}$. For both FedProx and SCAFFOLD, $\beta=0$ (no server momentum) yielded the best performance. For FedAvg, Mime, and MimeLite $\beta=0.9$ was the fastest. For FedProx, the regularization parameter $\mu$ was searched over $[0.1,0.5,1]$ and $\mu=0.1$ had highest test accuracy.

## C. 2 Description of large scale experiments

We perform 4 tasks over 3 datasets: i) On the EMNIST62 dataset [12] we run a convex multi-class (62 classes) logistic regression model, and ii) a convolution model with two CNN layers and two dense layers and dropout. iii) On the Shakespeare dataset, we train a single layer LSTM model with state size of 256 and embedding size of 8 to predict the next character [43]. iv) Finally, on the StackOVERFLOW dataset [16], we train a next word prediction language model with embedding size of 96, a LSTM layer of size 670, and a vocabulary size of 1000. In all cases we report the top-1 test accuracy in our experiments.

All datasets use the metadata indicating the original authors to separate them into multiple clients yielding naturally partitioned datasets. Table 3 summarizes the statistics about the different datasets. Note that the average number of rounds a client participates in (computed as sampled clients $\times$ number of rounds/number of clients) provides an indication of how much of the training data is seen with Shakespeare being closest to the cross-silo setting and StackOverflow representing the most cross-device in nature.

Table 3: Details about the datasets used and experiment setting.

|  | EMNIST62 | SHAKESPEARE | STACKOVERFLOW |
| :--- | :--- | :--- | :--- |
| Clients | 3,400 | 715 | 342,477 |
| Examples | 671,585 | 16,068 | $135,818,730$ |
| Batch size | 10 | 10 | 10 |
| Number of local epochs | 1 | 1 | 1 |
| Total number of rounds | 1000 | 1000 | 1000 |
| Avg. rounds each client participates | 5.9 | 28 | 0.15 |

We use Tensorflow federated datasets [60] to generate the datasets. Our federated learning simulation code is written in FedJAX $[52,53]$ and is open-sourced at github. cm/google/fedjax (see documentation). Black and white was reversed in EMNIST62 (i.e. subtracted from 1) to make them similar to MNIST. The preprocessing for Shakespeare and StackOverflow datasets exactly matches that of [49].

Table 4: Effective number of sampled clients.

|  | Total Comm. | EMNIST62 | SHAKESPEARE | STACKOvERFLOW |
| :--- | :--- | :--- | :--- | :--- |
| FedAvg | $2 \times$ | 20 | 20 | 50 |
| MimeLiteMom | $5 \times$ | 8 | 8 | 20 |
| MimeLiteAdagrad | $5 \times$ | 8 | 8 | 20 |
| MimeLiteAdam | $6 \times$ | 6 | 6 | 16 |
| MimeMom | $6 \times$ | 6 | 6 | 16 |
| MimeAdagrad | $6 \times$ | 6 | 6 | 16 |
| MimeAdam | $7 \times$ | 5 | 5 | 14 |

## C. 3 Practicality of experiments

In the experiments we only cared about the number of communication rounds, ignoring that MIME actually needs twice the number of bits per round and that the SERVER-ONLY methods have a much smaller computational requirement. This is standard in the federated learning setting as introduced by [43] and is justified because most of the time in cross-device FL is spent in establishing connections with devices rather than performing useful work such as communication or computation. In other words, latency and not bandwidth or computation are critical in cross device FL. However, one can certainly envision cases where this is not true. Incorporating communication compression strategies [58, 3, 33, 64] or client-model compression strategies [9, 19, 22] into our MIME framework can potentially address such issues and are important future research directions.
Regarding the algorithms evaluated, we chose not to include MVR as a base optimizer. This is because it is not a popular choice is practice even in the centralized setting, and serves more as a theoretical stand in to explain the benefit of the simpler SGD with momentum algorithm. Hence, we wouldn't expect MimeMVR to perform better than MimeMom. In general, our goal was to "mimic" centralized methods - methods which have better empirical performance (momentum and Adam) we showed also perform well in the federated setting when combined with Mime, and similarly methods which have better theoretical rates (MVR) have good rates with Mime as well.
Further, as we noted previously, we believe both the datasets and the tasks being studied here are close to real world settings since they contain natural heterogeneity. We now discuss our choice of other parameters in the experiment setup (number of training rounds, sampled clients, batch-size, etc.) Each round of federated learning takes 3 mins in the real world and is relatively independent of the size of communication [7] implying that training 1000 rounds takes 2 days even for small models. In contrast, running a centralized simulation takes about 15 mins. This underscores the importance of ensuring that the algorithms for federated learning converge in as few rounds as possible, as well as have very easy to set default hyper-parameters. Thus, in our experimental setup we keep all parameters other than the learning rate to their default values. In practice, this learning rate can be set by set using a small centralized dataset on the server (as in [23]). Thus, it is crucial for federated frameworks to be able to translate algorithms which work well in centralized settings directly to the federated setting without additinal hyper-parameter tuning. The choice of batch size being 10 was made both keeping in mind the limited memory available to each client as well as to match prior work. Finally, while we limit ourselves to sampling 20-50 workers per round due to computational constraints, in real world FL thousands of devices are often available for training simultaneously each round [7]. They also note that the probability of each of these devices being available has clear patterns and is far from uniform sampling. Conducting a large scale experimental study which mimics these alternate forms of heterogeneity is an important direction for future work.

## C. 4 Hyperparameter search

We run two hyper-parameter sweeps in our experiments: first a light setup which is reported in the main paper, and one we believe reflects the real world performance, and second a heavy tuning setting to showcase the performance of the methods as we vary the hyper-parameters.

Light-sweep setting (9×). For all Momentum methods, we pick momentum $\beta=0.9$. For Adam methods, we fix $\beta_{1}=0.9$ and $\beta_{2}=0.99$, and $\varepsilon_{0}=1 \times 10^{-7}$. For Adagrad we use the default initialization value of 0.1 and use $\varepsilon_{0}=1 \times 10^{-7}$. None of the algorithms use weight decay, clipping
etc. The learning rate is then tuned to obtain the best test accuracy. For all experiments, unless explicitly mentioned otherwise, the learning rate is searched over a grid ( $9 \times$ ):
$\eta \in\left[1,1 \times 10^{-0.5}, 1 \times 10^{-1}, 1 \times 10^{-1.5}, 1 \times 10^{-2}, 1 \times 10^{-2.5}, 1 \times 10^{-3}, 1 \times 10^{-3.5}, 1 \times 10^{-4}\right]$.

The server learning rate for all methods is kept at its default value of 1 .

Heavy-sweep setting (567×). For all Momentum methods, we pick momentum $\beta=0.9$. For Adam methods, we fix $\beta_{1}=0.9$ and $\beta_{2}=0.99$. For Adagrad we use the default initialization value of 0.1 . None of the algorithms use weight decay, clipping etc. The learning rate is then tuned to obtain the best test accuracy.
For all experiments, unless explicitly mentioned otherwise, the client learning rate is searched over a grid $(9 \times)$ :
$\eta_{\text {client }} \in\left[1,1 \times 10^{-0.5}, 1 \times 10^{-1}, 1 \times 10^{-1.5}, 1 \times 10^{-2}, 1 \times 10^{-2.5}, 1 \times 10^{-3}, 1 \times 10^{-3.5}, 1 \times 10^{-4}\right]$.

Further, we also search for the server learning rate is searched over a grid ( $9 \times$ ):
$\eta_{\text {server }} \in\left[1 \times 10^{1}, 1 \times 10^{0.5}, 1,1 \times 10^{-0.5}, 1 \times 10^{-1}, 1 \times 10^{-1.5}, 1 \times 10^{-2}, 1 \times 10^{-2.5}, 1 \times 10^{-3}\right]$.

Finally, for the adaptive methods such as Adam and Adagrad, we also tune the $\varepsilon_{0}$ parameter over a $\operatorname{grid}(7 \times)$ :

$$
\varepsilon_{0} \in\left[1,1 \times 10^{-1}, 1 \times 10^{-2}, 1 \times 10^{-3}, 1 \times 10^{-4}, 1 \times 10^{-5}, 1 \times 10^{-6}, 1 \times 10^{-7}\right]
$$

## C. 5 Comparison with previous results

As far as we are aware, [49] is the only prior work which conducts a systematic experimental study of federated learning algorithms over multiple realistic datasets. The algorithms comparable across the two works (e.g. FedAvgSGD, FedAvgMom, and FedAvgAdam) have qualitatively similar performance except with one exception: FedAvgAdam consistently underperforms FedAvgMom. This difference, as we show later, is because FedAvgAdam does not work with the default choices of hyper-parameters such as $\epsilon$ and requires additional tuning. As we explain in Section C.3, we chose to keep these parameters to the default values of their centralized counterparts to compare methods in a 'low-tuning' setting. We also point that while FedAvgAdam struggles to perform in this setup, MimeAdam and MimeLiteAdam are very stable and even often outperform their SGD counterparts.

## C. 6 Additional algorithmic details

Table 5: Decomposing base algorithms into a parameter update $(\mathcal{U})$ and statistics tracking $(\mathcal{V})$.

| Algorithm | Tracked statistics $\boldsymbol{s}$ | Update $\operatorname{step} \mathcal{U}$ | Tracking step $\mathcal{V}$ |
| :--- | :--- | :--- | :--- |
| SGD | - | $\boldsymbol{x}-\eta \boldsymbol{g}$ | - |
| SGDm/Mom | $\boldsymbol{m}$ | $\boldsymbol{x}-\eta((1-\beta) \boldsymbol{g}+\beta \boldsymbol{m})$ | $\boldsymbol{m}=(1-\beta) \boldsymbol{g}+\beta \boldsymbol{m}$ |
| AdaGrad | $\boldsymbol{v}$ | $\boldsymbol{x}-\frac{\eta}{\epsilon+\sqrt{\boldsymbol{v}}} \boldsymbol{g}$ | $\boldsymbol{v}=\boldsymbol{g}^{2}+\boldsymbol{v}$ |
| Adam | $\boldsymbol{m}, \boldsymbol{v}$ | $\boldsymbol{x}-\frac{\eta}{\epsilon+\sqrt{\boldsymbol{v}}}\left(\left(1-\beta_{1}\right) \boldsymbol{g}+\beta_{1} \boldsymbol{m}\right)$ | $\boldsymbol{m}=\left(1-\beta_{1}\right) \boldsymbol{g}+\beta_{1} \boldsymbol{m}$ <br> $\boldsymbol{v}=\left(1-\beta_{2}\right) \boldsymbol{g}^{2}+\beta_{2} \boldsymbol{v}$ |

## D Additional Adam experiments



Figure 3: Mime, MimeLite, FedAvg, and Loc-Mime with Adam using 10 local epochs, run on EMNIST62 and a 2 hidden layer (300u-100) MLP. (Left) Mime and MimeLite are nearly identical and outperform FedAvg. (Right) Locally adapting Adam state slows down convergence and makes it more unstable. Both these results are consistent with the earlier momentum results.

## E Stability of methods to hyper-parameters



Figure 4: Stability of adaptive methods with varying server learning: FedAvg (left), Mime (middle) and MimeLite (right) with Adam (top) and Adagrad (bottom) as base algorithms are run on EMNIST62 with CNN. For each value of server learning rate ( $y$-axis) and $\varepsilon_{0}$ ( $x$-axis), the client learning rate was tuned over the $9 \times$ grid and the accuracy reported. The red box highlights the default configuration in a centralized setting. We see that FedAvgAdam is very sensitive to the server learning rate and $\varepsilon_{0}$, performing poorly in the default centralized parameter regimes. Mime and MimeLite acheive their best performance with the centralized parameters. This justifies our claim that Mime and MimeLite can adapt any centralized method with the same hyper-parameters and only require tuning of a single learning rate. This, we believe, is crucial for real world deployment.


Figure 5: Stability of non-adaptive methods with varying server learning: FedAvg, Mime and MimeLite with SGD and momentum ( $\beta=0.9$ ) as base algorithms are run on EMNIST62 with CNN. For each value of server learning rate, the client learning rate was tuned over the $9 \times$ grid. The momentum methods are more insensitive to the server learning rate than the SGD methods. Server learning rate of 1 (default value) seems to work well for all methods.

## F Technicalities

We examine some additional definitions and introduce some technical lemmas.

## F. 1 Assumptions and definitions

We make precise a few definitions and explain some of their implications. We first discuss the two assumptions on the dissimilarity between the gradients (A1) and the Hessians (A2). Loosely, these two quantities are an extension of the concepts of variance and smoothness which occur in centralized SGD analysis to the federated learning setting. Just as the variance and smoothness are completely orthogonal concepts, we can have settings where $G^{2}$ (gradient dissimilarity) is large while $\delta$ (Hessian dissimilarity) is small, or vice-versa.

Our assumption about the bound on the $G$ gradient dissimilarity can easily be extended to $(G, B)$ gradient dissimilarity used by [33]:

$$
\begin{equation*}
\mathbb{E}_{i}\left\|\nabla f_{i}(\boldsymbol{x})\right\|^{2} \leq G^{2}+B^{2}\|\nabla f(\boldsymbol{x})\|^{2} . \tag{7}
\end{equation*}
$$

All the proofs in the paper extend in a straightforward manner to the above weaker notion. Since this notion does not present any novel technical challenge, we omit it in the rest of the proofs. Note however that the above weaker notion can potentially capture the fact that by increasing the model capacity, we can reduce $G$. In the extreme case, by taking a sufficiently over-parameterized model, it is possible to make $G=0$ in certain settings [63]. However, this comes both at a cost of increased resource requirements (i.e. higher memory and compute requirements per step) but can also result in other constants increasing (e.g. $B$ and $L$ ).

The second crucial definition we use in this work is that of $\delta$ bounded Hessian dissimilarity (A2). This has been used previously in the analyses of distributed [54, 6, 51] and federated learning [32], but has been restricted to quadratics. Here, we show how to extend both the notion as well as the analysis to general smooth functions. The main manner we will use this assumption is in Lemma 3 to claim that for any $\boldsymbol{x}$ and $\boldsymbol{y}$ the following holds:

$$
\begin{equation*}
\mathbb{E}\left\|\nabla f_{i}(\boldsymbol{y} ; \zeta)-\nabla f_{i}(\boldsymbol{x} ; \zeta)+\nabla f(\boldsymbol{x})-\nabla f(\boldsymbol{y})\right\|^{2} \leq \delta^{2}\|\boldsymbol{y}-\boldsymbol{x}\|^{2} . \tag{8}
\end{equation*}
$$

Here the expectation is over the choice of client $i$. To understand what the above condition means, it is illuminating to define $\Psi_{i}(\boldsymbol{z})=f_{i}(\boldsymbol{z} ; \zeta)-f(\boldsymbol{z})$. Then, we can rewrite (A2) and (8) respectively as

$$
\left\|\nabla^{2} \Psi_{i}(\boldsymbol{z})\right\| \leq \delta \quad \text { and } \quad \mathbb{E}\left\|\nabla \Psi_{i}(\boldsymbol{y})-\nabla \Psi_{i}(\boldsymbol{x})\right\|^{2} \leq \delta^{2}\|\boldsymbol{y}-\boldsymbol{x}\|^{2}
$$

Thus (8) and (A2) are both different notions of smoothness of $\Psi_{i}(\boldsymbol{x})$ (formal definition of smoothness will follow soon). The latter definition closely matches the notion of squared-smoothness used by [5] and is a promising relaxation of (A2). However, we run into some technical issues since in our case the variable $\boldsymbol{y}$ can also be a random variable and depend on the choice of the client $i$. Extending
our results to this weaker notion of Hessian-similarity and proving tight non-convex lower bounds is an exciting theoretical challenge.
Finally note that if the functions $f_{i}(\boldsymbol{x} ; \zeta)$ are assumed to be smooth as in $[54,6,32]$, then $\Psi_{i}((\boldsymbol{x})$ is $2 L$-smooth. Thus, we always have that $\delta \leq 2 L$. But, as shown in [54], it is possible to have $\delta \ll L$ if the data distribution amongst the clients is similar. Further, the lower bound from [6] proves that Hessian-similarity is the crucial quantity capturing the number of rounds of communication required for distributed/federated optimization.

We next define the terms smoothness and strong-convexity which we repeatedly use in the paper.
(A2*) $f_{i}$ is almost surely $\mathbf{L}$-smooth and satisfies:

$$
\begin{equation*}
\left\|\nabla f_{i}(\boldsymbol{x} ; \zeta)-\nabla f_{i}(\boldsymbol{y} ; \zeta)\right\| \leq L\|\boldsymbol{x}-\boldsymbol{y}\|, \text { for any } \boldsymbol{x}, \boldsymbol{y} \tag{9}
\end{equation*}
$$

The assumption (A2*) also implies the following quadratic upper bound on $f_{i}$

$$
\begin{equation*}
f_{i}(\boldsymbol{y}) \leq f_{i}(\boldsymbol{x})+\left\langle\nabla f_{i}(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\right\rangle+\frac{L}{2}\|\boldsymbol{y}-\boldsymbol{x}\|^{2} \tag{10}
\end{equation*}
$$

Further, if $f_{i}$ is twice-differentiable, (A2*) implies that $\left\|\nabla^{2} f_{i}(\boldsymbol{x} ; \zeta)\right\| \leq \beta$ for any $\boldsymbol{x}$.
(A3) We assume that the intra-client gradient variance is bounded by $\sigma^{2}$. For any client $i$, the following holds almost surely at any fixed $\boldsymbol{x}$ :

$$
\mathbb{E}_{\zeta_{i}}\left[\nabla f_{i}(\boldsymbol{x} ; \zeta)\right]=\nabla f_{i}(\boldsymbol{x}), \quad \text { and } \quad \mathbb{E}_{\zeta_{i}}\left\|\nabla f_{i}(\boldsymbol{x} ; \zeta)-\nabla f_{i}(\boldsymbol{x})\right\|^{2} \leq \sigma^{2}
$$

Note that we expect the intra-client variance to be smaller than inter-client variance and so typically $\sigma^{2} \leq G^{2}$.
(A4) $f$ satisfies the $\mu$-PL inequality [29] for $\mu>0$ if:

$$
\|\nabla f(\boldsymbol{x})\|^{2} \geq 2 \mu\left(f(\boldsymbol{x})-f^{\star}\right)
$$

Note that PL-inequality is much weaker than the standard notion of strong-convexity, and in fact is even satisfied by some non-convex functions [29].

## F. 2 Some technical lemmas

Now we cover some technical lemmas which are useful for computations later on. First, we state a relaxed triangle inequality true for the squared $\ell_{2}$ norm.
Lemma 1 (relaxed triangle inequality). Let $\left\{\boldsymbol{v}_{1}, \ldots, \boldsymbol{v}_{\tau}\right\}$ be $\tau$ vectors in $\mathbb{R}^{d}$. Then the following are true:

1. $\left\|\boldsymbol{v}_{i}+\boldsymbol{v}_{j}\right\|^{2} \leq(1+c)\left\|\boldsymbol{v}_{i}\right\|^{2}+\left(1+\frac{1}{c}\right)\left\|\boldsymbol{v}_{j}\right\|^{2}$ for any $c>0$, and
2. $\left\|\sum_{i=1}^{\tau} \boldsymbol{v}_{i}\right\|^{2} \leq \tau \sum_{i=1}^{\tau}\left\|\boldsymbol{v}_{i}\right\|^{2}$.

Proof. The proof of the first statement for any $c>0$ follows from the identity:

$$
\left\|\boldsymbol{v}_{i}+\boldsymbol{v}_{j}\right\|^{2}=(1+c)\left\|\boldsymbol{v}_{i}\right\|^{2}+\left(1+\frac{1}{c}\right)\left\|\boldsymbol{v}_{j}\right\|^{2}-\left\|\sqrt{c} \boldsymbol{v}_{i}+\frac{1}{\sqrt{c}} \boldsymbol{v}_{j}\right\|^{2} .
$$

For the second inequality, we use the convexity of $\boldsymbol{x} \rightarrow\|\boldsymbol{x}\|^{2}$ and Jensen's inequality

$$
\left\|\frac{1}{\tau} \sum_{i=1}^{\tau} \boldsymbol{v}_{i}\right\|^{2} \leq \frac{1}{\tau} \sum_{i=1}^{\tau}\left\|\boldsymbol{v}_{i}\right\|^{2}
$$

Next we state an elementary lemma about expectations of norms of random vectors.
Lemma 2 (separating mean and variance). Let $\left\{\Xi_{1}, \ldots, \Xi_{\tau}\right\}$ be $\tau$ random variables in $\mathbb{R}^{d}$ which are not necessarily independent. First suppose that their mean is $\mathbb{E}\left[\Xi_{i}\right]=\xi_{i}$ and variance is bounded as $\mathbb{E}\left[\left\|\Xi_{i}-\xi_{i}\right\|^{2}\right] \leq \sigma^{2}$. Then, the following holds

$$
\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}\right\|^{2}\right] \leq\left\|\sum_{i=1}^{\tau} \xi_{i}\right\|^{2}+\tau^{2} \sigma^{2}
$$

Now instead suppose that their conditional mean is $\mathbb{E}\left[\Xi_{i} \mid \Xi_{i-1}, \ldots \Xi_{1}\right]=\xi_{i}$ i.e. the variables $\left\{\Xi_{i}-\right.$ $\left.\xi_{i}\right\}$ form a martingale difference sequence, and the variance is bounded by $\mathbb{E}\left[\left\|\Xi_{i}-\xi_{i}\right\|^{2}\right] \leq \sigma^{2}$ as before. Then we can show the tighter bound

$$
\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}\right\|^{2}\right] \leq 2\left\|\sum_{i=1}^{\tau} \xi_{i}\right\|^{2}+2 \tau \sigma^{2}
$$

Proof. For any random variable $X, \mathbb{E}\left[X^{2}\right]=\mathbb{E}\left[(X-\mathbb{E}[X])^{2}\right]+(\mathbb{E}[X])^{2}$ implying

$$
\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}\right\|^{2}\right]=\left\|\sum_{i=1}^{\tau} \xi_{i}\right\|^{2}+\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}-\xi_{i}\right\|^{2}\right]
$$

Expanding the above expression using relaxed triangle inequality (Lemma 1) proves the first claim:

$$
\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}-\xi_{i}\right\|^{2}\right] \leq \tau \sum_{i=1}^{\tau} \mathbb{E}\left[\left\|\Xi_{i}-\xi_{i}\right\|^{2}\right] \leq \tau^{2} \sigma^{2}
$$

For the second statement, $\xi_{i}$ is not deterministic and depends on $\Xi_{i-1}, \ldots, \Xi_{1}$. Hence we have to resort to the cruder relaxed triangle inequality to claim

$$
\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}\right\|^{2}\right] \leq 2\left\|\sum_{i=1}^{\tau} \xi_{i}\right\|^{2}+2 \mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}-\xi_{i}\right\|^{2}\right]
$$

and then use the tighter expansion of the second term:

$$
\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_{i}-\xi_{i}\right\|^{2}\right]=\sum_{i, j} \mathbb{E}\left[\left(\Xi_{i}-\xi_{i}\right)^{\top}\left(\Xi_{j}-\xi_{j}\right)\right]=\sum_{i} \mathbb{E}\left[\left\|\Xi_{i}-\xi_{i}\right\|^{2}\right] \leq \tau \sigma^{2}
$$

The cross terms in the above expression have zero mean since $\left\{\Xi_{i}-\xi_{i}\right\}$ form a martingale difference sequence.

## F. 3 Properties of functions with bounded Hessian dissimilarity

We now study two lemmas which hold for any functions which satisfy (A2) and (A3). The first is closely related to the notion of smoothness (A2*).
Lemma 3 (similarity). The following holds for any two functions $f_{i}(\cdot)$ and $f(\cdot)$ satisfying (A2) and (A3), and any $\boldsymbol{x}, \boldsymbol{y}$ :

$$
\left\|\nabla f_{i}(\boldsymbol{y} ; \zeta)-\nabla f_{i}(\boldsymbol{x} ; \zeta)+\nabla f(\boldsymbol{x})-\nabla f(\boldsymbol{y})\right\|^{2} \leq \delta^{2}\|\boldsymbol{y}-\boldsymbol{x}\|^{2}
$$

Proof. Consider the function $\Psi(\boldsymbol{z}):=f_{i}(\boldsymbol{z} ; \zeta)-f(\boldsymbol{z})$. By the assumption (A2), we know that $\left\|\nabla^{2} \Psi(\boldsymbol{z})\right\| \leq \delta$ for all $\boldsymbol{z}$ i.e. $\Psi$ is $\delta$-smooth. By standard arguments based on taking limits [45], this implies that

$$
\|\nabla \Psi(\boldsymbol{y})-\nabla \Psi(\boldsymbol{x})\| \leq \delta\|\boldsymbol{y}-\boldsymbol{x}\|
$$

Plugging back the definition of $\Psi$ into the above inequality proves the lemma.
Next, we see how weakly-convex functions satisfy a weaker notion of "averaging does not hurt". This is used to get a handle on the effect of averaging of parameters in FedAvg.
Lemma 4 (averaging). Suppose $f$ is $\delta$-weakly convex. Then, for any $\gamma \geq \delta$, and a sequence of parameters $\left\{\boldsymbol{y}_{i}\right\}_{i \in \mathcal{S}}$ and $\boldsymbol{x}$ :

$$
\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} f\left(\boldsymbol{y}_{i}\right)+\frac{\gamma}{2}\left\|\boldsymbol{x}-\boldsymbol{y}_{i}\right\|^{2} \geq f(\overline{\boldsymbol{y}})+\frac{\gamma}{2}\|\boldsymbol{x}-\overline{\boldsymbol{y}}\|^{2}, \text { where } \overline{\boldsymbol{y}}:=\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \boldsymbol{y}_{i}
$$

Proof. Since $f$ is $\delta$-weakly convex, $\Phi(\boldsymbol{z}):=f(\boldsymbol{z})+\frac{\gamma}{2}\|\boldsymbol{z}-\boldsymbol{x}\|^{2}$ is convex. This proves the claim since $\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \Phi\left(\boldsymbol{y}_{i}\right) \geq \Phi(\overline{\boldsymbol{y}})$ by convexity.

## G Convergence with a generic base optimizer

Let us rewrite the Mime and MimeLite updates using notation convenient for analysis. In each round $t$, we sample clients $\mathcal{S}^{t}$ such that $\left|\mathcal{S}^{t}\right|=S$. The server communicates the server parameters $\boldsymbol{x}^{t-1}$ as well as the average gradient across the sampled clients $\boldsymbol{c}^{t}$ defined as

$$
\begin{equation*}
\boldsymbol{c}^{t}=\frac{1}{S} \sum_{i \in \mathcal{S}^{t}} \nabla f_{i}\left(\boldsymbol{x}^{t-1}\right) \tag{11}
\end{equation*}
$$

Note that computing $\boldsymbol{c}^{t}$ (required only by Mime but not by MimeLite) itself requires additional communication. In this proof, we do not make any assumption on how $\boldsymbol{c}^{t}$ is computed as long as it is unbiased and is computed over $S$ clients. In particular, it can either be computed on the sampled $\mathcal{S}^{t}$ or a different set of an independent sampled clients $\tilde{\mathcal{S}}^{t}$.
Then each client $i \in \mathcal{S}^{t}$ makes a copy $\boldsymbol{y}_{i, 0}^{t}=\boldsymbol{x}^{t-1}$ and perform $K$ local client updates. In each local client update $k \in[K]$, the client samples a dataset $\zeta_{i, k}^{t}$ and

$$
\begin{array}{rlr}
\boldsymbol{y}_{i, k}^{t} & =\boldsymbol{y}_{i, k-1}^{t}-\eta \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x}^{t-1} ; \zeta_{i, k}^{t}\right)+\boldsymbol{c}^{t} ; \boldsymbol{s}^{t-1}\right) \quad \text { (Mime client update) } \\
& =\boldsymbol{y}_{i, k-1}^{t}-\eta \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right) ; \boldsymbol{s}^{t-1}\right) . & \text { (MimeLite client update) }
\end{array}
$$

After $K$ such local updates, the server then aggregates the new client parameters as

$$
\begin{aligned}
\boldsymbol{x}^{t} & =\frac{1}{S} \sum_{i \in \mathcal{S}^{t}} \boldsymbol{y}_{i, K}^{t} & & \text { (Update server parameters) } \\
\boldsymbol{s}^{t} & =\mathcal{V}\left(\boldsymbol{c}^{t}, \boldsymbol{s}^{t-1}\right) . & & \text { (Update server statistics) }
\end{aligned}
$$

## G. 1 Proof of Theorem I (generic reduction)

Computing server update.
Lemma 5 (Deviation from central update.). For a linear updater $\mathcal{U}$ the server update for Mime can be written as

$$
\boldsymbol{x}^{t}=\boldsymbol{x}^{t-1}-\tilde{\eta} \mathcal{U}\left(\frac{1}{S} \sum_{i} \nabla f_{i}(\boldsymbol{x})+\boldsymbol{e}^{t} ; \boldsymbol{s}^{t-1}\right)
$$

and for MimeLite is becomes

$$
\boldsymbol{x}^{t}=\boldsymbol{x}^{t-1}-\tilde{\eta} \mathcal{U}\left(\frac{1}{K S} \sum_{i, k} \nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}\right)+\boldsymbol{e}^{t} ; \boldsymbol{s}^{t-1}\right)
$$

for $\tilde{\eta}:=K \eta$. The error is defined as $\boldsymbol{e}^{t}=\frac{1}{K S} \sum_{i, k}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1} ; \zeta_{i, k}\right)-\nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}\right)\right)$
Proof. Because the updater $\mathcal{U}$ is linear in its first parameter, we can rewrite the update to the server for MimeLite as

$$
\begin{aligned}
\boldsymbol{x}^{t}-\boldsymbol{x}^{t-1} & =\frac{1}{S} \sum_{i \in \mathcal{S}^{t}} \sum_{k=1}^{K}-\eta \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right) ; \boldsymbol{s}^{t-1}\right) \\
& =\eta K \mathcal{U}\left(\frac{1}{K S} \sum_{i, k} \nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right) ; \boldsymbol{s}^{t-1}\right)
\end{aligned}
$$

We drop the dependence on $t$ when obvious from context and $i$ by default sums over $\mathcal{S}^{t}$ and $k$ over $[K]$ by default. Using our definition of $\boldsymbol{e}^{t}$ we have

$$
\begin{aligned}
\boldsymbol{x}^{t}-\boldsymbol{x}^{t-1} & =\eta K \mathcal{U}\left(\frac{1}{K S} \sum_{i, k} \nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right) ; \boldsymbol{s}^{t-1}\right) \\
& =\tilde{\eta} \mathcal{U}\left(\frac{1}{K S} \sum_{i, k} \nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}\right)+\boldsymbol{e}^{t} ; \boldsymbol{s}^{t-1}\right) .
\end{aligned}
$$

Now let us examine the update of Mime. Again assuming $K$ is a multiple of epoch, we have $\sum_{i, k} \nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}^{t}\right)=K \sum_{i} \nabla f_{i}(\boldsymbol{x})=K S \boldsymbol{x}$. Hence,

$$
\begin{aligned}
\boldsymbol{x}^{t}-\boldsymbol{x}^{t-1} & =\frac{1}{S} \sum_{i \in \mathcal{S}^{t}} \sum_{k=1}^{K}-\eta \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}^{t}\right)+\boldsymbol{c} ; \boldsymbol{s}^{t-1}\right) \\
& =\eta K \mathcal{U}\left(\boldsymbol{c}+\boldsymbol{e}^{t} ; \boldsymbol{s}^{t-1}\right) \\
& =\eta K \mathcal{U}\left(\frac{1}{S} \sum_{i} \nabla f_{i}(\boldsymbol{x})+\boldsymbol{e}^{t} ; \boldsymbol{s}^{t-1}\right) .
\end{aligned}
$$

Thus we showed the lemma for both Mime and MimeLite.
Lemma 6 (Defining error). For $\boldsymbol{e}^{t}$ defined in Lemma 5, assuming all functions $f_{i}(\cdot, \zeta)$ are Lsmooth, we have

$$
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \leq L^{2} \mathcal{E}_{K}^{t}, \text { where } \mathcal{E}_{K}^{t}:=\frac{1}{K S} \sum_{i, k} \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2}
$$

Proof. Using the smoothness of the individual functions and the definition of $\boldsymbol{e}^{t}$,

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} & =\mathbb{E}\left\|\frac{1}{K S} \sum_{i, k}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1} ; \zeta_{i, k}\right)-\nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}\right)\right)\right\|^{2} \\
& \leq \frac{1}{K S} \sum_{i, k} \mathbb{E}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}\right)-\nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}\right)\right\|^{2} \leq L^{2} \mathcal{E}_{K}^{t}
\end{aligned}
$$

Henceforth, we will call $\mathcal{E}_{K}^{t}$ as the error, or as the client-drift following [32].
Bounding error in MimeLite. Now we will try bound the client drift $\mathcal{E}^{t}$ for MimeLite.
Lemma 7 (MimeLite error). Suppose that all functions $f_{i}(\cdot, \zeta)$ are $L$-smooth (A2*), $\sigma^{2}$ variance (A3), and (A1) is satisfied, and the updater $\mathcal{U}$ has B-Lipschitz updates. Then using step-size $\tilde{\eta} \leq$ $\frac{1}{2 B L}$,

$$
\frac{1}{18 B^{2} \tilde{\eta}^{2}} \mathcal{E}^{K} \leq \mathbb{E}\|\nabla f(\boldsymbol{x})\|^{2}+G^{2}+\frac{\sigma^{2}}{2 K}
$$

Proof. For $K=1$, we have $\mathbb{E}\left\|\boldsymbol{y}_{i, 1}-\boldsymbol{x}\right\|^{2} \leq B^{2} \eta^{2}\left(G^{2}+\sigma^{2}\right)+B^{2} \eta^{2} \mathbb{E}\|\nabla f(\boldsymbol{x})\|^{2}$. The lemma is easily shown to be true. Assuming $K \geq 2$ henceforth, and starting from the client update of MimeLite we have

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{y}_{i, k}-\boldsymbol{x}\right\|^{2} & =\mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\eta \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right) ; \boldsymbol{s}^{t-1}\right)-\boldsymbol{x}\right\|^{2} \\
& \leq \mathbb{E} \| \boldsymbol{y}_{i, k-1}-\eta \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \boldsymbol{s}^{t-1}\right)-\boldsymbol{x} \|^{2}+B^{2} \eta^{2} \sigma^{2}\right. \\
& \leq\left(1+\frac{1}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2}+K \eta^{2} \mathbb{E} \| \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \boldsymbol{s}^{t-1}\right) \|^{2}+B^{2} \eta^{2} \sigma^{2}\right. \\
& \leq\left(1+\frac{1}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2}+K B^{2} \eta^{2} \mathbb{E}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}\right) \pm \nabla f_{i}(\boldsymbol{x})\right\|^{2}+B^{2} \eta^{2} \sigma^{2} \\
& \leq\left(1+\frac{1}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2} \\
& \leq\left(1+\frac{2}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2}+2 K B^{2} \eta^{2} \mathbb{E}\|\nabla f(\boldsymbol{x})\|^{2}+2 K B^{2} \eta^{2} G^{2}+B^{2} \eta^{2} \sigma^{2}
\end{aligned}
$$

Here, we used the condition on our step size that $\tilde{\eta}=K \eta \leq \frac{1}{2 L B}$, which implies that $2 K B^{2} L^{2} \eta^{2} \leq$ $\frac{1}{K-1}$. Unrolling this recursion, we have

$$
\mathbb{E}\left\|\boldsymbol{y}_{i, k}-\boldsymbol{x}\right\|^{2} \leq\left(2 K B^{2} \eta^{2} \mathbb{E}\|\nabla f(\boldsymbol{x})\|^{2}+2 K B^{2} \eta^{2} G^{2}+B^{2} \eta^{2} \sigma^{2}\right) \sum_{k=1}^{K}\left(1+\frac{2}{K-1}\right)^{k}
$$

Note that $\left(1+\frac{2}{K-1}\right)^{k} \leq 9$. Averaging then over $k$ and $i$, we get

$$
\mathcal{E}_{K}^{t} \leq 18 K^{2} B^{2} \eta^{2} \mathbb{E}\|\nabla f(\boldsymbol{x})\|^{2}+18 K^{2} B^{2} \eta^{2} G^{2}+9 K B^{2} \eta^{2} \sigma^{2} .
$$

Finally, recalling that $\tilde{\eta}=K \eta$ finishes the lemma.

Bounding error in Mime. Next we will try bound the client drift $\mathcal{E}^{t}$ for Mime. The additional SVRG correction term used in Mime improves the bound on the error.
Lemma 8 (Mime Error). Suppose that all functions $f_{i}(\cdot, \zeta)$ are $L$-smooth (A2*), $\sigma^{2}$ variance (A3), and (A1) is satisfied, and the updater $\mathcal{U}$ has B-Lipschitz, updates. Then using step-size $\tilde{\eta} \leq \frac{1}{2 B L}$,

$$
\mathcal{E}^{K} \leq 18 B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\frac{1}{S} \sum_{i} \nabla_{i} f(\boldsymbol{x})\right\|^{2}
$$

Proof. For $K=1$, the Mime update loos like

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{y}_{i, 1}-\boldsymbol{x}\right\|^{2} & =\eta^{2} \mathbb{E}\left\|\mathcal{U}\left(\boldsymbol{c} ; \boldsymbol{s}^{t-1}\right)\right\|^{2} \\
& \leq \eta^{2} B^{2} \mathbb{E}\|\boldsymbol{c}\|^{2}
\end{aligned}
$$

Assuming $K \geq 2$ henceforth, and starting from the client update of Mime we have

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{y}_{i, k}-\boldsymbol{x}\right\|^{2} & =\mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\eta \mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}^{t}\right)+\boldsymbol{c}^{t} ; \boldsymbol{s}^{t-1}\right)-\boldsymbol{x}\right\|^{2} \\
\leq & \left(1+\frac{1}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2} \\
& +K \eta^{2} \mathbb{E}\left\|\mathcal{U}\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}^{t}\right)+\boldsymbol{c}^{t} ; \boldsymbol{s}^{t-1}\right)\right\|^{2} \\
\leq & \left(1+\frac{1}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2}+K \eta^{2} B^{2} \mathbb{E}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i, k-1} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}^{t}\right)+\boldsymbol{c}^{t}\right\|^{2} \\
\leq & \left(1+\frac{1}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2} \\
\leq & \left(1+\frac{1}{K-1}+2 K \eta^{2} B^{2} L^{2}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2}+2 K \eta^{2} B^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2} \\
\leq & \left(1+\frac{2}{K-1}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}-\boldsymbol{x}\right\|^{2}+2 K \eta^{2} B^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2}
\end{aligned}
$$

Here, we used the condition on our step size that $\tilde{\eta}=K \eta \leq \frac{1}{2 L B}$, which implies that $2 K B^{2} L^{2} \eta^{2} \leq$ $\frac{1}{K-1}$. Unrolling this recursion, we have

$$
\mathbb{E}\left\|\boldsymbol{y}_{i, k}-\boldsymbol{x}\right\|^{2} \leq 2 K B^{2} \eta^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2} \sum_{k=1}^{K}\left(1+\frac{2}{K-1}\right)^{k} \leq 18 K^{2} B^{2} \eta^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2}
$$

Note that $\left(1+\frac{2}{K-1}\right)^{k} \leq 9$. Averaging then over $k$ and $i$, recalling that $\tilde{\eta}=K \eta$ get

$$
\mathcal{E}_{K}^{t} \leq 18 B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2}
$$

## Putting it together (Theorem I).

Lemma 9. The updates of Mime and MimeLite for $\boldsymbol{g}^{t}$ satisfying $\mathbb{E}\left[\boldsymbol{g}^{t}\right]=\nabla f\left(\boldsymbol{x}^{t-1}\right)$, and we have for $\tilde{\eta} \leq \frac{1}{2 B L}$

$$
\begin{aligned}
\boldsymbol{x}^{t} & =\boldsymbol{x}^{t-1}-\tilde{\eta} \mathcal{U}\left(\boldsymbol{c}^{t}+\boldsymbol{e}^{t} ; \boldsymbol{s}^{t-1}\right) \\
\boldsymbol{s}^{t} & =\mathcal{V}\left(\boldsymbol{c}^{t} ; \boldsymbol{s}^{t-1}\right)
\end{aligned}
$$

Where, we have

$$
\frac{1}{18 B^{2} L^{2} \tilde{\eta}^{2}} \mathbb{E}_{t}\left\|\boldsymbol{e}_{t}\right\|^{2} \leq \begin{cases}\mathbb{E}\left\|\boldsymbol{c}_{t}\right\|^{2} & \text { Mime } \\ \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t}\right)\right\|^{2}+G^{2}+\frac{\sigma^{2}}{2 K} & \text { MimELite } .\end{cases}
$$

Proof. Now, combining Lemmas 5, 6, shows that running Mime or MimeLite is equivalent to

$$
\begin{aligned}
\boldsymbol{x}^{t} & =\boldsymbol{x}^{t-1}-\tilde{\eta} \mathcal{U}\left(\boldsymbol{g}^{t}+\boldsymbol{e}^{t} ; \boldsymbol{s}^{t-1}\right) \\
\boldsymbol{s}^{t} & =\mathcal{V}\left(\boldsymbol{g}^{t} ; \boldsymbol{s}^{t-1}\right)
\end{aligned}
$$

where for Mime we use

$$
\boldsymbol{g}_{\text {Mime }}^{t}=\frac{1}{S} \sum_{i} \nabla f_{i}(\boldsymbol{x}) \text { with } \mathbb{E}\left[\boldsymbol{g}_{\text {Mime }}^{t}\right]=\nabla f\left(\boldsymbol{x}^{t-1}\right) \text { and } \mathbb{E}\left\|\boldsymbol{g}_{\text {Mime }}^{t}-\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq \frac{G^{2}}{S}
$$

and for MimeLite we use
$\boldsymbol{g}_{\text {MimeLite }}^{t}=\frac{1}{K S} \sum_{i, k} \nabla f_{i}\left(\boldsymbol{x} ; \zeta_{i, k}\right)$ with $\mathbb{E}\left[\boldsymbol{g}_{\text {Mime }}^{t}\right]=\nabla f\left(\boldsymbol{x}^{t-1}\right)$ and $\mathbb{E}\left\|\boldsymbol{g}_{\text {Mime }}^{t}-\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq \frac{G^{2}}{S}+\frac{\sigma^{2}}{K S}$.
This shows the first part of the theorem. For the second part of the theorem, using the bound from Lemma 8 for Mime,

$$
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\| \leq L^{2} \mathcal{E}_{K}^{t} \leq 18 L^{2} B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2}
$$

For MimeLite, we will instead use the bound from Lemma 7,

$$
\mathbb{E}\left\|\boldsymbol{e}_{\text {MimeLite }}^{t}\right\| \leq L^{2} \mathcal{E}_{K}^{t}+\frac{\sigma^{2}}{K S} \leq 18 L^{2} B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t}\right)\right\|^{2}+18 L^{2} B^{2} \tilde{\eta}^{2} G^{2}+\frac{9 L^{2} B^{2} \tilde{\eta}^{2} \sigma^{2}}{K}+\frac{\sigma^{2}}{K S}
$$

Note that the Lemma we proved here is slightly stronger than the theorem in the main section (up to constants which were suppressed).

## G. 2 Convergence of MimeSGD and MimeLiteSGD (Corollary II)

Theorem I shows that Mime and MimeLite mimic a centralized algorithm quite closely up to error $\mathcal{O}\left(\tilde{\eta}^{2}\right)$. Then, analyzing the sensitivity of the base algorithm to such perturbation yields specific rates of convergence. We perform such an analysis using SGD as our base optimizer.

Properties of SGD as the base optimizer:

- $s^{t}$ is empty i.e. there are no global statistics used.
- $\mathcal{U}\left(\boldsymbol{g} ; \boldsymbol{s}^{t-1}\right)=\boldsymbol{g}$ for any $\boldsymbol{g}$ and $B=1$.

With this in mind, we proceed.
Lemma 10 (Progress in one round). Given that $f$ is $L$-smooth, and for any step-size $\tilde{\eta} \leq \frac{1}{2(B+2) L}$ for $B \geq 1$ we have

$$
f\left(\boldsymbol{x}^{t}\right) \leq f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{4} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\tilde{\eta} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S}
$$

Proof. Starting from the update equation and the smoothness of $f$, we have

$$
\begin{aligned}
\mathbb{E} f\left(\boldsymbol{x}^{t}\right) & \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)+\mathbb{E}\left\langle\nabla f\left(\boldsymbol{x}^{t-1}\right), \boldsymbol{x}^{t}-\boldsymbol{x}^{t-1}\right\rangle+\frac{L}{2} \mathbb{E}\left\|\boldsymbol{x}^{t}-\boldsymbol{x}^{t-1}\right\|^{2} \\
& =\mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\tilde{\eta} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\tilde{\eta}\left\langle\nabla f\left(\boldsymbol{x}^{t-1}\right), \boldsymbol{e}^{t}\right\rangle+\frac{L \tilde{\eta}^{2}}{2} \mathbb{E}\left\|\boldsymbol{c}^{t}+\boldsymbol{e}^{t}\right\|^{2} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta}}{2}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{2 L \tilde{\eta}^{2}}{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2}+\frac{2 L \tilde{\eta}^{2}}{2} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\left(\frac{\tilde{\eta}}{2}-\frac{2 L \tilde{\eta}^{2}}{2}\right) \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\left(L \tilde{\eta}^{2}+\frac{\tilde{\eta}}{2}\right) \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{2 L \tilde{\eta}^{2} G^{2}}{2 S} .
\end{aligned}
$$

Using the bound on the step size that $\tilde{\eta} \leq \frac{1}{4 L}$ yields the lemma.

One round progress for MimeSGD. Next, we specialize the convergence rate for Mime.
Lemma 11. Suppose $f$ is a L-smooth function satisfying PL-inequality for $\mu \geq 0$ ( $\mu=0$ corresponds to the general case). Running MimeSGD for $\tilde{\eta} \leq \frac{1}{12 B L}$ satisfies

$$
\frac{\tilde{\eta}}{16} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq\left(1-\frac{\mu \tilde{\eta}}{8}\right)\left(f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}\right)-\left(f\left(\boldsymbol{x}^{t}\right)-f^{\star}\right)+\frac{3 L \tilde{\eta}^{2} G^{2}}{S} .
$$

Proof. Recall from Lemma 9 that for Mime,

$$
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \leq 18 L^{2} B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2} \leq 18 L^{2} B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{18 L^{2} B^{2} \tilde{\eta}^{2} G^{2}}{S}
$$

Combining this with Lemma 10 yields the following progress for Mime

$$
\begin{aligned}
f\left(\boldsymbol{x}^{t}\right) & \leq f\left(\boldsymbol{x}^{t-1}\right)-\left(\frac{\tilde{\eta}}{4}-18 L^{2} B^{2} \tilde{\eta}^{3}\right) \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\left(L \tilde{\eta}^{2}+18 L^{2} B^{2} \tilde{\eta}^{3}\right) G^{2}}{S} \\
& \leq f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{8} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{3 L \tilde{\eta}^{2} G^{2}}{S}
\end{aligned}
$$

Here, we used the bound on the step size that $\tilde{\eta} \leq \frac{1}{12 L B}$ implies $18 L^{2} B^{2} \tilde{\eta}^{2} \leq \frac{1}{8}$. Now using PL-inequality, we can write

$$
f\left(\boldsymbol{x}^{t}\right)-f^{\star} \leq f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}-\frac{\mu \tilde{\eta}}{8}\left(f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}\right)-\frac{\tilde{\eta}}{16} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{3 L \tilde{\eta}^{2} G^{2}}{S}
$$

This yields the lemma.

We are now ready to derive the convergence rate.

Convergence rate of MimeSGD on general non-convex functions. Set $\mu=0$ in Lemma 11 and sum over $t$

$$
\begin{aligned}
\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} & \leq \frac{16\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{\tilde{\eta} T}+\frac{48 L \tilde{\eta} G^{2}}{S} \\
& \leq 16 \sqrt{\frac{3 L G^{2}\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{S T}}+\frac{192 B L\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{T}
\end{aligned}
$$

The final step used a step-size of $\tilde{\eta}=\min \left(\frac{1}{12 B L}, \frac{1}{4 L}, \sqrt{\frac{S\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{3 L T G^{2}}}\right)$. Here, we used $\boldsymbol{x}^{\text {out }}=\boldsymbol{x}^{\tau}$ where $\tau$ is uniformly at random chosen in $[T]$.

Convergence rate of MimeSGD on PL-inequality. Multiply Lemma 11 by $\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t}$ and sum over $t$

$$
\begin{aligned}
\sum_{t=1}^{T}\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq & \sum_{t=1}^{T}\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-(t-1)} \frac{16\left(f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}\right)}{\tilde{\eta}} \\
& -\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t} \frac{16\left(f\left(\boldsymbol{x}^{t}\right)-f^{\star}\right)}{\tilde{\eta}}+\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t} \frac{48 L \tilde{\eta} G^{2}}{S} \\
\leq & \left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T} \frac{16\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{\tilde{\eta}}+\sum_{t=1}^{T}\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t} \frac{48 L \tilde{\eta} G^{2}}{S}
\end{aligned}
$$

Output $\boldsymbol{x}^{\text {out }}=\boldsymbol{x}^{\tau}$ where $\tau$ is chosen with probability proportional to $\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t}$. Then, this yields

$$
\mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{\text {out }}\right)\right\|^{2} \leq\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T} \frac{16\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{\tilde{\eta}}+\frac{48 L \tilde{\eta} G^{2}}{S} \leq \tilde{\mathcal{O}}\left(\frac{\sigma^{2}}{\mu T}+L\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right) \exp \left(-\frac{\mu T}{12 B L}\right)\right)
$$

Using an appropriate step-size $\tilde{\eta}$ yields the final rate (see Lemma 1 of [32]).
One round progress for MimeLiteSGD. Next, we specialize the convergence rate for MimeLite.
Lemma 12. Suppose $f$ is a L-smooth function satisfying PL-inequality for $\mu \geq 0$ ( $\mu=0$ corresponds to the general case). Running MimeLiteSGD for $\tilde{\eta} \leq \frac{1}{12 B L}$ satisfies

$$
\frac{\tilde{\eta}}{16} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq\left(1-\frac{\mu \tilde{\eta}}{8}\right)\left(f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}\right)-\left(f\left(\boldsymbol{x}^{t}\right)-f^{\star}\right)+\frac{L \tilde{\eta}^{2} G^{2}}{S}+18 L^{2} B^{2} \tilde{\eta}^{3}\left(G^{2}+\sigma^{2} / K\right)
$$

Proof. Recall from Lemma 9 that,

$$
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \leq 18 L^{2} B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2} \leq 18 L^{2} B^{2} \tilde{\eta}^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+18 L^{2} B^{2} \tilde{\eta}^{2} G^{2}+\frac{9 L^{2} B^{2} \tilde{\eta}^{2} \sigma^{2}}{K}
$$

Combining this with Lemma 10 yields the following progress for Mime

$$
\begin{aligned}
f\left(\boldsymbol{x}^{t}\right) & \leq f\left(\boldsymbol{x}^{t-1}\right)-\left(\frac{\tilde{\eta}}{4}-18 L^{2} B^{2} \tilde{\eta}^{3}\right) \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S}+18 L^{2} B^{2} \tilde{\eta}^{3}\left(G^{2}+\sigma^{2} / K\right) \\
& \leq f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{8} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S}+18 L^{2} B^{2} \tilde{\eta}^{3}\left(G^{2}+\sigma^{2} / K\right)
\end{aligned}
$$

Here, we used the bound on the step size that $\tilde{\eta} \leq \frac{1}{12 L B}$ implies $18 L^{2} B^{2} \tilde{\eta}^{2} \leq \frac{1}{8}$. Now using PL-inequality, we can write

$$
\begin{aligned}
f\left(\boldsymbol{x}^{t}\right)-f^{\star}- & \left(f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}\right) \leq \\
& -\frac{\mu \tilde{\eta}}{8}\left(f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}\right)-\frac{\tilde{\eta}}{16} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S}+18 L^{2} B^{2} \tilde{\eta}^{3}\left(G^{2}+\sigma^{2} / K\right)
\end{aligned}
$$

This yields the lemma.
We are now ready to derive the convergence rate.
Convergence rate of MimeLiteSGD on general non-convex functions. Define $\tilde{G}^{2}=G^{2}+$ $\sigma^{2} / K$. Set $\mu=0$ in Lemma 12 and sum over $t$

$$
\begin{aligned}
& \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq \frac{16\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{\tilde{\eta} T}+\frac{16 L \tilde{\eta} G^{2}}{S}+288 L^{2} B^{2} \tilde{\eta}^{2} \tilde{G}^{2} \\
& \leq 16 \sqrt{\frac{L G^{2}\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{S T}}+84\left(\frac{L \tilde{G}\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{T}\right)^{2 / 3}+\frac{192 B L\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{T}
\end{aligned}
$$

The final step used an appropriate step-size of $\tilde{\eta}$, see Lemma 2 of [32]. Here, we used $\boldsymbol{x}^{\text {out }}=\boldsymbol{x}^{\tau}$ where $\tau$ is uniformly at random chosen in $[T]$. Finally note that if $K \geq \frac{\sigma^{2}}{G^{2}}$, then $\tilde{G}^{2} \leq 2 G^{2}$.

Convergence rate of MimeLiteSGD on PL-inequality. Multiply Lemma 12 by $\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t}$ and sum over $t$

$$
\begin{aligned}
& \sum_{t=1}^{T}\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq \sum_{t=1}^{T}\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-(t-1)} \frac{16\left(f\left(\boldsymbol{x}^{t-1}\right)-f^{\star}\right)}{\tilde{\eta}} \\
&-\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t} \frac{16\left(f\left(\boldsymbol{x}^{t}\right)-f^{\star}\right)}{\tilde{\eta}} \\
&+\sum_{t=1}^{T}\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t}\left(\frac{16 L \tilde{\eta} G^{2}}{S}+288 L^{2} B^{2} \tilde{\eta}^{2} \tilde{G}^{2}\right) \\
& \leq\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T} \frac{16\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{\tilde{\eta}} \\
&+\sum_{t=1}^{T}\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t}\left(\frac{16 L \tilde{\eta} G^{2}}{S}+288 L^{2} B^{2} \tilde{\eta}^{2} \tilde{G}^{2}\right)
\end{aligned}
$$

Output $\boldsymbol{x}^{\text {out }}=\boldsymbol{x}^{\tau}$ where $\tau$ is chosen with probability proportional to $\left(1-\frac{\mu \tilde{\eta}}{8}\right)^{T-t}$. Then, this yields with appropriate step-size $\tilde{\eta}$ yields the final rate (see Lemma 1 of [32]).

$$
\mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{\text {out }}\right)\right\|^{2} \leq \tilde{\mathcal{O}}\left(\frac{\sigma^{2}}{\mu T}+\frac{L^{2} \tilde{G}^{2}}{\mu^{2} T^{2}}+L\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right) \exp \left(-\frac{\mu T}{12 B L}\right)\right)
$$

## G. 3 Convergence of MimeAdam and MimeLiteAdam (Corollary III)

We will largely follow the convergence analysis of [75] for the analysis of Adam. A crucial difference between their setting and ours is that in our algorithm we use the global statistics (second order moment) corresponding to $t-1$ i.e. $\sqrt{\boldsymbol{v}^{t-1}}$ instead of $\sqrt{\boldsymbol{v}^{t}}$ where the $\sqrt{ }$ - operator is applied element wise. Practically, this does not make a significant difference since the discount (momentum) factor for the second momentum is very large. Theoretically however, this difference simplifies our proof significantly removing otherwise hard to handle stochastic dependencies.
In this section, we will use Adam as our base optimizer with $\varepsilon_{0}>0$ parameter for stability and $\beta_{1}=0$ (i.e. RMSProp). This is identical to the setting in the centralized algorithm analyzed by [75]. The properties of our base optimizer are then:

- $\boldsymbol{s}^{t}=\boldsymbol{v}^{t}$ which is a running average estimate of the second moment and satisfies $\boldsymbol{v}^{t}>0$.
- $\mathcal{U}\left(\boldsymbol{g} ; \boldsymbol{v}^{t-1}\right)=\frac{\boldsymbol{g}}{\sqrt{\boldsymbol{v}^{t-1}}+\varepsilon_{0}}$ for any $\boldsymbol{g}$. This update for any $\boldsymbol{v}^{t-1}$ is $B$-Lipschitz for $B=\frac{1}{\varepsilon_{0}}$.

In this sub-section, all operations on vectors (multiplication, division, addition, comparison) are applied element-wise with appropriate broad-casting.

## One round progress of Adam.

Lemma 13 (Effective step-sizes). Suppose that $\left|\nabla_{j} f_{i}(\boldsymbol{x})\right| \leq H$. Then Adam has effective step-sizes

$$
\frac{1}{H+\varepsilon_{0}} \boldsymbol{g} \leq \mathcal{U}\left(\boldsymbol{g} ; \boldsymbol{v}^{t-1}\right) \leq \frac{1}{\varepsilon_{0}} \boldsymbol{g} .
$$

Proof. Recall that $\boldsymbol{v}^{t}=\beta_{2} \boldsymbol{v}^{t-1}+\left(1-\beta_{2}\right)\left(\boldsymbol{c}^{t}\right)^{2}$ starting from $\boldsymbol{v}^{0}=0$. Thus for any $t \geq 0$, we have $\boldsymbol{v}^{t} \geq 0$ and hence $\sqrt{\boldsymbol{v}^{t-1}}+\varepsilon_{0} \geq \varepsilon_{0}$. For the other side, recall that $\boldsymbol{v}^{t}$ is updated with centralized stochastic gradients $\boldsymbol{c}^{t}=\frac{1}{S} \sum_{i} \nabla f_{i}(\boldsymbol{x})$.

$$
\left[\boldsymbol{c}^{t}\right]_{j}=\frac{1}{S} \sum_{i}\left[\nabla f_{i}(\boldsymbol{x})\right]_{j} \leq H
$$

Further,

$$
\left[\boldsymbol{v}^{t}\right]_{j}=\beta_{2}\left[\boldsymbol{v}^{t-1}\right]_{j}+\left(1-\beta_{2}\right)\left[\boldsymbol{c}^{t}\right]_{j}^{2} \leq \beta_{2}\left[\boldsymbol{v}^{t-1}\right]_{j}+\left(1-\beta_{2}\right) H^{2} \leq H^{2}
$$

Hence $\sqrt{\boldsymbol{v}^{t-1}+\varepsilon_{0}} \leq H+\varepsilon_{0}$.

Lemma 14 (One round progress). For one round of Adam with error $e^{t}$ in the update $\mathcal{U}$ and using $c^{t}$ for update $\mathcal{V}$, we have

$$
\mathbb{E} f\left(\boldsymbol{x}^{t}\right) \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta}\left(\left(H+\varepsilon_{0}\right)+\varepsilon_{0} /\left(H+\varepsilon_{0}\right)\right)}{2 \varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}}
$$

Proof. Starting from Lemma 13 and the smoothness of $f$, we have

$$
\begin{aligned}
\mathbb{E} f\left(\boldsymbol{x}^{t}\right) & \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\tilde{\eta} \mathbb{E}\left\langle\nabla f\left(\boldsymbol{x}^{t-1}\right), \mathbb{E}_{t}\left[\mathcal{U}\left(\boldsymbol{c}^{t}+\boldsymbol{e}^{t}\right)\right]\right\rangle+\frac{L \tilde{\eta}^{2}}{2} \mathbb{E}\left\|\mathcal{U}\left(\boldsymbol{c}^{t}+\boldsymbol{e}^{t} ; \boldsymbol{v}^{t-1}\right)\right\|^{2} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\tilde{\eta} \mathbb{E}\left\langle\nabla f\left(\boldsymbol{x}^{t-1}\right), \mathbb{E}_{t}\left[\frac{\boldsymbol{c}^{t}+\boldsymbol{e}^{t}}{\sqrt{\boldsymbol{v}^{t-1}}+\varepsilon_{0}}\right]\right\rangle+\frac{L \tilde{\eta}^{2}}{2} \mathbb{E}\left\|\mathcal{U}\left(\boldsymbol{c}^{t}+\boldsymbol{e}^{t} ; \boldsymbol{v}^{t-1}\right)\right\|^{2} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\tilde{\eta} \mathbb{E}\left\langle\nabla f\left(\boldsymbol{x}^{t-1}\right),\left[\frac{\nabla f\left(\boldsymbol{x}^{t-1}\right)+\boldsymbol{e}^{t}}{\sqrt{\boldsymbol{v}^{t-1}}+\varepsilon_{0}}\right]\right\rangle+\frac{L \tilde{\eta}^{2}}{2 \varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{c}^{t}+\boldsymbol{e}^{t}\right\|^{2} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{H+\varepsilon_{0}}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}-\tilde{\eta} \mathbb{E}\left\langle\nabla f\left(\boldsymbol{x}^{t-1}\right), \frac{\boldsymbol{e}^{t}}{\sqrt{\boldsymbol{v}^{t-1}+\varepsilon_{0}}}\right\rangle+\frac{L \tilde{\eta}^{2}}{2 \varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{c}^{t}+\boldsymbol{e}^{t}\right\|^{2} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{2\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta}\left(H+\varepsilon_{0}\right)}{2} \mathbb{E}\left\|\frac{\boldsymbol{e}^{t}}{\sqrt{\boldsymbol{v}^{t-1}}+\varepsilon_{0}}\right\|^{2}+\frac{L \tilde{\eta}^{2}}{2 \varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{c}^{t}+\boldsymbol{e}^{t}\right\|^{2} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\left(\frac{\tilde{\eta}}{2\left(H+\varepsilon_{0}\right)}-\frac{L \tilde{\eta}^{2}}{\varepsilon_{0}^{2}}\right)\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta}\left(H+\varepsilon_{0}\right)+2 L \tilde{\eta}^{2}}{2 \varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta}\left(\left(H+\varepsilon_{0}\right)+\varepsilon_{0} /\left(H+\varepsilon_{0}\right)\right)}{2 \varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}}
\end{aligned}
$$

Here we used our bound on the step-size that $\tilde{\eta} \leq \frac{\varepsilon_{0}}{4 L\left(H+\varepsilon_{0}\right)}$.

## Convergence of MimeAdam.

Lemma 15. Suppose that assumptions A1-(A3) hold and further $\left|\nabla_{j} f_{i}(\boldsymbol{x})\right| \leq H$. Then, running MimeAdam with step-size $\tilde{\eta} \leq \frac{\varepsilon_{0}^{2}}{12 L\left(H+\varepsilon_{0}\right)}$, we have

$$
\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq \frac{96 L\left(H+\varepsilon_{0}\right)^{2}\left(f\left(\boldsymbol{x}_{0}\right)-f^{\star}\right)}{\varepsilon_{0}^{2} T}+\frac{2 G^{2}}{S}
$$

Combining Lemma 14 with the bound on $\boldsymbol{e}^{t}$ from Lemma 9 we get,

$$
\begin{aligned}
& \mathbb{E} f\left(\boldsymbol{x}^{t}\right) \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta}\left(\left(H+\varepsilon_{0}\right)+\varepsilon_{0} /\left(H+\varepsilon_{0}\right)\right)}{2 \varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right) \\
&- \frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{9 L^{2} \tilde{\eta}^{3}\left(\left(H+\varepsilon_{0}\right)+\varepsilon_{0} /\left(H+\varepsilon_{0}\right)\right)}{\varepsilon_{0}^{4}} \mathbb{E}\left\|\boldsymbol{c}^{t}\right\|^{2} \\
&+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\left(\frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}-\frac{9 L^{2} \tilde{\eta}^{3}\left(\left(H+\varepsilon_{0}\right)+\varepsilon_{0} /\left(H+\varepsilon_{0}\right)\right)}{\varepsilon_{0}^{4}}\right)\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \\
&+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}}+\frac{9 L^{2} \tilde{\eta}^{3}\left(\left(H+\varepsilon_{0}\right)+\varepsilon_{0} /\left(H+\varepsilon_{0}\right)\right) G^{2}}{S \varepsilon_{0}^{4}} \\
& \leq+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}}+\frac{18 L^{2} \tilde{\eta}^{3}\left(H+\varepsilon_{0}\right) G^{2}}{S \varepsilon_{0}^{4}} \\
& \leq-\left(\frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}-\frac{18 L^{2} \tilde{\eta}^{3}\left(H+\varepsilon_{0}\right)}{\varepsilon_{0}^{4}}\right)\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \\
& \leq
\end{aligned}
$$

To simplify computations, here we assumed we assumed $\left(H+\varepsilon_{0}\right)^{2} \geq \varepsilon_{0}$ without loss of generality. If this is not true, we can replace $H$ by $\max \left(H, \sqrt{\varepsilon_{0}}-\varepsilon_{0}\right)$. Assuming $\tilde{\eta} \leq \frac{\varepsilon_{0}^{2}}{12 L\left(H+\varepsilon_{0}\right)}$, we have $\frac{18 L^{2} \tilde{\eta}^{2}\left(H+\varepsilon_{0}\right)}{\varepsilon_{0}^{4}} \leq \frac{1}{8\left(H+\varepsilon_{0}\right)}$. Rearranging the terms and substituting the bounds on the step-size yields the lemma.

## Convergence of MimeLiteAdam.

Lemma 16. Suppose that assumptions A1-(A3) hold and further $\left|\nabla_{j} f_{i}(\boldsymbol{x})\right| \leq H$. Then, running MimeLiteAdam with step-size $\tilde{\eta} \leq \frac{\varepsilon_{0}^{2}}{12 L \sqrt{S}\left(H+\varepsilon_{0}\right)}$, we have for $\tilde{G}^{2}:=G^{2}+\sigma^{2} / K$,

$$
\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \leq \frac{96 L \sqrt{S}\left(H+\varepsilon_{0}\right)^{2}\left(f\left(\boldsymbol{x}_{0}\right)-f^{\star}\right)}{\varepsilon_{0}^{2} T}+\frac{2 \tilde{G}^{2}}{S}
$$

Combining Lemma 14 with the bound on $\boldsymbol{e}^{t}$ from Lemma 9 we get for $\tilde{G}^{2}:=G^{2}+\sigma^{2} / K$,

$$
\begin{aligned}
\mathbb{E} f\left(\boldsymbol{x}^{t}\right) \leq & \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta}\left(H+\varepsilon_{0}\right)}{\varepsilon_{0}^{2}} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}} \\
\leq & \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{4\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2} \\
& +\frac{18 L^{2} \tilde{\eta}^{3}\left(H+\varepsilon_{0}\right)}{\varepsilon_{0}^{4}} \mathbb{E}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{18 L^{2} \tilde{\eta}^{3}\left(H+\varepsilon_{0}\right)\left(\tilde{G}^{2}\right)}{\varepsilon_{0}^{4}}+\frac{L \tilde{\eta}^{2} G^{2}}{S \varepsilon_{0}^{2}} \\
\leq & \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)-\frac{\tilde{\eta}}{8\left(H+\varepsilon_{0}\right)}\left\|\nabla f\left(\boldsymbol{x}^{t-1}\right)\right\|^{2}+\frac{\tilde{\eta} \tilde{G}^{2}}{4 S\left(H+\varepsilon_{0}\right)}
\end{aligned}
$$

Again as before to simplify computations, here we assumed $\left(H+\varepsilon_{0}\right)^{2} \geq \varepsilon_{0}$ without loss of generality. If this is not true, we can replace $H$ by $\max \left(H, \sqrt{\varepsilon_{0}}-\varepsilon_{0}\right)$. Assuming $\tilde{\eta} \leq \frac{\varepsilon_{0}^{2}}{12 L\left(H+\varepsilon_{0}\right) \sqrt{S}}$, we have $\frac{18 L^{2} \tilde{\eta}^{2}\left(H+\varepsilon_{0}\right)}{\varepsilon_{0}^{4}} \leq \frac{1}{8 S\left(H+\varepsilon_{0}\right)}$. Rearranging the terms and substituting the bounds on the step-size yields the lemma.

## H Circumventing server-only lower bounds

In this section we see how to use momentum based variance reduction $[14,62]$ to reduce the variance of the updates and improve convergence. It should be noted that MVR does not exactly fit the MIME framework (BASEOPT) since it requires computing gradients at two points on the same batch. However, it is straightforward to extend the idea of Mime to MVR as we will now do. We use MVR as a theoretical justification for why the usual momentum works well in practice. An interesting future direction would be to adapt the algorithm and analysis of [13], which does fit the framework of Mime.

For the sake of convenience, we summarize the notation used in the proof in a table.
Table 6: Summary of all notation used in the MVR proofs

| $\sigma^{2}, G^{2}$, and $\delta$ | intra-client gradient, inter-client gradient, and inter-client Hessian variance |
| :---: | :--- |
| $\eta, a$ | step-size, $(1-\beta)$ momentum parameters |
| $T, t$ | total number, index of communication rounds |
| $K, k$ | total number, index of client local update steps |
| $\mathcal{S}^{t}, S$, and $i$ | sampled set, size, and index of clients in round $t$ |
| $\boldsymbol{x}^{t}$ | aggregated server model after round $t$ |
| $\boldsymbol{m}^{t}$ | server momentum computed after round $t$ |
| $\boldsymbol{c}^{t}$ | control variate of server after round $t$ (only MiME) |
| $\boldsymbol{y}_{i, k}^{t}$ | model parameters of $i$ th client in round $t$ after step $k$ |
| $\zeta_{i, k}^{t}$ | mini-batch data used by $i$ th client in round $t$ and step $k$ |
| $\boldsymbol{d}_{i, k}^{t}$ | parameter update by $i$ th client in round $t$, step $k$ |
| $\boldsymbol{e}^{t}$ | error in momentum $\boldsymbol{m}^{t}-\nabla f\left(\boldsymbol{x}^{t-1}\right)$ |
| $\Delta_{i, k}^{t}, \Delta^{t-1}$ | $\mathbb{E}\left\\|\boldsymbol{y}_{i, k}^{t}-\boldsymbol{x}^{t-2}\right\\|^{2}, \mathbb{E}\left\\|\boldsymbol{x}^{t-1}-\boldsymbol{x}^{t-2}\right\\|^{2}=\Delta_{i, 0}^{t}$ |

## H. 1 Algorithm descriptions

Now, we formally describe the Mime MVR and MimeLite MVR algorithms. In each round $t$, we sample clients $\mathcal{S}^{t}$ such that $\left|\mathcal{S}^{t}\right|=S$. The server communicates the server parameters $\boldsymbol{x}^{t-1}$, the past parameters $\boldsymbol{x}^{t-2}$, and the momentum $\boldsymbol{m}^{t-1}$ term. Mime additionally uses a control variate $\boldsymbol{c}^{t-1}$ as we describe next.

Control variate in Mime. Mime uses an additional control variate $\boldsymbol{c}^{t-1}$ to reduce the variance.

$$
\begin{equation*}
\boldsymbol{c}^{t-1}=\frac{1}{S} \sum_{i \in \mathcal{S}^{t}} \nabla f_{i}\left(\boldsymbol{x}^{t-2}\right) \tag{12}
\end{equation*}
$$

Note that both $\boldsymbol{c}^{t-1}$ and $\boldsymbol{m}^{t-1}$ use gradients and parameters from previous rounds (different from the previous section). A naive implementation of this method requires two steps of communication per round to implement this algorithm. Alternatively, we can reserve some clients in the previous round for computing $\boldsymbol{c}^{t-1}$ which can then be used in the current round, removing the need for two steps of communication. In particular, it can be computed on a different set of an independent sampled clients $\tilde{\mathcal{S}}^{t-1}$. In fact, all our theoretical results hold even if we use a single client to perform the local updates and the rest of clients are used only to compute $\boldsymbol{c}^{t-1}$ each round.

Local client updates. Then each client $i \in \mathcal{S}^{t}$ makes a copy $\boldsymbol{y}_{i, 0}^{t}=\boldsymbol{x}^{t-1}$ and perform $K$ local client updates. In each local client update $k \in[K]$, the client samples a dataset $\zeta_{i, k}^{t}$. Mime performs the following update:

$$
\begin{align*}
& \boldsymbol{y}_{i, k}^{t}=\boldsymbol{y}_{i, k-1}^{t}-\eta \boldsymbol{d}_{i, k}^{t}, \text { where } \\
& \begin{aligned}
& \boldsymbol{d}_{i, k}^{t}=a\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x}^{t-1} ; \zeta_{i, k}^{t}\right)+\boldsymbol{c}^{t-1}\right)+(1-a) \boldsymbol{m}^{t-1} \\
& \quad+(1-a)\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x}^{t-1} ; \zeta_{i, k}^{t}\right)\right)
\end{aligned} \tag{13}
\end{align*}
$$

MimeLite on the other hand uses a very similar but simpler update scheme which does not rely on $c^{t-1}$ :

$$
\begin{align*}
& \boldsymbol{y}_{i, k}^{t}=\boldsymbol{y}_{i, k-1}^{t}-\eta \boldsymbol{d}_{i, k}^{t}, \text { where } \\
& \begin{aligned}
\boldsymbol{d}_{i, k}^{t}=a \nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ;\right. & \left.; \zeta_{i, k}^{t}\right)+(1-a) \boldsymbol{m}^{t-1} \\
& +(1-a)\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x}^{t-1} ; \zeta_{i, k}^{t}\right)\right)
\end{aligned} \tag{14}
\end{align*}
$$

Server updates. After $K$ such local updates, the server then aggregates the new client parameters as

$$
\begin{equation*}
\boldsymbol{x}^{t}=\frac{1}{S} \sum_{j \in \mathcal{S}^{t}} \boldsymbol{y}_{j, K}^{t} \tag{15}
\end{equation*}
$$

The momentum term is updated at the end of the round for $a \geq 0$ as

$$
\begin{equation*}
\boldsymbol{m}^{t}=\underbrace{a\left(\frac{1}{S} \sum_{j \in \mathcal{S}^{t}} \nabla f_{j}\left(\boldsymbol{x}^{t-1}\right)\right)+(1-a) \boldsymbol{m}^{t-1}}_{\text {Mom }}+\underbrace{(1-a)\left(\frac{1}{S} \sum_{j \in \mathcal{S}^{t}} \nabla f_{j}\left(\boldsymbol{x}^{t-1}\right)-\nabla f_{j}\left(\boldsymbol{x}^{t-2}\right)\right)}_{\text {correction }} . \tag{16}
\end{equation*}
$$

As we can see, the momentum update of MVR can be broken down into the usual Mom update, and a correction. Intuitively, this correction term is very small since $f_{i}$ is smooth and $\boldsymbol{x}^{t-1} \approx \boldsymbol{x}^{t-2}$. Another way of looking at the update (16) is to note that if all functions are identical i.e. $f_{j}=f_{k}$ for any $j, k$, then (16) just becomes the usual gradient descent. Thus MimeMVR tries to maintain an exponential moving average of only the variance terms, reducing its bias. We refer to [14] for more detailed explanation of MVR.

## H. 2 Bias in updates

The main difference in MimeMVR from the centralized versions of [62,14] is the additional local steps which are biased. In particular, for $k \geq 1$ the expected gradient $\mathbb{E}\left[\nabla f_{i}\left(\boldsymbol{y}_{i, k}^{t}\right)\right] \neq \nabla f\left(\boldsymbol{y}_{i, k}^{t}\right)$ because $\boldsymbol{y}_{i, k}^{t}$ also depends on the sample $i$. This bias is in fact the underlying cause of client drift and controlling it is a crucial step for our analysis.

Lemma 17 (Mime bias). For any values of $\boldsymbol{x}$ and $\boldsymbol{y}_{i}$ where $\boldsymbol{y}_{i}$ may depend on $i$, the following holds for any client $i$ almost surely given that (A1) and (A2) hold:

$$
\mathbb{E}_{\mathcal{S}, \zeta}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)+\frac{1}{|\mathcal{S}|} \sum_{j \in \mathcal{S}} \nabla f_{j}(\boldsymbol{x})-\nabla f_{i}(\boldsymbol{x} ; \zeta) \quad-\quad \nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2} \leq 2 \delta^{2} \mathbb{E}_{\mathcal{S}}\left\|\boldsymbol{y}_{i}-\boldsymbol{x}\right\|^{2}+\frac{2 G^{2}}{S}
$$

Proof. We can separate the noise from the rest of the terms and expand as

$$
\begin{aligned}
\mathbb{E}_{\zeta, \mathcal{S}} \| & \nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)+\frac{1}{|\mathcal{S}|} \sum_{j \in \mathcal{S}} \nabla f_{j}(\boldsymbol{x})-\nabla f_{i}(\boldsymbol{x} ; \zeta)-\nabla f\left(\boldsymbol{y}_{i}\right) \|^{2} \\
& \leq 2 \mathbb{E}_{\mathcal{S}}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)+\nabla f(\boldsymbol{x})-\nabla f_{i}(\boldsymbol{x} ; \zeta)-\nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2}+2 \mathbb{E}_{\mathcal{S}}\left\|\frac{1}{|\mathcal{S}|} \sum_{j \in \mathcal{S}} \nabla f_{j}(\boldsymbol{x})-\nabla f(\boldsymbol{x})\right\|^{2} \\
& \leq 2 \mathbb{E}_{\mathcal{S}}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)+\nabla f(\boldsymbol{x})-\nabla f_{i}(\boldsymbol{x} ; \zeta)-\nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2}+\frac{2 G^{2}}{S} \\
& \leq 2 \mathbb{E}_{\mathcal{S}} \delta^{2}\left\|\boldsymbol{y}_{i}-\boldsymbol{x}\right\|^{2}+\frac{2 G^{2}}{S}
\end{aligned}
$$

The first inequality used Young's inequality, the second used (A1), and the last used (A2) in the form of Lemma 3.

We can perform a similar analysis of the bias of local updates encountered by MimeLite.
Lemma 18 (MimeLite bias). For any values of $\boldsymbol{x}$ and $\boldsymbol{y}_{i}$ where $\boldsymbol{y}_{i}$ may depend on i, the following holds for any client $i$ randomly chosen from $\mathcal{C}$ given that (A1), (A2) and (A3) hold:

$$
\mathbb{E}_{i, \zeta}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right) \quad-\quad \nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2} \leq 2 \delta^{2} \mathbb{E}_{i}\left\|\boldsymbol{y}_{i}-\boldsymbol{x}\right\|^{2}+2 G^{2}+\sigma^{2}
$$

Proof. We can separate the noise from the rest of the terms and expand as

$$
\begin{aligned}
\mathbb{E}_{\zeta, i}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)-\nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2} & =\mathbb{E}_{\zeta, i}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right) \pm \nabla f_{i}(\boldsymbol{x}) \pm \nabla f(\boldsymbol{x})-\nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2} \\
& \leq \mathbb{E}_{i}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i}\right) \pm \nabla f_{i}(\boldsymbol{x}) \pm \nabla f(\boldsymbol{x})-\nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2}+\sigma^{2} \\
\leq & 2 \mathbb{E}_{i}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i}\right)+\nabla f(\boldsymbol{x})-\nabla f_{i}(\boldsymbol{x})-\nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2} \\
& \quad+2 \mathbb{E}_{i}\left\|\nabla f_{i}(\boldsymbol{x})-\nabla f(\boldsymbol{x})\right\|^{2}+\sigma^{2} \\
& \leq 2 \mathbb{E}_{i}\left\|\nabla f_{i}\left(\boldsymbol{y}_{i}\right)+\nabla f(\boldsymbol{x})-\nabla f_{i}(\boldsymbol{x})-\nabla f\left(\boldsymbol{y}_{i}\right)\right\|^{2}+2 G^{2}+\sigma^{2} \\
\leq & 2 \delta^{2} \mathbb{E}_{i}\left\|\boldsymbol{y}_{i}-\boldsymbol{x}\right\|^{2}+2 G^{2}+\sigma^{2} .
\end{aligned}
$$

The first inequality used (A3), the second used Young's inequality, the third used (A1), and the last used (A2) in the form of Lemma 3.

Note that the bias for MimeLite is very similar to that of Mime, except that Mime has dependence of $\frac{G^{2}}{S}$, whereas MimeLite has $G^{2}+\sigma^{2}$. Hence, the rate of convergence of MimeLite will depend on $G^{2}$ wheras Mime will have the optimal dependency of $G^{2} / S$. Hence, in the rest of the proof, we will consider only Mime and simply replace $G^{2} / S$ with $\left(G^{2}+\sigma^{2}\right)$ to obtain the corresponding results for MimeLite.

## H. 3 Change in each client update

Client update variance. Now we examine the variance of our update in each local step $\boldsymbol{d}_{i, k}^{t}$.
Lemma 19. For the client update (13), given (A1) and (A2), the following holds for any $a \in[0,1]$ where $\boldsymbol{e}^{t}:=\boldsymbol{m}^{t}-\nabla f\left(\boldsymbol{x}^{t-1}\right)$ and $\Delta_{i, k}^{t}:=\mathbb{E}\left\|\boldsymbol{y}_{i, k}^{t}-\boldsymbol{x}^{t-2}\right\|^{2}$ :

$$
\mathbb{E}\left\|\boldsymbol{d}_{i, k}^{t}-\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \leq 3 \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+3 \delta^{2} \Delta_{i, k-1}^{t}+\frac{3 a^{2} G^{2}}{S}
$$

Proof. Starting from the client update (13), we can rewrite it as

$$
\begin{aligned}
\boldsymbol{d}_{i, k}^{t}-\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)= & (1-a) \boldsymbol{e}^{t-1} \\
& \left.+\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x}^{t-2} ; \zeta_{i, k}^{t}\right)\right)-\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)+\nabla f\left(\boldsymbol{x}^{t-2}\right)\right) \\
& +a\left(\frac{1}{S} \sum_{j \in \mathcal{S}^{t}} \nabla f_{j}\left(\boldsymbol{x}^{t-2}\right)-\nabla f\left(\boldsymbol{x}^{t-2}\right)\right)
\end{aligned}
$$

We can use the relaxed triangle inequality Lemma 1 to claim

$$
\begin{aligned}
\mathbb{E} \| \boldsymbol{d}_{i, k}^{t}-\nabla & f\left(\boldsymbol{y}_{i, k-1}^{t}\right) \|^{2} \\
= & 3(1-a)^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2} \\
& +3(1-a)^{2}\left\|\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x}^{t-2} ; \zeta_{i, k}^{t}\right)\right)-\left(\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)-\nabla f\left(\boldsymbol{x}^{t-2}\right)\right)\right\|^{2} \\
& +3 a^{2}\left\|\frac{1}{S} \sum_{j \in \mathcal{S}^{t}} \nabla f_{j}\left(\boldsymbol{x}^{t-2}\right)-\nabla f\left(\boldsymbol{x}^{t-2}\right)\right\|^{2} \\
\leq & 3 \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+3 \delta^{2}\left\|\boldsymbol{y}_{i, k-1}^{t}-\boldsymbol{x}^{t-2}\right\|^{2}+\frac{3 a^{2} G^{2}}{S}
\end{aligned}
$$

The last inequality used the Hessian similarity Lemma 3 to bound the second term and the heterogeneity bound (A1) to bound the last term. Also, $(1-a)^{2} \leq 1$ since $a \in[0,1]$.

Distance moved in each step. We show that the distance moved by a client in each step during the client update can be controlled.
Lemma 20. For MimeMVR updates (13) with $\eta \leq \frac{1}{6 K \delta}$ and given (A1) and (A2), the following holds

$$
\Delta_{i, k}^{t} \leq\left(1+\frac{1}{K}\right) \Delta_{i, k-1}^{t}+18 \eta^{2} K a^{2} \frac{G^{2}}{S}+18 \eta^{2} K \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+6 \eta^{2} K\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}
$$

where we define $\Delta_{i, k}^{t}:=\mathbb{E}\left\|\boldsymbol{y}_{i, k}^{t}-\boldsymbol{x}^{t-2}\right\|^{2}$.
Proof. Starting from the MimeMVR update (13) and the relaxed triangle inequality with $c=2 K$,

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{y}_{i, k}^{t}-\boldsymbol{x}^{t-2}\right\|^{2}= & \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\eta \boldsymbol{d}_{i, k}^{t}-\boldsymbol{x}^{t-2}\right\|^{2} \\
\leq & \left(1+\frac{1}{2 K}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\boldsymbol{x}^{t-2}\right\|^{2}+(2 K+1) \eta^{2} \mathbb{E}\left\|\boldsymbol{d}_{i, k}^{t}\right\|^{2} \\
\leq & \left(1+\frac{1}{2 K}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\boldsymbol{x}^{t-2}\right\|^{2}+6 K \eta^{2} \mathbb{E}\left\|\boldsymbol{d}_{i, k}^{t}-\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
& +6 K \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
\leq & \left(1+\frac{1}{2 K}+18 K \eta^{2} \delta^{2}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\boldsymbol{x}^{t-2}\right\|^{2} \\
& +18 K \eta^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{18 K \eta^{2} a^{2} G^{2}}{S}+6 K \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} .
\end{aligned}
$$

The last inequality used the update variance bound Lemma 19. We can simplify the expression further since $\eta \leq \frac{1}{6 K \delta}$ implies $18 K \eta^{2} \delta^{2} \leq \frac{1}{2 K}$.

Progress in one step. Now we can compute the progress made in each step.
Lemma 21. For any client update step with step size $\eta \leq \min \left(\frac{1}{L}, \frac{1}{192 \delta K}\right)$ and given that (A1), (A2) hold, we have

$$
\begin{aligned}
\mathbb{E} f\left(\boldsymbol{y}_{i, k}^{t}\right)+\delta\left(1+\frac{2}{K}\right)^{K-k} \Delta_{i, k}^{t} \leq \mathbb{E} f\left(\boldsymbol{y}_{i, k-1}^{t}\right) & +\delta\left(1+\frac{2}{K}\right)^{K-(k-1)} \Delta_{i, k-1}^{t} \\
& -\frac{\eta}{4} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}+3 \eta \mathbb{E}\left\|e^{t-1}\right\|^{2}+\frac{3 \eta a^{2} G^{2}}{S}
\end{aligned}
$$

Proof. The assumption that $f$ is $L$-smooth implies a quadratic upper bound (10).

$$
\begin{aligned}
f\left(\boldsymbol{y}_{i, k}^{t}\right)-f\left(\boldsymbol{y}_{i, k-1}^{t}\right) & \leq-\eta\left\langle\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right), \boldsymbol{d}_{i, k}^{t}\right\rangle+\frac{L \eta^{2}}{2}\left\|\boldsymbol{d}_{i, k}^{t}\right\|^{2} \\
& =-\frac{\eta}{2}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}+\frac{L \eta^{2}-\eta}{2}\left\|\boldsymbol{d}_{i, k}^{t}\right\|^{2}+\frac{\eta}{2}\left\|\boldsymbol{d}_{i, k}^{t}-\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}
\end{aligned}
$$

The second equality used the fact that for any $a, b,-2 a b=(a-b)^{2}-a^{2}-b^{2}$. The second term can be removed since $\eta \leq \frac{1}{L}$. Taking expectation on both sides and using the update variance bound Lemma 19,

$$
\begin{aligned}
& \mathbb{E} f\left(\boldsymbol{y}_{i, k}^{t}\right)-\mathbb{E} f\left(\boldsymbol{y}_{i, k-1}^{t}\right) \leq-\frac{\eta}{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}+\frac{3 \eta a^{2} G^{2}}{2 S} \\
&+\frac{3 \eta}{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 \eta \delta^{2}}{2} \Delta_{i, k-1}^{t} \\
& \leq-\frac{\eta}{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}+\frac{3 \eta a^{2} G^{2}}{2 S} \\
&+\frac{3 \eta}{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 \eta \delta^{2}}{2} \Delta_{i, k-1}^{t}
\end{aligned}
$$

Multiplying the distance bound Lemma 20 by $\delta\left(1+\frac{2}{K}\right)^{K-k}$. Note that for any $K \geq 1$ and $k \in[K]$, we have $1 \leq\left(1+\frac{2}{K}\right)^{K-k} \leq 8$. Then we get

$$
\begin{aligned}
& \delta\left(1+\frac{2}{K}\right)^{K-k} \Delta_{i, k}^{t} \leq \delta\left(1+\frac{2}{K}\right)^{K-k}( \left(1+\frac{1}{K}\right) \Delta_{i, k-1}^{t}+18 \eta^{2} K a^{2} \frac{G^{2}}{S} \\
&\left.+18 \eta^{2} K \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+6 \eta^{2} K\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}\right) \\
& \leq \delta\left(1+\frac{2}{K}\right)^{K-(k-1)} \Delta_{i, k-1}^{t}-\frac{\delta}{K}\left(1+\frac{2}{K}\right)^{K-k} \Delta_{i, k-1}^{t} \\
&+48 \eta^{2} \delta K \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}+\frac{144 \eta^{2} \delta K a^{2} G^{2}}{S}+144 \eta^{2} \delta K \mathbb{E}\left\|e^{t-1}\right\|^{2} \\
& \leq \delta\left(1+\frac{2}{K}\right)^{K-(k-1)} \Delta_{i, k-1}^{t}-\frac{\delta}{K} \Delta_{i, k-1}^{t}+48 \eta^{2} \delta K \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
&+\frac{144 \eta^{2} \delta K a^{2} G^{2}}{S}+144 \eta^{2} \delta K \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2} .
\end{aligned}
$$

Adding these two inequalities together yields

$$
\begin{aligned}
\mathbb{E} f\left(\boldsymbol{y}_{i, k}^{t}\right)+\delta\left(1+\frac{2}{K}\right)^{K-k} \Delta_{i, k}^{t} \leq \mathbb{E} f & \left(\boldsymbol{y}_{i, k-1}^{t}\right)+\delta\left(1+\frac{2}{K}\right)^{K-(k-1)} \Delta_{i, k-1}^{t} \\
& -\left(\frac{\eta}{2}-48 \eta^{2} \delta K\right) \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
& +\left(\frac{3 \eta}{2}+144 \eta^{2} \delta K\right) \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2} \\
& +\left(\frac{3 \eta}{2}++144 \eta^{2} \delta K\right) \frac{a^{2} G^{2}}{S}
\end{aligned}
$$

Using our bound on the step-size that $\eta \leq \frac{1}{192 \delta K}$ implies that $\eta \delta K \leq \frac{1}{48 * 4}$.

## H. 4 Change in each round

We now see how the quantities we defined change across rounds.

## Distance moved in a round.

Lemma 22. For MimeMVR updates (13) with $\eta \leq \frac{1}{6 K \delta}$ and given (A1) and (A2), the following holds

$$
\Delta^{t} \leq 54 K^{2} \eta^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{54 K^{2} \eta^{2} a^{2} G^{2}}{S}+\frac{1}{K S} \sum_{i, k} 18 K^{2} \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}
$$

where we define $\Delta^{t}:=\mathbb{E}\left\|\boldsymbol{x}^{t}-\boldsymbol{x}^{t-1}\right\|^{2}$.
Proof. Starting from the MimeMVR update (13) and following the proof of Lemma 20,

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{y}_{i, k}^{t}-\boldsymbol{x}^{t-1}\right\|^{2}= & \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\eta \boldsymbol{d}_{i, k}^{t}-\boldsymbol{x}^{t-1}\right\|^{2} \\
\leq & \left(1+\frac{1}{2 K}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\boldsymbol{x}^{t-1}\right\|^{2}+(2 K+1) \eta^{2} \mathbb{E}\left\|\boldsymbol{d}_{i, k}^{t}\right\|^{2} \\
\leq & \left(1+\frac{1}{2 K}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\boldsymbol{x}^{t-1}\right\|^{2}+6 K \eta^{2} \mathbb{E}\left\|\boldsymbol{d}_{i, k}^{t}-\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
& +6 K \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
\leq & \left(1+\frac{1}{K}\right) \mathbb{E}\left\|\boldsymbol{y}_{i, k-1}^{t}-\boldsymbol{x}^{t-1}\right\|^{2} \\
& +18 K \eta^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{18 K \eta^{2} a^{2} G^{2}}{S}+6 K \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}
\end{aligned}
$$

Note that $\boldsymbol{x}^{t}=\frac{1}{S} \sum_{i \in \mathcal{S}} \boldsymbol{y}_{i, K}^{t}$ and so,

$$
\begin{aligned}
& \mathbb{E}\left\|\boldsymbol{x}^{t}-\boldsymbol{x}^{t-1}\right\|^{2} \\
& \quad \leq \frac{1}{S} \sum_{i \in \mathcal{S}} \mathbb{E}\left\|\boldsymbol{y}_{i, K}^{t}-\boldsymbol{x}^{t-1}\right\|^{2} \\
& \quad \leq \frac{1}{S} \sum_{i \in \mathcal{S}} \sum_{k}\left(18 K \eta^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{18 K \eta^{2} a^{2} G^{2}}{S}+6 K \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}\right)\left(1+\frac{1}{K}\right)^{K-k} \\
& \quad \leq 54 K^{2} \eta^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{54 K^{2} \eta^{2} a^{2} G^{2}}{S}+\frac{1}{K S} \sum_{i, k} 18 K^{2} \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} .
\end{aligned}
$$

Here we used the inequality that for all $k,\left(1+\frac{1}{K}\right)^{K-k} \leq 3$.
Server momentum variance. We compute the error of the server momentum $\boldsymbol{m}^{t-1}$ defined as $\boldsymbol{e}^{t}=\boldsymbol{m}^{t}-\nabla f\left(\boldsymbol{x}^{t-1}\right)$. Its expected norm can be bounded as follows.
Lemma 23. For the momentum update (16), given (A1) and (A2), the following holds for any $\eta \leq \frac{1}{51 \delta K}$ and $1 \geq a \geq 2592 K^{2} \delta^{2} \eta^{2}$,

$$
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \leq\left(1-\frac{23 a}{24}\right) \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 a^{2} G^{2}}{S}+\frac{1}{K S} \sum_{i, k} 36 K^{2} \delta^{2} \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}
$$

Proof. Starting from the momentum update (16),

$$
\begin{aligned}
\boldsymbol{e}^{t}= & (1-a) \boldsymbol{e}^{t-1} \\
& +(1-a)\left(\frac{1}{S} \sum_{j \in \mathcal{S}^{t}}\left(\nabla f_{j}\left(\boldsymbol{x}^{t-1}\right)-\nabla f_{j}\left(\boldsymbol{x}^{t-2}\right)\right)-\nabla f\left(\boldsymbol{x}^{t-1}\right)+\nabla f\left(\boldsymbol{x}^{t-2}\right)\right) \\
& +a\left(\frac{1}{S} \sum_{j \in \mathcal{S}^{t}}\left(\nabla f_{j}\left(\boldsymbol{x}^{t-1}\right)-\nabla f\left(\boldsymbol{x}^{t-1}\right)\right) .\right.
\end{aligned}
$$

Now, the term $e^{t-1}$ does not have any information from round $t$ and hence is statistically independent of the rest of the terms. Further, the rest of the terms have mean 0 . Hence, we can separate out the zero mean noise terms from the $e^{t-1}$ following Lemma 2 and then the relaxed triangle inequality Lemma 1 to claim

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \leq & (1-a)^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2} \\
& +2(1-a)^{2}\left\|\frac{1}{S} \sum_{j \in \mathcal{S}^{t}}\left(\nabla f_{j}\left(\boldsymbol{x}^{t-1}\right)-\nabla f_{j}\left(\boldsymbol{x}^{t-2}\right)\right)-\nabla f\left(\boldsymbol{x}^{t-1}\right)+\nabla f\left(\boldsymbol{x}^{t-2}\right)\right\|^{2} \\
& +2 a^{2} \| \frac{1}{S} \sum_{j \in \mathcal{S}^{t}}\left(\nabla f_{j}\left(\boldsymbol{x}^{t-1}\right)-\nabla f\left(\boldsymbol{x}^{t-1}\right) \|^{2}\right. \\
\leq & (1-a)^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+2(1-a)^{2} \delta^{2}\left\|\boldsymbol{x}^{t-1}-\boldsymbol{x}^{t-2}\right\|^{2}+\frac{2 a^{2} G^{2}}{S}
\end{aligned}
$$

The inequality used the Hessian similarity Lemma 3 to bound the second term and the heterogeneity bound (A1) to bound the last term. Finally, note that $(1-a)^{2} \leq(1-a) \leq 1$ for $a \in[0,1]$. We can continue by bounding $\Delta^{t-1}$ using Lemma 22 .

$$
\begin{aligned}
\mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \leq & (1-a) \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+2 \delta^{2} \Delta^{t-1}+\frac{2 a^{2} G^{2}}{S} \\
\leq & (1-a) \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{2 a^{2} G^{2}}{S} \\
& +108 K^{2} \delta^{2} \eta^{2} \mathbb{E}\left\|e^{t-1}\right\|^{2}+\frac{108 K^{2} \delta^{2} \eta^{2} a^{2} G^{2}}{S}+\frac{1}{K S} \sum_{i, k} 36 K^{2} \delta^{2} \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
\leq & \left(1-\frac{23 a}{24}\right) \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 a^{2} G^{2}}{S}+\frac{1}{K S} \sum_{i, k} 36 K^{2} \delta^{2} \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} .
\end{aligned}
$$

The last step used our bound on the momentum parameter that $1 \geq a \geq 2592 \eta^{2} \delta^{2} K^{2}$. Note that $\eta \leq \frac{1}{51 \delta K}$ ensures that this set is non-empty.

Progress in one round. Finally, we can compute the progress made in a round. Note that we need a technical condition that $f$ is $\delta$-weakly convex. However, this is only needed because we insist on running the algorithm on $S$ clients in parallel and then averaging their weights-the averaging requires weak convexity to ensure that the loss doesn't blow up. It has been experimentally observed in [43] that with the right initialization, averaging of the parameters does not increase the loss value and so weak convexity within this region might be vaalid. Finally note that if we instead simply run the local updates on a single chosen client with all the rest only being used to compute $\boldsymbol{c}^{t-1}$, we will retain all convergence rates without needing weak-convexity.
Lemma 24. For any round of MimeMVR with step size $\eta \leq \min \left(\frac{1}{L}, \frac{1}{864 \delta K}\right)$ and momentum parameter $a \geq 912 \eta^{2} \delta^{2} K^{2}$. Then, given that (A1)-(A2) hold and $f$ is $\delta$-weakly convex, we have

$$
\frac{\eta}{24 K S} \sum_{k \in[K], j \in \mathcal{S}^{t}} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \leq \Phi^{t-1}-\Phi^{t}+\frac{17 \eta a \delta^{2} K^{2} G^{2}}{S}
$$

where we define the sequence

$$
\Phi^{t}:=\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)-f^{\star}\right]+\frac{96 \eta}{23 a} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{8 \delta}{K} \Delta^{t} .
$$

Proof. We start by summing over the progress in single client updates as in Lemma 21

$$
\begin{aligned}
& \sum_{k \in[K]} \frac{\eta}{4} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, 0}^{t}\right)\right\|^{2} \leq \mathbb{E} f\left(\boldsymbol{y}_{i, 0}^{t}\right)+\delta\left(1+\frac{2}{K}\right)^{K} \Delta_{i, 0}^{t} \\
&-\mathbb{E} f\left(\boldsymbol{y}_{i, K}^{t}\right)-\delta \Delta_{i, K}^{t} \\
&+3 \eta K \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 \eta K a^{2} G^{2}}{S} \\
& \leq \mathbb{E} f\left(\boldsymbol{y}_{i, 0}^{t}\right)+8 \delta \Delta_{i, 0}^{t}-\mathbb{E} f\left(\boldsymbol{y}_{i, K}^{t}\right)-\delta \Delta_{i, K}^{t} \\
&+3 \eta K \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 \eta K a^{2} G^{2}}{S} \\
& \leq \mathbb{E} f\left(\boldsymbol{x}^{t-1}\right)+8 \delta \Delta^{t-1}-\mathbb{E} f\left(\boldsymbol{y}_{i, K}^{t}\right)-\delta \Delta_{i, K}^{t} \\
&+3 \eta K \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 \eta K a^{2} G^{2}}{S}
\end{aligned}
$$

Recall that $\Delta_{i, k}^{t}=\mathbb{E}\left\|\boldsymbol{y}_{i, k}^{t}-\boldsymbol{x}^{t-2}\right\|^{2}$ and $\boldsymbol{y}_{i, 0}^{t}=\boldsymbol{x}^{t-1}$. This gives the last step above, making $\Delta_{i, 0}^{t}=\Delta^{t-1}$. Then by the averaging Lemma 4, we have

$$
\begin{aligned}
\frac{1}{S} \sum_{j \in \mathcal{S}^{t}} \mathbb{E}\left[f\left(\boldsymbol{y}_{j, K}^{t}\right)\right]+\delta \Delta_{j, K}^{t} & =\frac{1}{S} \sum_{j \in \mathcal{S}} \mathbb{E}\left[f\left(\boldsymbol{y}_{j, K}^{t}\right)\right]+\delta \mathbb{E}\left\|\boldsymbol{x}^{t-2}-\boldsymbol{y}_{j, K}^{t}\right\|^{2} \\
& \geq \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)\right]+\delta \mathbb{E}\left\|\boldsymbol{x}^{t-2}-\boldsymbol{x}^{t}\right\|^{2} .
\end{aligned}
$$

So by averaging our inequality over the sampled clients, and diving our summation over the updates by $K$, we get

$$
\begin{aligned}
& \frac{\eta}{4 K S} \sum_{k \in[K], j \in \mathcal{S}^{t}} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
& \leq \frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t-1}\right)\right]+3 \eta \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{8 \delta}{K} \Delta^{t-1}-\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)\right]+\frac{3 \eta a^{2} G^{2}}{S}
\end{aligned}
$$

We can use the bound on $\Delta_{t}$ from Lemma 22 to proceed as

$$
\begin{aligned}
\frac{\eta}{4 K S} \sum_{k \in[K], j \in \mathcal{S}^{t}} \mathbb{E} \| \nabla & f\left(\boldsymbol{y}_{i, k-1}^{t}\right) \|^{2} \\
\leq & \frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t-1}\right)\right]-\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)\right]+3 \eta \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{3 \eta a^{2} G^{2}}{S} \\
& +\frac{8 \delta}{K} \Delta^{t-1}-\frac{8 \delta}{K} \Delta^{t} \\
& +432 K \delta \eta^{2} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{432 K \delta \eta^{2} a^{2} G^{2}}{S}+\frac{1}{K S} \sum_{i, k} 144 K \delta \eta^{2} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
\leq & \frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t-1}\right)\right]-\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)\right]+4 \eta \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{4 \eta a^{2} G^{2}}{S} \\
& +\frac{8 \delta}{K} \Delta^{t-1}-\frac{8 \delta}{K} \Delta^{t}+\frac{\eta}{6 K S} \sum_{i, k} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2}
\end{aligned}
$$

The last step used the bound on the step size that $\eta \leq \frac{1}{864 \delta K}$. Now, multiplying the error bound Lemma 23 by $\frac{96 \eta}{23 a}$ gives

$$
\frac{96 \eta}{23 a} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \leq \frac{4 * 24 \eta}{23 a}\left(1-\frac{23 a}{24}\right) \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}+\frac{13 \eta a G^{2}}{S}+\frac{1}{K S} \sum_{i, k} \frac{38 K^{2} \delta^{2} \eta^{3}}{a} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} .
$$

Adding this to the previously obtained bound yields

$$
\begin{aligned}
\frac{\eta}{4 K S} \sum_{k \in[K], j \in \mathcal{S}^{t}} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \leq\left(\frac{1}{6}\right. & \left.+\frac{38 K^{2} \delta^{2} \eta^{2}}{a}\right) \frac{\eta}{K S} \sum_{k \in[K], j \in \mathcal{S}^{t}} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \\
& +\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t-1}\right)\right]-\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)\right] \\
& +\frac{96 \eta}{23 a} \mathbb{E}\left\|\boldsymbol{e}^{t-1}\right\|^{2}-\frac{96 \eta}{23 a} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \\
& +\frac{8 \delta}{K} \Delta^{t-1}-\frac{8 \delta}{K} \Delta^{t} \\
& -\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)\right]-\frac{4 \eta}{a} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2} \\
& +\left(13 \eta a+3 \eta a^{2}\right) \frac{G^{2}}{S}
\end{aligned}
$$

Since $a \geq 912 \eta^{2} K^{2} \delta^{2}$, we have $\frac{1}{4}-\left(\frac{1}{6}-\frac{38 K^{2} \delta^{2} \eta^{2}}{a}\right) \geq \frac{1}{24}$. Using this proves the lemma.

## H. 5 Final convergence rates

Theorem $\mathbf{V}$ (Convergence of MimeMVR). Let us run MimeMVR with step size $\eta=$ $\min \left(\frac{1}{L}, \frac{1}{864 \delta K},\left(\frac{S\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{6936 K^{3} T \delta^{2} G^{2}}\right)^{1 / 3}\right)$ and momentum parameter $a=\max \left(1536 \eta^{2} \delta^{2} K^{2}, \frac{1}{T}\right)$.
Then, given that (A1) and (A2) hold, we have

$$
\frac{1}{K S T} \sum_{t \in[T]} \sum_{k \in[K]} \sum_{j \in \mathcal{S}^{t}} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \leq \mathcal{O}\left(\left(\frac{\delta^{2} G^{2} F}{S T^{2}}\right)^{1 / 3}+\frac{G^{2}}{S T}+\frac{(L+\delta K) F}{K T}\right)
$$

where we define $F:=f\left(\boldsymbol{x}^{0}\right)-f^{\star}$.
Proof. Unroll the one round progress Lemma 24 and average over $T$ rounds to get

$$
\frac{1}{K S T} \sum_{t \in[T]} \sum_{k \in[K]} \sum_{j \in \mathcal{S}^{t}} \mathbb{E}\left\|f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \leq \frac{24\left(\Phi^{0}-\Phi^{T}\right)}{\eta T}+\frac{408 a G^{2}}{S}
$$

Recall that we defined

$$
\Phi^{t}:=\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{t}\right)-f^{\star}\right]+\frac{96 \eta}{23 a} \mathbb{E}\left\|\boldsymbol{e}^{t}\right\|^{2}+\frac{8 \delta}{K} \Delta^{t} .
$$

Hence, $\Phi^{T} \geq 0$. Further, note that by definition $\Delta^{0}=0$ and $\mathbb{E}\left\|\boldsymbol{e}_{0}\right\|^{2}:=\mathbb{E}\left\|\boldsymbol{m}^{0}-\nabla f\left(\boldsymbol{x}^{0}\right)\right\|^{2}$. [14] show that by using time-varying step sizes, it is possible to directly control the error $\boldsymbol{e}_{0}$. Alternatively, [62] use a large initial accumulation for the momentum term. For the sake of simplicity, we will follow the latter approach. It is straightforward to extend our techniques to the time-varying step-size case as well but with additional proof complexity. Note that either way, the total complexity only changes by a factor of 2 . Suppose that we run the algorithm for $2 T$ rounds wherein for the first $T$ rounds, we simply compute $\boldsymbol{m}^{0}=\frac{1}{T_{0} S} \sum_{t=1}^{T_{0}} \sum_{j \in \mathcal{S}^{t}} \nabla f_{j}\left(\boldsymbol{x}^{0}\right)$. With this, we have $\boldsymbol{e}_{0}=\mathbb{E}\left\|\boldsymbol{m}^{0}-\nabla f\left(\boldsymbol{x}^{0}\right)\right\|^{2} \leq \frac{G^{2}}{S T}$. Thus, we have for the first round $t=1$

$$
\Phi^{0}=\frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right]+\frac{96 \eta}{23 a} \mathbb{E}\left\|\boldsymbol{e}^{0}\right\|^{2} \leq \frac{1}{K} \mathbb{E}\left[f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right]+\frac{96 \eta G^{2}}{23 a T S} .
$$

Together, this gives

$$
\frac{1}{K S T} \sum_{t \in[T]} \sum_{k \in[K]} \sum_{i \in \mathcal{S}^{t}} \mathbb{E}\left\|f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \leq \frac{24\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{\eta K T}+\frac{96 G^{2}}{a T^{2} S}+\frac{408 a G^{2}}{S}
$$

The above equation holds for any choice of $\eta \leq \min \left(\frac{1}{L}, \frac{1}{864 \delta K}\right)$ and momentum parameter $a \geq$ $912 \eta^{2} \delta^{2} K^{2}$. Set the momentum parameter as

$$
a=\max \left(912 \eta^{2} \delta^{2} K^{2}, \frac{1}{T}\right)
$$

With this choice, we can simplify the rate of convergence as

$$
\frac{24\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{\eta K T}+\frac{96 G^{2}}{T S}+\frac{166464 \eta^{2} \delta^{2} K^{2} G^{2}}{S}+\frac{408 G^{2}}{S T}
$$

Now let us pick

$$
\eta=\min \left(\frac{1}{L}, \frac{1}{864 \delta K},\left(\frac{S\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{6936 K^{3} T \delta^{2} G^{2}}\right)^{1 / 3}\right)
$$

For this combination of step size $\eta$ and $a$, the rate simplifies to

$$
\frac{504 G^{2}}{T S}+916\left(\frac{\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right) \delta^{2} G^{2}}{S T^{2}}\right)^{1 / 3}+\frac{24(L+864 \delta K)\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{K T}
$$

This finishes the proof of the theorem.
Theorem VI (Convergence of MimeLiteMVR). Let us run MimeLiteMVR with step size $\eta=$ $\min \left(\frac{1}{L}, \frac{1}{864 \delta K},\left(\frac{\left(f\left(\boldsymbol{x}^{0}\right)-f^{\star}\right)}{6936 K^{3} T \delta^{2}\left(G^{2}+\sigma^{2}\right)}\right)^{1 / 3}\right)$ and momentum parameter $a=\max \left(1536 \eta^{2} \delta^{2} K^{2}, \frac{1}{T}\right)$.
Then, given that (A1) and (A2*) hold, we have
$\frac{1}{K S T} \sum_{t \in[T]} \sum_{k \in[K]} \sum_{j \in \mathcal{S}^{t}} \mathbb{E}\left\|\nabla f\left(\boldsymbol{y}_{i, k-1}^{t}\right)\right\|^{2} \leq \mathcal{O}\left(\left(\frac{\delta^{2}\left(G^{2}+\sigma^{2}\right) F}{T^{2}}\right)^{1 / 3}+\frac{G^{2}+\sigma^{2}}{T}+\frac{(L+\delta K) F}{K T}\right)$,
where we define $F:=f\left(\boldsymbol{x}^{0}\right)-f^{\star}$.
Proof. The proof for MimeLiteMVR is identical to that of MimeMVR, except that as noted in Lemma 18, the $\frac{G^{2}}{S}$ term in Mime gets replaced by $\left(G^{2}+\sigma^{2}\right)$ everywhere. Note that MimeLiteMVR (Lemma 18) requires a weaker Hessian variance condition of $\left\|\nabla^{2} f_{i}(\boldsymbol{x})-\nabla^{2} f(\boldsymbol{x})\right\| \leq \delta$ as opposed to MimeMVR which needs $\left\|\nabla^{2} f_{i}(\boldsymbol{x} ; \zeta)-\nabla^{2} f(\boldsymbol{x})\right\| \leq \delta$.

Note that the final convergence rates of MimeMVR and MimeLiteMVR both include the intermediate client parameters. To implement this algorithm would require additional communication where in each round a random client parameter in $\left\{\boldsymbol{y}_{i, 1}^{t}, \ldots, \boldsymbol{y}_{i, K}^{t}\right\}$ is communicated to the server. However, as is common in non-convex stochastic analysis, we expect the last iterate to converge at a similar rate as well in practice.

## I Algorithm pseudocodes

```
Algorithm 2 FedAvg framework
    input: initial \(\boldsymbol{x}\) and \(\boldsymbol{s}\), server learning rate \(\eta_{g}\), client learning rate \(\eta_{l}\), and base optimizer \(\mathcal{B}=\)
    \((\mathcal{U}, \mathcal{V})\)
    for each round \(t=1, \cdots, T\) do
        sample subset \(\mathcal{S}\) of clients
        communicate \(\boldsymbol{x}\) to all clients \(i \in \mathcal{S}\)
        on client \(i \in \mathcal{S}\) in parallel do
            initialize local model \(\boldsymbol{y}_{i} \leftarrow \boldsymbol{x}\)
            for \(k=1, \cdots, K\) do
                sample mini-batch \(\zeta\) from local data
                    update \(\boldsymbol{y}_{i} \leftarrow \boldsymbol{y}_{i}-\eta_{l} \nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)\)
            end for
            communicate \(\boldsymbol{y}_{i}\)
        end on client
        compute aggregate pseudo-gradient \(\boldsymbol{g} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}}\left(\boldsymbol{x}-\boldsymbol{y}_{i}\right)\)
        \(\boldsymbol{x} \leftarrow \eta_{g} \mathcal{U}(\boldsymbol{g}, \boldsymbol{s})\) (update server parameters)
        \(\boldsymbol{s} \leftarrow \mathcal{V}(\boldsymbol{g}, \boldsymbol{s})\) (update optimizer state)
    end for
```

```
Algorithm 3 MimeMom and MimeLiteMom
    input: initial \(\boldsymbol{x}\), and hyperparameters \(\eta, \beta\). optional \(\eta_{g}(\) default \(=1)\)
    initialize \(\boldsymbol{m} \leftarrow 0, \boldsymbol{c} \leftarrow 0\)
    for each round \(t=1, \cdots, T\) do
        sample subset \(\mathcal{S}\) of clients
        communicate \((\boldsymbol{x}, \boldsymbol{s}=\boldsymbol{m})\) and \(\boldsymbol{c}\) (only Mime) to all clients \(i \in \mathcal{S}\)
        on client \(i \in \mathcal{S}\) in parallel do
            initialize local model \(\boldsymbol{y}_{i} \leftarrow \boldsymbol{x}\)
            for \(k=1, \cdots, K\) do
                sample mini-batch \(\zeta\) from local data
                \(\boldsymbol{g}_{i} \leftarrow \nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)-\nabla f_{i}(\boldsymbol{x} ; \zeta)+\boldsymbol{c}\) (Mime)
                    \(\boldsymbol{g}_{i} \leftarrow \nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)\) (MimeLite)
                    update using server momentum \(\boldsymbol{y}_{i} \leftarrow \boldsymbol{y}_{i}-\eta\left(\left(1-\beta_{1}\right) \boldsymbol{g}_{i}+\beta_{1} \boldsymbol{m}\right)\)
            end for
            compute full local-batch gradient \(\nabla f_{i}(\boldsymbol{x})\)
            communicate \(\left(\boldsymbol{y}_{i}, \nabla f_{i}(\boldsymbol{x})\right)\)
        end on client
        compute \(\boldsymbol{c} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \nabla f_{i}(\boldsymbol{x})\)
        \(\boldsymbol{m} \leftarrow\left(\left(1-\beta_{1}\right) \boldsymbol{c}+\beta_{1} \boldsymbol{m}\right)\) (update server momentum)
        \(\boldsymbol{x} \leftarrow \boldsymbol{x}-\eta_{g} \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}}\left(\boldsymbol{x}-\boldsymbol{y}_{i}\right)\) (update server parameters)
    end for
```

Algorithm 4 MimeAdam and MimeLiteAdam
input: initial $\boldsymbol{x}$, and hyperparameters $\eta, \beta_{1}, \beta_{2}, \varepsilon_{0}$. optional $\eta_{g}($ default $=1)$
initialize $\boldsymbol{m} \leftarrow 0, \boldsymbol{v} \leftarrow 0, \boldsymbol{c} \leftarrow 0$
for each round $t=1, \cdots, T$ do
sample subset $\mathcal{S}$ of clients
communicate $(\boldsymbol{x}, \boldsymbol{s}=(\boldsymbol{m}, \boldsymbol{v}))$ and $\boldsymbol{c}$ (only Mime) to all clients $i \in \mathcal{S}$
on client $i \in \mathcal{S}$ in parallel do
initialize local model $\boldsymbol{y}_{i} \leftarrow \boldsymbol{x}$
for $k=1, \cdots, K$ do
sample mini-batch $\zeta$ from local data
$\boldsymbol{g}_{i} \leftarrow \nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)-\nabla f_{i}(\boldsymbol{x} ; \zeta)+\boldsymbol{c}$ (Mime)
$\boldsymbol{g}_{i} \leftarrow \nabla f_{i}\left(\boldsymbol{y}_{i} ; \zeta\right)$ (MimeLite)
update $\boldsymbol{y}_{i} \leftarrow \boldsymbol{y}_{i}-\eta\left(\left(1-\beta_{1}\right) \boldsymbol{g}_{i}+\beta_{1} \boldsymbol{m}\right) /\left(\left(\sqrt{\boldsymbol{v}}+\varepsilon_{0}\right)\left(1-\beta_{1}^{t}\right)\right)$
end for
compute full local-batch gradient $\nabla f_{i}(\boldsymbol{x})$
communicate $\left(\boldsymbol{y}_{i}, \nabla f_{i}(\boldsymbol{x})\right)$
end on client
compute $\boldsymbol{c} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \nabla f_{i}(\boldsymbol{x})$
$\boldsymbol{m} \leftarrow\left(\left(1-\beta_{1}\right) \boldsymbol{c}+\beta_{1} \boldsymbol{m}\right) /\left(1-\beta_{1}^{t}\right)$
$\boldsymbol{v} \leftarrow\left(\left(1-\beta_{2}\right) \boldsymbol{c}^{2}+\beta_{2} \boldsymbol{v}\right) /\left(1-\beta_{2}^{t}\right)$
$\boldsymbol{x} \leftarrow \boldsymbol{x}-\eta_{g} \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}}\left(\boldsymbol{x}-\boldsymbol{y}_{i}\right)$ (update server parameters)
end for

```
Algorithm 5 MimeMVR pseudocode
    input: initial \(\boldsymbol{x}^{0}\), learning rate \(\eta\)
    initialize \(\boldsymbol{c}^{0} \leftarrow 0, \boldsymbol{m}^{0} \leftarrow 0\)
    for each round \(t=1, \cdots, T\) do
        sample subset \(\mathcal{S}\) of clients
        communicate \(\boldsymbol{x}^{t-1}, \boldsymbol{x}^{t-2}, m^{t-1}, c^{t-1}\) to all clients \(i \in \mathcal{S}\)
        on client \(i \in \mathcal{S}\) in parallel do
            initialize local model \(\boldsymbol{y}_{i, 0}^{t} \leftarrow \boldsymbol{x}^{t-1}\)
            for \(k=1, \cdots, K\) do
                sample mini-batch \(\zeta_{i, k}^{t}\) from local data
            compute SVRG gradient \(\boldsymbol{g}_{i, k}^{t} \leftarrow \nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\nabla f_{i}\left(\boldsymbol{x}^{t-1} ; \zeta_{i, k}^{t}\right)+\boldsymbol{c}^{t-1}\)
            compute corrected momentum \(\boldsymbol{d}_{i, k}^{t} \leftarrow a \boldsymbol{g}_{i, k}^{t}+(1-a) \boldsymbol{m}^{t-1}+(1-a)\left(\nabla f_{i}\left(\boldsymbol{y}_{i, k-1}^{t} ; \zeta_{i, k}^{t}\right)-\right.\)
            \(\left.\nabla f_{i}\left(\boldsymbol{x}^{t-1} ; \zeta_{i, k}^{t}\right)\right)\)
            update \(\boldsymbol{y}_{i, k}^{t}=\boldsymbol{y}_{i, k-1}^{t}-\eta \boldsymbol{d}_{i, k}^{t}\)
            end for
            compute full local-batch gradients \(\nabla f_{i}\left(\boldsymbol{x}^{t-1}\right), \nabla f_{i}\left(\boldsymbol{x}^{t-2}\right)\)
            communicate \(\left(\boldsymbol{y}_{i, K}^{t}, \nabla f_{i}\left(\boldsymbol{x}^{t-1}\right), \nabla f_{i}\left(\boldsymbol{x}^{t-2}\right)\right)\)
        end on client
        compute new aggregate pseudo-gradient \(\boldsymbol{c}^{t} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \nabla f_{i}\left(\boldsymbol{x}^{t-1}\right)\)
        compute old aggregate pseudo-gradient \(\tilde{\boldsymbol{c}}^{t} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \nabla f_{i}\left(\boldsymbol{x}^{t-2}\right)\)
        update server momentum \(\boldsymbol{m}^{t} \leftarrow a \boldsymbol{c}^{t}+(1-a) \boldsymbol{m}^{t-1}+(1-a)\left(\boldsymbol{c}^{t}-\tilde{\boldsymbol{c}}^{t}\right)\)
        update server parameters \(\boldsymbol{x}^{t} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \boldsymbol{y}_{i, K}^{t}\)
    end for
```

