

# Network Adaptive Federated Learning: Congestion and Lossy Compression

Parikshit Hegde

*Electrical and Computer Engineering*  
*The University of Texas at Austin*  
Austin, Texas, USA  
hegde@utexas.edu

Gustavo de Veciana

*Electrical and Computer Engineering*  
*The University of Texas at Austin*  
Austin, Texas, USA  
gustavo@ece.utexas.edu

Aryan Mokhtari

*Electrical and Computer Engineering*  
*The University of Texas at Austin*  
Austin, Texas, USA  
mokhtari@austin.utexas.edu

**Abstract**—In order to achieve the dual goals of privacy and learning across distributed data, Federated Learning (FL) systems rely on frequent exchanges of large files (model updates) between a set of clients and the server. As such FL systems are exposed to, or indeed the cause of, congestion across a wide set of network resources. Lossy compression can be used to reduce the size of exchanged files and associated delays, at the cost of adding noise to model updates. By judiciously adapting clients’ compression to varying network congestion, an FL application can reduce wall clock training time. To that end, we propose a Network Adaptive Compression (NAC-FL) policy, which dynamically varies the client’s lossy compression choices to network congestion variations. We prove, under appropriate assumptions, that NAC-FL is asymptotically optimal in terms of directly minimizing the expected wall clock training time. Further, we show via simulation that NAC-FL achieves robust performance improvements with higher gains in settings with positively correlated delays across time.

**Index Terms**—federated learning, rate adaptation, resilience

## I. INTRODUCTION

Communication costs and delays of sending model updates from clients to the server are a known bottleneck in training Federated Learning (FL) systems [1]–[4]. Two common techniques used to alleviate this issue are: 1) *local computations* where clients perform several local steps before communicating with the server, and 2) *(lossy) compression* where clients communicate quantized/compressed updates to the server. The eventual end goal of these approaches is to minimize the *wall clock time* for convergence of the training algorithm (hereon referred to as FL algorithm) by reducing the amount of data communicated from clients to the server.

To this end, several works have analyzed the relationship between compression, local computations and the number of rounds needed by FL algorithms to converge [5]–[12]. However, these works ignore the impact of changing network congestion, *both across clients and across time*, on the wall

This material is based upon work of Hegde and de Veciana supported by the National Science Foundation (NSF) under grant No. 2148224 and is supported in part by funds from OUSD R&E, NIST, and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program and the WNCG/6G@UT industrial affiliates. The work of Mokhtari is supported in part by the NSF AI Institute for Future Edge Networks and Distributed Intelligence (AI-EDGE) via NSF grant 2112471, the Machine Learning Lab (MLL) at UT Austin, and the Wireless Networking and Communications Group (WNCG) Industrial Affiliates Program.

clock time to converge. For instance, a client may choose a high degree of compression when it sees high network congestion, while a client seeing lower congestion may *opportunistically* choose not to compress as much. In this work, we ask the following question: “Can we design a policy that adapts the amount of compression across clients and time according to changing network conditions in order to optimize the wall clock time?” To answer this question, we first characterize the impact that changing network congestion and an adaptive compression policy have on the wall clock time. Second, we propose the Network Adaptive Compression for Federated Learning (NAC-FL) policy that judiciously chooses compression levels based on network congestion to minimize the wall clock time. Crucially, NAC-FL does not rely on the prior knowledge of the distribution of network congestion. Instead, it learns to optimize its compression decisions on-the-fly based on the congestion seen by clients.

NAC-FL works in an opportunistic manner by adaptively choosing high or low amounts of compression across clients and across time based on low or high network congestion. It further considers two effects that compression has on the wall clock time. First, with increasing amount of compression, the FL algorithm would require more communication rounds to converge, as the server receives “noisier”, and hence inaccurate, model updates. Second, with higher degrees of compression, the duration of each round would decrease as a smaller model update is communicated. Since the wall clock time is affected by both the number of rounds and the duration of each round (it is effectively the product of the two quantities), a policy for choosing compression levels should consider these jointly. Fig. 1 provides an illustrative visualization. Hence, NAC-FL aims to find the “sweet-spot” compression levels over time varying network congestion.

**Contributions.** We propose a general framework to study how to best adapt compression of client model updates. Assuming a stationary Markov model for the underlying network congestion state, we show that optimal policies are state dependent and characterize the expected stopping time for convergence to a predefined model accuracy.

This characterization provides the underlying insight for our proposed NAC-FL policy. To our knowledge this is the first




Fig. 1: Illustration of how compression level affects round duration, number of rounds and wall clock time.

policy for compression that adapts to the stochastic variations of the underlying network congestion process. Under appropriate assumptions on the FL algorithm and underlying network congestion and delays, we provide a proof of the asymptotic optimality of NAC-FL in terms of minimizing the mean time until the convergence criterion is met. To our knowledge this is the first theoretical result of this type.

Finally we demonstrate via simulation the performance gains and robustness of NAC-FL vs alternative fixed compression and/or fixed error per round policies. We explore a variety of models for network congestion, finding that in particular NAC-FL excels in the practically relevant setting where the network sees positive correlations in the network congestion across time.

### A. Related Work

Perhaps the most related papers to our work are [13]–[17] which explored adaptive compression schemes for FL settings. In [13]–[15] the authors propose adapting compression to network congestion. In these works, the algorithm to select compression has a per round budget, e.g., a budget on delay (or compression error) per round, and possibly heterogeneous compression levels are chosen across the clients based on the current network congestion to minimize the compression error (or delay) for the round. These works exploit the diversity of network congestion across the clients, but *not across time*. Meanwhile [16], [17] have observed that using a higher amount of compression at the start and gradually reducing compression through time may improve the wall clock time. Our proposed policy is novel in that it learns how to best exploit congestion variation across clients and across time to optimize the wall clock time.

Another line of work that aims to reduce the overall communication cost is client sampling [18]–[21], where at each round, only a subset of the clients are chosen to participate. The authors of [21] propose a client sampling and power control policy that adapts to time varying channels of clients sharing a single base station and optimizes a proxy for wall clock time. Overall we view lossy compression and client sampling as alternative approaches geared at addressing communication bottlenecks. A study of how to jointly adapt lossy compression and client sampling to changing network congestion is left for future work.

### B. Paper Organization

In Section II, we introduce our system model. In Section III, we propose our NAC-FL algorithm for lossy compression and under appropriate assumptions prove it is asymptotically optimal. Section IV is devoted to exploring the method for several problem instances and in particular for various models for the underlying network congestion in terms of correlation across clients and time. In Section V, we comment on the practical aspects of estimating the file transfer delay of clients when deploying NAC-FL. Finally, in Section VI, we close the paper with some concluding remarks.

**Notation.** *Throughout this document, unless otherwise mentioned, quantities denoted with lowercase letters correspond to constants, and uppercase letters correspond to random variables. Bold symbols correspond to vectors, and regular symbols indicate scalars. For example,  $\mathbf{x}$  is a constant vector,  $\mathbf{X}$  is a random vector,  $x$  is a constant scalar, and  $X$  is a random scalar/variable. Lowercase and uppercase forms of the same letter correspond to constant and random variable notions of the same quantity. A sequence indexed by  $n$  will be denoted as  $(x^n)_n$ .*

## II. MODEL SETUP

In this paper, we focus on a federated architecture, where a server aims to find a model that performs well with respect to the data of a group of  $m$  clients, and in which nodes exchange updates based on their local information with only the server. More precisely, suppose the loss function associated with client  $j$  is denoted by  $f_j(\mathbf{w})$ , where  $\mathbf{w}$  represents the weights of the model, e.g., the weights of a neural network. The goal is to find the model that minimizes the average loss across clients

$$f(\mathbf{w}) = \frac{1}{m} \sum_{j=1}^m f_j(\mathbf{w}).$$

The FL algorithm proceeds in rounds. Each round consists of two stages: (i) a local stage in which each client updates the most recent model received from the server via gradient-based updates based on its local data and (ii), an aggregation stage in which the server updates the global model by aggregating the local updates received from clients. We shall let  $\mathbf{w}^n$  denote the global model at the server at round  $n$ . Further, we let  $\tau^n$  denote the total number of local steps (such as gradient steps) that each client performs at round  $n$ , and let  $\mathbf{w}_j^{\tau_j^n, n}$  denote the resulting local model at node  $j$ .

In this paper, we are interested in the setting where each client sends a compressed version  $\tilde{\mathbf{g}}_{Q_j}^n$  of its local model  $\mathbf{w}_j^{\tau_j^n, n}$  to the server using a lossy compression algorithm (or, *compressor*)  $\mathcal{Q}(\cdot, \cdot)$ . The compressor accepts a vector  $\mathbf{x}$  and a parameter  $q \in [0, q_{\max}]$  indicating the amount of compression with the maximum value being  $q_{\max}$ , and outputs  $\tilde{\mathbf{X}} = \mathcal{Q}(\mathbf{x}, q)$  which is an approximation of  $\mathbf{x}$ , but has a decreased file size as compared to  $\mathbf{x}$ .  $\tilde{\mathbf{X}}$  is capitalized to highlight that the compressor  $\mathcal{Q}(\cdot, \cdot)$  may use randomness in its compression. We shall denote by  $q_j^n$  the compression

parameter used by client  $j$  for round  $n$ , and denote by  $\mathbf{q}^n \triangleq (q_j^n)_{j=1}^m$  the vector of parameters used by the clients in round  $n$ . After receiving updates from all the clients, the server aggregates the compressed local models and produces the next global model  $\mathbf{w}^{n+1}$ .

Given a target tolerance  $\varepsilon > 0$ , the goal of FL is to generate a sequence of global models until on some round  $r_\varepsilon$  a prespecified *stopping criterion* is first met, e.g., the norm of the global loss function gradient is at most  $\varepsilon$ , i.e.,  $\|\nabla f(\mathbf{w}^{r_\varepsilon})\| \leq \varepsilon$ . Our goal is to find an adaptive compression policy that dynamically adapts to the possibly time varying network states such that the target accuracy is achieved with a minimum overall wall clock time.

We formalize the *overall wall clock time*, denoted  $t_\varepsilon$ , required to achieve the target accuracy as follows. The duration  $d(\tau^n, \mathbf{q}^n, \mathbf{c}^n)$  of a round  $n$  depends on:

- $\tau^n$ , the number of local computations performed by clients which we will assume to be the same across clients;
- $\mathbf{q}^n$ , an  $m$  dimensional vector of clients' compression parameters ;
- $\mathbf{c}^n$ , the *network state* which models network congestion and is assumed to be an element of a finite set  $\mathcal{C}$ .

This allows some flexibility, e.g., the round's duration may depend on the max delay to deliver the model update from clients to server, or the sum of the delays if clients share a single resource in TDMA (Time Division Multiple Access) fashion. The total wall clock time is then given by

$$t_\varepsilon = \sum_{n=1}^{r_\varepsilon} d(\tau^n, \mathbf{q}^n, \mathbf{c}^n). \quad (1)$$

In our system model, the sequence of network states,  $(\mathbf{c}^n)_n$ , is assumed to be exogenous, i.e., not be controlled by the server or the clients nor their choices of  $\tau^n$  and  $\mathbf{q}^n$ . The delays associated with the server multicasting global models to clients are assumed to be exogeneous i.e., can not be controlled by the FL server/clients and are not compressed, whence are not part of the model. Still, in this work, based on observing the network state we will devise an approach to select the clients compression parameters so as to minimize the wall clock time. As discussed in Section V, in practice observation of the network state may involve light weight in band estimation by probing delays of message bits as they are delivered in a given round.

A policy for choosing compression parameters is called a *state dependent stationary policy* if it can be expressed as a function  $\pi$  of the current network state, i.e.,  $\mathbf{q}^n = \pi(\mathbf{c}^n)$  for all rounds  $n \in \mathbb{N}$ . Such a policy will be referred to simply as policy  $\pi$ . Given a random sequence of network states,  $(\mathbf{C}^n)_n$ , let  $R_\varepsilon^\pi$  be the random variable denoting the minimum number of rounds needed to converge to error tolerance  $\varepsilon$  under policy  $\pi$ . Then, the corresponding wall clock time, denoted by  $T_\varepsilon^\pi$ , is expressed as,

$$T_\varepsilon^\pi = \sum_{n=1}^{R_\varepsilon^\pi} d(\tau^n, \pi(\mathbf{C}^n), \mathbf{C}^n).$$

### III. NETWORK ADAPTIVE COMPRESSION FOR FEDERATED LEARNING (NAC-FL)

Our approach to designing a policy to adapt clients' compression parameters centers on recognizing that the expected wall clock time can be broken up into a product of the *expected number of rounds*  $r_\varepsilon$  needed to converge to an error tolerance  $\varepsilon$  and the *average duration of each round*  $\hat{d}$ . We start by characterizing the relationship between  $r_\varepsilon$ ,  $\hat{d}$ , and the sequence of selected quantization parameters  $(\mathbf{q}^n)_n$  and network states  $(\mathbf{c}^n)_n$  for a given FL algorithm.

Below we state an assumption relating  $r_\varepsilon$  to  $(\mathbf{q}^n)_n$ . To that end we introduce a strictly increasing, continuous and bounded scalar function  $h_\varepsilon : [0, q_{\max}] \rightarrow \mathbb{R}^+$  of compression parameter  $q$  and an associated vector function  $\mathbf{h}_\varepsilon : [0, q_{\max}]^m \rightarrow \mathbb{R}_+^m$  of a compression vector  $\mathbf{q}$  where  $\mathbf{h}_{\varepsilon,j}(\mathbf{q}) = h_\varepsilon(q_j)$ . We let  $\mathbf{h}_\varepsilon^{-1}$  denote the inverse of this vector function.

**Assumption 1.** *For a given FL algorithm there exists a strictly increasing, continuous and bounded function  $h_\varepsilon(q)$  and norm  $\|\cdot\|$  such that given a sequence of compression parameters  $(\mathbf{q}^n)_n$ , the FL algorithm has reached the desired error tolerance  $\varepsilon$  by round  $r$  if and only if,*

$$r > \frac{1}{r} \sum_{n=1}^r \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|$$

for some norm.

The above assumption implies that the expected number of rounds can be written as the average of an increasing function of the sequence of selected quantization parameters. Roughly speaking, given a lossy compression policy that generates a stationary parameter sequence  $(\mathbf{Q}^n)_n$  whose marginal distribution is the same as the random vector  $\mathbf{Q}$ , the above criterion means that the expected number of rounds to converge to the desired error tolerance is approximately  $\mathbb{E}[\|\mathbf{h}_\varepsilon(\mathbf{Q})\|]$ .

This is a general condition that is motivated by convergence bounds of several FL algorithms with compression, including, [5], [8], [11]. In particular in Appendix A, we illustrate this motivation for an extension of the FedCOM algorithm [11], when  $q$  indicates the normalized-variance introduced by the compressor, the scalar function is  $h_\varepsilon(q) = O(\sqrt{q+1}/\varepsilon)$  and the norm is the  $L_2$  norm.

**Assumption 2.** *For any sequence of compression parameters  $(\mathbf{q}^n)_n$  the minimum number of rounds  $r_\varepsilon$  needed to converge to an error tolerance  $\varepsilon$  is such that  $r_\varepsilon = \Theta(1/\text{poly}(\varepsilon))$ , where  $\text{poly}(\varepsilon)$  denotes a polynomial of  $\varepsilon$ .*

Assumption 2 is a natural assumption for gradient based optimization algorithms. It requires the convergence guarantees for the FL algorithm to be such that when we require a more accurate solution, the number of required communication rounds grows. This argument indeed holds even for the settings that we do not exchange compressed signals.

We also make the following additional assumption about the round duration function.




Fig. 2: Illustration of a round duration as a function of compression parameter  $q$  for a fixed local computation  $\tau$  and network state  $c$ .

**Assumption 3.** Given a network state  $c$ , number of local computations  $\tau$ , and compression parameters  $\mathbf{q} = \mathbf{h}_\varepsilon^{-1}(\mathbf{r})$ , the round duration  $d(\tau, \mathbf{q}, c) = d(\tau, \mathbf{h}_\varepsilon^{-1}(\mathbf{r}), c)$  is bounded, convex in  $\mathbf{r}$  and decreasing in every coordinate of  $\mathbf{r}$ .

In Assumption 3, the round duration being decreasing in  $\mathbf{r}$  is reasonable, since we expect more rounds as well as smaller file sizes with higher compression. The convexity is motivated by the notion that we use a “good compressor” as illustrated next. Consulting Fig. 2, for any two parameters  $q_1, q_2$  and  $0 < \alpha < 1$ , a new time-sharing compressor  $\mathcal{Q}'$  may be derived which outputs  $\mathcal{Q}(\mathbf{x}, q_1)$  with probability  $\alpha$  and outputs  $\mathcal{Q}(\mathbf{x}, q_2)$  with probability  $(1-\alpha)$ . This compressor has expected round duration  $\alpha d(\tau, q_1, c) + (1-\alpha)d(\tau, q_2, c)$ . And, in certain cases, its compression parameter is  $q_\alpha = \alpha q_1 + (1-\alpha)q_2$  (such as when the stochastic quantizer parameterized by its normalized variance [5] is used). If  $\mathcal{Q}$  is a “good compressor”, then its round duration,  $d(\tau, q_\alpha, c)$ , should be lower compared to that of the simple time-shared compressor,  $\alpha d(\tau, q_1, c) + (1-\alpha)d(\tau, q_2, c)$ . Therefore, the convexity of the round duration function is a reasonable assumption for “good compressors” (considering  $h_\varepsilon(q) \propto q$  for simplicity).

**Assumption 4.** The sequence of network states  $(C^n)_n$  forms an irreducible aperiodic stationary Markov Chain on a finite state space  $\mathcal{C}$  with invariant distribution  $\mu$ .

Assumption 4 is a natural assumption made to facilitate the analysis of algorithms (see e.g., [22]).

#### A. Expected Wall Clock Time Formulation

Given the above mentioned assumptions, we are now ready to introduce the proposed framework. We begin by showing that we need only consider state dependent stationary policies for choosing compression parameters when optimizing the overall wall clock time.

**Lemma 1.** Under Assumptions 1-4 there exists a state dependent stationary policy to select compression parameters which

is asymptotically optimal in terms of minimizing the wall clock time to reach a desired error tolerance of  $\varepsilon$  as  $\varepsilon \rightarrow 0$ .

The proof of Lemma 1 depends on two critical observations. First, since by Assumption 2 the number of rounds needed to converge grows large as  $\varepsilon \rightarrow 0$ , one can expect the empirical distribution of the network states modelled by the finite state Markov Chain to concentrate around the invariant prior to the stopping time. Second, due to the convexity of the round duration function in Assumption 3, given a sequence of network states there exists a state dependent stationary policy that is near optimal and depends solely on the empirical distribution of the sequence. The proof is in Appendix C.

Here, we will focus on the setting where  $\varepsilon$  is small, hence by Lemma 1, we only need to consider state dependent stationary policies,  $\mathbf{q}^n = \boldsymbol{\pi}(c^n)$ .

**Lemma 2.** Under Assumptions 1-4 and a fixed number of local computations per round  $\tau$ , for every  $\delta > 0$ , there exists an  $\varepsilon_{th} > 0$  such that, for all  $\varepsilon < \varepsilon_{th}$  and any state-dependent stationary policy  $\boldsymbol{\pi}$ , the expected wall clock time is bounded as,

$$1 - \delta \leq \frac{\mathbb{E}[T_\varepsilon^\boldsymbol{\pi}]}{\mathbb{E}[\|\mathbf{h}_\varepsilon(\boldsymbol{\pi}(C))\|] \mathbb{E}[d(\tau, \boldsymbol{\pi}(C), C)]} \leq 1 + \delta, \quad (2)$$

where,  $C$  denotes a random variable whose distributions is  $\mu$  (see Assumption 3).

Lemma 2 is proved in Appendix D. Define,

$$\hat{t}_\varepsilon^\boldsymbol{\pi} \triangleq \mathbb{E}[\|\mathbf{h}_\varepsilon(\boldsymbol{\pi}(C))\|] \mathbb{E}[d(\tau, \boldsymbol{\pi}(C), C)]. \quad (3)$$

Due to Lemma 2, for small enough  $\varepsilon$ ,  $\hat{t}_\varepsilon^\boldsymbol{\pi}$  provides an accurate approximation for  $\mathbb{E}[T_\varepsilon^\boldsymbol{\pi}]$ . Therefore, from here onwards, we shall assume implicitly that a small  $\varepsilon$  is considered and focus on finding a policy to optimize  $\hat{t}_\varepsilon^\boldsymbol{\pi}$ .

Suppose the distribution of  $C$  is known. Then, one could compute expected wall clock time as given in (3) for any state dependent stationary policy  $\boldsymbol{\pi}$ . In this case, we could determine an optimal policy  $\boldsymbol{\pi}^*$  by solving the optimization problem,

$$\min_{\boldsymbol{\pi} \in \mathcal{Q}_{m|c}} \hat{t}_\varepsilon^\boldsymbol{\pi} = \mathbb{E}[\|\mathbf{h}_\varepsilon(\boldsymbol{\pi}(C))\|] \mathbb{E}[d(\tau, \boldsymbol{\pi}(C), C)], \quad (4)$$

where  $\mathcal{Q}_{m|c}$  is the set of all state-dependent stationary policies.

Alas, in practice, we often cannot directly solve the above problem, as the distribution of  $C$  is unknown. Hence, below, we propose a stochastic approximation like algorithm that achieves the optimal wall clock time of  $\boldsymbol{\pi}^*$  asymptotically.

#### B. NAC-FL: Informal Description

The idea underlying NAC-FL is to keep running estimates for  $\mathbb{E}[\|\mathbf{h}_\varepsilon(\mathbf{Q})\|]$  and  $\mathbb{E}[d(\tau, \mathbf{Q}, C)]$  i.e.,

$$\hat{r}_\varepsilon^n = \frac{1}{n} \sum_{k=1}^n \|\mathbf{h}_\varepsilon(\mathbf{q}^{(k)})\|, \quad \hat{d}^n = \frac{1}{n} \sum_{k=1}^n d(\tau, \mathbf{q}^{(k)}, c^{(k)}).$$

Given a network state of  $c^{n+1}$  at round  $n+1$ , and, a possible choice for compression parameters  $\mathbf{q}$ , the running averages would be updated as follows,

$$\begin{aligned}\hat{r}_\varepsilon^{n+1} &= \frac{n}{n+1}\hat{r}_\varepsilon^n + \frac{1}{n+1}\|\mathbf{h}_\varepsilon(\mathbf{q})\|, \\ \hat{d}^{n+1} &= \frac{n}{n+1}\hat{d}^n + \frac{1}{n+1}d(\tau, \mathbf{q}, c^{n+1}).\end{aligned}\quad (5)$$

As seen in (3), to minimize the wall clock time one should minimize  $\hat{r}_\varepsilon^{n+1}\hat{d}^{n+1}$ , which can be expanded as,

$$\begin{aligned}\hat{r}_\varepsilon^{n+1}\hat{d}^{n+1} &= \frac{n}{(n+1)^2}\left[r_\varepsilon^n d(\tau, \mathbf{q}, c^{n+1}) + \hat{d}^n \|\mathbf{h}_\varepsilon(\mathbf{q})\|\right] \\ &\quad + \frac{n^2}{(n+1)^2}r_\varepsilon^n \hat{d}^n + O\left(\frac{1}{(n+1)^2}\right).\end{aligned}$$

Given the fact that  $\hat{r}_\varepsilon^n$  and  $\hat{d}^n$  are constants, and neglecting the term  $O(1/(n+1)^2)$ , an optimal choice for  $\mathbf{q}^{n+1}$  is

$$\mathbf{q}^{n+1} = \underset{\mathbf{q}}{\operatorname{argmin}} \quad \hat{r}_\varepsilon^n d(\tau, \mathbf{q}, c^{n+1}) + \hat{d}^n \|\mathbf{h}_\varepsilon(\mathbf{q})\|. \quad (6)$$

The NAC-FL policy is summarized in Algorithm 1. To retrieve policy informally described above the tunable parameters  $(\beta_n)_n$  and  $\alpha$  should be set to  $\beta_n = \frac{1}{n}$  and  $\alpha = 1$ .

Consider two possible network states  $c$  and  $c'$  at a round  $n$ . If the delay under state  $c$  is higher compared to  $c'$  for any compression parameters, then NAC-FL would choose a higher compression amount  $\mathbf{q}$  for state  $c$  compared to compression amount  $\mathbf{q}'$  for state  $c'$ , i.e.,  $\mathbf{q} > \mathbf{q}'$  elementwise. This may be concluded from the selection policy of (6), and noting that  $r_\varepsilon(\mathbf{q})$  is increasing in  $\mathbf{q}$  (Assumption 1), and  $d(\tau, \mathbf{q}, c)$  is decreasing in  $\mathbf{q}$  (Assumption 3).

Observe that since the estimates  $\hat{r}_\varepsilon^n$  and  $\hat{d}^n$  will initially change across rounds, NAC-FL may choose different compression parameters in two rounds for which the network was in the same state, i.e., NAC-FL is not a state-dependent stationary policy. Still, we will show NAC-FL is asymptotically near optimal. To develop this result we shall next present NAC-FL in a more formal manner.

---

#### Algorithm 1: NAC-FL

---

**Input** : Initialization:  $\hat{r}_\varepsilon^{(0)}, \hat{d}^{(0)}$ ; step size schedule  $\{\beta_n\}_{n=1}^\infty$ ; parameter  $\alpha$ .

- 1 **for**  $n = 1, \dots$ , **until termination do**
- 2     Server observes network state  $c^n$ ;
- 3      $\mathbf{q}^n =$   
 $\quad \underset{\mathbf{q}}{\operatorname{argmin}} \quad \alpha \hat{r}_\varepsilon^{(n-1)} d(\tau, \mathbf{q}, c^n) + \hat{d}^{(n-1)} \|\mathbf{h}_\varepsilon(\mathbf{q})\|;$
- 4      $\hat{r}_\varepsilon^n = (1 - \beta_n)\hat{r}_\varepsilon^{(n-1)} + \beta_n \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|$ ;
- 5      $\hat{d}^n = (1 - \beta_n)\hat{d}^{(n-1)} + \beta_n d(\tau, \mathbf{q}^n, c^n)$ ;
- 6 **end**

---

### C. NAC-FL: Formal Description

Our NAC-FL approach is also inspired by the Frank-Wolfe Algorithm [23]. We start by reformulating the optimization

program in (4). Denote by set  $V_\varepsilon$  all possible pairs of expectations  $(\hat{r}_\varepsilon, \hat{d})$ ,

$$\begin{aligned}V_\varepsilon = \left\{ (\hat{r}_\varepsilon, \hat{d}) : \exists \boldsymbol{\pi} \in \mathcal{Q}_{m|c|} \text{ s.t. } \hat{r}_\varepsilon = \mathbb{E}[\|\mathbf{h}_\varepsilon(\boldsymbol{\pi}(C))\|], \right. \\ \left. \hat{d} = \mathbb{E}[d(\tau, \boldsymbol{\pi}(C), C)] \right\}.\end{aligned}\quad (7)$$

Using the set  $V_\varepsilon$ , and denoting  $H(r, d) \triangleq rd$ , we may write the optimization (4) characterizing the optimal policy  $\boldsymbol{\pi}^*$  as

$$\min_{\hat{r}_\varepsilon, \hat{d}} \{H(\hat{r}_\varepsilon, \hat{d}) : (\hat{r}_\varepsilon, \hat{d}) \in V_\varepsilon\}. \quad (8)$$

In this case, from a point  $(\hat{r}_\varepsilon^n, \hat{d}^n)$ , the Frank-Wolfe update would be given as,

$$\begin{aligned}(\hat{r}_\varepsilon, \hat{d}) &= \underset{(r, d) \in V_\varepsilon}{\operatorname{argmin}} \quad \nabla H(\hat{r}_\varepsilon^n, \hat{d}^n)^\top \begin{pmatrix} r \\ d \end{pmatrix}, \\ \hat{r}_\varepsilon^{n+1} &= (1 - \beta)\hat{r}_\varepsilon^n + \beta\hat{r}_\varepsilon, \\ \hat{d}^{n+1} &= (1 - \beta)\hat{d}^n + \beta\hat{d}.\end{aligned}\quad (9)$$

The gradient  $\nabla H(\hat{r}_\varepsilon, \hat{d})$  is,  $\nabla H(\hat{r}_\varepsilon, \hat{d}) = \begin{pmatrix} \hat{d} & \hat{r}_\varepsilon \end{pmatrix}^\top$ .  $V_\varepsilon$  is a set of feasible averages of  $\hat{r}_\varepsilon$  and  $\hat{d}$ . Therefore, at round  $(n+1)$ , not all the pairs  $(r, d) \in V_\varepsilon$  may be achievable. Hence, NAC-FL approximates equation (9) as,

$$\mathbf{q}^{n+1} = \underset{\mathbf{q}}{\operatorname{argmin}} \quad \hat{r}_\varepsilon^n d(\tau, \mathbf{q}, c^{n+1}) + \hat{d}^n \|\mathbf{h}_\varepsilon(\mathbf{q})\|.$$

We have thus retrieved our proposed NAC-FL algorithm based on the Frank-Wolfe update, with one difference. The above derivation suggests the use of a fixed step-size  $\beta$  at all rounds while the previously derived algorithm used a decaying the step-size  $\beta_n = 1/n$ . In our simulations, we will embrace the latter.

The following assumption is required to show the asymptotic optimality of NAC-FL. A state dependent stationary policy  $\boldsymbol{\pi}$  maps from a domain of finite size  $|\mathcal{C}|$ , to a range positive-real vectors of dimension  $m$ . Therefore, the policy may be represented by a positive-real vector,  $\boldsymbol{\pi}$ , of dimension  $m|\mathcal{C}|$ . Further, a vector  $\mathbf{r}^\boldsymbol{\pi}$  may be obtained by applying  $h_\varepsilon(\cdot)$  elementwise to the policy vector  $\boldsymbol{\pi}$ ,  $\mathbf{r}^\boldsymbol{\pi} \triangleq \mathbf{h}_\varepsilon(\boldsymbol{\pi})$ . This representation is used in the following assumption.

**Assumption 5.** *The objective function  $\hat{t}_\varepsilon^\boldsymbol{\pi}$  of the optimization problem in (4) is a strictly quasiconvex function in  $\boldsymbol{\pi}$  in the following sense,*

$$\mathbf{r}^{\boldsymbol{\pi}^\top} (\nabla_{\mathbf{r}^\boldsymbol{\pi}} \hat{t}_\varepsilon^\boldsymbol{\pi}) = 0 \implies \mathbf{r}^{\boldsymbol{\pi}^\top} (\nabla_{\mathbf{r}^\boldsymbol{\pi}}^2 \hat{t}_\varepsilon^\boldsymbol{\pi}) \mathbf{r}^\boldsymbol{\pi} > 0. \quad (10)$$

Assumption 5 ensures that there is a unique state dependent stationary policy  $\boldsymbol{\pi}^*$  which optimizes (4). We have observed that the considered network model, compression model and the  $\|\mathbf{h}_\varepsilon(\mathbf{q})\|$  function associated with the FedCOM algorithm indeed satisfy this assumption.

Next we shall establish an optimality property for NAC-FL. To that end we shall consider executing NAC-FL without termination with  $\beta_n = \beta$  for all  $n$  and let  $(\mathbf{Q}_\beta^n)_n, \hat{R}_{\varepsilon, \beta}^n$  and

$\hat{D}_\beta^n$  be the corresponding sequence of compression parameters and the associated estimates.

**Theorem 1.** *Let  $\pi^*$  be the solution and  $\hat{t}_\varepsilon^*$  the minimum of the optimization problem in (4). If Assumptions 1-5 hold, then there exists a positive sequence  $(\beta_i)_{i=1}^\infty$  with  $\beta_i \rightarrow 0$  as  $i \rightarrow \infty$ , such that for every  $\rho > 0$ , there exists a threshold  $n_{th}(\rho)$  such that,*

$$\lim_{i \rightarrow \infty} \sup_{n \geq n_{th}(\rho)/\beta_i} P \left( \left\| \begin{pmatrix} \hat{R}_{\varepsilon, \beta_i}^n - \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|] \\ \hat{D}_{\beta_i}^n - \mathbb{E}[d(\tau, \pi^*(C), C)] \end{pmatrix} \right\| > \rho \right) = 0,$$

The proof of Theorem 1 is included in Appendix B.

**Remark 1.** *Theorem 1 should be interpreted with some subtlety. Say the desired error-tolerance  $\varepsilon$  is very small such that the number of rounds needed to converge under any compression policy is such that  $r_\varepsilon \gg n_{th}(\rho)/\beta$ . Then, based on Theorem 1, one can show that NAC-FL compression choices will be near optimal after  $n_{th}(\rho)/\beta$  rounds. Thereafter, since  $r_\varepsilon$  is large, NAC-FL will make near optimal choices for long enough leading to a near optimal expected wall clock time.*

We further remark on the meaning of the asymptotic result in the context of minimizing the wall clock time. In applications that require a very low error-tolerance  $\varepsilon$ , one needs to have a large number (i.e., in the asymptotic region) of communication rounds  $r_\varepsilon$  for convergence. Therefore, even though the wall clock time obtained by using NAC-FL may be large in this setting, it is near-optimal compared to other methods of choosing compression parameters.

#### IV. SIMULATION

In this section, we present our simulation results. We begin by describing additional model details used in our simulations.

##### A. Additional Model Details

1) *Compression Model:* We shall use the *stochastic quantizer* in [5] which we will denote as  $\mathcal{Q}_q(\cdot, b)$ . The quantizer has a parameter  $b \in \{1, \dots, 32\}$  corresponding to the number of bits used to represent each co-ordinate, in addition to the bit used to denote signs. When input a vector  $\mathbf{x}$ , it outputs,

$$\mathcal{Q}_q(\mathbf{x}, b) = \|\mathbf{x}\|_\infty \text{sign}(\mathbf{x})\zeta(\mathbf{x}, b) \quad (11)$$

where  $\text{sign}(\mathbf{x})$  is the element-wise sign operator and where the function  $\zeta(\mathbf{x}, b)$  uniformly quantizes each co-ordinate amongst  $2^b - 1$  levels between 0 and 1. That is, if  $x_i/\|\mathbf{x}\|_\infty \in \left[ \frac{l}{2^b-1}, \frac{l+1}{2^b-1} \right)$ , then it is quantized as,

$$\zeta_i(\mathbf{x}, b) = \begin{cases} \frac{l+1}{2^b-1}, & \text{with prob. } \frac{|x_i|}{\|\mathbf{x}\|_\infty} (2^b - 1) - l, \\ \frac{l}{2^b-1}, & \text{otherwise.} \end{cases}$$

When  $\mathbf{x}$  is quantized to  $b$ -bits per co-ordinate, its file size is given by the function,  $s(b) = \|\mathbf{x}\|_0 (b + 1) + 32$  bits. Here, the zero-norm,  $\|\mathbf{x}\|_0$ , gives the length of the vector, the 1 indicates the bit used to denote the sign, and the 32 bits are

for a floating point number denoting the norm,  $\|\mathbf{x}\|_\infty$ . Finally, if client  $j$  uses the parameter  $b_j$ , then the vector of parameters used by the clients is denoted as,  $\mathbf{b} = (b_j)_{j=1}^m$ .

2) *Network Congestion Model:* For purposes of evaluating the performance of various algorithms over different types of network congestion we propose the following general, albeit idealized, model. We let  $\mathbf{C}^n$  be a  $m$  dimensional random vector denoting the *Bit Transmission Delay* (BTD) for clients during round  $n$ . We further let  $\mathbf{C}^n = \exp(\mathbf{Z}^n)$  i.e., coordinate-wise exponentiation of an  $m$  dimensional first order autoregressive process given by  $(\mathbf{Z}_i)_{i=0}^\infty$  where  $\mathbf{Z}_0 = \mathbf{0}$ , where

$$\mathbf{Z}^n = A \mathbf{Z}^{(n-1)} + \mathbf{E}^n, \quad \forall n \geq 1, \quad (12)$$

where  $A$  is an  $m \times m$  deterministic matrix, and  $\mathbf{E}^n \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$  are i.i.d.,  $m$  dimensional normal random vectors. Different correlations across time and clients may be modelled by varying  $A$ ,  $\boldsymbol{\mu}$  and  $\Sigma$ . The marginal distributions of  $\mathbf{C}^n$  are thus log-normal but can be correlated in different ways based on the underlying autoregressive process. In particular:

**Homogeneous Independent:** the parameters are set to  $A = 0$ ,  $\boldsymbol{\mu} = \mathbf{1}$ , and  $\Sigma = \sigma^2 I$ . This results in a process which is independent and identically distributed across clients and time.

**Heterogeneous Independent:** the parameters are set to  $A = 0$ ,  $\mu_i = 0$  for  $i \in \{1, \dots, 5\}$  and  $\mu_i = 2$  for  $i \in \{6, \dots, 10\}$ , and  $\Sigma = I$ . This results in a process which is independent across clients and time, with the BTD being lower for the first 5 clients compared to the rest.

**Perfectly correlated:** the parameters are set to  $A$  such that  $A_{i,j} = \frac{a}{m}$  where  $a \in (0, 1)$ ,  $\boldsymbol{\mu} = \mathbf{0}$ , and  $\Sigma$  such that  $\Sigma_{i,j} = \sigma^2 = 1$ . This results in a process where all clients see the same positively correlated time-varying delays.

**Partially correlated:** the parameters are set to  $A$  such that  $A_{i,j} = \frac{a}{m}$ ,  $\boldsymbol{\mu} = \mathbf{0}$ , and  $\Sigma$  such that  $\Sigma_{i,i} = 1$  and  $\Sigma_{i,j} = 1/2$  for  $i \neq j$ . This results in a process where delays are positively correlated across clients and time.

3) *Model for Round Durations:* We will model the duration of a round as the maximum across clients' delays, i.e.,

$$d(\tau, \mathbf{b}, c) = \max_j [\theta\tau + c_j s(b_j)],$$

where  $\theta$  represents the compute time per local computation, and  $c_j s(b_j)$  the BTD of client  $j$  times the size of the client  $j$ 's file capturing the time taken to communicate its update. For simplicity we will set  $\theta = 0$ .

4) *Compression Level Choice Policies:* We compare NAC-FL to the following policies,

a) *Fixed Bit:* Here, a number  $b$  is fixed, and all the clients use the stochastic quantizer  $\mathcal{Q}_q(\mathbf{x}, b)$  from (11) with the parameter  $b$ . We present results for  $b \in \{1, 2, 3\}$ , as we didn't notice a performance improvement for larger parameters in our experiments.

b) *Fixed Error:* This method was suggested in [13] and is parameterized by a number  $q$ . At each round  $n$ , the parameters  $\mathbf{b}^n$  of the stochastic quantizers are such that the average normalized-variance  $\bar{q}^n$  (see equation (15)) is smaller than

$q$ , and the duration of the round  $d(\tau, \mathbf{q}^n, \mathbf{c}^n)$  is minimized. We fix  $q = 5.25$  in all our experiments after finding it to be performing well across different settings.

5) *Machine Learning Model*: We consider  $m = 10$  clients. We consider the MNIST dataset [24] which may be distributed homogeneously or heterogeneously amongst the clients. Since data is heterogeneous across clients in most FL applications, we consider the heterogenous data case. That is, each client has data corresponding to 1 unique label. The MNIST dataset has 60,000 training samples, 10,000 test samples and 10 labels. The clients and the server aim to train a fully connected neural network with the architecture (784, 250, 10) with the sigmoid activation for the hidden layer. The learning rate is initialized to  $\eta_0 = 0.07$ , and is decayed by a factor 0.9 every 10 rounds. The aggregation rate and local computations per round are fixed throughout the training to  $\gamma = 1$  and  $\tau = 2$  respectively. As for the parameters of the NAC-FL policy, we set  $\beta_n = \frac{1}{n}$ , and  $\alpha = 2$ .

We measure the performance of the global model using the following,

a) *Training Loss*: The training loss of the global model is the empirical cross entropy loss across the entire set of training samples.

b) *Test Accuracy*: The test accuracy is measured over all the test samples. Here, in some experiments, we run 20 simulations with different random seeds, and report the mean, 90th percentile and 10th percentile times to reach a test accuracy of 90%. The 90th and 10th percentile scores are reported to capture the variation in performance across the 20 simulations. We also report a *gain* metric, which is sample mean of the time gained to reach 90% accuracy by NAC-FL compared to another policy reported in percentage. For instance, let  $x_i, y_i$  be the times under NAC-FL and another policy for a random seed  $i$ , then the gain is  $100 * \left( \frac{\sum_{i=1}^{20} y_i / x_i - 1}{20} \right)$ .

## B. Simulation Results

1) *Homogeneous Independent BTD*: We simulated over  $\sigma^2 \in \{1, 2, 3\}$  in order to study the change in performance over increasing variance. We observe that in all the cases, NAC-FL and the Fixed Error policy have very similar performance across all the considered statistics. This is because the Fixed Error policy was designed to operate well in the i.i.d., network delay case. However, both NAC-FL and Fixed Error policy perform better than all the Fixed Bit policies according to all the statistics across all the considered parameters. Moreover, we observed that the gap in the performance to Fixed Bit policies increased with increasing variance. For instance, the gain of the best Fixed Bit policy increased from 145% to 250% when the variance was increased from 1 to 3, while the gain of the worst fixed bit policy increased from 314% to 881%. This is as expected because both NAC-FL and Fixed Error policy adapt to the heterogenous delay of clients at any given time. Surprisingly, NAC-FL lagged behind Fixed Error policy in some metrics, but it performed better in terms of the gain metric in all the 3 cases, with the gain over Fixed Error policy ranging from 1% to 8%.

$\sigma^2$		1 bit	2 bits	3 bits	Fixed Error	NAC-FL
1	Mean	6.31	3.82	4.15	<b>1.58</b>	1.60
	90th	6.95	4.72	5.00	<b>1.86</b>	2.05
	10th	5.63	3.20	3.38	1.20	<b>1.14</b>
	Gain	314%	145%	168%	3%	-
2	Mean	54.8	32.5	34.9	12.5	<b>12.2</b>
	90th	70.6	44.7	43.1	<b>19.0</b>	20.8
	10th	42.5	19.2	21.0	6.26	<b>5.82</b>
	Gain	522%	216%	240%	8%	-
3	Mean	799	430	458	<b>165</b>	168
	90th	1430	752	665	<b>318</b>	320
	10th	418	157	148	<b>46.2</b>	57.9
	Gain	881%	270%	250%	1%	-

TABLE I: Performance comparison of policies with homogeneous independent BTD in terms of the mean, 90th percentile and 10th percentile times to reach 90% test accuracy under the different policies, and their average sample-path gain compared to NAC-FL. All the numbers represented are in  $10^7$  seconds.

2) *Heterogeneous Independent BTD*: We considered this case since the first 5 clients would have consistently worse delay, NAC-FL and the Fixed Error policy would consistently compress the updates of those clients heavily. Since the data distribution is heterogeneous, it may be possible heavy compression of updates from specific clients throughout the training may hurt the performance. On the other hand, the Fixed Bit policies use the same amount of compression across all clients equally irrespective of their delays. Still, we observed that NAC-FL and the Fixed Error policy perform better than the Fixed Bit policies as can be seen in Table II. In fact, performance in terms of the gain metric is very comparable to the i.i.d., network delay case with  $\sigma^2 = 1$  in Table I.

	1 bit	2 bits	3 bits	Fixed Error	NAC-FL
Mean	9.49	5.85	6.46	2.49	<b>2.48</b>
90th	11.5	7.16	8.09	<b>3.48</b>	3.54
10th	8.30	4.37	4.98	1.74	<b>1.54</b>
Gain	319%	146%	173%	4%	-

TABLE II: Performance comparison of policies with heterogeneous independent BTD. The numbers shown are the mean, 90th percentile and 10th percentile times to reach 90% test accuracy under the different policies, and their average sample-path gain compared to NAC-FL. All the numbers represented are in  $10^8$  seconds.

3) *Perfectly Correlated BTD*: We will demonstrate that NAC-FL performs better than Fixed Error and Fixed Bit policies under increasing correlated delay across time since they are not designed to optimize the wall clock time under this case.

To study the variation of network delay across rounds, consider the marginal auto-regressive process of 1 client which may be represented by the following scalar autoregressive process,

$$Z^n = a'Z^{(n-1)} + E^n, \quad (13)$$

where  $E^n \sim \mathcal{N}(0, 1)$ . We define metric called *asymptotic*

variance, denoted  $\sigma_\infty^2$ , which is designed to capture the variance, and long and short term correlations of a random process,

$$\sigma_\infty^2 \triangleq \lim_{n \rightarrow \infty} \frac{\mathbb{E} \left[ (Z^{(1)} + \dots + Z^{(n)})^2 \right]}{n}. \quad (14)$$

For the autoregressive process in (13), it may be computed to be,  $\sigma_\infty^2 = 1/(1 - a')^2$ .

Table III shows the performance of the different policies under varying asymptotic variance of the marginals. We observe that in addition to beating the baseline fixed bit policies on all the metrics, the NAC-FL performs better than the Fixed Error policy in most metrics as well. Considering the gain metric, we observe gain of 13% over the Fixed Error policy for low asymptotic variance of  $\sigma_\infty^2 = 1.56$ , and is as large as 27% for higher asymptotic variance of  $\sigma_\infty^2 = 4$ . Notably, in terms of the 10th percentile time to reach 90% accuracy, the Fixed Error policy required 40%, 23% and 32% more time compared to NAC-FL in the  $\sigma_\infty^2=1.56, 4$  and 16 cases respectively.

$\sigma_\infty^2$		1 bit	2 bits	3 bits	Fixed Error	NAC-FL
1.56	Mean	5.14	3.04	3.47	2.21	<b>2.11</b>
	90th	5.94	3.65	4.43	<b>2.66</b>	3.32
	10th	3.88	2.38	2.18	1.43	<b>1.02</b>
	Gain	191%	58%	75%	13%	-
4	Mean	5.82	3.49	4.03	2.47	<b>2.23</b>
	90th	7.43	4.77	6.28	<b>3.94</b>	4.00
	10th	3.88	2.22	1.98	1.21	<b>0.981</b>
	Gain	252%	82%	107%	27%	-
16	Mean	8.42	5.19	6.15	3.75	<b>3.36</b>
	90th	12.8	10.3	13.4	7.94	<b>7.2</b>
	10th	4.34	1.40	1.67	1.15	<b>0.87</b>
	Gain	316%	72%	98%	21%	-

TABLE III: Performance comparison of policies with perfectly correlated BTD in terms of the mean, 90th percentile and 10th percentile times to reach 90% test accuracy under the different policies, and their average sample-path gain compared to NAC-FL. All the numbers represented are in  $10^7$  seconds.

4) *Partially Correlated BTD*: In Table IV, we show results for the partially correlated BTD case with asymptotic variance  $\sigma_\infty^2 = 4$ . We consider this case to demonstrate that NAC-FL is effective with positive (but, not 100%) correlation across clients as well. Indeed, we observe NAC-FL performing better compared to all the other policies across all the considered metrics, with a gain of 10% over the Fixed Error policy, and 129% over the best fixed bit policy. Notably, in terms of the 10th percentile and 90th percentile metrics, NAC-FL outperformed Fixed Error policy by 30% and 15% respectively.

Figure 3 contains sample path plots of Training Loss and Accuracy vs Wall Clock Time for the independent homogeneous ( $\sigma^2 = 2$ ), heterogeneous and perfectly correlated ( $\sigma_\infty^2 = 4$ ) BTD cases. Both accuracy and loss plots for NAC-FL and Fixed Error are overlapping in the independent homogeneous and heterogeneous BTD cases, as expected. However, in the perfectly correlated BTD case, NAC-FL dominates the performance of Fixed Error policy.

In summary, we observe that NAC-FL’s performance is robust under a range of network models considered. NAC-FL

	1 bit	2 bits	3 bits	Fixed Error	NAC-FL
Mean	13.6	8.33	9.51	4.22	<b>3.83</b>
90th	15.9	10.5	13.9	6.24	<b>5.46</b>
10th	9.51	5.47	5.80	2.64	<b>2.02</b>
Gain	307%	129%	159%	10%	-

TABLE IV: Performance comparison of policies with partially correlated BTD in terms of the mean, 90th percentile and 10th percentile times to reach 90% test accuracy under the different policies, and their average sample-path gain compared to NAC-FL. All the numbers represented are in  $10^7$  seconds.

vastly outperformed the baseline Fixed Bit policies in all the network models. The performance of NAC-FL was observed to be similar to that of Fixed Error policy in the independent BTD setting, albeit, it outperformed Fixed Error policy in terms of the gain metric under all the network models. Notably, the gap between NAC-FL and Fixed Error policy was observed to be noticeably high in the perfectly and partially correlated BTD settings, where NAC-FL was able to adapt to positive correlations of BTD across time, whereas Fixed Error could not.

## V. NAC-FL IN PRACTICE

In this section we briefly comment on some practical aspects underlying estimating model update delays. This involves estimating the network’s current average BTD to each client. A simple approach to doing so is to observe that for the stochastic quantizer described in Section IV-A1, clients always send the vector of signs of their updates, no matter what are the bits per coordinate that will be chosen. So, as the clients send their signs, the server may probe the delay characteristics to estimate the BTD of clients without having to request vacuous (non update related) bits to do so. It may then use these estimates to perform the optimization in (6) for the round.

## VI. CONCLUSION

Due to their distributed character FL algorithms are exposed to congestion across a potentially large number of network resources, whence one might say they are exposed to network congestion and variability at scale. Building adaptive algorithms that minimize the impact of time varying congestion across clients presents a significant challenge, particularly when the aim is to directly optimize the expected wall clock time. NAC-FL exemplifies a new class of robust algorithms to optimally adapt clients’ lossy compression. This paper further provides the technical roadmap to formalizing and showing asymptotic optimality for such algorithms.

## REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [2] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.



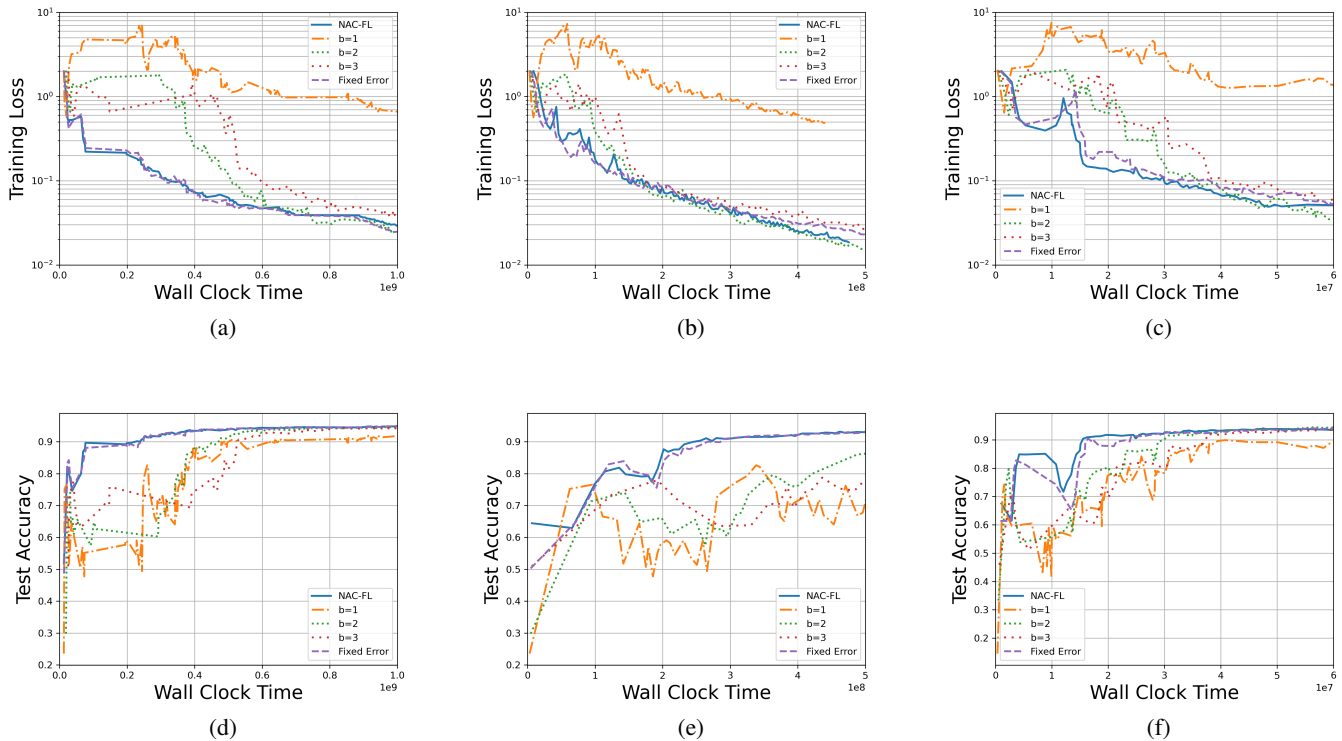


Fig. 3: Plots of Training Loss and Test Accuracy vs Wall Clock time on different network models. Figures (a) and (d) correspond to homogeneous independent BTM case ( $\sigma^2 = 2$ ), Figures (b) and (e) correspond to the heterogeneous independent BTM case, and Figures (c) and (f) correspond to the perfectly correlated BTM case ( $\sigma_\infty^2 = 4$ ).

[3] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, “Advances and open problems in federated learning,” *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[4] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.

[5] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, “QSGD: Communication-efficient SGD via gradient quantization and encoding,” *Advances in neural information processing systems*, vol. 30, 2017.

[6] D. Basu, D. Data, C. Karakus, and S. Diggavi, “Qsparse-local-SGD: Distributed SGD with quantization, sparsification and local computations,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[7] S. U. Stich, “Local SGD converges fast and communicates little,” *arXiv preprint arXiv:1805.09767*, 2018.

[8] S. U. Stich, J.-B. Cordonnier, and M. Jaggi, “Sparsified SGD with memory,” *Advances in Neural Information Processing Systems*, vol. 31, 2018.

[9] S. Ghadimi and G. Lan, “Stochastic first-and zeroth-order methods for nonconvex stochastic programming,” *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2341–2368, 2013.

[10] A. Reiszadeh, A. Mokhtari, H. Hassani, A. Jadbabaie, and R. Pedarsani, “Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2021–2031.

[11] F. Haddadpour, M. M. Kamani, A. Mokhtari, and M. Mahdavi, “Federated learning with compression: Unified analysis and sharp guarantees,” *arXiv preprint arXiv:2007.01154*, 2020.

[12] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, “Robust and communication-efficient federated learning from non-iid data,” *IEEE transactions on neural networks and learning systems*, vol. 31, no. 9, pp. 3400–3413, 2019.

[13] P. S. Bouzinis, P. D. Diamantoulakis, and G. K. Karagiannidis, “Wireless quantized federated learning: A joint computation and communication design,” *arXiv preprint arXiv:2203.05878*, 2022.

[14] X. Zhang, X. Zhu, J. Wang, H. Yan, H. Chen, and W. Bao, “Federated learning with adaptive communication compression under dynamic bandwidth and unreliable networks,” *Information Sciences*, vol. 540, pp. 242–262, 2020.

[15] H. Sun, X. Ma, and R. Q. Hu, “Adaptive federated learning with gradient compression in uplink NOMA,” *IEEE Transactions on Vehicular Technology*, vol. 69, no. 12, pp. 16 325–16 329, 2020.

[16] D. Jhunjunwala, A. Gadhikar, G. Joshi, and Y. C. Eldar, “Adaptive quantization of model updates for communication-efficient federated learning,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 3110–3114.

[17] R. Hönl, Y. Zhao, and R. Mullins, “DAdaQuant: Doubly-adaptive quantization for communication-efficient federated learning,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 8852–8866.

[18] W. Chen, S. Horvath, and P. Richtárik, “Optimal client sampling for federated learning,” *arXiv preprint arXiv:2010.13723*, 2020.

[19] M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, “A joint learning and communications framework for federated learning over wireless networks,” *IEEE Transactions on Wireless Communications*, vol. 20, no. 1, pp. 269–283, 2020.

[20] M. Ribero and H. Vikalo, “Communication-efficient federated learning via optimal client sampling,” *arXiv preprint arXiv:2007.15197*, 2020.

[21] J. Perazzone, S. Wang, M. Ji, and K. S. Chan, “Communication-efficient device scheduling for federated learning using stochastic optimization,” in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 1449–1458.

[22] A. L. Stolyar, “On the asymptotic optimality of the gradient scheduling algorithm for multiuser throughput allocation,” *Operations research*, vol. 53, no. 1, pp. 12–25, 2005.

[23] M. Frank and P. Wolfe, “An algorithm for quadratic programming,” *Naval research logistics quarterly*, vol. 3, no. 1-2, pp. 95–110, 1956.

[24] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.

- [25] G. Sparling, "Honors calculus notes." [Online]. Available: <http://www.math.pitt.edu/~sparling/23014/23014notes6/node24.html>
- [26] P. Billingsley, *Convergence of probability measures*. John Wiley & Sons, 2013.
- [27] S. N. Ethier and T. G. Kurtz, *Markov processes: characterization and convergence*. John Wiley & Sons, 2009.
- [28] A. El Gamal and Y.-H. Kim, *Network information theory*. Cambridge university press, 2011.
- [29] K.-M. Chung, H. Lam, Z. Liu, and M. Mitzenmacher, "Chernoff-Hoeffding bounds for markov chains: Generalized and simplified," *arXiv preprint arXiv:1201.0559*, 2012.
- [30] D. A. Levin and Y. Peres, *Markov chains and mixing times*. American Mathematical Soc., 2017, vol. 107.

APPENDIX A  
FEDERATED LEARNING WITH ADAPTIVE COMPRESSION (FLAC)

In this section, we consider a variant of the FedCOM algorithm [11], which we will call FedCOM-V. FedCOM is based on fixing a quantization parameter throughout run of the FL algorithm. On the other hand, FedCOM-V allows for an arbitrary sequence of quantization parameters  $(q^n)_n$ , in order to account for adaptive compression policies such as NAC-FL. FedCOM-V is presented in Algorithm 2.

---

**Algorithm 2:** FedCOM-V

---

**Input** : number of local computations schedule  $(\tau_n)_{n=1}^\infty$ , local learning rate schedule  $(\eta_n)_{n=1}^\infty$ , adaptively chosen global learning rate schedule  $(\gamma_n)_{n=1}^\infty$ , adaptively chosen number of rounds  $r$ , initial global model  $w^1$ .

**Result:**  $w^{r+1}$ : Final model

```

1 for  $n = 1, \dots, r$  do
2   for each client  $j \in [m]$  do
3     Set  $w_j^{1,n} = w^n$ ;
4     for  $a = 1, \dots, \tau_n$  do
5       Sample a minibatch  $\mathcal{Z}_j^{a,n}$  and compute  $\tilde{g}_j^{a,n} \triangleq \nabla f(w_j^{a,n}; \mathcal{Z}_j^{a,n})$ ;
6        $w_j^{a+1,n} = w_j^{a,n} - \eta_n \tilde{g}_j^{a,n}$ ;
7     end
8     Device sends  $\tilde{g}_{Qj}^n = \mathcal{Q}((w^n - w_j^{\tau_n+1,n})/\eta_n, q_j^n)$  back to the server;
9   end
10  Server computes,  $\tilde{g}_Q^n = \frac{1}{m} \sum_{j=1}^m \tilde{g}_{Qj}^n$ ;
11  Server computes  $w^{n+1} = w^n - \eta_n \gamma_n \tilde{g}_Q^n$  and broadcasts to all devices;
12 end

```

---

In order to study the convergence properties of FedCOM-V, we make the following standard assumptions.

**Assumption 6** (Smoothness and Lower Boundedness). *The objective function  $f(\cdot)$  is differentiable and  $L$ -smooth. That is,  $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\|$ , for every  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ . Moreover, the optimal value of  $f$  is lower bounded,  $f^* = \min_w f(w) > -\infty$ .*

**Assumption 7** (Bounded Variance). *For all clients  $j \in [m]$  and rounds  $n$  and local step  $a$ , we can sample an independent minibatch  $\mathcal{Z}_j^{a,n}$  of size  $|\mathcal{Z}_j^{a,n}| = b$  and compute an unbiased stochastic gradient  $\tilde{g}_j^{a,n} = \nabla f(w; \mathcal{Z}_j^{a,n})$ , i.e.,  $\mathbb{E}_{\mathcal{Z}_j^{a,n}}[\tilde{g}_j] = \nabla f(w_j^{a,n})$ . Moreover, the variance is bounded by a constant  $\sigma^2$ , i.e.,  $\mathbb{E}_{\mathcal{Z}_j^{a,n}}[\|\tilde{g}_j^{a,n} - \nabla f(w_j^{a,n})\|^2] \leq \sigma^2$ .*

**Assumption 8** (Compression Model). *The output of the compressor  $\mathcal{Q}(\mathbf{x}, q)$  is an unbiased estimator of  $\mathbf{x}$ , i.e.,  $\mathbb{E}[\mathcal{Q}(\mathbf{x}, q)|\mathbf{x}] = \mathbf{x}$ , and, its variance is bounded as,  $\mathbb{E}[\|\mathcal{Q}(\mathbf{x}, q) - \mathbf{x}\|^2 | \mathbf{x}] \leq q \|\mathbf{x}\|^2$ .*

We denote the maximum normalized-variance by  $q_{\max}$  and the average normalized-variance used at round  $n$  by

$$\bar{q}^n = \frac{1}{m} \sum_{j=1}^m q_j^n. \quad (15)$$

The following Theorem states the relationship between  $(q^n)_n$ ,  $\varepsilon$  and  $r_\varepsilon$  and is proved in Appendix F.

**Theorem 2.** *Let Algorithm 2 be run with a sequence of compressors such that the average normalized-variance at round  $n$  is  $\bar{Q}^n$ . Further, assume that the sequence  $(\bar{Q}^n)_n$  forms a stationary process with the stationary distribution represented by a random variable  $Q$ . To obtain  $\mathbb{E}[\|\nabla f(w)\|^2] \leq \varepsilon$ , we can choose,*

$$r_\varepsilon = O\left(\log(1/\varepsilon) \frac{\mathbb{E}[\sqrt{Q+1}]}{\varepsilon}\right), \quad \tau^n = O(n).$$

The upper bound on  $r_\varepsilon$  in Theorem 2 provides a justification for Assumption 1 with  $h_\varepsilon(q) = O(\sqrt{q+1})$ . Here,  $\tau^n$  is a function of  $n$ , but for the purposes of NAC-FL we may use the average of  $\tau^{(1)}$  to  $\tau^{(r_\varepsilon)}$  in the expression of the duration function. One may obtain a similar expression for other popular FL algorithms [5], [8].

APPENDIX B  
PROOF OF THEOREM 1

In this section we show that NAC-FL converges to the optimal solution asymptotically as  $\beta \downarrow 0$ .

In order to consider the effect of  $\beta \downarrow 0$  on NAC-FL estimates  $\hat{R}_\varepsilon^n$  and  $\hat{D}^n$  in (9), we shall denote these as  $\hat{R}_{\varepsilon,\beta}^n$  and  $\hat{D}_\beta^n$  respectively. Let  $\text{conv}(V_\varepsilon)$  be the convex hull of the set  $V_\varepsilon$  defined in (7). Recall the positive sequence  $(\beta_i)_i$  with  $\beta_i \rightarrow 0$  from the statement of Theorem 1. Letting  $\mathbf{X}_\beta^n \triangleq (\hat{R}_{\varepsilon,\beta}^n \hat{D}_\beta^n)^\top$ , and  $H(\mathbf{x}) \triangleq x_1 x_2$  over the domain  $\mathbb{R}_+^2$ , we have the following result.

**Proposition B.1.** *Let the initialization  $\mathbf{X}_\beta^0$  be equal to  $\mathbf{x}^0 \in \mathbb{R}_+^2$  almost surely for any  $0 < \beta < 1$ , then, for any  $s > 0$ ,  $\lim_{i \rightarrow \infty} \mathbf{X}_{\beta_i}^{\lfloor s/\beta_i \rfloor}$  exists, is almost surely deterministic and denoted as  $\mathbf{x}(s) \triangleq \lim_{i \rightarrow \infty} \mathbf{X}_{\beta_i}^{\lfloor s/\beta_i \rfloor}$ . Further, for any  $\mathbf{x}^0 \in \mathbb{R}_+^2$ ,  $\mathbf{x}(s)$  obeys the following differential equation,*

$$\begin{aligned} \mathbf{x}(0) &= \mathbf{x}^0, \\ \dot{\mathbf{x}}(s) &= \mathbf{v}(s) - \mathbf{x}(s), \quad s > 0, \\ \mathbf{v}(s) &= \underset{\mathbf{v} \in \text{conv}(V_\varepsilon)}{\text{argmin}} \quad \nabla H(\mathbf{x}(s))^\top \mathbf{v}, \quad s > 0. \end{aligned} \tag{16}$$

The proof of Proposition B.1 is very similar to that of the main result of [22]. For the sake of completeness, we briefly show the proof at the end of this section. From hereon, (16) will be referred to as the Fluid-Frank-Wolfe (FFW) process.

**Proposition B.2.** *Under Assumption 5, the FFW process in (16) has a unique fixed point  $\mathbf{x}^* \in \text{conv}(V_\varepsilon)$  such that,*

$$\mathbf{x}^* = \underset{\mathbf{x} \in \text{conv}(V_\varepsilon)}{\text{argmin}} \quad \nabla H(\mathbf{x}^*)^\top \mathbf{x}.$$

Moreover,  $\mathbf{x}^* \in V_\varepsilon$ , and  $\mathbf{x}^*$  is the minimizer of  $H$  over the set  $V_\varepsilon$ .

Proposition B.2 is proved in Appendix E.

Denote,  $G(\mathbf{x}) = \min_{\mathbf{v} \in \text{conv}(V_\varepsilon)} \nabla H(\mathbf{x})^\top (\mathbf{v} - \mathbf{x})$ . Since,  $\nabla H$  is a continuous function of  $\mathbf{x}$ ,  $G(\mathbf{x})$  is a continuous function of  $\mathbf{x}$  as well.

As a consequence of Proposition B.2, there exists a unique point  $\mathbf{x}^* \in \text{conv}(V)_\varepsilon$  such that  $G(\mathbf{x}^*) = 0$ . For all other  $\mathbf{x} \in \text{conv}(V_\varepsilon)$ ,  $G(\mathbf{x}) < 0$ . We will, in fact, prove a stronger result that  $G(\mathbf{x})$  is bounded away from 0 for points that are a distance away from  $\mathbf{x}^*$ .

**Claim 1:** for any  $\omega > 0$ , there exists a  $\xi > 0$  such that if  $\mathbf{x} \in \text{conv}(V_\varepsilon)$  and  $\|\mathbf{x} - \mathbf{x}^*\| \geq \omega$ , then  $G(\mathbf{x}) < -\xi$ .

*Proof.* We prove this claim by contradiction. Suppose there exists an  $\omega > 0$  such that for all  $\xi > 0$ , the set,

$$\begin{aligned} \mathcal{X}^\xi &\triangleq \{\mathbf{x}^\xi : \mathbf{x}^\xi \in \text{conv}(V_\varepsilon), \|\mathbf{x}^\xi - \mathbf{x}^*\| \geq \omega \text{ and } G(\mathbf{x}^\xi) \geq -\xi\}, \\ &= \text{conv}(V_\varepsilon) \cap \{\mathbf{x}^\xi : \|\mathbf{x}^\xi - \mathbf{x}^*\| \geq \omega\} \cap \{\mathbf{x}^\xi : G(\mathbf{x}^\xi) \geq -\xi\}, \end{aligned}$$

is non-empty.

$\text{conv}(V_\varepsilon)$  is a compact set because it is the convex hull of a compact set,  $V_\varepsilon$ . Further, the sets  $\{\mathbf{x}^\xi : \|\mathbf{x}^\xi - \mathbf{x}^*\| \geq \omega\}$  and  $\{\mathbf{x}^\xi : G(\mathbf{x}^\xi) \geq -\xi\}$  are also closed because they are the pre-image of continuous functions over closed sets. Therefore,  $\mathcal{X}^\xi$  is a closed set since it is the intersection of three closed sets. Further, it is also bounded because  $\text{conv}(V_\varepsilon)$  is bounded. Therefore,  $\mathcal{X}^\xi$  is a compact set.

Consider  $\xi_1 > \xi_2 > 0$ . Since,  $G(\mathbf{x}) \geq -\xi_2$  implies that  $G(\mathbf{x}) \geq -\xi_1$ , we have that  $\mathcal{X}^{\xi_1} \supset \mathcal{X}^{\xi_2}$ . Consider a decreasing sequence  $(\xi_i)_{i \in \mathbb{N}}$  with  $\lim_{i \rightarrow \infty} \xi_i = 0$ . Then,  $(\mathcal{X}^{\xi_i})_{i \in \mathbb{N}}$  is a decreasing sequence of compact and non-empty sets. We know that a decreasing sequence of non empty compact sets has a limit, and the limit is non-empty [25]. Therefore,

$$\mathcal{X}^0 \triangleq \bigcap_{i=1}^{\infty} \mathcal{X}^{\xi_i},$$

exists and is non-empty. Since  $\xi_i \downarrow 0$ , this means that any  $\mathbf{x} \in \mathcal{X}^0$  satisfies  $G(\mathbf{x}) = 0$ . Since  $\|\mathbf{x} - \mathbf{x}^*\| \geq \omega$  and  $\mathbf{x} \in \text{conv}(V_\varepsilon)$  for any  $\mathbf{x} \in \mathcal{X}^0$ , this is a contradiction to the fact that  $\mathbf{x}^*$  is a unique point in  $\text{conv}(V_\varepsilon)$  with  $G(\mathbf{x}) = 0$ . Therefore, there must exist some  $\xi > 0$  for which  $\mathcal{X}^\xi$  is empty.  $\square$

Next we proceed to study the asymptotic convergence of the process  $\mathbf{x}(\cdot)$ . Note that since  $V_\varepsilon$  is apriori unknown, the initialization  $\mathbf{x}^0$  may not be in the set  $V_\varepsilon$ . Nevertheless, the FFW process  $\mathbf{x}(\cdot)$  eventually reaches the set  $\text{conv}(V_\varepsilon)$ . In order to formalize this, let  $\text{conv}^\zeta(V_\varepsilon)$  denote the  $\zeta$ -thickening of the set  $\text{conv}(V_\varepsilon)$ ,

$$\text{conv}^\zeta(V_\varepsilon) = \{\mathbf{y} : \exists \mathbf{x} \in \text{conv}(V_\varepsilon) \text{ such that } \|\mathbf{y} - \mathbf{x}\|_2 \leq \zeta\}.$$

**Proposition B.3.** Consider the FFW process defined in (16). For every  $\zeta > 0$ , there exists an  $s_\zeta > 0$  such that,  $\mathbf{x}(s) \in \text{conv}^\zeta(V_\varepsilon)$  for all  $s > s_\zeta$ .

The proof is the same as that of Corollary 2 in [22].

**Claim 2:**  $\mathbf{x}(s) \rightarrow \mathbf{x}^*$  as  $s \rightarrow \infty$ .

*Proof.* First we prove that  $\liminf_{s \rightarrow \infty} \|\mathbf{x}(s) - \mathbf{x}^*\| = 0$  by contradiction. As a contradiction assume that there exists an  $\omega > 0$  and  $s^\omega > 0$  such that  $\|\mathbf{x}(s) - \mathbf{x}^*\| > \omega$  for all  $s > s^\omega$ .

Let  $\xi > 0$  be the constant according to Claim 1 which ensures that  $G(\mathbf{x}) < -\xi$  for all  $\mathbf{x}$  in  $\text{conv}(V_\varepsilon)$  satisfying  $\|\mathbf{x} - \mathbf{x}^*\| > \omega$ . Moreover, due to continuity of  $G(\cdot)$ , there exists a  $\xi' > 0$  and a small enough  $\zeta > 0$  such that  $G(\mathbf{x}) < -\xi'$  for all  $\mathbf{x}$  in  $\text{conv}^\zeta(V_\varepsilon)$  that satisfy  $\|\mathbf{x} - \mathbf{x}^*\| \geq \omega$ . Due to Proposition B.3, there exists a constant  $s_\zeta > 0$  such that  $\mathbf{x}(s) \in \text{conv}^\zeta(V_\varepsilon)$  for all  $s > s_\zeta$ .

Define,  $s_*^\omega = s^\omega + s_\zeta + (H(\mathbf{x}(s_\zeta + s^\omega)) + 1)/\xi'$ . Then,

$$\begin{aligned} H(\mathbf{x}(s_*^\omega)) &= H(\mathbf{x}(s_\zeta + s^\omega)) + \int_{s_\zeta + s^\omega}^{s_*^\omega} dH(\mathbf{x}(s)), \\ &= H(\mathbf{x}(s_\zeta + s^\omega)) + \int_{s_\zeta + s^\omega}^{s_*^\omega} \nabla H(\mathbf{x}(s))^\top \dot{\mathbf{x}}(s) ds, \\ &= H(\mathbf{x}(s_\zeta + s^\omega)) + \int_{s_\zeta + s^\omega}^{s_*^\omega} \nabla H(\mathbf{x}(s))^\top (\mathbf{v}(s) - \mathbf{x}(s)) ds, \\ &= H(\mathbf{x}(s_\zeta + s^\omega)) + \int_{s_\zeta + s^\omega}^{s_*^\omega} G(\mathbf{x}(s)) ds, \\ &< H(\mathbf{x}(s_\zeta + s^\omega)) + \int_{s_\zeta + s^\omega}^{s_*^\omega} -\xi' ds, \\ &= H(\mathbf{x}(s_\zeta + s^\omega)) - H(\mathbf{x}(s_\zeta + s^\omega)) - 1 < 0. \end{aligned}$$

Since,  $H$  is a positive function, this is a contradiction. Therefore, there exists a time  $s > s^\omega + s_\zeta$  such that  $\|\mathbf{x}(s) - \mathbf{x}^*\| < \omega$ . Since this is true for every  $\omega > 0$  and  $s^\omega > 0$ , we have proved that  $\liminf_{s \rightarrow \infty} \|\mathbf{x}(s) - \mathbf{x}^*\| = 0$ .

Next we prove that  $\lim_{s \rightarrow \infty} \mathbf{x}(s) = \mathbf{x}^*$ . Define,

$$H^\omega = \max_{\substack{\mathbf{x} \in \text{conv}^\zeta(V_\varepsilon) \\ \|\mathbf{x} - \mathbf{x}^*\| \leq \omega}} H(\mathbf{x}).$$

Since  $\liminf_{s \rightarrow \infty} \|\mathbf{x}(s) - \mathbf{x}^*\| = 0$ , there exists a constant  $s_{th}^\omega > s_\zeta$  such that  $\|\mathbf{x}(s_{th}^\omega) - \mathbf{x}^*\| \leq \omega$ . Due to Proposition B.3, for all  $s > s_{th}^\omega$ , we have  $\mathbf{x}(s) \in \text{conv}^\zeta(V_\varepsilon)$ . Therefore, if for any  $s > s_{th}^\omega$ ,  $H(\mathbf{x}(s)) > H^\omega$  is true, then  $\mathbf{x}(s)$  satisfies  $\mathbf{x}(s) \in \text{conv}^\zeta(V_\varepsilon)$  and  $\|\mathbf{x}(s) - \mathbf{x}^*\| > \omega$ . Therefore, due to Claim 1 at all such points, the gradient satisfies,  $dH(\mathbf{x}(s))/ds = G(\mathbf{x}(s)) < 0$ . This implies that  $H(\mathbf{x}(s)) \leq H^\omega$  for all  $s > s_{th}^\omega$ .

Moreover, by the continuity of  $H(\cdot)$ ,  $H^\omega \rightarrow H(\mathbf{x}^*)$  as  $\omega \downarrow 0$ . And, by definition of the minimum  $\mathbf{x}^*$ ,  $H(\mathbf{x}(s)) \geq H(\mathbf{x}^*)$  for any  $s > 0$ . Therefore, by the Sandwich Theorem,  $\lim_{s \rightarrow \infty} H(\mathbf{x}(s)) = H(\mathbf{x}^*)$ .

Further, by the continuity of  $H(\cdot)$  and the uniqueness of the minimum  $\mathbf{x}^*$ ,  $\lim_{s \rightarrow \infty} H(\mathbf{x}(s)) = H(\mathbf{x}^*)$  implies that  $\lim_{s \rightarrow \infty} \mathbf{x}(s) = \mathbf{x}^*$ .  $\square$

Claim 2 proves that the Fluid-Frank-Wolfe process converges to the optimal solution  $\mathbf{x}^*$  asymptotically. In particular, for any  $\rho > 0$ , there exists an  $n_{th}(\rho) > 0$  such that,

$$\sup_{s > n_{th}(\rho)} \|\mathbf{x}(s) - \mathbf{x}^*\| \leq \rho.$$

Denote,  $\mathbf{x}_\beta(s) = \mathbf{X}_\beta^{\lfloor s/\beta \rfloor}$ . Then, since the functions converge as follows,  $(\mathbf{x}_{\beta_i}) \rightarrow \mathbf{x}$  as  $i \rightarrow \infty$ , from the Continuous Mapping Theorem [26, Theorem 2.7], we have,

$$\lim_{i \rightarrow \infty} \sup_{s > n_{th}(\rho)} P(\|\mathbf{x}_{\beta_i}(s) - \mathbf{x}^*\| > \rho) = 0.$$

The above implies the Theorem statement,

$$\lim_{i \rightarrow \infty} \sup_{n > n_{th}(\rho)/\beta_i} P(\|\mathbf{X}_{\beta_i}^n - \mathbf{x}^*\| > \rho) = 0.$$

A. Proof of Proposition B.1

Define the ‘‘scaled process’’ as,  $\mathbf{x}_\beta(s) \triangleq \mathbf{X}_\beta^{\lfloor s/\beta \rfloor}$ . Denote  $D_{\mathbb{R}^2}[0, \infty)$  as the set of functions with domain  $[0, \infty)$ , range  $\mathbb{R}^2$ , and which are right continuous with left limits. Observe that  $\mathbf{x}_\beta$  has sample paths in  $D_{\mathbb{R}^2}[0, \infty)$  for any  $0 < \beta < 1$ .

Denote,  $\mathbf{V}_\beta^n \triangleq \begin{bmatrix} \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\| \\ d(\tau, \mathbf{q}^n, \mathbf{c}) \end{bmatrix}$ , which is the *action* taken by the NAC-FL algorithm (Algorithm 1) at round  $n$ , and  $\mathbf{v}_\beta(s) \triangleq \mathbf{V}_\beta^{\lfloor s/\beta \rfloor}$ . Defining,

$$K = \max \left( \mathbf{x}^0, \max_{q \in [0, q_{\max}], \mathbf{C} \in \mathcal{C}} \left\| \begin{bmatrix} \|\mathbf{h}_\varepsilon(\mathbf{q})\| \\ d(\tau, \mathbf{q}, \mathbf{C}) \end{bmatrix} \right\| \right),$$

by the update rule of NAC-FL,  $\mathbf{x}_\beta(s) = (1 - \beta) \mathbf{x}_\beta(s - \beta) + \beta \mathbf{v}_\beta(s - \beta)$ , we have  $\mathbf{x}_\beta(s) \leq K$  for any  $0 < \beta < 1$  and  $s > 0$ . Further, rearranging the NAC-FL update rule as,

$$\mathbf{x}_\beta(s) - \mathbf{x}_\beta(s - \beta) = \beta (\mathbf{v}_\beta(s - \beta) - \mathbf{x}_\beta(s - \beta))$$

we obtain,  $\|\mathbf{x}_\beta(s) - \mathbf{x}_\beta(s - \beta)\| \leq 2\beta K$ . More generally, for any  $s_1, s_2 > 0$ , we have,

$$\|\mathbf{x}_\beta(s_1) - \mathbf{x}_\beta(s_2)\| \leq 2K \max(\beta, |s_1 - s_2|).$$

This implies the ‘‘asymptotic Lipschitz’’ property,

$$\lim_{\beta \rightarrow 0} \|\mathbf{x}_\beta(s_1) - \mathbf{x}_\beta(s_2)\| \leq 2K |s_1 - s_2|, \quad \forall s_1, s_2 > 0.$$

Then, by Corollary 7.4 in Chapter 3 of [27], the set of stochastic processes  $\{\mathbf{x}_\beta(\cdot)\}_{0 < \beta < 1}$  is *relatively compact*. Therefore, there exists a sequence  $(\beta_i)_i$  with  $\beta_i \rightarrow 0$  as  $i \rightarrow \infty$  such that  $\mathbf{x}_{\beta_i}(\cdot) \rightarrow \mathbf{x}(\cdot)$  as  $i \rightarrow \infty$  for some stochastic process  $\mathbf{x}(\cdot)$  with sample paths in  $D_{\mathbb{R}^2}[0, \infty)$ .

Next, we need to prove that  $\mathbf{x}(\cdot)$  behaves according to (16). To do so, observe that due to the ‘‘continuity property’’ (i.e.,  $\|\mathbf{X}_\beta^n - \mathbf{X}_\beta^{n-1}\| \leq 2K\beta$ ), for any  $\delta > 0$ , there exists a small enough  $\beta > 0$  and  $\Delta > 0$  such that, for any integer  $n$  in the range  $[s/\beta, (s + \Delta)/\beta]$ , we have,

$$|(\nabla H(\mathbf{X}_\beta^n))^\top \mathbf{V}_\beta^n - Y_{\mathbf{C}^n}^*| \leq \delta,$$

where,

$$Y_{\mathbf{C}}^* \triangleq \min_{q \in [0, q_{\max}]} (\nabla H(\mathbf{x}_\beta(s)))^\top \begin{bmatrix} \|\mathbf{h}_\varepsilon(\mathbf{q})\| \\ d(\tau, \mathbf{q}, \mathbf{C}) \end{bmatrix}, \quad \mathbf{C} \in \mathcal{C}.$$

The above equations say that the optimal action at any round in the considered range is very close to the optimal action at the start of the range, for an appropriate selection of parameters. Summing across  $n$  in the range  $[s/\beta, (s + \Delta)/\beta]$  we obtain,

$$\left| \sum_{s/\beta \leq n \leq (s+\Delta)/\beta} (\nabla H(\mathbf{X}_\beta^n))^\top \mathbf{V}_\beta^n - \sum_{s/\beta \leq n \leq (s+\Delta)/\beta} Y_{\mathbf{C}^n}^* \right| \leq \delta \Delta / \beta.$$

Multiplying by  $\beta$  on both sides, from the definition of the scaled process, we have,

$$\left| \int_s^{s+\Delta} (\nabla H(\mathbf{x}_\beta(\xi)))^\top \mathbf{v}_\beta(\xi) d\xi - \sum_{s/\beta \leq n \leq (s+\Delta)/\beta} \beta Y_{\mathbf{C}^n}^* \right| \leq \delta \Delta.$$

From the Law of Large Numbers for Markov Chains, we have  $\lim_{\beta \rightarrow 0} \sum_{s/\beta \leq n \leq (s+\Delta)/\beta} \beta Y_{\mathbf{C}^n}^* = \Delta \sum_{\mathbf{C} \in \mathcal{C}} \mu(\mathbf{C}) Y_{\mathbf{C}}^*$ . Similar to the convergence of  $\mathbf{x}_\beta$  shown above, one can prove convergence of  $\mathbf{v}_\beta$  to a process  $\mathbf{v}$ . Therefore, taking limit  $i \rightarrow \infty$  along the sequence  $(\beta_i)_i$ , we get,

$$\left| \int_s^{s+\Delta} (\nabla H(\mathbf{x}(\xi)))^\top \mathbf{v}(\xi) d\xi - \Delta \sum_{\mathbf{C} \in \mathcal{C}} \mu(\mathbf{C}) Y_{\mathbf{C}}^* \right| \leq \delta \Delta.$$

Observe that  $\sum_{\mathbf{C} \in \mathcal{C}} \mu(\mathbf{C}) Y_{\mathbf{C}}^* = \min_{\mathbf{v} \in \text{conv}(V_\varepsilon)} (\nabla H(\mathbf{x}(s)))^\top \mathbf{v}$ . Therefore, by choosing a  $\Delta$  small enough, we get,

$$\left| (\nabla H(\mathbf{x}(s)))^\top \mathbf{v}(s) - \min_{\mathbf{v} \in \text{conv}(V_\varepsilon)} (\nabla H(\mathbf{x}(s)))^\top \mathbf{v} \right| \leq \delta.$$

Since  $\delta$  can also be chosen arbitrarily small, we have,

$$(\nabla H(\mathbf{x}(s)))^\top \mathbf{v}(s) = \min_{\mathbf{v} \in \text{conv}(V_\varepsilon)} (\nabla H(\mathbf{x}(s)))^\top \mathbf{v}.$$

APPENDIX C  
PROOF OF LEMMA 1

In this section we show that a state-dependent stationary policy asymptotically optimizes the wall clock time. To do so, we first define the notion of a *type* for sequences of network states and compression parameters. Then, we show that for a given network state sequence, a policy for choosing compression parameters which depends on the sequence type optimizes the wall clock type. Finally, because the type asymptotically concentrates for markov processes, we show that a state-dependent stationary policy asymptotically optimizes the wall clock time.

We start by defining the notion of an empirical distribution, called *type*, and its associated expectation and conditional expectations.

**Definition 1** (Type). *The type of a finite sequence,  $x^{[r]} \triangleq (x^n)_{n=1}^r$  with elements in domain  $\mathcal{X}$ , is a function,  $\hat{p}(\cdot; x^{[r]}): \mathcal{X} \rightarrow [0, 1]$ , defined as,*

$$\hat{p}(x; x^{[r]}) = \frac{\sum_{n=1}^r \mathbb{1}(x^n = x)}{r}, \quad \forall x \in \mathcal{X},$$

where  $\mathbb{1}(x = y) = 1$  if  $x = y$ , and 0 otherwise.

Similarly, the *conditional type* and the *joint type* are defined as follows.

**Definition 2** (Joint Type and Conditional Type). *The joint type of two finite sequences,  $x^{[r]}$  and  $y^{[r]}$  with domains  $\mathcal{X}$  and  $\mathcal{Y}$  respectively, is a function,  $\hat{p}(\cdot; x^{[r]}, y^{[r]}): \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ , defined as,*

$$\hat{p}(x, y; x^{[r]}, y^{[r]}) = \frac{\sum_{n=1}^r \mathbb{1}(x^n = x, y^n = y)}{r}, \quad \forall x \in \mathcal{X}, y \in \mathcal{Y}.$$

The *conditional type*  $\hat{p}(\cdot; x^{[r]}, y^{[r]}): \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$  is defined as,

$$\begin{aligned} \hat{p}(x|y; x^{[r]}, y^{[r]}) &= \frac{\sum_{n=1}^r \mathbb{1}(x^n = x, y^n = y)}{\sum_{n=1}^r \mathbb{1}(y^n = y)}, \quad \forall x \in \mathcal{X}, y \in \mathcal{Y} \text{ such that } \hat{p}(y; y^{[r]}) > 0, \\ &= \frac{\hat{p}(x, y; x^{[r]}, y^{[r]})}{\hat{p}(y; y^{[r]})}. \end{aligned}$$

Then, the expectation and conditional expectation with respect to the type may be defined as follows.

**Definition 3** (Expectation and Conditional Expectation). *The expectation of a non-negative function  $g: \mathcal{X} \rightarrow \mathbb{R}^+$  with respect to type  $\hat{p}(\cdot; x^{[r]})$  is defined as<sup>1</sup>,*

$$\hat{\mathbb{E}}[g(X); x^{[r]}] \triangleq \sum_{x \in \mathcal{X}} g(x) \hat{p}(x; x^{[r]}),$$

where  $X$  denotes a random variable with distribution  $\hat{p}(x; x^{[r]})$ . Similarly, the conditional expectation of a non-negative function  $l: \mathcal{X} \rightarrow \mathbb{R}^+$  with respect to the type  $\hat{p}(\cdot; x^{[r]}, y^{[r]})$  is defined as,

$$\hat{\mathbb{E}}[l(X)|Y = y; x^{[r]}, y^{[r]}] \triangleq \sum_{x \in \mathcal{X}} l(x) \hat{p}(x|y; x^{[r]}, y^{[r]}), \quad \forall y \in \mathcal{Y} \text{ such that } \hat{p}(y; y^{[r]}) > 0,$$

where the random variable pair  $(X, Y)$  has joint distribution  $\hat{p}(x, y; x^{[r]}, y^{[r]})$ .

**Proposition C.1.** *Suppose Assumptions 1 and 3 hold, and let  $(c^n)_n$  denote an observed sequence of network states and  $(q^n)$  denote a sequence of compression parameters. Then, for any positive integer  $r$  and positive  $\varepsilon$  with the associated function  $h_\varepsilon(\cdot)$  defined in Assumption 1, there exists a sequence dependent, state dependent stationary policy  $\pi$  such that,*

$$\sum_{n=1}^r \|h_\varepsilon(q^n)\| \geq \sum_{n=1}^r \|h_\varepsilon(\pi(c^n))\|, \quad (17)$$

and,

$$\sum_{n=1}^r d(\tau, q^n, c^n) \geq \sum_{n=1}^r d(\tau, \pi(c^n), c^n). \quad (18)$$

<sup>1</sup>If  $\mathcal{X}$  is uncountably infinite, then,  $\sum_{x \in \mathcal{X}} g(x) \triangleq \sup \{ \sum_{x \in \mathcal{F}} g(x) : \mathcal{F} \subset \mathcal{X}, \mathcal{F} \text{ is finite} \}$ .

*Proof of Proposition C.1.* Given the sequence  $(\mathbf{q}^n, c^n)_{n=1}^r$ , we obtain the joint type  $p(\cdot, \cdot; \mathbf{q}^{[r]}, c^{[r]})$ . Thus, one may interpret the sequence as given an observed network state  $c$ , the policy plays the compression parameters  $\mathbf{q}$  with probability  $\hat{p}(\mathbf{q}|c; \mathbf{q}^{[r]}, c^{[r]})$ . Define the state-dependent stationary policy  $\boldsymbol{\pi}$  as playing the conditional mean (w.r.t., the function  $\mathbf{h}_\varepsilon$ ) given any network state  $c$ . That is,

$$\boldsymbol{\pi}(c) = \mathbf{h}_\varepsilon^{-1} \left( \hat{\mathbb{E}} \left[ \mathbf{h}_\varepsilon(\mathbf{Q}) \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right] \right), \quad \forall c \in \mathcal{C} \text{ such that } \hat{p}(c; c^{[r]}) > 0. \quad (19)$$

Such a choice for  $\boldsymbol{\pi}(c)$  always exists because  $\mathbf{h}_\varepsilon(\cdot)$  is continuous, bounded and strictly increasing coordinate-wise applied function which implies that the inverse operator of  $\mathbf{h}_\varepsilon(\cdot)$  is well-defined. Note that  $\mathbf{h}_\varepsilon(\boldsymbol{\pi}(c)) = \hat{\mathbb{E}} \left[ \mathbf{h}_\varepsilon(\mathbf{Q}) \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right]$ . So, due to the convexity of  $\|\cdot\|$ ,

$$\|\mathbf{h}_\varepsilon(\boldsymbol{\pi}(c))\| \leq \hat{\mathbb{E}} \left[ \|\mathbf{h}_\varepsilon(\mathbf{Q})\| \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right], \quad \forall c \in \mathcal{C}. \quad (20)$$

Then,

$$\begin{aligned} \sum_{n=1}^r \|\mathbf{h}_\varepsilon(\boldsymbol{\pi}(c^n))\| &= r \hat{\mathbb{E}} \left[ \|\mathbf{h}_\varepsilon(\boldsymbol{\pi}(C))\| ; c^{[r]} \right], \\ &\stackrel{(a)}{\leq} r \hat{\mathbb{E}} \left[ \hat{\mathbb{E}} \left[ \|\mathbf{h}_\varepsilon(\mathbf{Q})\| \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right] ; c^{[r]} \right], \\ &\stackrel{(b)}{=} r \hat{\mathbb{E}} \left[ \|\mathbf{h}_\varepsilon(\mathbf{Q})\| ; \mathbf{q}^{[r]} \right], \\ &= \sum_{n=1}^r \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|, \end{aligned} \quad (21)$$

where (a) follows from (20), and (b) follows from the Tower-rule of expectations.

Next we bound  $d(\tau, \boldsymbol{\pi}(c), c)$ . By the definition of  $\boldsymbol{\pi}$  in (19), for all  $c \in \mathcal{C}$ ,

$$\begin{aligned} d(\tau, \boldsymbol{\pi}(c), c) &= d(\tau, \mathbf{h}_\varepsilon^{-1}(\mathbf{h}_\varepsilon(\boldsymbol{\pi}(c))), c), \\ &\stackrel{(a)}{=} d\left(\tau, \mathbf{h}_\varepsilon^{-1} \left( \hat{\mathbb{E}} \left[ \mathbf{h}_\varepsilon(\mathbf{Q}) \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right] \right), c\right), \\ &\stackrel{(b)}{\leq} \hat{\mathbb{E}} \left[ d(\tau, \mathbf{h}_\varepsilon^{-1}(\mathbf{h}_\varepsilon(\mathbf{Q})), c) \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right], \\ &= \hat{\mathbb{E}} \left[ d(\tau, \mathbf{Q}, c) \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right], \end{aligned} \quad (22)$$

where (a) follows from the definition of policy  $\boldsymbol{\pi}$ , and (b) follows from the convexity of  $d(\tau, \mathbf{h}_\varepsilon^{-1}(\cdot), c)$  (Assumption 3). (22) is analogous to (20). Therefore, we may repeat the same calculation in (21) for the delay,

$$\begin{aligned} \sum_{n=1}^r d(\tau, \boldsymbol{\pi}(c^n), c^n) &= r \hat{\mathbb{E}} \left[ d(\tau, \boldsymbol{\pi}(C), C) ; c^{[r]} \right], \\ &\leq r \hat{\mathbb{E}} \left[ \hat{\mathbb{E}} \left[ d(\tau, \mathbf{Q}, C) \mid C = c; \mathbf{q}^{[r]}, c^{[r]} \right] ; c^{[r]} \right], \\ &= r \hat{\mathbb{E}} \left[ d(\tau, \mathbf{Q}, C) ; \mathbf{q}^{[r]}, c^{[r]} \right], \\ &= \sum_{n=1}^r d(\tau, \mathbf{q}^n, c^n), \end{aligned}$$

□

Equation (17) in Proposition C.1 states that if the FL algorithm with a sequence of compression parameters  $(\mathbf{q}^n)_n$  has reached an error tolerance of  $\varepsilon$  by round  $r$ , then, under Assumption 1, it has also reached error tolerance  $\varepsilon$  under sequence of compression parameters  $(\boldsymbol{\pi}(c^n))_n$  by around  $r$ . Moreover, (18) states that  $(\boldsymbol{\pi}(c^n))_n$  takes lesser amount of time up to round  $r$  compared to  $(\mathbf{q}^n)_n$ . However, this construction of  $\boldsymbol{\pi}$  is sequence dependent. More specifically, it is dependent on the type  $\hat{p}(\cdot; c^{[r]})$  of the network state sequence observed. In order to prove Lemma 1, we need to construct a state-dependent but sequence-independent stationary policy that is near-optimal in minimizing the expected wall clock time. Therefore, in the following, we first define the notion of a *typical set* and show in Proposition C.2 that the type of an observed network state sequence concentrates around its mean with high probability.

**Definition 4** (Typical Set). *For a distribution  $p$  on network sets, a typical set with parameters  $(r, \nu)$ , is defined as,*

$$\mathcal{T}_\nu^r(p) \triangleq \{c^r : |\hat{p}(c|c^r) - p(c)| \leq \nu p(c), \text{ for all } c \in \mathcal{C}\}.$$



We will use the following result called the *Typical Averaging Lemma* for typical sets [28, Section 2.4].

**Lemma C.1.** *Let  $c^r \in \mathcal{T}_\nu^r(p)$ . Then, for any non-negative function  $g : \mathcal{C} \rightarrow \mathbb{R}^+$ ,*

$$(1 - \nu) \mathbb{E}[g(C)] \leq \frac{1}{r} \sum_{n=1}^r g(c^n) \leq (1 + \nu) \mathbb{E}[g(C)],$$

where  $C$  is a random variable with distribution  $p$ .

Next, due to ergodicity of stationary Markov chains, we have the following proposition which is proved at the end of this section.

**Proposition C.2.** *Let Assumption 4 be true. Then, there exist positive constants  $\kappa_1$  and  $\kappa_2$  such that, for every  $\nu > 0$ , and  $r \in \mathbb{N}$ ,*

$$P\left(\exists r' \geq r \text{ such that } C^{r'} \notin \mathcal{T}_\nu^{r'}(\mu)\right) \leq \kappa_1 \exp(-\kappa_2 \nu^2 r).$$

Lemma C.1 will be used to argue that if two network-state sequences have similar types (they belong to  $\mathcal{T}_\nu^r(\mu)$ ), then a state dependent stationary policy  $\pi$  will have a similar expected wall clock to converge under both sequences. Then, Proposition C.2 will be used to argue that one observes a typical network state sequence with high probability. We proceed to prove this formally in the following proof of Lemma 1.

*Proof of Lemma 1.* Denote  $h_\varepsilon^{\min}$  and  $h_\varepsilon^{\max}$  as the minimum and maximum of the bounded function  $h_\varepsilon(\cdot)$ . Then, Assumption 1 implies that the number of rounds needed to converge to an error tolerance  $\varepsilon$  under any sequence of compression parameters is bounded between  $h_\varepsilon^{\min}$  and  $h_\varepsilon^{\max}$ .

For a positive  $\nu$ , let  $\varepsilon$  be small enough such that  $h_\varepsilon^{\min} > 2(1 + \nu)/\nu$ . First, we will consider network state sequences which are typical with respect to  $\nu$ . Specifically, we consider a sequence  $(c^n)_n$  such that  $c^r$  belongs to  $\mathcal{T}_\nu^r(p)$  for every  $r > h_\varepsilon^{\min}$ .

Due to Proposition C.1, there exists a state-dependent stationary policy that optimizes the wall clock time to reach error tolerance  $\varepsilon$  with respect to the sequence  $(c^n)_n$ . Let  $\pi'$  represent this policy, and  $r_\varepsilon^{\pi'}$  be the minimum number of rounds taken to converge by  $\pi'$ . Then, the wall clock time for  $\pi'$  can be lower bounded as,

$$\begin{aligned} \sum_{n=1}^{r_\varepsilon^{\pi'}} d(\tau, \pi'(c^n), c^n) &= \left( \frac{1}{r_\varepsilon^{\pi'}} \sum_{n=1}^{r_\varepsilon^{\pi'}} d(\tau, \pi'(c^n), c^n) \right) r_\varepsilon^{\pi'}, \\ &\stackrel{(a)}{\geq} (1 - \nu) \mathbb{E}[d(\tau, \pi'(C), C)] r_\varepsilon^{\pi'}, \\ &\stackrel{(b)}{\geq} (1 - \nu) \mathbb{E}[d(\tau, \pi'(C), C)] \left( \frac{1}{r_\varepsilon^{\pi'}} \sum_{n=1}^{r_\varepsilon^{\pi'}} \|\mathbf{h}_\varepsilon(\pi'(c^n))\| \right), \\ &\stackrel{(c)}{\geq} (1 - \nu)^2 \mathbb{E}[d(\tau, \pi'(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi'(C))\|], \\ &\stackrel{(d)}{\geq} (1 - \nu)^2 \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|]. \end{aligned} \tag{23}$$

(a) and (c) follow from Lemma C.1, (b) follows from Assumption 1, and (d) follows by the following definition,

$$\pi^* = \arg \min_{\pi \in \Pi} \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|].$$

Performing a similar calculation for  $\pi^*$ ,

$$\begin{aligned} \sum_{n=1}^{r_\varepsilon^{\pi^*}} d(\tau, \pi^*(c^n), c^n) &= \left( \frac{1}{r_\varepsilon^{\pi^*}} \sum_{n=1}^{r_\varepsilon^{\pi^*}} d(\tau, \pi^*(c^n), c^n) \right) r_\varepsilon^{\pi^*}, \\ &\stackrel{(a)}{\leq} (1 + \nu) \mathbb{E}[d(\tau, \pi^*(C), C)] r_\varepsilon^{\pi^*}, \\ &\stackrel{(b)}{\leq} (1 + \nu)^2 \mathbb{E}[d(\tau, \pi^*(C), C)] \left( \frac{1}{r_\varepsilon^{\pi^*}} \sum_{n=1}^{r_\varepsilon^{\pi^*}} \|\mathbf{h}_\varepsilon(\pi^*(c^n))\| \right), \\ &\stackrel{(c)}{\leq} (1 + \nu)^3 \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|]. \end{aligned} \tag{24}$$

(a) and (c) follow from Lemma C.1 and (b) follows from Proposition C.3 proved at the end of this section.

The expected wall clock time to reach error tolerance  $\varepsilon$  under the state-dependent stationary policy  $\pi^*$  can be upper bounded as,

$$\begin{aligned} \mathbb{E} [T_\varepsilon^{\pi^*}] &\stackrel{(a)}{\leq} P\left(C^{R_\varepsilon^{\pi^*}} \in \mathcal{T}_\nu^{R_\varepsilon^{\pi^*}}(\mu)\right) (1+\nu)^3 \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|] + P\left(C^{R_\varepsilon^{\pi^*}} \notin \mathcal{T}_\nu^{R_\varepsilon^{\pi^*}}(\mu)\right) h_\varepsilon^{\max} d^{\max}, \\ &\stackrel{(b)}{\leq} (1+\nu)^3 \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|] + P\left(\exists r' \geq h_\varepsilon^{\min}, C^{r'} \notin \mathcal{T}_\nu^{r'}(\mu)\right) h_\varepsilon^{\max} d^{\max}, \\ &\stackrel{(c)}{\leq} (1+\nu)^3 \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|] + \kappa_1 \exp(-\kappa_2 \nu^2 h_\varepsilon^{\min}) h_\varepsilon^{\max} d^{\max}, \end{aligned} \quad (25)$$

where, (a) follows by using (25) for typical sequences, and upper bounding  $R_\varepsilon^{\pi^*}$  by  $h_\varepsilon^{\max}$  and the round duration by  $d^{\max}$  for non-typical sequences. (b) follows by upper bounding  $P\left(C^{R_\varepsilon^{\pi^*}} \in \mathcal{T}_\nu^{R_\varepsilon^{\pi^*}}(\mu)\right)$  by 1, and upper bounding  $P\left(C^{R_\varepsilon^{\pi^*}} \notin \mathcal{T}_\nu^{R_\varepsilon^{\pi^*}}(\mu)\right)$  by  $P\left(\exists r' \geq h_\varepsilon^{\min}, C^{r'} \notin \mathcal{T}_\nu^{r'}(\mu)\right)$  because, almost surely,  $R_\varepsilon^{\pi^*} \geq h_\varepsilon^{\min}$ . (c) follows from Proposition C.2.

Let  $T_\varepsilon^*$  be the random variable representing the wall-clock time to reach error tolerance  $\varepsilon$  when one uses the optimal sample-path sequence dependent policy on random network-state sequence  $(C^n)_n$ . Then, denoting  $R_\varepsilon^*$  as the corresponding random variable denoting the number of rounds needed to reach error tolerance  $\varepsilon$ , we have,

$$\begin{aligned} \mathbb{E} [T_\varepsilon^*] &\stackrel{(a)}{\geq} P\left(C^{R_\varepsilon^*} \in \mathcal{T}_\nu^{R_\varepsilon^*}(\mu)\right) (1-\nu)^2 \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|], \\ &\stackrel{(b)}{\geq} P\left(\forall r' \geq h_\varepsilon^{\min}, C^{r'} \in \mathcal{T}_\nu^{r'}(\mu)\right) (1-\nu)^2 \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|], \\ &\stackrel{(c)}{\geq} (1-\nu)^2 (1 - \kappa_1 \exp(-\kappa_2 \nu^2 h_\varepsilon^{\min})) \mathbb{E}[d(\tau, \pi^*(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi^*(C))\|]. \end{aligned} \quad (26)$$

(a) follows due to (23), and (b) follows since, almost surely,  $R_\varepsilon^* \geq h_\varepsilon^{\min}$ . (c) follows from Proposition C.2.

Due to Assumption 2, both  $h_\varepsilon^{\min}$  and  $h_\varepsilon^{\max}$  are  $\Theta(1/\text{poly}(\varepsilon))$ . And,  $\nu$  can be made as small as desired. Therefore, from (25) and (26),

$$\mathbb{E} [T_\varepsilon^{\pi^*}] \rightarrow \mathbb{E} [T_\varepsilon^*] \quad \text{as } \varepsilon \rightarrow 0.$$

□

**Proposition C.3.** *Under Assumptions 1 and 2, for  $\nu > 0$ , if  $\varepsilon$  is small enough such that  $h_\varepsilon^{\min} > 2(1+\nu)/\nu$ , then, for any state-dependent stationary policy  $\pi$ , the minimum number of rounds  $r_\varepsilon^\pi$  to reach an error tolerance  $\varepsilon$  is such that,*

$$r_\varepsilon^\pi \leq (1+\nu) \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|.$$

*Proof.* Due to Assumption 1,  $r_\varepsilon^\pi$  satisfies,

$$r_\varepsilon^\pi > \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|.$$

Moreover, since  $r_\varepsilon^\pi$  is the earliest round at which error tolerance  $\varepsilon$  is reached, due to Assumption 1, round  $r_\varepsilon^\pi - 1$  satisfies,

$$\begin{aligned} r_\varepsilon^\pi - 1 &\leq \frac{1}{r_\varepsilon^\pi - 1} \sum_{n=1}^{r_\varepsilon^\pi - 1} \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|, \\ &\leq \frac{1}{r_\varepsilon^\pi - 1} \sum_{n=1}^{r_\varepsilon^\pi} \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|, \\ &= \left(\frac{r_\varepsilon^\pi}{r_\varepsilon^\pi - 1}\right) \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|, \\ \implies r_\varepsilon^\pi &\leq \left(\frac{r_\varepsilon^\pi}{r_\varepsilon^\pi - 1}\right)^2 \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} \|\mathbf{h}_\varepsilon(\mathbf{q}^n)\|. \end{aligned} \quad (27)$$

Now, we upper bound  $r_\varepsilon^\pi / (r_\varepsilon^\pi - 1)$ ,

$$\begin{aligned} \frac{r_\varepsilon^\pi}{r_\varepsilon^\pi - 1} &= 1 + \frac{1}{r_\varepsilon^\pi - 1}, \\ &\stackrel{(a)}{<} 1 + \frac{\nu}{2 + \nu}, \\ &= \frac{1 + \nu}{1 + \nu/2}, \\ &\stackrel{(b)}{\leq} \sqrt{1 + \nu}. \end{aligned} \tag{28}$$

(a) follows by rearranging the assumption  $r_\varepsilon^\pi > 2(1 + \nu)/\nu$  to obtain  $r_\varepsilon^\pi - 1 > (2 + \nu)/\nu$ . (b) follows since  $1 + \nu/2 \geq \sqrt{1 + \nu}$  for  $\nu > 0$ . Substituting (28) in (27) completes the result.  $\square$

In order to prove Proposition C.2, we use Theorem 3.1 from [29] which we re-state below for clarity.

**Theorem 3.** *Let Assumption 4 hold. Define  $r_{\text{mix}}$  as the  $1/8$  mixing time<sup>2</sup> of the Markov chain  $(C^n)_n$  and  $f : \mathcal{C} \rightarrow [0, 1]$  be a function. Let,  $\mu_f \triangleq \mathbb{E} [f(C^{(1)})]$ . Then, there exists a constant  $\kappa_c$  such that for every  $0 \leq \nu \leq 1$  and  $r \in \mathbb{N}$ ,*

$$P \left( \left| \frac{1}{r} \sum_{n=1}^r f(C^n) - \mu_f \right| \geq \nu \mu_f \right) \leq \kappa_c \exp \left( -\frac{\nu^2 \mu_f r}{72 r_{\text{mix}}} \right).$$

*Proof of Proposition C.2.* For some  $c \in \mathcal{C}$ , define  $f(c') = \mathbb{1}(c' = c)$ . Then, due to Theorem 3, there exist a constant  $\kappa_c$  such that,

$$P \left( \left| \frac{1}{r'} \sum_{n=1}^{r'} \mathbb{1}(C^n = c) - \mu(c) \right| \geq \nu \mu(c) \right) \leq \kappa_c \exp \left( -\frac{\nu^2 \mu(c) r'}{72 r_{\text{mix}}} \right). \tag{29}$$

Denote,  $\kappa \triangleq \sum_{c \in \mathcal{C}} \kappa_c$  and  $\mu_{\min} = \min_{c \in \mathcal{C}} \mu(c)$ . Since the Markov chain is irreducible,  $\mu_{\min} > 0$ . Then, using (29) and taking a union bound over all  $c \in \mathcal{C}$ , we obtain,

$$P \left( C^{r'} \notin \mathcal{T}_\nu^{r'}(\mu) \right) \leq \kappa \exp \left( -\frac{\nu^2 \mu_{\min} r'}{72 r_{\text{mix}}} \right).$$

Define  $\kappa_2 \triangleq \mu_{\min} / (72 r_{\text{mix}})$ . Taking a further union bound over all  $r' \geq r$ ,

$$P \left( \exists r' \geq r \text{ such that } C^{r'} \notin \mathcal{T}_\nu^{r'}(\mu) \right) \leq \frac{\kappa}{1 - \exp(-\kappa_2 \nu^2)} \exp(-\kappa_2 \nu^2 r)$$

Defining  $\kappa_1 \triangleq \kappa / (1 - \exp(-\kappa_2 \nu^2))$  completes the proof.  $\square$

<sup>2</sup>Denoting  $M$  as the transition-matrix of the Markov chain, and  $\mu$  as its stationary distribution,  $r_{\text{mix}} \triangleq \max_\psi \{r : \|M^r \psi - \mu\|_{TV} \leq 1/8\}$ , where  $\|\cdot\|_{TV}$  denotes the TV-norm. Due to Theorem 4.9 of [30],  $r_{\text{mix}}$  is finite for an aperiodic and irreducible Markov chain.

APPENDIX D  
PROOF OF LEMMA 2

Here we prove Lemma 2 which states that the expected wall clock is approximately equal to the product of the expected number of rounds and the expected round duration asymptotically.

The proof is very similar to the proof of Lemma 1 in Appendix C. As such, we will use the notation introduced in Appendix C.

Denote  $h_\varepsilon^{\min}$  and  $h_\varepsilon^{\max}$  as the minimum and maximum of the bounded function  $h_\varepsilon(\cdot)$ . And, let  $d^{\min}$  and  $d^{\max}$  denote the minimum and maximum of the positive, bounded function  $d(\cdot, \cdot, \cdot)$ .

Similar to the proof of Lemma 1, we consider network state sequences which are typical. In order to define the parameters for the typical set, for any given  $\delta > 0$ , we choose a small enough  $\varepsilon_{th} > 0$  and  $\delta' > 0$  such that,

- 1)  $\varepsilon_{th}$  is small enough such that  $h_{\varepsilon_{th}}^{\min} > 2(1 + \delta')/\delta'$ .
- 2)  $\delta' > 0$  is such that for all  $0 < \varepsilon \leq \varepsilon_{th}$ ,

$$(1 - \delta) < (1 - \delta')^2 \left( 1 - \kappa_1 \exp \left( -\kappa_2 (\delta')^2 h_\varepsilon^{\min} \right) \right),$$

and,

$$\max \left\{ (1 + \delta')^3, 1 + \frac{\kappa_1 \exp \left( -\kappa_2 (\delta')^2 h_\varepsilon^{\min} \right) h_\varepsilon^{\max} d^{\max}}{h_\varepsilon^{\min} d^{\min}} \right\} < (1 + \delta/2).$$

Such a choice of  $\delta'$  is possible because  $h_\varepsilon^{\min} = \Theta(1/\text{poly}(\varepsilon))$ , and  $\exp(-\kappa_2 (\delta')^2 h_\varepsilon^{\min}) h_\varepsilon^{\max} / h_\varepsilon^{\min} = \exp(-\Omega(\delta'^2 / \text{poly}(\varepsilon)))$ .

We consider a sequence  $(c^n)_n$  such that  $c^r \in \mathcal{T}_{\delta'}^r(\mu)$  for every  $r \geq h_{\varepsilon_{th}}^{\min}$ .

For a policy  $\pi$ , let  $r_\varepsilon^\pi$  denote the minimum number of rounds needed to converge to error tolerance  $\varepsilon < \varepsilon_{th}$  given network state sequence  $(c^n)_n$ . Then, the wall clock time for  $\pi$  can be lower bounded as,

$$\begin{aligned} \sum_{n=1}^{r_\varepsilon^\pi} d(\tau, \pi(c^n), c^n) &= \left( \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} d(\tau, \pi(c^n), c^n) \right) r_\varepsilon^\pi, \\ &\stackrel{(a)}{\geq} (1 - \delta') \mathbb{E} [d(\tau, \pi(C), C)] r_\varepsilon^\pi, \\ &\stackrel{(b)}{\geq} (1 - \delta') \mathbb{E} [d(\tau, \pi(C), C)] \left( \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} \|\mathbf{h}_\varepsilon(\pi(c^n))\| \right), \\ &\stackrel{(c)}{\geq} (1 - \delta')^2 \mathbb{E} [d(\tau, \pi(C), C)] \mathbb{E} [\|\mathbf{h}_\varepsilon(\pi(C))\|], \end{aligned} \tag{30}$$

(a) and (c) follow from Lemma C.1 since  $r_\varepsilon^\pi > h_{\varepsilon_{th}}^{\min}$ , and (b) follows from Assumption 1.

Performing a similar calculation for the upper bound,

$$\begin{aligned} \sum_{n=1}^{r_\varepsilon^\pi} d(\tau, \pi(c^n), c^n) &= \left( \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} d(\tau, \pi(c^n), c^n) \right) r_\varepsilon^\pi, \\ &\stackrel{(a)}{\leq} (1 + \delta') \mathbb{E} [d(\tau, \pi(C), C)] r_\varepsilon^\pi, \\ &\stackrel{(b)}{\leq} (1 + \delta')^2 \mathbb{E} [d(\tau, \pi(C), C)] \left( \frac{1}{r_\varepsilon^\pi} \sum_{n=1}^{r_\varepsilon^\pi} \|\mathbf{h}_\varepsilon(\pi(c^n))\| \right), \\ &\stackrel{(c)}{\leq} (1 + \delta')^3 \mathbb{E} [d(\tau, \pi(C), C)] \mathbb{E} [\|\mathbf{h}_\varepsilon(\pi(C))\|]. \end{aligned} \tag{31}$$

(a) and (c) follow from Lemma C.1, and (b) follows from Proposition C.3.

Let  $R_\varepsilon^\pi$  be the random variable denoting the minimum number of rounds needed to converge to an error tolerance  $\varepsilon < \varepsilon_{th}$  when policy  $\pi$  is used on the sequence  $(C^n)_n$ . The expected wall clock times of  $\pi$  can be lower bounded as,

$$\begin{aligned}
\mathbb{E}[T_\varepsilon^\pi] &\geq P\left(C^{R_\varepsilon^\pi} \in \mathcal{T}_{\delta'}^{R_\varepsilon^\pi}(\mu)\right) (1 - \delta')^2 \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|], \\
&\stackrel{(a)}{\geq} P\left(\forall r' \geq h_\varepsilon^{\min}, C^{r'} \in \mathcal{T}_{\nu'}^{r'}(\mu)\right) (1 - \delta')^2 \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|] \\
&\stackrel{(b)}{\geq} (1 - \delta')^2 \left(1 - \kappa_1 \exp\left(-\kappa_2 \delta'^2 h_\varepsilon^{\min}\right)\right) \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|] \\
&\stackrel{(c)}{\geq} (1 - \delta) \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|].
\end{aligned} \tag{32}$$

(a) follows since  $R_\varepsilon^\pi \geq h_\varepsilon^{\min}$  almost surely, (b) follows from Proposition C.2 and (c) follows from the choice of  $\delta'$ .

The expected wall clock times of  $\pi$  can be upper bounded as,

$$\begin{aligned}
\mathbb{E}[T_\varepsilon^\pi] &\stackrel{(a)}{\leq} P\left(C^{R_\varepsilon^\pi} \in \mathcal{T}_{\delta'}^{R_\varepsilon^\pi}(\mu)\right) (1 + \delta')^3 \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|] + P\left(C^{R_\varepsilon^\pi} \notin \mathcal{T}_{\delta'}^{R_\varepsilon^\pi}(\mu)\right) h_\varepsilon^{\max} d^{\max}, \\
&\stackrel{(b)}{\leq} (1 + \delta')^3 \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|] + P\left(\exists r' \geq h_\varepsilon^{\min} C^{r'} \notin \mathcal{T}_{\delta'}^{r'}(\mu)\right) h_\varepsilon^{\max} d^{\max}, \\
&\stackrel{(c)}{\leq} (1 + \delta')^3 \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|] + \kappa_1 \exp\left(-\kappa_2 \delta'^2 h_\varepsilon^{\min}\right) h_\varepsilon^{\max} d^{\max}, \\
&\stackrel{(d)}{\leq} (1 + \delta) \mathbb{E}[d(\tau, \pi(C), C)] \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|].
\end{aligned} \tag{33}$$

(a) follows by using (31) to upper bound the wall clock time for typical sequences, and using the worst-case upper bound  $h_\varepsilon^{\max} d^{\max}$  for non-typical sequences. (b) follows because, almost surely,  $R_\varepsilon^\pi \geq h_\varepsilon^{\min}$ . (c) follows from Proposition C.2. (d) follows from the choice of  $\delta'$ .

(32) and (33), jointly, conclude the proof.

APPENDIX E  
PROOF OF PROPOSITION B.2

In this section, we prove Proposition B.2 which states that the Fluid-Frank-Wolfe (FFW) process has a unique stationary point in the set  $\text{conv}(V_\varepsilon)$ . It further states that the stationary point lies in the set  $V_\varepsilon$ . In order to demonstrate the main arguments of the proof, we will first consider the case with a single client ( $m = 1$ ) and a single network state,  $\mathcal{C} = \{c\}$ . Later, in Subsection B, we will generalize these arguments to complete the proof of Proposition B.2.

*A. Warmup: 1 Client and 1 Network State (1C1NS) Case*

Recall that the set  $V_\varepsilon$  is the set of pairs of achievable expected number of rounds  $\hat{r}_\varepsilon$  and expected round duration  $\hat{d}$  of state-dependent stationary policy. Here, we observe that  $\text{conv}(V_\varepsilon)$  may be interpreted as the corresponding feasibility set for (possibly random) stationary policies. We start by making this observation precise. In the 1C1NS case, a stationary policy may be represented by a one-dimensional random variable  $\Pi$  that denotes the possibly randomly selected compression parameter. The space of possible stationary policies is denoted by the set of distributions  $\mathcal{Q}_1$ ,

$$\mathcal{Q}_1 = \{f_\Pi : f_\Pi \text{ is a distribution over } [0, q_{\max}]\}.$$

Under a policy corresponding to  $\Pi$  and a small enough error tolerance  $\varepsilon$ , due to Lemma 2, the expression for the expected wall clock time for stationary policies is given by,

$$\mathbb{E}[T_\varepsilon^\Pi] \approx \hat{t}_\varepsilon^\Pi = \mathbb{E}[h_\varepsilon(\Pi)] \mathbb{E}[d(\tau, \Pi, c)],$$

where  $T_\varepsilon^\Pi$  is the wall clock time to reach error tolerance  $\varepsilon$  under compression parameter  $\Pi$ . Then, the feasible set  $\text{conv}(V_\varepsilon)$  for stationary policies may be written for this case as,

$$\text{conv}(V_\varepsilon) = \left\{ (\hat{r}_\varepsilon, \hat{d}) : \exists f_\Pi \in \mathcal{Q}_1 \text{ s.t., } \Pi \sim f_\Pi \text{ satisfies } \hat{r}_\varepsilon = \mathbb{E}[h_\varepsilon(\Pi)], \hat{d} = \mathbb{E}[d(\tau, \Pi, c)] \right\}.$$

Here, stationary policies may be separated into two categories,

- deterministic policies: here,  $\Pi$  is a constant.
- stochastic policies: here,  $\Pi$  is non-deterministic random variable.

Our aim is to prove that there exists a unique fixed-point of the FFW update in  $\text{conv}(V_\varepsilon)$ . As we will see, this will prove Proposition B.2 for the special case of 1 client and 1 network state. Before delving into the proof of this statement, we state a couple of properties of  $\text{conv}(V_\varepsilon)$  which are useful in proving the existence of a unique fixed-point of the FFW update.

**Proposition E.4.** *For any  $h_\varepsilon^{\min} \leq h \leq h_\varepsilon^{\max}$ , there exists a deterministic policy that minimizes  $\hat{t}_\varepsilon^\Pi$  given a constraint  $\mathbb{E}[h_\varepsilon(\Pi)] = h$ .*

*Proof.* As a contradiction, assume that there is no deterministic policy that minimizes the expected wall clock time given a constraint  $\mathbb{E}[h_\varepsilon(\Pi)] = h$ . Let  $\Pi^*$  be a stochastic policy that minimizes the expected wall clock time with the constraint  $\mathbb{E}[h_\varepsilon(\Pi)] = h$ . Then, consider an alternate policy with deterministic compression parameter  $\pi$  chosen as,  $h_\varepsilon(\pi) = \mathbb{E}[h_\varepsilon(\Pi^*)]$ . Such a  $\pi$  exists due to the Intermediate Value Theorem since  $h_\varepsilon(\cdot)$  is a continuous function. In this case, by the strict convexity of the duration function assumed in Assumption 3, we have,

$$\mathbb{E}[d(\tau, \pi, c)] < \mathbb{E}[d(\tau, \Pi^*, c)].$$

Therefore, the relation between their expected wall clock times is,

$$\hat{t}_\varepsilon^\pi = \mathbb{E}[h_\varepsilon(\pi)] \mathbb{E}[d(\tau, \pi, c)] < \mathbb{E}[h_\varepsilon(\Pi^*)] \mathbb{E}[d(\tau, \Pi^*, c)] = \hat{t}_\varepsilon^{\Pi^*}.$$

This is a contradiction to the assumption that a stochastic policy minimizes the expected wall clock time given the constraint.  $\square$

For notational brevity, denote,

$$\bar{d}(\hat{r}_\varepsilon) = d(\tau, h_\varepsilon^{-1}(\hat{r}_\varepsilon), c).$$

Recall that we may denote points  $(\hat{r}_\varepsilon, \hat{d}) \in \text{conv}(V_\varepsilon)$  by a two-dimensional vector  $\mathbf{x} = (\hat{r}_\varepsilon, \hat{d})^\top$ . Also, recall the function  $H(\mathbf{x}) = x_1 x_2$ . The following proposition states several equivalent ways of describing a point  $\mathbf{x}$  in the set  $V_\varepsilon$ .

**Proposition E.5.** *The following statements are equivalent*

- I.  $\mathbf{x} \in \text{conv}(V_\varepsilon)$  is of the form  $(\hat{r}_\varepsilon, \bar{d}(\hat{r}_\varepsilon))$  for some  $h_\varepsilon^{\min} \leq \hat{r}_\varepsilon \leq h_\varepsilon^{\max}$ .
- II.  $\mathbf{x} \in \text{conv}(V_\varepsilon)$  is such that  $\alpha \mathbf{x} \notin \text{conv}(V_\varepsilon)$  for any  $0 < \alpha < 1$ .
- III.  $\mathbf{x} = (\hat{r}_\varepsilon, \hat{d}) \in \text{conv}(V_\varepsilon)$  is such that,  $\hat{d} = \min\{d' : (\hat{r}_\varepsilon, d') \in \text{conv}(V_\varepsilon)\}$ .
- IV.  $\mathbf{x} = (\hat{r}_\varepsilon, \hat{d}) \in \text{conv}(V_\varepsilon)$  is such that,  $\hat{r}_\varepsilon \hat{d} = \min\{\hat{r}_\varepsilon d' : (\hat{r}_\varepsilon, d') \in \text{conv}(V_\varepsilon)\}$ .




Fig. 4: Feasibility set for stationary policies  $\text{conv}(V_\varepsilon)$ . In the IC1NS case, a point  $(\hat{r}_\varepsilon, \hat{d}) \in \text{conv}(V_\varepsilon)$  corresponds to a policy  $\Pi$  such that,  $\hat{r}_\varepsilon = \mathbb{E}[h_\varepsilon(\Pi)]$  and  $\hat{d} = \mathbb{E}[d(\tau, \Pi, c)]$ . The function  $\bar{d}(\hat{r}_\varepsilon)$  is represented by the blue curve.

V.  $x$  is an extreme point<sup>3</sup> of  $\text{conv}(V_\varepsilon)$ .

*Proof Sketch.* The equivalences may be inferred from the structure of the feasibility set  $\text{conv}(V_\varepsilon)$  as shown in Fig. 4. We briefly describe the arguments required to prove the equivalences. I  $\iff$  III because  $(\hat{r}_\varepsilon, \bar{d}_\varepsilon)$  corresponds to a deterministic policy by definition, and, due to part b, a deterministic policy minimizes the wall clock time  $\hat{r}_\varepsilon \hat{d}$  amongst the set of policies  $\{f_\Pi \in \mathcal{Q}_1 : \text{s.t.}, \Pi \sim f_\Pi \text{ satisfies } \mathbb{E}[h_\varepsilon(\Pi)] = \hat{r}_\varepsilon\}$ . I, III  $\iff$  II because  $\bar{d}(\hat{r}_\varepsilon)$  is a strictly decreasing function. I  $\iff$  V because  $\bar{d}$  is a strictly convex function. III and IV are trivially equivalent.  $\square$

From here on, we will call points of the form described in Proposition E.5 as extreme points, and all other points in  $\text{conv}(V_\varepsilon)$  as non-extreme points. Note that in the IC1NS case, the set of extreme points is equal to the set  $V_\varepsilon$ . However, we refrain from using this fact here because in the general case of multiple clients and multiple network states, this is no longer true.

Next, we prove the existence of a unique fixed point of the FFW update in two steps. First, we show that a non-extreme point cannot be a fixed point of the FFW update.

**Proposition E.6.** *If a point  $x \in \text{conv}(V_\varepsilon)$  is not an extreme point of  $\text{conv}(V_\varepsilon)$ , then it is not a fixed-point of the FFW update.*

*Proof.* Due Proposition E.5,  $x$  being a non-extreme point implies that there exists a constant  $0 < \alpha < 1$  such that  $\alpha x \in \text{conv}(V_\varepsilon)$ .

Observe that  $\nabla H(x) = (x_2 \ x_1)^\top$ . Since,  $h_\varepsilon(\cdot)$  and  $d(\cdot)$  are non-negative functions, we have that  $\nabla H(x)$  has non-negative entries for any  $x \in V_\varepsilon$ .

Recall that  $x$  is a fixed-point of the FFW update if and only if  $x = \arg \min_{x' \in \text{conv}(V_\varepsilon)} \nabla H(x)^\top x'$ . Due to the elementwise non-negativity of  $\nabla H$ ,  $\nabla H(x)^\top (\alpha x) < \nabla H(x)^\top x$ . Since  $\alpha x \in \text{conv}(V_\varepsilon)$ ,  $x$  is not a fixed point of the FFW-update.  $\square$

Due to Proposition E.6, we focus on only extreme points in the next step. Before proving the existence of a unique fixed point of the FFW update amongst the set of extreme points, we state a result about an equivalent description of a fixed point. For this purpose, define  $\hat{t}(\hat{r}_\varepsilon) = \hat{r}_\varepsilon \bar{d}(\hat{r}_\varepsilon)$ .

**Proposition E.7.** *Under Assumption 5, an extreme-point  $x = (\hat{r}_\varepsilon, \bar{d}(\hat{r}_\varepsilon))^\top$  with  $h_\varepsilon^{\min} < \hat{r}_\varepsilon < \hat{h}_\varepsilon^{\max}$  is a fixed point for the FFW-update if and only if  $\hat{t}'(\hat{r}_\varepsilon) = 0$ .*

*Proof.* First, as a contradiction, assume that there exists an  $\hat{r}_\varepsilon$  with  $\hat{t}'(\hat{r}_\varepsilon) = 0$ , but  $x = (\hat{r}_\varepsilon, \bar{d}(\hat{r}_\varepsilon))^\top$  is not a fixed point of the FFW update. Recall that  $x$  is a fixed point of the FFW update if  $x = \arg \min_{x' \in \text{conv}(V_\varepsilon)} \nabla H(x)^\top x'$ . Therefore, this implies that there exists a different point  $z = (\hat{r}', \bar{d}(\hat{r}'))^\top$  in  $\text{conv}(V_\varepsilon)$  such that,

$$\nabla H(x)^\top (z - x) \leq 0.$$

Refer to Fig. 5 for an illustration of the following argument. Due to strict convexity of the curve  $\bar{d}(\cdot)$ , there exists a point  $w = (\hat{r}, \bar{d}(\hat{r}))^\top$  with  $\hat{r}$  being in between  $\hat{r}_\varepsilon$  and  $\hat{r}'$  such that  $\nabla H(x)^\top (w - x) = -\xi < 0$ . Then, for any  $0 < \theta < 1$ , denote

<sup>3</sup>A point  $x \in \text{conv}(V_\varepsilon)$  is called an *extreme point* if it cannot be written as the convex combination of two other points in  $\text{conv}(V_\varepsilon)$ .




Fig. 5: Illustration of the construction for a part of the proof of Proposition E.7 and E.14. Specifically, proof of the statement,  $\hat{t}'(\hat{r}_\varepsilon) = 0$  implies a fixed point for the FW-update.

$\mathbf{w}_\theta^c = (1 - \theta)\mathbf{x} + \theta\mathbf{w}$  and  $\mathbf{w}_\theta = ((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}, \bar{d}((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}))$ . It is easily verified that  $\nabla H(\mathbf{x})^\top(\mathbf{w}_\theta^c - \mathbf{x}) = -\xi\theta$ . And, due to the convexity of  $\bar{d}(\cdot)$ ,  $\nabla H(\mathbf{x})^\top(\mathbf{w}_\theta - \mathbf{x}) \leq \nabla H(\mathbf{x})^\top(\mathbf{w}_\theta^c - \mathbf{x})$ . That is,

$$\nabla H(\mathbf{x})^\top(\mathbf{w}_\theta - \mathbf{x}) \leq -\xi\theta, \quad (34)$$

By a Taylor Series expansion of  $H$ , we have  $\hat{t}((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}) = \hat{t}(\hat{r}_\varepsilon) + \nabla H(\mathbf{x})^\top(\mathbf{w}_\theta - \mathbf{x}) + O(\theta^2)$ . Therefore, for small enough  $\theta$ , (34) implies that  $\hat{t}((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}) < \hat{t}(\hat{r}_\varepsilon)$ . This is a contradiction to Assumption 5 which states that  $\hat{t}''(\hat{r}_\varepsilon) > 0$ .

Conversely, consider for contradiction that  $\mathbf{x} = (\hat{r}_\varepsilon, \bar{d}(\hat{r}_\varepsilon))^\top$  is a fixed point of the FFW update, but  $\hat{t}'(\hat{r}_\varepsilon) \neq 0$ . Define  $\hat{r}^\delta = \hat{r}_\varepsilon + \delta$ , and  $\mathbf{z}^\delta = (\hat{r}^\delta, \bar{d}(\hat{r}^\delta))^\top$ . Then, for small enough  $\delta > 0$ , since  $\hat{t}'(\hat{r}_\varepsilon) \neq 0$ , we have that

$$\begin{aligned} \hat{t}(\hat{r}^\delta) &< \hat{t}(\hat{r}_\varepsilon) - \Omega(\delta), \\ \text{(OR)} \\ \hat{t}(\hat{r}^{-\delta}) &< \hat{t}(\hat{r}_\varepsilon) - \Omega(\delta). \end{aligned} \quad (35)$$

Again, by Taylor series expansion,

$$\begin{aligned} H(\mathbf{z}^\delta) &= H(\mathbf{x}) + \nabla H(\mathbf{x})(\mathbf{z}^\delta - \mathbf{x}) + O(\delta^2), \\ \iff \hat{t}(\hat{r}^\delta) &= \hat{t}(\hat{r}_\varepsilon) + \nabla H(\mathbf{x})(\mathbf{z}^\delta - \mathbf{x}) + O(\delta^2). \end{aligned} \quad (36)$$

Comparing (35) and (36), we have that  $\nabla H(\mathbf{x})(\mathbf{z}^\delta - \mathbf{x}) < 0$  or  $\nabla H(\mathbf{x})(\mathbf{z}^{-\delta} - \mathbf{x}) < 0$ . This is a contradiction to the premise that  $\mathbf{x}$  is a fixed point for the FFW update.  $\square$

**Remark 2.**  $\mathbf{x} = (h_\varepsilon^{\min}, \bar{d}(h_\varepsilon^{\min}))$  is a fixed point of the FFW update if  $t'(h_\varepsilon^{\min}) \geq 0$ . Similarly,  $\mathbf{x} = (h_\varepsilon^{\max}, \bar{d}(h_\varepsilon^{\max}))$  is a fixed point of the FFW update if  $t'(h_\varepsilon^{\max}) \leq 0$ . The proof of these statements is very similar to that of Proposition E.5. We don't prove these statements for the ICINS case for succinctness. But, we prove it rigorously in the multiple client and multiple network state case.

**Proposition E.8.** For the ICINS case, there exists a unique fixed point  $\mathbf{x}$  of the FFW update in  $\text{conv}(V_\varepsilon)$ .

*Proof.* Due to Proposition E.6, a non-extreme point of  $\text{conv}(V_\varepsilon)$  cannot be a fixed point of the FFW update.

Proposition E.7 and Remark 2 jointly characterize the fixed point of the FFW update amongst the set of extreme points in terms of  $\hat{t}(\cdot)$ . In particular, they imply that a point  $\mathbf{x} = (\hat{r}_\varepsilon, \bar{d}(\hat{r}_\varepsilon))$  is a fixed-point if and only if  $\hat{r}_\varepsilon$  is a local-minimum of  $\hat{t}(\cdot)$  in the domain  $h_\varepsilon^{\min} \leq \hat{r}_\varepsilon \leq h_\varepsilon^{\max}$  (since  $\hat{t}(\cdot)$  is quasiconvex, it does not have any local maximum in the interior of its domain). Since  $\hat{t}$  is a strictly quasiconvex function (Assumption 5) on a bounded domain, it has a unique local minimum. Therefore, the FFW update has a unique fixed point among the set of extreme points of  $\text{conv}(V_\varepsilon)$ .  $\square$

### B. Multiple Clients and Multiple Network States (MCMNS) Case

This section contains the complete proof of Proposition B.2. The argument is a generalization of what we saw in the previous subsection to the case with multiple clients and multiple network states. We start by introducing notation to describe policies in this general setting.



We may denote the policy  $\pi$  as a function,  $\pi : \mathcal{C} \rightarrow [0, q_{\max}]^m$ , or as a vector  $\underline{\pi}$  of dimension  $m|\mathcal{C}|$ . In specific, enumerating the elements of  $\mathcal{C}$  as,  $\mathcal{C} = \{c_1, c_2, \dots, c_{|\mathcal{C}|}\}$ , the vector  $\underline{\pi}$  is represented as,

$$\underline{\pi} = (\pi_1(c_1), \dots, \pi_m(c_1), \pi_1(c_2), \dots, \pi_m(c_2), \dots, \pi_1(c_{|\mathcal{C}|}), \dots, \pi_m(c_{|\mathcal{C}|}))^\top,$$

where  $\pi_i(c)$  indicates the  $i^{\text{th}}$  entry of  $m$ -dimensional vector  $\pi(c)$ .

Next, we explain how to define the quantities,  $\mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|]$  and  $\mathbb{E}[d(\tau, \pi(C), C)]$  in terms of the vector representation of  $\pi$ . Denote  $\mathbf{e}^i$  as an  $m|\mathcal{C}|$  dimensional vector with,

$$\mathbf{e}_j^i = \begin{cases} 1 & , \text{ if } (i-1)|\mathcal{C}| \leq j < i|\mathcal{C}|, \\ 0 & , \text{ otherwise.} \end{cases}$$

Recall that  $\mu$  denotes the stationary distribution of the Markov chain of network states. Then, for a deterministic policy  $\pi$ , we may write,

$$\tilde{r}_\varepsilon(\underline{\pi}) \triangleq \mathbb{E}[\|\mathbf{h}_\varepsilon(\pi(C))\|] = \sum_{i=1}^{|\mathcal{C}|} \mu(c_i) \|\mathbf{h}_\varepsilon(\underline{\pi}) \odot \mathbf{e}^i\|, \quad (37)$$

$$\tilde{d}(\underline{\pi}) \triangleq \mathbb{E}[d(\tau, \pi(C), C)] = \sum_{i=1}^{|\mathcal{C}|} \mu(c_i) d\left(\tau, \underline{\pi}_{(i-1)|\mathcal{C}|}^{i|\mathcal{C}|-1}, c_i\right), \quad (38)$$

where  $\odot$  denotes the elementwise product of two vectors, and  $\underline{\pi}_j^k$  denotes the vector  $(\pi_j, \pi_{j+1}, \dots, \pi_k)^\top$ .

Similar to the 1C1NS case,  $\text{conv}(V_\varepsilon)$  may be interpreted as the set of achievable pairs of expected rounds and expected round duration of (possibly, stochastic) stationary policies. Therefore, a stationary policy may be represented by a random vector,  $\underline{\mathbf{II}}$ . Then, we will denote the set of all (deterministic and stochastic) stationary policies as,

$$\mathcal{Q}_{m|\mathcal{C}|} = \left\{ \underline{f}_{\underline{\mathbf{II}}} : \underline{f}_{\underline{\mathbf{II}}} \text{ is a } m|\mathcal{C}| \text{ dimensional distribution over } [0, q_{\max}]^{m|\mathcal{C}|} \right\}.$$

The feasible set  $\text{conv}(V_\varepsilon)$  may be defined as,

$$\text{conv}(V_\varepsilon) = \left\{ (\hat{r}_\varepsilon, \hat{d}) : \exists \underline{f}_{\underline{\mathbf{II}}} \in \mathcal{Q}_{m|\mathcal{C}|} \text{ s.t., } \underline{\mathbf{II}} \sim \underline{f}_{\underline{\mathbf{II}}} \text{ satisfies } \hat{r}_\varepsilon = \mathbb{E}[\tilde{r}_\varepsilon(\underline{\mathbf{II}})], \quad \hat{d} = \mathbb{E}[\tilde{d}(\underline{\mathbf{II}})] \right\}.$$

Next, we prove an analogous result of Proposition E.4.

**Proposition E.9.** *For any  $h_\varepsilon^{\min} \leq h \leq h_\varepsilon^{\max}$ , there exists a deterministic policy  $\underline{\pi}$  that minimizes  $\mathbb{E}[\tilde{d}(\underline{\mathbf{II}})]$  over the constrained domain  $\{\underline{\mathbf{II}} \in \mathcal{Q}_{m|\mathcal{C}|} : \mathbb{E}[\tilde{r}_\varepsilon(\underline{\mathbf{II}})] = h\}$ .*

*Proof.* The proof is similar to that of Proposition E.4.

As a contradiction, assume that there is no deterministic policy that minimizes the expected wall clock time given a constraint  $\mathbb{E}[\tilde{r}_\varepsilon(\underline{\mathbf{II}})] = h$ . Let  $\underline{\mathbf{II}}^*$  be a stochastic policy that minimizes the expected wall clock time with the constraint  $\mathbb{E}[\tilde{r}_\varepsilon(\underline{\mathbf{II}})] = h$ . Then, consider an alternate policy with deterministic compression parameters  $\underline{\pi}$  chosen as,  $\mathbf{h}_\varepsilon(\underline{\pi}) = \mathbb{E}[\mathbf{h}_\varepsilon(\underline{\mathbf{II}}^*)]$ . Such a  $\underline{\pi}$  exists due to the Intermediate Value Theorem since  $h_\varepsilon(\cdot)$  is a continuous function. In this case, by the strict convexity of the duration function assumed in Assumption 3, we have,

$$\mathbb{E}[\tilde{d}(\underline{\pi})] < \mathbb{E}[\tilde{d}(\underline{\mathbf{II}}^*)].$$

By the convexity of the norm operator,  $\tilde{r}_\varepsilon(\underline{\pi}) \leq h$ . Notice that  $\tilde{r}_\varepsilon(\cdot)$  is increasing in every co-ordinate, whereas  $\tilde{d}(\cdot)$  is decreasing in every co-ordinate. Therefore, for any  $\underline{\pi}' \geq \underline{\pi}$  (elementwise inequality) with  $\tilde{r}_\varepsilon(\underline{\pi}') = h$ , we have,

$$\mathbb{E}[\tilde{d}(\underline{\pi}')] \leq \mathbb{E}[\tilde{d}(\underline{\pi})] < \mathbb{E}[\tilde{d}(\underline{\mathbf{II}}^*)].$$

This is a contradiction to the assumption that a stochastic policy minimizes the expected wall clock time given the constraint.  $\square$

At this point in the proof of the 1C1NS case, we defined a function  $\bar{d}(\hat{r}_\varepsilon)$  which was the round duration of the deterministic policy whose number of rounds for convergence was  $\hat{r}_\varepsilon$ . In the MCMNS case, since there could be multiple deterministic policies corresponding to a rounds for convergence  $\hat{r}_\varepsilon$ , we define  $\bar{d}(\hat{r}_\varepsilon)$  with respect to the policy that minimizes the round duration.

$$\bar{d}(r_\varepsilon) = \min_{\underline{\pi} : \tilde{r}_\varepsilon(\underline{\pi}) = r_\varepsilon} \tilde{d}(\underline{\pi}). \quad (39)$$

In the rest of the proof, we need to use the fact that  $\bar{d}(\hat{r}_\varepsilon)$  is strictly convex, and that  $\hat{t}(\hat{r}_\varepsilon) \triangleq \hat{r}_\varepsilon \bar{d}(\hat{r}_\varepsilon)$  is strictly quasiconvex. In the 1C1NS case, these facts were a direct consequence of Assumptions 3 and 5 because  $\bar{d}(\hat{r}_\varepsilon)$  was simply  $d(\tau, h_\varepsilon^{-1}(\hat{r}_\varepsilon), c)$ . For the MCMNS case we prove these results in the following two propositions.

**Proposition E.10.** *Under Assumptions 1 and 3,  $\bar{d}(r_\varepsilon)$  is decreasing and strictly convex in  $r_\varepsilon$ .*

*Proof.* From Assumption 1, recall that since  $h_\varepsilon(\cdot)$  is a strictly increasing, continuous function, it has an inverse  $h_\varepsilon^{-1}(\cdot)$ . Denote by,  $\mathbf{h}_\varepsilon^{-1} : [h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|} \rightarrow [0, q_{\max}]^{m|C|}$ , the function that outputs a vector obtained by applying  $h_\varepsilon^{-1}(\cdot)$  elementwise to the input vector.

$\bar{d}(r_\varepsilon)$  from (39) may be redefined as,

$$\bar{d}(r_\varepsilon) = \min_{\substack{\mathbf{r} : \mathbf{r} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|} \\ \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})) = r_\varepsilon}} \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})).$$

Due to Assumption 1,  $\tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}))$  is an increasing function in every element of  $\mathbf{r}$ , and, due to Assumption 3,  $\tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}))$  is a decreasing function in every element of  $\mathbf{r}$ . Therefore, we can further reformulate  $\bar{d}(r_\varepsilon)$  as,

$$\bar{d}(r_\varepsilon) = \min_{\substack{\mathbf{r} : \mathbf{r} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|} \\ \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})) \leq r_\varepsilon}} \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})). \quad (40)$$

$\bar{d}(r_\varepsilon)$  is decreasing because the feasibility set in the minimization problem in (40) is an monotonically increasing set with increasing  $\hat{r}_\varepsilon$ .

Consider two points,  $r_{\varepsilon,1}, r_{\varepsilon,2} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]$ , and let  $\mathbf{r}_{\varepsilon,1}$  and  $\mathbf{r}_{\varepsilon,2}$  be their corresponding minimizers according to (40),

$$\begin{aligned} \mathbf{r}_{\varepsilon,1} &\triangleq \arg \min_{\substack{\mathbf{r} : \mathbf{r} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|} \\ \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})) \leq r_{\varepsilon,1}}} \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})), \\ \mathbf{r}_{\varepsilon,2} &\triangleq \arg \min_{\substack{\mathbf{r} : \mathbf{r} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|} \\ \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})) \leq r_{\varepsilon,2}}} \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})). \end{aligned} \quad (41)$$

Let  $0 < \theta < 1$ ,  $r_{\varepsilon,\theta} = \theta r_{\varepsilon,1} + (1-\theta)r_{\varepsilon,2}$  and  $\mathbf{r}_{\varepsilon,\theta} = \theta \mathbf{r}_{\varepsilon,1} + (1-\theta) \mathbf{r}_{\varepsilon,2}$ . By the strict convexity of  $\tilde{d}(\mathbf{h}_\varepsilon^{-1}(\cdot))$  as considered in Assumption 3,

$$\tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,\theta})) < \theta \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,1})) + (1-\theta) \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,2})). \quad (42)$$

From the definition of  $\tilde{r}_\varepsilon(\cdot)$  in (37),

$$\begin{aligned} \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,\theta})) &= \sum_{i=1}^{|C|} \mu(c_i) \left\| \mathbf{r}_{\varepsilon,\theta} \odot \mathbf{e}^{(i)} \right\|, \\ &\stackrel{(a)}{\leq} \theta \sum_{i=1}^{|C|} \mu(c_i) \left\| \mathbf{r}_{\varepsilon,1} \odot \mathbf{e}^{(i)} \right\| + (1-\theta) \sum_{i=1}^{|C|} \mu(c_i) \left\| \mathbf{r}_{\varepsilon,2} \odot \mathbf{e}^{(i)} \right\|, \\ &\stackrel{(b)}{=} \theta \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,1})) + (1-\theta) \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,2})), \\ &\stackrel{(c)}{\leq} \theta r_{\varepsilon,1} + (1-\theta) r_{\varepsilon,2}, \\ &\stackrel{(d)}{=} r_{\varepsilon,\theta}. \end{aligned} \quad (43)$$

(a) follows from the convexity of the norm operator, and (b), (c) and (d) follow by definition.

Due to (43),  $\mathbf{r}_{\varepsilon,\theta}$  is a feasible point in the constraint set of the minimization problem in (40) evaluated at  $r_{\varepsilon,\theta}$ . Therefore,

$$\begin{aligned} \bar{d}(r_{\varepsilon,\theta}) &\leq \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,\theta})), \\ &\stackrel{(a)}{<} \theta \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,1})) + (1-\theta) \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}_{\varepsilon,2})), \\ &\stackrel{(b)}{=} \theta \bar{d}(r_{\varepsilon,1}) + (1-\theta) \bar{d}(r_{\varepsilon,2}). \end{aligned}$$

(a) follows from (42), and (b) follows by definition. Since this is true for any  $r_{\varepsilon,1}, r_{\varepsilon,2} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]$ , and  $0 < \theta < 1$ ,  $\bar{d}(\cdot)$  is strictly convex.  $\square$

However, unlike the 1 client and 1 network state case,  $\bar{d}(r_\varepsilon)$  may not be differentiable. But, since it is convex, it has left-derivative and right-derivative functions denoted as  $\bar{d}'_L(r_\varepsilon)$  and  $\bar{d}'_R(r_\varepsilon)$  respectively.

$$\bar{d}'_L(r_\varepsilon) = \begin{cases} \lim_{\delta \downarrow 0} \frac{\bar{d}(r_\varepsilon) - \bar{d}(r_\varepsilon - \delta)}{\delta}, & r_\varepsilon > h_{\min}(\varepsilon) \\ -\infty, & r_\varepsilon = h_{\min}(\varepsilon), \end{cases}$$

$$\bar{d}'_R(r_\varepsilon) = \begin{cases} \lim_{\delta \downarrow 0} \frac{\bar{d}(r_\varepsilon + \delta) - \bar{d}(r_\varepsilon)}{\delta}, & r_\varepsilon < h_{\max}(\varepsilon) \\ 0, & r_\varepsilon = h_{\max}(\varepsilon). \end{cases}$$

Similar to the case of 1 client and 1 network state, given a constraint,  $\tilde{r}_\varepsilon(\boldsymbol{\pi}) = r_\varepsilon$ , the optimal expected wall clock time may be expressed as  $\hat{t}(r_\varepsilon) = r_\varepsilon \bar{d}(r_\varepsilon)$ . Since  $\bar{d}(r_\varepsilon)$  has left and right derivatives everywhere, so does  $\hat{t}(r_\varepsilon)$ . Denote them by  $\hat{t}'_L(r_\varepsilon)$  and  $\hat{t}'_R(r_\varepsilon)$  respectively.

**Proposition E.11.**  $\hat{t}(r_\varepsilon)$  is strictly quasiconvex. That is, for any  $r_{\varepsilon,1}, r_{\varepsilon,2} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]$ , and  $0 < \theta < 1$ ,

$$\hat{t}(\theta r_{\varepsilon,1} + (1 - \theta)r_{\varepsilon,2}) < \max\{\hat{t}(r_{\varepsilon,1}), \hat{t}(r_{\varepsilon,2})\}.$$

*Proof.* Here, we reuse the definitions introduced in the proof of Proposition E.10.

Consider two points  $r_{\varepsilon,1}, r_{\varepsilon,2} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]$ . Let  $\mathbf{r}_{\varepsilon,1}$  and  $\mathbf{r}_{\varepsilon,2}$  be defined as in (41). Consider,  $0 < \theta < 1$ . Since  $\tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}))$  is continuous in  $\mathbf{r}$ , by the Intermediate Value Theorem, there exists a  $0 < \delta < 1$  such that  $\mathbf{r}^\delta = \delta \mathbf{r}_{\varepsilon,1} + (1 - \delta) \mathbf{r}_{\varepsilon,2}$  has  $\tilde{r}_\varepsilon(\mathbf{r}^\delta) = \theta r_{\varepsilon,1} + (1 - \theta)r_{\varepsilon,2}$ .

From the strict quasiconvexity of the wall clock time as in Assumption 5, we have,

$$\tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}^\delta)) \bar{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}^\delta)) < \max\{\hat{t}(r_{\varepsilon,1}), \hat{t}(r_{\varepsilon,2})\}.$$

By definition,

$$\hat{t}(\theta r_{\varepsilon,1} + (1 - \theta)r_{\varepsilon,2}) \leq \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}^\delta)) \hat{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}^\delta)).$$

Therefore,

$$\hat{t}(\theta r_{\varepsilon,1} + (1 - \theta)r_{\varepsilon,2}) < \max\{\hat{t}(r_{\varepsilon,1}), \hat{t}(r_{\varepsilon,2})\}.$$

□

At this point in the proof of the 1C1NS case, we stated some equivalent descriptions of a point  $\mathbf{x} \in V_\varepsilon$ . The description is the same for the MCMNS case as well, which we restate here for clarity.

**Proposition E.12.** *The following statements are equivalent*

- I.  $\mathbf{x} \in \text{conv}(V_\varepsilon)$  is of the form  $(\hat{r}_\varepsilon, \bar{d}(\hat{r}_\varepsilon))$  for some  $h_\varepsilon^{\min} \leq \hat{r}_\varepsilon \leq h_\varepsilon^{\max}$ .
- II.  $\mathbf{x} \in \text{conv}(V_\varepsilon)$  is such that  $\alpha \mathbf{x} \notin \text{conv}(V_\varepsilon)$  for any  $0 < \alpha < 1$ .
- III.  $\mathbf{x} = (\hat{r}_\varepsilon, \hat{d}) \in \text{conv}(V_\varepsilon)$  is such that,  $\hat{d} = \min\{d' : (\hat{r}_\varepsilon, d') \in \text{conv}(V_\varepsilon)\}$ .
- IV.  $\mathbf{x} = (\hat{r}_\varepsilon, \hat{d}) \in \text{conv}(V_\varepsilon)$  is such that,  $\hat{r}_\varepsilon \hat{d} = \min\{\hat{r}_\varepsilon d' : (\hat{r}_\varepsilon, d') \in \text{conv}(V_\varepsilon)\}$ .
- V.  $\mathbf{x}$  is an extreme point of  $\text{conv}(V_\varepsilon)$ .

*Proof Sketch.* The only difference in the proof from that of Proposition E.5 is the equivalence I  $\iff$  III. Here, I  $\iff$  III follows from the definition of  $\bar{d}(\cdot)$  in (39). □

At this point in the 1C1NS case, we showed that a non-extreme point of  $\text{conv}(V_\varepsilon)$  cannot be a fixed point of the FFW update. The result is the same in the MCMNS case. We restate the result for clarity, but skip the proof as it is the same as that of Proposition E.6.

**Proposition E.13.** *If a point  $\mathbf{x} \in \text{conv}(V_\varepsilon)$  is not an extreme point of  $\text{conv}(V_\varepsilon)$ , then it is not a fixed-point of the FFW update.*

Next, similar to the 1C1NS case, we show necessary and sufficient condition for an extreme-point of  $\text{conv}(V_\varepsilon)$  to be a fixed-point of the FFW update.

**Proposition E.14.** *A point  $\mathbf{x} = (r_\varepsilon, \bar{d}(r_\varepsilon))^\top$  is a fixed point for the FFW update if and only if  $\hat{t}'_L(r_\varepsilon) \leq 0$  and  $\hat{t}'_R(r_\varepsilon) \geq 0$ .*

*Proof.* The proof is very similar to the proof of Proposition E.7, but with some extra care because, here,  $\hat{t}(\cdot)$  may not be differentiable.

Consider a point  $\mathbf{x} = (r_\varepsilon, \bar{d}(r_\varepsilon))^\top$  which is not a fixed point of the FFW update. This implies that there exists another point,  $\mathbf{z} = (r, \bar{d}(r))^\top$ , such that,

$$\nabla H(\mathbf{x})^\top (\mathbf{z} - \mathbf{x}) \leq 0.$$

Refer to Fig. 5 for an illustration of the following argument. Due to strict convexity of the curve  $\bar{d}(\cdot)$  (Proposition E.10), there exists a point  $\mathbf{w} = (\hat{r}', \bar{d}(\hat{r}'))^\top$  with  $\hat{r}$  being in between  $\hat{r}_\varepsilon$  and  $\hat{r}'$  such that  $\nabla H(\mathbf{x})^\top (\mathbf{w} - \mathbf{x}) = -\xi < 0$ . Then, for any  $0 < \theta < 1$ , consider  $\mathbf{w}_\theta^c = (1 - \theta)\mathbf{x} + \theta\mathbf{w}$  and  $\mathbf{w}_\theta = ((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}', \bar{d}((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}'))$ . It is easily verified that  $\nabla H(\mathbf{x})^\top (\mathbf{w}_\theta^c - \mathbf{x}) = -\xi\theta$ . And, due to the convexity of  $\bar{d}(\cdot)$ ,  $\nabla H(\mathbf{x})^\top (\mathbf{w}_\theta - \mathbf{x}) \leq \nabla H(\mathbf{x})^\top (\mathbf{w}_\theta^c - \mathbf{x})$ . That is,

$$\nabla H(\mathbf{x})^\top (\mathbf{w}_\theta - \mathbf{x}) \leq -\xi\theta, \quad (44)$$

for some positive constant  $\xi$ .

By a Taylor Series expansion of  $H$ , we have  $\hat{t}((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}') = \hat{t}(\hat{r}_\varepsilon) + \nabla H(\mathbf{x})^\top (\mathbf{w}_\theta - \mathbf{x}) + O(\theta^2)$ . Therefore, due to (44), for small enough  $\theta$ , we have  $\hat{t}((1 - \theta)\hat{r}_\varepsilon + \theta\hat{r}') < \hat{t}(\hat{r}_\varepsilon) - \xi'\theta$ , where  $\xi'$  is some positive constant. This, in turn, implies that either  $\hat{t}'_L(\hat{r}_\varepsilon) > 0$  or  $\hat{t}'_R(\hat{r}_\varepsilon) < 0$ .

Conversely, consider for contradiction that  $\mathbf{x} = (r_\varepsilon, \bar{d}(r_\varepsilon))^\top$  is a fixed point of the FFW update, but  $\hat{t}'_L(r_\varepsilon) > 0$  or  $\hat{t}'_R(r_\varepsilon) < 0$ . Define  $r^\delta = r_\varepsilon + \delta$ , and  $\mathbf{z}^\delta = (r^\delta, \bar{d}(r^\delta))^\top$ . Then, for small enough  $\delta > 0$ , we have that

$$\begin{aligned} \hat{t}(r^\delta) &< \hat{t}(r_\varepsilon) - \Omega(\delta), \\ \text{(OR)} \\ \hat{t}(r^{-\delta}) &< \hat{t}(r_\varepsilon) - \Omega(\delta). \end{aligned} \quad (45)$$

Again, by Taylor series expansion,

$$\hat{t}(r^\delta) = \hat{t}(r_\varepsilon) + \nabla H(\mathbf{x})(\mathbf{z}^\delta - \mathbf{x}) + O(\delta^2). \quad (46)$$

Comparing (45) and (46), we have that  $\nabla H(\mathbf{x})(\mathbf{z}^\delta - \mathbf{x}) < 0$  or  $\nabla H(\mathbf{x})(\mathbf{z}^{-\delta} - \mathbf{x}) < 0$ . This is a contradiction to the premise that  $\mathbf{x}$  is a fixed point for the FFW update.  $\square$

Next, similar to the ICINS case, we use the equivalence derived in Proposition E.12 to prove that the fixed point of the FFW update is unique.

**Proposition E.15.** *There exists a unique point  $\mathbf{x} = (r, \bar{d}(r))^\top$  with  $h_\varepsilon^{\min} \leq r \leq h_\varepsilon^{\max}$  such that  $\hat{t}'_L(r) \leq 0$  and  $\hat{t}'_R(r) \geq 0$ .*

*Proof.* Denote,

$$\mathbf{r}(r_\varepsilon) = \arg \min_{\substack{\mathbf{r} : \mathbf{r} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|} \\ \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})) = r_\varepsilon}} \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})).$$

Also, denote,

$$\tilde{t}_\varepsilon(\mathbf{r}) = \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r})) \tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}))$$

We use the following claim which is proved at the end of this section.

**Claim E.1.** *If  $\hat{t}'_L(r_\varepsilon) \leq 0$  and  $\hat{t}'_R(r_\varepsilon) \geq 0$ , then,*

$$\nabla_{\mathbf{r}} (\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon))) = 0,$$

*or, the left and right derivatives of  $\mathbf{r}(\cdot)$  evaluated at  $r_\varepsilon$ ,  $\mathbf{r}'_L(r_\varepsilon)$  and  $\mathbf{r}'_R(r_\varepsilon)$ , exist, and,*

$$\begin{aligned} \mathbf{r}'_L(r_\varepsilon)^\top \nabla_{\mathbf{r}} (\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon))) &\leq 0, \\ \text{(AND)} \end{aligned}$$

$$\mathbf{r}'_R(r_\varepsilon)^\top \nabla_{\mathbf{r}} (\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon))) \geq 0.$$

In either case, Assumption 5 ensures that for all small enough  $\delta > 0$ , we have  $\hat{t}(r_\varepsilon - \delta) > \hat{t}(r_\varepsilon)$  and  $\hat{t}(r_\varepsilon + \delta) > \hat{t}(r_\varepsilon)$ . That is, if  $\hat{t}'_L(r_\varepsilon) \leq 0$ , then for all small enough  $\delta > 0$ ,  $\hat{t}(r_\varepsilon - \delta) > \hat{t}(r_\varepsilon)$  (whenever  $r_\varepsilon - \delta$  is in  $[h_\varepsilon^{\min}, h_\varepsilon^{\max}]$ ). Similarly, if  $\hat{t}'_R(r_\varepsilon) \geq 0$ , then for small enough  $\delta > 0$ ,  $\hat{t}(r_\varepsilon + \delta) > \hat{t}(r_\varepsilon)$  (whenever  $r_\varepsilon + \delta$  is in  $[h_\varepsilon^{\min}, h_\varepsilon^{\max}]$ ). Therefore,  $r_\varepsilon$  is a strict local minimum of  $\hat{t}$ .

Since  $\hat{t}(\cdot)$  is strictly quasiconvex (Proposition E.11), there can be at most one point  $r_\varepsilon$  which is a strict local minimum of  $\hat{t}(\cdot)$ .

Moreover, since  $\hat{t}$  is a continuous function over a closed and bounded set, it attains a local minimum over its domain. And,  $\hat{t}'_L(r_\varepsilon) \leq 0$  and  $\hat{t}'_R(r_\varepsilon) \geq 0$  is necessary condition for a local minimum. Therefore, there exists at least one point  $r_\varepsilon$  in the domain such that  $\hat{t}'_L(r_\varepsilon) \leq 0$  and  $\hat{t}'_R(r_\varepsilon) \geq 0$ .  $\square$

*Proof of Claim E.1.* Recall the notation,  $\tilde{t}_\varepsilon(\mathbf{r}) = \tilde{r}_\varepsilon(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}))\tilde{d}(\mathbf{h}_\varepsilon^{-1}(\mathbf{r}))$ .

a) *Case 1:* : First, consider the case that  $\mathbf{r}(r_\varepsilon)$  is in the interior of the domain  $[h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|}$ . In this case we will show that,  $\nabla\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon)) = 0$ . As a contradiction, assume that  $\nabla\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon)) \neq 0$ .

Let  $L(r_\varepsilon) \triangleq \{\mathbf{r} : \tilde{r}(\mathbf{r}) = r_\varepsilon\}$ . Since,  $\mathbf{r}(r_\varepsilon)$  is the minimizer of  $\tilde{t}_\varepsilon(\mathbf{r})$  over the set  $L(r_\varepsilon)$ ,  $\nabla\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon))$  has to be normal to  $L(r_\varepsilon)$  at the point  $\mathbf{r}(r_\varepsilon)$ . Let  $\mathbf{n}$  denote the normal to the set  $L$  at point  $r_\varepsilon$ . Then, one of the following is true,

$$\begin{aligned}\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon) + \delta\mathbf{n}) &= \tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon)) - \Theta(\delta), \\ (OR) \\ \tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon) - \delta\mathbf{n}) &= \tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon)) - \Theta(\delta).\end{aligned}$$

Then, by the definition of  $t(\cdot)$ , one of the following is true,

$$\begin{aligned}t(r_\varepsilon + \delta) &= t(r_\varepsilon) - \Omega(\delta), \\ (OR) \\ t(r_\varepsilon - \delta) &= t(r_\varepsilon) - \Omega(\delta).\end{aligned}$$

This implies that either  $t'_L(r_\varepsilon) > 0$  or  $t'_R(r_\varepsilon) > 0$ . This is a contradiction. Therefore,  $\nabla\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon)) = 0$ , for all  $\mathbf{r}(r_\varepsilon)$  in the interior of  $[h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|}$ .

b) *Case 2:* : Consider the case that  $\mathbf{r}(r_\varepsilon)$  is on the boundary of  $[h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|}$  and  $\nabla\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon)) \neq 0$ .

Here, the left derivative exists and its direction is expressed as,

$$\mathbf{r}'_L(r_\varepsilon) \propto \arg \max_{\mathbf{s} : \mathbf{r}(r_\varepsilon) + \mathbf{s} \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]^{m|C|}} \frac{\mathbf{s}^\top \nabla\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon))}{\|\mathbf{s}\|_2}.$$

Since,  $t'_L(r_\varepsilon) \leq 0$ , we obtain  $\mathbf{r}'_L(r_\varepsilon)^\top \nabla\tilde{t}_\varepsilon(\mathbf{r}(r_\varepsilon)) \leq 0$ .

A similar argument holds for the right derivative as well.  $\square$

Finally, we summarize how the results in this section prove Proposition B.2.

*Proof Summary of Proposition B.2.* Due to Proposition E.13, a non-extreme point of  $\text{conv}(V_\varepsilon)$  is not a fixed-point of the FFW update. Further, Proposition E.12 shows that an equivalent condition for an extreme point  $\mathbf{x} = (\hat{r}_\varepsilon, \tilde{d}(\hat{r}_\varepsilon))$  of  $\text{conv}(V_\varepsilon)$  being a fixed point of the FFW update is that  $t'_L(\hat{r}_\varepsilon) \leq 0$  and  $t'_R(\hat{r}_\varepsilon) \geq 0$ . Then, in Proposition E.15 we showed that there is a unique point which satisfies  $t'_L(\hat{r}_\varepsilon) \leq 0$  and  $t'_R(\hat{r}_\varepsilon) \geq 0$ . Therefore, this proves that there is a unique fixed point  $\mathbf{x}^*$  for the FFW update in  $\text{conv}(V_\varepsilon)$ .

Further, since  $\mathbf{x}^*$  is an extreme point of  $\text{conv}(V_\varepsilon)$ , and  $\text{conv}(V_\varepsilon)$  is the convex hull of the set  $V_\varepsilon$ ,  $\mathbf{x}^*$  lies in the set  $V_\varepsilon$ .

Finally, to prove that  $\mathbf{x}^*$  is the minimizer of  $H(\cdot)$  over the set  $V_\varepsilon$ , we make the following observation which is a consequence of IV in Proposition E.12,

$$\min_{\mathbf{x} \in V_\varepsilon} H(\mathbf{x}) = \min_{r \in [h_\varepsilon^{\min}, h_\varepsilon^{\max}]} \hat{t}(r).$$

Recall from Proposition E.12 that  $\mathbf{x}^* = (r^*, \tilde{d}(r^*))$  is such that  $t'_L(r^*) \leq 0$  and  $t'_R(r^*) \geq 0$ . Further, since  $\hat{t}$  is strictly quasiconvex (Proposition E.11),  $r^*$  is the global minimizer of  $\hat{t}$ . Therefore,  $\mathbf{x}^*$  is the minimizer of  $H(\cdot)$  over the set  $V_\varepsilon$ .  $\square$

APPENDIX F  
PROOF OF THEOREM 2

In this section, we prove Theorem 2 which bounds the convergence of the FedCOM-V Algorithm shown in Algorithm 2. Moreover, we explicitly state the choice of local learning rate schedule  $(\eta^n)$ , and global learning rate schedule  $(\gamma^n)$  required to achieve the convergence rate stated in Theorem 2.

In the rest of the section, we violate our convention by sometimes using lower case letters instead of capital letters to denote random variables/vectors in order to stay consistent with the notation used in FL literature.

We start by defining sigma-algebras and filtrations associated with the probability space.

**Remark 3.** Let  $\sigma(X)$  denote the sigma-algebra generated by the random variable  $X$ . Let  $\sigma(X_1, \dots, X_n)$  denote the sigma-algebra generated by the set of random variables  $X_1$  to  $X_n$ . Similarly, let  $\sigma(\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_n)$  denote the smallest sigma-algebra containing the sigma-algebras  $\mathcal{F}_1$  to  $\mathcal{F}_n$ .

First, we describe the sigma-algebra associated with the network state process. Recalling that  $C^n$  denotes the network state at round  $n$ , denote,

$$\mathcal{F}^C \triangleq \sigma(C^n : n \geq 1).$$

We remark that although the compression parameters  $(\mathbf{q}^n)_n$  may be random vectors dependent across various rounds  $n$ , their randomness only depends on the network states under the NAC-FL policy as well as under other baseline policies we consider in this paper. More precisely,  $\mathbf{q}^n$  is measurable in  $\mathcal{F}^C$  for all rounds  $n$ , and is therefore not dependent on the updates of the FedCOM-V algorithm. Moreover, the aim in this proof is to study the convergence of the FedCOM-V algorithm for arbitrary choices of  $(\mathbf{q}^n)_n$ . So, in this section, all expectations will be conditioned on  $\mathcal{F}^C$ , and therefore,  $q_j^n$ 's will be treated as arbitrary, but known, constants in this Section.

Next, we describe the sigma-algebras associated with the stochastic gradients. Recall that at round  $n$ , by local-step  $b$ , client  $j$  has sampled mini-batches  $(\mathcal{Z}_j^{a,n})_{a=1}^b$  to compute the stochastic-gradients. So, we denote the associated sigma-algebra across all clients as,

$$\mathcal{D}_{b,n} \triangleq \sigma(\mathcal{Z}_j^{a,n} : j \in [m], a \in [b]), \quad b \in [\tau_n],$$

The sigma-algebra associated with the compressors used in round  $n$  is denoted as,

$$\mathcal{Q}_n \triangleq \sigma(\mathcal{Q}(\cdot, q_j^n) : j \in [m]).$$

Finally, we describe the filtration for the entire system. Since, in this Section, we condition all events on the knowledge of the network states, the filtration is initialized as,

$$\mathcal{F}_0 = \mathcal{F}^C.$$

At the  $b^{\text{th}}$  local step of round  $n$ , the filtration includes the knowledge of all previous rounds and the stochastic-gradients up to the  $b^{\text{th}}$  local step,

$$\mathcal{F}_{b,n} = \sigma(\mathcal{F}_{n-1}, \mathcal{D}_{b,n}), \quad b \in [\tau_n], n \geq 1,$$

And, finally, at the end of round  $n$ , the filtration includes the knowledge of all previous rounds, the stochastic gradients and compressors of round  $n$ ,

$$\mathcal{F}_n = \sigma(\mathcal{F}_{n-1}, \mathcal{D}_{\tau_n, n}, \mathcal{Q}_n), \quad n \geq 1.$$

Next we recap the FedCOM-V Algorithm and introduce further notation used in the proof.

At the start of round  $n$ , client  $j$  receives the global model  $\mathbf{w}^n$  from the server, which it initializes as  $\mathbf{w}_j^{1,n}$ . Then, at local-step  $b$ , it samples a mini-batch  $\mathcal{Z}_j^{b,n}$  and computes the stochastic gradient,  $\tilde{\mathbf{g}}_j^{b,n} \triangleq \nabla f(\mathbf{w}_j^{b,n}, \mathcal{Z}_j^{b,n})$ , while performing the local model update as,  $\mathbf{w}_j^{b+1,n} = \mathbf{w}_j^{b,n} - \eta_n \tilde{\mathbf{g}}_j^{b,n}$ . At this point, we remark that there are two sources of randomness involved in the evaluation of a stochastic gradient at a local step  $b$  of round  $n$ . One is from the model  $\mathbf{w}_j^{b,n}$  at which the gradient is evaluated, which is itself obtained by stochastic gradient and compressed aggregation updates of previous rounds and local steps (i.e.,  $\mathbf{w}_j^{b,n}$  is measurable in  $\mathcal{F}_{b-1,n}$ ). The second source of randomness is from the mini-batch,  $\mathcal{Z}_j^{b,n}$ , used to compute the stochastic gradient  $\tilde{\mathbf{g}}_j^{b,n}$ . So,  $\tilde{\mathbf{g}}_j^{b,n}$  is measurable in  $\mathcal{F}_{b,n}$ .

Additionally, for the analysis, we will define the ‘‘true-gradient’’ at the local model  $\mathbf{w}_j^{b,n}$  as,  $\mathbf{g}_j^{b,n} \triangleq \nabla f(\mathbf{w}_j^{b,n})$ . Observe that the true gradient at local step  $b$  of round  $n$  is itself a random vector, as it is evaluated at  $\mathbf{w}_j^{b,n}$ . But, it is independent of the mini-batch,  $\mathcal{Z}_j^{b,n}$ , sampled at that step. Therefore,  $\mathbf{g}_j^{b,n}$  is measurable in  $\mathcal{F}_{b-1,n}$ . Refer to Figure 6 for an illustration of this process.

After  $\tau_n$  local computations, the client computes its ‘‘pre-compressed’’ update,  $\tilde{\mathbf{g}}_j^n \triangleq \sum_{b=1}^{\tau_n} \tilde{\mathbf{g}}_j^{b,n}$ , which can also be expressed as,  $\tilde{\mathbf{g}}_j^n = (\mathbf{w}^n - \mathbf{w}_j^{\tau_n+1,n})/\eta_n$ . Next, the client sends the compressed message,  $\tilde{\mathbf{g}}_{Q_j}^n = \mathcal{Q}(\tilde{\mathbf{g}}_j^n, q_j^n)$ , to the server.




Fig. 6: Illustration of the local steps at a client.

The server aggregates the compressed messages received from the clients as,  $\tilde{\mathbf{g}}_Q^n = 1/m \sum_{j=1}^m \tilde{\mathbf{g}}_{Qj}^n$ , and performs the update  $\mathbf{w}^{n+1} = \mathbf{w}^n - \eta_n \gamma_n \tilde{\mathbf{g}}_Q^n$ .

For the purpose of analysis, define  $\tilde{\mathbf{g}}^n \triangleq 1/m \sum_{j=1}^m \tilde{\mathbf{g}}_j^n$ , which may be interpreted as the message aggregated at the server had the clients not used any compression. Further, define the “true-gradient” analogies of  $\tilde{\mathbf{g}}_j^n$  and  $\tilde{\mathbf{g}}^n$  as,  $\mathbf{g}_j^n = \sum_{b=1}^{\tau_n} \mathbf{g}_j^{b,n}$  and  $\mathbf{g}^n = 1/m \sum_{j=1}^m \mathbf{g}_j^n$ .  $\mathbf{g}^n$  and  $\mathbf{g}_j^n$  are random variables since their components,  $\mathbf{g}_j^{b,n}$ s, are evaluated at models which are obtained by stochastic gradient updates.

**Remark 4.** *The presence of a tilde and the subscript  $Q$ , such as in  $\tilde{\mathbf{g}}_Q^n$ , will indicate that vector is both compressed and has stochastic gradient components. The presence of just a tilde, such as in  $\tilde{\mathbf{g}}_j^n$ , will indicate that vector (or its components) has two sources of randomness: one from the model at which it (or its components) is evaluated, and second from the mini-batch using which it (or its components) is evaluated. The absence of both the tilde and subscript  $Q$ , such as in  $\mathbf{g}_j^n$ , will indicate that the vector (or its components) has one source of randomness, which is from the model at which it (or its components) is evaluated.*

Also, denote the average noise-variance across clients per round as  $\bar{q}^n$ :

$$\bar{q}^n = \frac{1}{m} \sum_{j=1}^m q_j^n.$$

We start by stating some results which will assist in proving Theorem 2. First is a lemma that bounds the distance between the sum of stochastic gradients at a client in a round to the sum of the true gradients across the local steps at the client in the round.

**Lemma F.2.**  $\mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^n - \mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right] \leq \tau_n \sigma^2$ .

*Proof.*

$$\begin{aligned}
\mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^n - \mathbf{g}_j^n\|^2 \mid \mathcal{F}_{n-1} \right] &\stackrel{(a)}{=} \mathbb{E} \left[ \left\| \sum_{b=1}^{\tau_n} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \right\|^2 \mid \mathcal{F}_{n-1} \right], \\
&\stackrel{(b)}{=} \mathbb{E} \left[ \mathbb{E} \left[ \left\| \sum_{b=1}^{\tau_n} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \right\|^2 \mid \mathcal{F}_{\tau_n-1,n} \right] \mid \mathcal{F}_{n-1} \right], \\
&\stackrel{(c)}{=} \mathbb{E} \left[ \mathbb{E} \left[ \left\| \sum_{b=1}^{\tau_n-1} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \right\|^2 \mid \mathcal{F}_{\tau_n-1,n} \right] + \mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^{\tau_n,n} - \mathbf{g}_j^{\tau_n,n}\|^2 \mid \mathcal{F}_{\tau_n-1,n} \right] \right. \\
&\quad \left. + 2 \mathbb{E} \left[ (\tilde{\mathbf{g}}_j^{\tau_n,n} - \mathbf{g}_j^{\tau_n,n})^\top \sum_{b=1}^{\tau_n-1} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \mid \mathcal{F}_{\tau_n-1,n} \right] \mid \mathcal{F}_{n-1} \right], \\
&\stackrel{(d)}{=} \mathbb{E} \left[ \left\| \sum_{b=1}^{\tau_n-1} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \right\|^2 + \mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^{\tau_n,n} - \mathbf{g}_j^{\tau_n,n}\|^2 \mid \mathcal{F}_{\tau_n-1,n} \right] \right. \\
&\quad \left. + 2 \mathbb{E} \left[ (\tilde{\mathbf{g}}_j^{\tau_n,n} - \mathbf{g}_j^{\tau_n,n})^\top \sum_{b=1}^{\tau_n-1} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \mid \mathcal{F}_{\tau_n-1,n} \right] \mid \mathcal{F}_{n-1} \right], \\
&\stackrel{(e)}{=} \mathbb{E} \left[ \left\| \sum_{b=1}^{\tau_n-1} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \right\|^2 + \mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^{\tau_n,n} - \mathbf{g}_j^{\tau_n,n}\|^2 \mid \mathcal{F}_{\tau_n-1,n} \right] \mid \mathcal{F}_{n-1} \right], \\
&\stackrel{(f)}{\leq} \mathbb{E} \left[ \left\| \sum_{b=1}^{\tau_n-1} (\tilde{\mathbf{g}}_j^{b,n} - \mathbf{g}_j^{b,n}) \right\|^2 \mid \mathcal{F}_{n-1} \right] + \sigma^2, \\
&\vdots \\
&\leq \tau_n \sigma^2
\end{aligned}$$

(a) follows by definition. (b) follows by law of iterated expectations. (c) follows by partially expanding the summation. (d) follows because all random vectors except  $\tilde{\mathbf{g}}_j^{\tau_n,n}$  are measurable in  $\mathcal{F}_{\tau_n-1,n}$ . (e) follows because  $\tilde{\mathbf{g}}_j^{\tau_n,n}$  is an unbiased estimate of  $\mathbf{g}_j^{\tau_n,n}$  (Assumption 7). (f) follows by Assumption 7 which bounds the variance of the stochastic gradients. Repeating steps (b)-(f)  $(\tau_n - 1)$  more times by taking internal conditional expectations w.r.t.,  $\mathcal{F}_{\tau_n-2,n}, \dots, \mathcal{F}_{1,n}, \mathcal{F}_{n-1}$  gives us the result.  $\square$

Next is a lemma comparing the expected norm of the (compressed and stochastic) gradient received by the server to the true gradients at all the local steps.

**Lemma F.3.** *Under Assumption 7 and 8,*

$$\mathbb{E} \left[ \|\tilde{\mathbf{g}}_Q^n\|^2 \mid \mathcal{F}_{n-1} \right] \leq \frac{2\tau_n}{m} \left( \frac{q_{\max}}{m} + 1 \right) \sum_{j=1}^m \sum_{b=1}^{\tau_n} \mathbb{E} \left[ \|\mathbf{g}_j^{b,n}\|^2 \mid \mathcal{F}_{n-1} \right] + (\bar{q}^n + 1) \frac{2\tau_n \sigma^2}{m}.$$



*Proof.*

$$\begin{aligned}
& \mathbb{E} \left[ \|\tilde{\mathbf{g}}_Q^n\|^2 \middle| \mathcal{F}_{n-1} \right] \\
& \stackrel{(a)}{=} \mathbb{E} \left[ \mathbb{E} \left[ \left\| \frac{1}{m} \sum_{j=1}^m \tilde{\mathbf{g}}_{Q,j}^n \right\|^2 \middle| \mathcal{F}_{\tau_n, n} \right] \middle| \mathcal{F}_{n-1} \right], \\
& \stackrel{(b)}{=} \mathbb{E} \left[ \mathbb{E} \left[ \left\| \frac{1}{m} \sum_{j=1}^m \tilde{\mathbf{g}}_{Q,j}^n - \frac{1}{m} \sum_{j=1}^m \mathbb{E} [\tilde{\mathbf{g}}_{Q,j}^n | \mathcal{F}_{\tau_n, n}] \right\|^2 \middle| \mathcal{F}_{n-1} \right] + \left\| \mathbb{E} \left[ \frac{1}{m} \sum_{j=1}^m \tilde{\mathbf{g}}_{Q,j}^n | \mathcal{F}_{\tau_n, n} \right] \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
& \stackrel{(c)}{=} \mathbb{E} \left[ \frac{1}{m^2} \sum_{j=1}^m \left( \mathbb{E} \left[ \|\tilde{\mathbf{g}}_{Q,j}^n - \tilde{\mathbf{g}}_j^n\|^2 \middle| \mathcal{F}_{\tau_n, n} \right] \right) + \left\| \frac{1}{m} \sum_{j=1}^m \tilde{\mathbf{g}}_j^n \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
& \stackrel{(d)}{\leq} \mathbb{E} \left[ \sum_{j=1}^m \frac{q_j^n}{m^2} \|\tilde{\mathbf{g}}_j^n\|^2 + \|\tilde{\mathbf{g}}^n\|^2 \middle| \mathcal{F}_{n-1} \right]. \tag{47}
\end{aligned}$$

(a) follows by Law of Iterated Expectations. (b) follows by the identity,  $\mathbb{E}[\|\mathbf{X}\|^2] = \mathbb{E}[\|\mathbf{X} - \mathbb{E}[\mathbf{X}]\|^2] + \|\mathbb{E}[\mathbf{X}]\|^2$  (this is the  $\mathbb{E}X^2 = \text{var}(X) + (\mathbb{E}X)^2$  identity applied to vectors). The first summation term of (c) is obtained by expanding out the first squared norm from (b) and observing that  $\tilde{\mathbf{g}}_{Q,j}^n - \tilde{\mathbf{g}}_j^n$  is a zero mean random vector, and independent across different  $j$ 's. The second term in (c) follows from the linearity of expectation and the unbiased property of the compressor (Assumption 8). (d) follows from Assumption 8, which bounds the noise introduced by the compressor.

Let's bound  $\mathbb{E}[\|\tilde{\mathbf{g}}^n\|^2 | \mathcal{F}_{n-1}]$ ,

$$\begin{aligned}
\mathbb{E} \left[ \|\tilde{\mathbf{g}}^n\|^2 \middle| \mathcal{F}_{n-1} \right] & \stackrel{(a)}{\leq} 2 \mathbb{E} \left[ \|\tilde{\mathbf{g}}^n - \mathbf{g}^n\|^2 \middle| \mathcal{F}_{n-1} \right] + 2 \mathbb{E} \left[ \|\mathbf{g}^n\|^2 \middle| \mathcal{F}_{n-1} \right], \\
& \stackrel{(b)}{=} 2 \mathbb{E} \left[ \left\| \frac{1}{m} \sum_{j=1}^m (\tilde{\mathbf{g}}_j^n - \mathbf{g}_j^n) \right\|^2 \middle| \mathcal{F}_{n-1} \right] + 2 \mathbb{E} \left[ \left\| \frac{1}{m} \sum_{j=1}^m \mathbf{g}_j^n \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
& \stackrel{(c)}{=} \frac{2}{m^2} \sum_{j=1}^m \mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^n - \mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right] + 2 \mathbb{E} \left[ \left\| \frac{1}{m} \sum_{j=1}^m \mathbf{g}_j^n \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
& \stackrel{(d)}{\leq} \frac{2}{m^2} \sum_{j=1}^m \mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^n - \mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right] + \frac{2}{m} \sum_{j=1}^m \mathbb{E} \left[ \|\mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right], \\
& \stackrel{(e)}{\leq} \frac{2\tau_n\sigma^2}{m} + \frac{2}{m} \sum_{j=1}^m \mathbb{E} \left[ \|\mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right]. \tag{48}
\end{aligned}$$

Above, (a) follows from the identity  $\|\mathbf{x}\|^2 \leq 2\|\mathbf{x} - \mathbf{y}\|^2 + 2\|\mathbf{y}\|^2$  (to prove the identity, observe that  $\|\cdot\|^2$  is a convex function, and use Jensen's inequality,  $\|\mathbf{x}/2\|^2 \leq 1/2\|\mathbf{x} - \mathbf{y}\|^2 + 1/2\|\mathbf{y}\|^2$ ), and (b) follows by definition. (c) is true because, given  $\mathcal{F}_{n-1}$ ,  $(\tilde{\mathbf{g}}_j^n - \mathbf{g}_j^n)$  is independent of  $(\tilde{\mathbf{g}}_k^n - \mathbf{g}_k^n)$  for  $k \neq j$ . (d) follows from Jensen's Inequality, and (e) is true by Lemma F.2.

Performing a similar calculation again,

$$\begin{aligned}
\mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right] & \leq 2 \mathbb{E} \left[ \|\tilde{\mathbf{g}}_j^n - \mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right] + 2 \mathbb{E} \left[ \|\mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right], \\
& \leq 2\tau_n\sigma^2 + 2 \mathbb{E} \left[ \|\mathbf{g}_j^n\|^2 \middle| \mathcal{F}_{n-1} \right]. \tag{49}
\end{aligned}$$

Using Jensen's inequality, we get,

$$\|\mathbf{g}_j^n\|^2 \leq \tau_n \sum_{b=1}^{\tau_n} \|\mathbf{g}_j^{b,n}\|^2. \tag{50}$$

Substituting (48), (49) and (50) in (47), we get the result,

$$\mathbb{E} \left[ \|\tilde{\mathbf{g}}_Q^n\|^2 \middle| \mathcal{F}_{n-1} \right] \leq \frac{2\tau_n}{m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \left( \frac{q_j^n}{m} + 1 \right) \mathbb{E} \left[ \|\mathbf{g}_j^{b,n}\|^2 \middle| \mathcal{F}_{n-1} \right] + \left( \sum_{j=1}^m \frac{q_j^n}{m} + 1 \right) \frac{2\tau_n\sigma^2}{m}. \tag{51}$$

□

The following lemma bounds the inner product between the true gradient evaluated at the global model at the start of a round and the approximate gradient received by the server.

**Lemma F.4.** *Under Assumption 6, the FedCOM-V updates follow,*

$$-\mathbb{E} \left[ \langle \nabla f(\mathbf{w}^n), \tilde{\mathbf{g}}_Q^n \rangle \middle| \mathcal{F}_{n-1} \right] \leq \frac{1}{2m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \left( -\|\nabla f(\mathbf{w}^n)\|^2 - \mathbb{E} \left[ \|\mathbf{g}_j^{b,n}\|^2 \middle| \mathcal{F}_{n-1} \right] + L^2 \mathbb{E} \left[ \|\mathbf{w}^n - \mathbf{w}_j^{b,n}\|^2 \middle| \mathcal{F}_{n-1} \right] \right).$$

*Proof.*

$$\begin{aligned} -\mathbb{E} \left[ \langle \nabla f(\mathbf{w}^n), \tilde{\mathbf{g}}_Q^n \rangle \middle| \mathcal{F}_{n-1} \right] &\stackrel{(a)}{=} -\mathbb{E} \left[ \langle \nabla f(\mathbf{w}^n), \mathbb{E} [\tilde{\mathbf{g}}_Q^n \middle| \mathcal{F}_{\tau_n, n}] \rangle \middle| \mathcal{F}_{n-1} \right], \\ &\stackrel{(b)}{=} -\mathbb{E} \left[ \langle \nabla f(\mathbf{w}^n), \tilde{\mathbf{g}}^n \rangle \middle| \mathcal{F}_{n-1} \right], \\ &\stackrel{(c)}{=} -\mathbb{E} \left[ \left\langle \nabla f(\mathbf{w}^n), \frac{1}{m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \tilde{\mathbf{g}}_j^{b,n} \right\rangle \middle| \mathcal{F}_{n-1} \right], \\ &\stackrel{(d)}{=} -\mathbb{E} \left[ \frac{1}{m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \langle \nabla f(\mathbf{w}^n), \mathbb{E} [\tilde{\mathbf{g}}_j^{b,n} \middle| \mathcal{F}_{b-1, n}] \rangle \middle| \mathcal{F}_{n-1} \right], \\ &\stackrel{(e)}{=} -\mathbb{E} \left[ \frac{1}{m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \langle \nabla f(\mathbf{w}^n), \mathbf{g}_j^{b,n} \rangle \middle| \mathcal{F}_{n-1} \right], \\ &\stackrel{(f)}{=} \mathbb{E} \left[ \frac{1}{2m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \left( -\|\nabla f(\mathbf{w}^n)\|^2 - \|\mathbf{g}_j^{b,n}\|^2 + \|\nabla f(\mathbf{w}^n) - \mathbf{g}_j^{b,n}\|^2 \right) \middle| \mathcal{F}_{n-1} \right], \\ &\stackrel{(g)}{\leq} \mathbb{E} \left[ \frac{1}{2m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \left( -\|\nabla f(\mathbf{w}^n)\|^2 - \|\mathbf{g}_j^{b,n}\|^2 + L^2 \|\mathbf{w}^n - \mathbf{w}_j^{b,n}\|^2 \right) \middle| \mathcal{F}_{n-1} \right], \end{aligned}$$

(a) follows by Law of Iterated Expectations since  $\nabla f(\mathbf{w}^n)$  is measurable in  $\mathcal{F}_{\tau_n, n}$ . (b) follows since  $\tilde{\mathbf{g}}_Q^n$  is an unbiased estimate of  $\tilde{\mathbf{g}}^n$  by Assumption 8. (c) follows from the definition of  $\tilde{\mathbf{g}}^n$ . (d) follows from the Law of Iterated Expectations. (e) follows from the unbiased property of the stochastic gradients as stated in Assumption 7. (f) follows from the relation  $2\langle \mathbf{x}, \mathbf{y} \rangle = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2$ . (g) follows from  $L$ -smoothness of Assumption 6 since  $\mathbf{g}_j^{b,n} = \nabla f(\mathbf{w}_j^{b,n})$ . □

The following Lemma bounds the distance between the global model at the start of a round to the local model at a local step of a client.

**Lemma F.5.** *Under Assumption 7, FedCOM-V updates follow,*

$$\mathbb{E} \left[ \|\mathbf{w}^n - \mathbf{w}_j^{b,n}\|^2 \middle| \mathcal{F}_{n-1} \right] \leq 2\eta^2 \tau_n \sum_{b=1}^{\tau_n} \mathbb{E} \left[ \|\mathbf{g}_j^{b,n}\|^2 \middle| \mathcal{F}_{n-1} \right] + 2\eta_n^2 \tau_n \sigma^2.$$

*Proof.*

$$\begin{aligned}
\mathbb{E} \left[ \left\| \mathbf{w}^n - \mathbf{w}_j^{b,n} \right\|^2 \middle| \mathcal{F}_{n-1} \right] &\stackrel{(a)}{=} \mathbb{E} \left[ \left\| \eta_n \sum_{a=1}^{b-1} \tilde{\mathbf{g}}_j^{a,n} \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
&\stackrel{(b)}{\leq} 2\eta_n^2 \mathbb{E} \left[ \left\| \sum_{a=1}^{b-1} (\tilde{\mathbf{g}}_j^{a,n} - \mathbf{g}_j^{a,n}) \right\|^2 \middle| \mathcal{F}_{n-1} \right] + 2\eta_n^2 \mathbb{E} \left[ \left\| \sum_{a=1}^{b-1} \mathbf{g}_j^{a,n} \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
&\stackrel{(c)}{\leq} 2\eta_n^2 (b-1)\sigma^2 + 2\eta_n^2 \mathbb{E} \left[ \left\| \sum_{a=1}^{b-1} \mathbf{g}_j^{a,n} \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
&\stackrel{(d)}{\leq} 2\eta_n^2 (b-1)\sigma^2 + 2\eta_n^2 (b-1) \sum_{a=1}^{b-1} \mathbb{E} \left[ \left\| \mathbf{g}_j^{a,n} \right\|^2 \middle| \mathcal{F}_{n-1} \right], \\
&\leq 2\eta_n^2 \tau_n \sigma^2 + 2\eta_n^2 \tau_n \sum_{a=1}^{\tau_n} \mathbb{E} \left[ \left\| \mathbf{g}_j^{a,n} \right\|^2 \middle| \mathcal{F}_{n-1} \right].
\end{aligned}$$

(a) follows from the local update rule of Algorithm 2. (b) follows from the identity,  $\|\mathbf{x}\|^2 \leq 2\|\mathbf{x} - \mathbf{y}\|^2 + 2\|\mathbf{y}\|^2$ . (c) follows from a similar calculation as in the proof of Lemma F.2. (d) follows from Jensen's inequality.  $\square$

We prove Theorem 2 in two steps. First, in Theorem 4, we bound the convergence rate of FedCOM-V for a general choice of learning rates and local computations. Then, in Theorem 5, we prove a more explicit form of Theorem 2 for a specific choice of learning rates and local computations.

**Theorem 4.** *Under Assumptions 6 to 8, if the local learning rates  $(\eta_n)$ , local computations  $(\tau_n)$  and global learning rates  $(\gamma_n)$  satisfy,*

$$1 \geq 2\tau_n^2 L^2 \eta_n^2 + 2 \left( \frac{q_{\max}}{m} + 1 \right) \eta_n \gamma_n L \tau_n, \quad \forall n,$$

*then, the FedCOM-V updates satisfy,*

$$\frac{\sum_{n=1}^r \eta_n \tau_n \gamma_n \mathbb{E} \left[ \left\| \nabla f(\mathbf{w}^n) \right\|^2 \middle| \mathcal{F}^C \right]}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} \leq \frac{2(f(\mathbf{w}^{(0)}) - f(\mathbf{w}^*))}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} + \frac{2L\sigma^2 \sum_{n=0}^{r-1} \eta_n^2 \tau_n \gamma_n^2 (\bar{q}^n + 1)}{m \sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} + 2L^2 \sigma^2 \frac{\sum_{n=0}^{r-1} \eta_n^3 \tau_n^2 \gamma_n}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n}.$$

*Proof.* Recall the global update rule,  $\mathbf{w}^{n+1} = \mathbf{w}^n - \eta_n \gamma_n \tilde{\mathbf{g}}_Q^n$ . From the  $L$ -smoothness of  $f(\cdot)$ , we can write,

$$f(\mathbf{w}^{n+1}) - f(\mathbf{w}^n) \leq -\eta_n \gamma_n \langle \nabla f(\mathbf{w}^n), \tilde{\mathbf{g}}_Q^n \rangle + \frac{\eta_n^2 \gamma_n^2 L}{2} \left\| \tilde{\mathbf{g}}_Q^n \right\|^2.$$

Now, we bound the conditional expectation,

$$\begin{aligned} \mathbb{E} [f(\mathbf{w}^{n+1}) - f(\mathbf{w}^n) | \mathcal{F}_{n-1}] &\leq -\eta_n \gamma_n \mathbb{E} [\langle \nabla f(\mathbf{w}^n), \tilde{\mathbf{g}}_Q^n \rangle | \mathcal{F}_{n-1}] + \frac{\eta_n^2 \gamma_n^2 L}{2} \mathbb{E} [\|\tilde{\mathbf{g}}_Q^n\|^2 | \mathcal{F}_{n-1}], \\ &\stackrel{(a)}{\leq} \frac{\eta_n \gamma_n}{2m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \left( -\|\nabla f(\mathbf{w}^n)\|^2 - \mathbb{E} [\|\mathbf{g}_j^{b,n}\|^2 | \mathcal{F}_{n-1}] + L^2 \mathbb{E} [\|\mathbf{w}^n - \mathbf{w}_j^{b,n}\|^2 | \mathcal{F}_{n-1}] \right) \\ &\quad + \frac{2\tau_n L \eta_n^2 \gamma_n^2}{2m} \left( \frac{q_{\max}}{m} + 1 \right) \sum_{j=1}^m \sum_{b=1}^{\tau_n} \mathbb{E} [\|\mathbf{g}_j^{b,n}\|^2 | \mathcal{F}_{n-1}] + (\bar{q}^n + 1) \frac{2\tau_n L \eta_n^2 \gamma_n^2 \sigma^2}{2m}, \end{aligned} \quad (52)$$

$$\begin{aligned} &\stackrel{(b)}{\leq} \frac{\eta_n \gamma_n}{2m} \sum_{j=1}^m \sum_{b=1}^{\tau_n} \left( -\|\nabla f(\mathbf{w}^n)\|^2 - \mathbb{E} [\|\mathbf{g}_j^{b,n}\|^2 | \mathcal{F}_{n-1}] + 2L^2 \eta_n^2 \tau_n \sum_{b'=1}^{\tau_n} \mathbb{E} [\|\mathbf{g}_j^{b',n}\|^2 | \mathcal{F}_{n-1}] \right) \\ &\quad + 2L^2 \eta_n^2 \tau_n \sigma^2 + \frac{\tau_n L \eta_n^2 \gamma_n^2}{m} \left( \frac{q_{\max}}{m} + 1 \right) \sum_{j=1}^m \sum_{b=1}^{\tau_n} \mathbb{E} [\|\mathbf{g}_j^{b,n}\|^2 | \mathcal{F}_{n-1}] + (\bar{q}^n + 1) \frac{\tau_n L \eta_n^2 \gamma_n^2 \sigma^2}{m}, \end{aligned} \quad (53)$$

$$\begin{aligned} &\stackrel{(c)}{\leq} -\frac{\eta_n \gamma_n \tau_n}{2} \|\nabla f(\mathbf{w}^n)\|^2 + \frac{L \tau_n \eta_n^2 \gamma_n}{m} (m L \tau_n \eta_n + \gamma_n (\bar{q}^n + 1)) \sigma^2 \\ &\quad - \frac{\eta_n \gamma_n}{2m} \left( 1 - 2L^2 \eta_n^2 \tau_n^2 - 2L \tau_n \eta_n \gamma_n \left( \frac{q_{\max}}{m} + 1 \right) \right) \sum_{j=1}^m \sum_{b=1}^{\tau_n} \mathbb{E} [\|\mathbf{g}_j^{b,n}\|^2 | \mathcal{F}_{n-1}], \\ &\stackrel{(d)}{\leq} -\frac{\eta_n \gamma_n \tau_n}{2} \|\nabla f(\mathbf{w}^n)\|^2 + \frac{L \tau_n \eta_n^2 \gamma_n}{m} (m L \tau_n \eta_n + \gamma_n (\bar{q}^n + 1)) \sigma^2, \end{aligned} \quad (54)$$

where, (a) is obtained by using Lemmas F.3 and F.4, and (b) is obtained by using Lemma F.5. (c) is a rearrangement of terms, and (d) follows from the premise of the theorem,

$$1 \geq 2\tau_n^2 L^2 \eta_n^2 + 2 \left( \frac{q_{\max}}{m} + 1 \right) \eta_n \gamma_n L \tau_n.$$

Taking an expectation, rearranging terms and summing up equation (54) for all the rounds  $r$ , we have a telescopic cancellation to get,

$$\begin{aligned} \frac{\sum_{n=1}^r \eta_n \tau_n \gamma_n \mathbb{E} [\|\nabla f(\mathbf{w}^n)\|^2 | \mathcal{F}^C]}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} &= \frac{2(f(\mathbf{w}^{(0)}) - f(\mathbf{w}^r))}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} + \frac{2L\sigma^2 \sum_{n=0}^{r-1} \eta_n^2 \tau_n \gamma_n^2 (\bar{q}^n + 1)}{m \sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} + 2L^2 \sigma^2 \frac{\sum_{n=0}^{r-1} \eta_n^3 \tau_n^2 \gamma_n}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n}, \\ &\leq \frac{2(f(\mathbf{w}^{(0)}) - f(\mathbf{w}^*))}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} + \frac{2L\sigma^2 \sum_{n=0}^{r-1} \eta_n^2 \tau_n \gamma_n^2 (\bar{q}^n + 1)}{m \sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n} + 2L^2 \sigma^2 \frac{\sum_{n=0}^{r-1} \eta_n^3 \tau_n^2 \gamma_n}{\sum_{n=0}^{r-1} \eta_n \tau_n \gamma_n}. \end{aligned} \quad (55)$$

□

**Theorem 5.** *If we choose the learning rates for round  $n$ ,  $\eta_n$  and  $\gamma_n$ , and the number of local computations  $\tau_n$  as,*

$$\eta_n = \frac{c_\eta}{Ln}, \quad \gamma_n = \frac{c_\gamma}{\sqrt{\bar{q}^n + 1}}, \quad \tau_n = \frac{n}{2c_\eta},$$

where,

$$c_\eta = 2 \left( \frac{L \Delta_f \sqrt{m}}{\sigma} \left( \frac{q_{\max}}{m} + 1 \right) \right)^2, \quad c_\gamma = \frac{1}{2 \left( \frac{q_{\max}}{m} + 1 \right)},$$

where,  $\Delta_f \triangleq \sqrt{2(f(\mathbf{w}^{(0)}) - f(\mathbf{w}^*))}/L$ , and, if FedCOM-V is run for  $r$  communication rounds such that,

$$\frac{r}{1 + \log r} \geq \max \left\{ \left( \frac{q_{\max}}{m} + 1 \right)^2 \frac{4L^2 \Delta_f^2 m \sqrt{q_{\max} + 1}}{\varepsilon}, \left( \frac{q_{\max}}{m} + 1 \right) \frac{12L^2 \Delta_f^2 \sigma \sum_{n=1}^r \sqrt{\bar{q}^n + 1}}{\varepsilon r} \right\}, \quad (56)$$

then we have,

$$\frac{\sum_{n=1}^r \eta_n \gamma_n \tau_n \mathbb{E} [\|f(\mathbf{w}^n)\|^2 | \mathcal{F}^C]}{\sum_{n=1}^r \eta_n \tau_n \gamma_n} \leq \varepsilon. \quad (57)$$

*Proof of Theorem 5.* This proof will use the result of Theorem 4. Therefore, we first check that the premise of Theorem 4 is satisfied. Due to the choice of  $\eta_n$ ,  $\tau_n$  and  $\gamma_n$ ,

$$\begin{aligned} 2\tau_n^2 L^2 \eta_n^2 + 2 \left( \frac{q_{\max}}{m} + 1 \right) \eta_n \gamma_n L \tau_n &= 1/2 + \left( \frac{q_{\max}}{m} + 1 \right) c_\gamma, \\ &= 1. \end{aligned}$$

Therefore, the following result of Theorem 4 holds true.

$$\frac{\sum_{n=1}^r \eta_n \tau_n \gamma_n \mathbb{E} \left[ \|\nabla f(\mathbf{w}^n)\|^2 \mid \mathcal{F}^C \right]}{\sum_{n=1}^r \eta_n \tau_n \gamma_n} \leq \underbrace{\frac{2(f(\mathbf{w}^0) - f(\mathbf{w}^*))}{\sum_{n=1}^r \eta_n \tau_n \gamma_n} + \frac{2L\sigma^2 \sum_{n=1}^r \eta_n^2 \tau_n \gamma_n^2 (\bar{q}^n + 1)}{m \sum_{n=1}^r \eta_n \tau_n \gamma_n}}_{\text{I}} + \underbrace{2L^2 \sigma^2 \frac{\sum_{n=1}^r \eta_n^3 \tau_n^2 \gamma_n}{\sum_{n=1}^r \eta_n \tau_n \gamma_n}}_{\text{II}}.$$

Consider I,

$$\begin{aligned} \text{I} &\stackrel{(a)}{=} \frac{4L(f(\mathbf{w}^{(0)}) - f(\mathbf{w}^*))}{c_\gamma \sum_{n=1}^r 1/\sqrt{\bar{q}^n + 1}} + \frac{2\sigma^2 c_\eta c_\gamma}{m} \frac{\sum_{n=1}^r 1/n}{\sum_{n=1}^r 1/\sqrt{\bar{q}^n + 1}}, \\ &\stackrel{(b)}{\leq} \frac{4L(f(\mathbf{w}^{(0)}) - f(\mathbf{w}^*))}{c_\gamma \sum_{n=1}^r 1/\sqrt{\bar{q}^n + 1}} + \frac{2\sigma^2 c_\eta c_\gamma}{m} \frac{1 + \log r}{\sum_{n=1}^r 1/\sqrt{\bar{q}^n + 1}}, \\ &\stackrel{(c)}{\leq} \frac{2(1 + \log r)}{\sum_{n=1}^r 1/\sqrt{\bar{q}^n + 1}} \left( \frac{L^2 \Delta_f^2}{c_\gamma} + \frac{\sigma^2 c_\eta c_\gamma}{m} \right), \\ &\stackrel{(d)}{=} \frac{1 + \log r}{\sum_{n=1}^r 1/\sqrt{\bar{q}^n + 1}} 6L^2 \Delta_f^2 \left( \frac{q_{\max}}{m} + 1 \right), \\ &\stackrel{(e)}{\leq} 6L^2 \Delta_f^2 \left( \frac{q_{\max}}{m} + 1 \right) \frac{1 + \log r}{r} \frac{\sum_{n=1}^r \sqrt{\bar{q}^n + 1}}{r}. \end{aligned} \tag{58}$$

(a) is obtained by substituting the expressions for  $\eta_n$ ,  $\gamma_n$  and  $\tau_n$ . (b) is obtained by the bound,  $\sum_{n=1}^r 1/n \leq 1 + \log r$ . (c) is a rearrangement of terms and the bound  $1 \leq 1 + \log r$ . (d) is obtained by substituting the expressions for  $c_\eta$  and  $c_\gamma$ . (e) is obtained using the result that the harmonic mean is smaller than the arithmetic mean.

Consider II,

$$\begin{aligned} \text{II} &= 2L^2 \sigma^2 \frac{\sum_{n=1}^r \eta_n^3 \tau_n^2 \gamma_n}{\sum_{n=1}^r \eta_n \tau_n \gamma_n}, \\ &= \sigma^2 c_\eta \frac{\sum_{n=1}^r 1/(n\sqrt{\bar{q}^n + 1})}{\sum_{n=1}^r 1/\sqrt{\bar{q}^n + 1}}, \\ &\leq \sigma^2 c_\eta \sqrt{q_{\max} + 1} \frac{\sum_{n=1}^r 1/n}{r}, \\ &\leq \sigma^2 c_\eta \sqrt{q_{\max} + 1} \frac{1 + \log r}{r}, \\ &= 2 \left( \frac{q_{\max}}{m} + 1 \right)^2 L^2 \Delta_f^2 m \sqrt{q_{\max} + 1} \frac{1 + \log r}{r}. \end{aligned} \tag{59}$$

Therefore, if we chose the total number of communication rounds  $r$  such that,

$$6L^2 \Delta_f^2 \left( \frac{q_{\max}}{m} + 1 \right) \frac{1 + \log r}{r} \frac{\sum_{n=1}^r \sqrt{\bar{q}^n + 1}}{r} \leq \frac{\varepsilon}{2}, \tag{60}$$

and,

$$2 \left( \frac{q_{\max}}{m} + 1 \right)^2 L^2 \Delta_f^2 m \sqrt{q_{\max} + 1} \frac{1 + \log r}{r} \leq \frac{\varepsilon}{2}, \tag{61}$$

then (57) is satisfied. (60) and (61) are simply a restatement of (56). This completes the proof.  $\square$

If the compression parameters  $(\mathbf{Q}^n)_n$  formed a stationary process with a stationary distribution according to a random variable  $\mathbf{Q}$ , then  $\sum_{n=1}^r \sqrt{\bar{Q}^n + 1}/r \rightarrow \mathbb{E}[\sqrt{\bar{Q} + 1}]$ . Therefore, Theorem 5 proves Theorem 2.