Using Predictions in Online Optimization: Looking Forward with an Eye on the Past*

Niangjun Chen[†], Joshua Comden[‡], Zhenhua Liu[‡], Anshul Gandhi[‡], Adam Wierman[†]

[†]California Institute of Technology, Pasadena, CA, E-mail: {ncchen, adamw}@caltech.edu

[‡]Stony Brook University, NY, E-mail: {joshua.comden, zhenhua.liu, anshul.gandhi}@stonybrook.edu

ABSTRACT

We consider online convex optimization (OCO) problems with switching costs and noisy predictions. While the design of online algorithms for OCO problems has received considerable attention, the design of algorithms in the context of noisy predictions is largely open. To this point, two promising algorithms have been proposed: Receding Horizon Control (RHC) and Averaging Fixed Horizon Control (AFHC). The comparison of these policies is largely open. AFHC has been shown to provide better worst-case performance, while RHC outperforms AFHC in many realistic settings. In this paper, we introduce a new class of policies, Committed Horizon Control (CHC), that generalizes both RHC and AFHC. We provide average-case analysis and concentration results for CHC policies, yielding the first analysis of RHC for OCO problems with noisy predictions. Further, we provide explicit results characterizing the optimal CHC policy as a function of properties of the prediction noise, e.g., variance and correlation structure. Our results provide a characterization of when AFHC outperforms RHC and vice versa, as well as when other CHC policies outperform both RHC and AFHC.

1. INTRODUCTION

In an online convex optimization (OCO) problem, an algorithm interacts with an environment in a sequence of rounds. In round t the algorithm chooses an action x_t from a convex decision/action space F , the environment reveals a convex cost function h_t , and the algorithm pays cost $h_t(x_t)$. The goal of the algorithm is to minimize cost over a horizon T .

OCO has a rich history, with applications in wide-ranging areas of computer science and beyond [54, 27, 21, 50, 33, 34, 35, 36, 37]. In recent years, OCO has seen considerable interest from applications in the networking and distributed systems communities. In particular, OCO has enabled novel designs for dynamic capacity planning, load shifting and demand response for data centers [29, 34, 35, 36, 40], geographical load balancing of internet-scale systems [33, 47],

electrical vehicle charging [15, 29], video streaming [41, 25], and thermal management of systems-on-chip [51, 52, 7].

Applications of OCO in the networking and distributed systems communities typically differ in two significant ways from the classical OCO literature: (i) actions incur *switching costs* and (ii) *noisy predictions* about the future are available.

Switching costs capture the cost that is incurred by systems when moving from one state to another. This is modeled by adding an extra term to the cost paid by the algorithm in each round, i.e., the cost becomes $h_t(x_t) + \beta \|x_t - x_{t-1}\|$, where $\|\cdot\|$ is a norm (often the one-norm), and $\beta \in \mathbb{R}^+$. This additional term models, e.g., the cost of turning servers on/off in dynamic capacity planning [29, 34, 35, 36, 40, 17] or the cost of changing a quality level in the case of video streaming [24, 25]. The addition of switching costs makes the algorithmic problem harder as it forces current actions to depend on beliefs about future cost functions.

Predictions are of great importance in networking and distributed systems. Despite the considerable noise that is often inherent in forecasts, predictions can be extremely useful. For example, predictions of future demand are critical in the case of dynamic capacity planning in data centers to ensure sufficient capacity [16, 30, 19, 29, 34, 35, 36, 40, 17]. Unfortunately, designing online algorithms that exploit noisy predictions is an open, challenging topic.

In this paper, we focus on *OCO problems that have both switching costs and noisy predictions*. While there is a significant literature on OCO problems with switching costs [13, 4, 20, 5, 33, 34, 6], there is much less work studying the impact of predictions [33, 13, 34]. Further, the analytic work that does focus on predictions typically assumes perfect lookahead, the lone exception being Chen et al. [13].

Two promising algorithms. Perhaps the most natural starting point for studying algorithms for OCO problems with switching costs and noisy predictions is the class of Model Predictive Control (MPC) algorithms. MPC is a prominent and widely-studied class of algorithms in the control theory community [8, 48, 31, 12, 18, 38], and much of the work studying predictions in OCO problems has focused on MPC and its variants, e.g., [33, 15, 14].

From these works, two promising algorithms have emerged: Receding Horizon Control (RHC) and Averaging Fixed Horizon Control (AFHC). (See Section 3 for formal definitions of these two algorithms.) Both algorithms use a prediction horizon/window of size w , but make decisions in very different ways. RHC considers, at each point in time, the predictions available in the current horizon, determines the trajectory of w actions that minimize the cost within that horizon, and then commits *only* the first action in that trajectory. By contrast, AFHC works by averaging the actions of multiple Fixed Horizon Control (FHC) algorithms, each of which work similarly to RHC but commit to *all* w actions in a given prediction horizon.

*This work is partially supported by the NSF through CNS-1464388, CNS-1464151, CNS-1319820, NETS-1518941 and an A*STAR NSS (PhD) scholarship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGMETRICS '16, June 14–18, 2016, Antibes Juan-Les-Pins, France.

© 2016 ACM. ISBN 978-1-4503-4266-7/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2896377.2901464>

Receding Horizon Control (RHC) has a long history in the control theory literature [8, 48, 31, 12, 18, 38], but was first studied analytically in the context of OCO in [33]. In [33], RHC was proven to have a competitive ratio (the ratio of the cost incurred by RHC to the cost incurred by the offline optimal algorithm) of $1 + O(1/w)$ in the one-dimensional setting, where w is the size of the prediction window. However, the competitive ratio of RHC is $1 + \Omega(1)$ in the general case, and thus does *not* decrease to one as the prediction window grows in the worst case; this is despite the fact that predictions are assumed to have no noise (the perfect lookahead model). To this point there is no analytic work characterizing the performance of RHC with noisy predictions.

The poor worst-case performance of RHC motivated the proposal of Averaging Fixed Horizon Control (AFHC) [33], which provides an interesting contrast. While RHC is entirely “forward looking”, AFHC keeps an “eye on the past” by respecting the actions of FHC algorithms in previous timesteps and thus avoids switching costs incurred by moving too quickly between actions. As a result, AFHC achieves a competitive ratio of $1 + O(1/w)$ in both single and multi-dimensional action spaces, under the assumption of perfect lookahead, [33]. Further, strong guarantees on the performance of AFHC have been established in the case of noisy predictions [13].

Surprisingly, while the competitive ratio of AFHC is smaller than that of RHC, RHC provides *better* performance than AFHC in many practical cases. Further, RHC is seemingly more resistant to prediction noise in many settings (see Figure 1 for an example), though no analytic results are known for this case. Thus, at this point, two promising algorithms have been proposed, but it is unclear in what settings each should be used and it is unclear if there are other algorithms that dominate these two proposals.

Contributions of this paper. The goal of this paper is to provide new insights into the design of algorithms for OCO problems with noisy predictions. In particular, our results highlight the importance of *commitment* in online algorithms, and the significant performance gains that can be achieved by tuning the *commitment level* of an algorithm as a function of structural properties of the prediction noise such as variance and correlation structure.

In terms of commitment, receding horizon control (RHC) and averaging fixed horizon control (AFHC) represent two extreme algorithm designs – RHC commits to only one action at a time whereas AFHC averages over algorithms that commit to actions spanning the whole prediction horizon. While the non-committal nature of RHC enables quick response to improved predictions, it makes RHC *susceptible to switching costs*. On the other hand, the cautious nature of AFHC averts switching costs but makes it entirely *dependent on the accuracy of predictions*.

Motivated by these deficiencies in existing algorithm design, we introduce a new class of policies, **Committed Horizon Control** (CHC), that allows for arbitrary levels of commitment and thus subsumes RHC and AFHC. We present both average-case analysis (Theorems 1 and 6) and concentration results (Theorems 7) for CHC policies. In doing so, we provide the *first analysis* of RHC with noisy predictions.

Our results demonstrate that intermediate levels of commitment can provide significant reductions in cost, to the tune of more than 50% (e.g., Figure 4(a), Figure 5(a) and Figure 6(a)). Further, our results also reveal the impact of correlation structure and variance of prediction noise on the optimal level of commitment, and provide simple guidelines on how to choose between RHC and AFHC.

These results are enabled by a key step in our proof that transforms the control strategy employed by the offline optimal algorithm, OPT to the strategy of CHC via a trajectory

of intermediate strategies. We exploit the structure of our algorithm at each intermediate step to bound the difference in costs; the sum of these costs over the entire transformation then gives us a bound on the difference in costs between OPT and CHC.

To summarize, this paper makes the following contributions to the literature on OCO with noisy predictions:

- We provide the first analysis of RHC for OCO problems with noisy predictions.
- We characterize when RHC/AFHC is better as a function of the correlation structure and variance of prediction noise.
- We introduce and analyze a new class of Committed Horizon Control (CHC) policies that generalizes AFHC and RHC.
- We highlight how the commitment level of a policy should be tuned depending on structural properties of prediction noise. By optimizing the level of commitment, CHC policies can achieve performance improvements of more than 50% over AFHC and RHC.

2. PROBLEM FORMULATION

We consider online convex optimization (OCO) problems with switching costs and noisy predictions. We first introduce OCO with switching costs (Section 2.1) and then describe the model of prediction noise (Section 2.2). Finally, we discuss the performance metric we consider in this paper – the competitive difference – and how it relates to common measures such as regret and competitive ratio (Section 2.3).

2.1 OCO with switching costs

An OCO problem with switching costs considers a convex, compact decision/action space $F \subset \mathbb{R}^n$ and a sequence of cost functions $\{h_1, h_2, \dots\}$, where each $h_t : F \rightarrow \mathbb{R}^+$ is convex, and F is a compact set.

At time t , the following sequence occurs: (i) the online algorithm first chooses an action, which is a vector $x_t \in F \subset \mathbb{R}^n$, (ii) the environment chooses a cost function h_t from a set $\mathcal{C}$, and (iii) the algorithm pays a stage cost $h_t(x_t)$ and a switching cost $\beta \|x_t - x_{t-1}\|$, where $\beta \in \mathbb{R}^+$, and $\|\cdot\|$ can be any norm in $\mathbb{R}^n$, and F is bounded in terms of this norm, i.e., $\|x - y\| \leq D$ for all $x, y \in F$.

Motivated by path planning and image labeling problems [45, 13, 42], we consider a variation of the above that uses a *parameterized* cost function $h_t(x_t) = h(x_t, y_t)$, where the parameter $y_t \in \mathbb{R}^m$ is the focus of prediction. This yields a total cost over T rounds of

$$\min_{x_t \in F} \sum_{t=1}^T h(x_t, y_t) + \beta \|x_t - x_{t-1}\|, \quad (1)$$

Note that prior work [13] considers only the case where a least-square penalty is paid each round, i.e., an online LASSO formulation with $h(x_t, y_t) = \frac{1}{2} \|y_t - Kx_t\|_2^2$. In this paper, we consider more general h . We impose that $h(x_t, y_t)$ is separately convex in both x_t and y_t along with the following smoothness criteria.

DEFINITION 1. A function h is α -Hölder continuous in the second argument for $\alpha \in \mathbb{R}^+$, i.e., for all $x \in F$, there exists $G \in \mathbb{R}^+$, such that

$$|h(x, y_1) - h(x, y_2)| \leq G \|y_1 - y_2\|_2^\alpha, \forall y_1, y_2.$$

G and α control the sensitivity of the cost function to a disturbance in y .

For this paper, we focus on $\alpha \leq 1$, since the only α -Hölder continuous function with $\alpha > 1$ is the constant function [2].

When $\alpha = 1$, h is G -Lipschitz in the second argument; if h is differentiable in the second argument, this is equivalent to $\|\partial_y h(x, y)\|_2 \leq G$, $\forall x, y$.

2.2 Modeling prediction noise

Predictions about the future play a crucial role in almost all online decision problems. However, while significant effort has gone into designing predictors, e.g., [53, 43, 44, 26], much less work has gone into *integrating* predictions efficiently into algorithms. This is, in part, due to a lack of tractable, practical models for prediction noise. As a result, most papers that study online decision making problems, such as OCO, use numerical simulations to evaluate the impact of prediction noise, e.g., [1, 3, 36, 40].

The papers that do consider analytic models often use either i.i.d. prediction noise or worst-case bounds on prediction errors for tractability. An exception is the recent work [13, 15] which introduces a model for prediction noise that captures three important features of real predictors: (i) it allows for correlations in prediction errors (both short range and long range); (ii) the quality of predictions decreases the further in the future we try to look ahead; and (iii) predictions about the future are refined as time passes. Further, [13] shows that it is tractable in the context of OCO. Thus, we adopt the model from [13] for this paper.

Specifically, throughout this paper we model prediction error via the following equation:

$$y_t - y_{t|\tau} = \sum_{s=\tau+1}^t f(t-s)e(s), \quad (2)$$

where $y_{t|\tau}$ is the prediction of y_t made at time $\tau < t$. This model characterizes prediction error as white noise being passed through a causal filter. In particular, the prediction error is a weighted linear combination of per-step noise terms $e(s)$ with weights $f(t-s)$, where f is a deterministic impulse response function. The noise terms $e(s)$ are assumed to be uncorrelated with mean zero and positive definite covariance R_e ; let $\sigma^2 = \text{tr}(R_e)$. Further, the impulse response function f is assumed to satisfy $f(0) = I$ and $f(t) = 0$ for $t < 0$.

Note that i.i.d. prediction noise can be recovered by imposing that $f(0) = I$ and $f(t) = 0$ for all $t \neq 0$. Further, the model can represent prediction errors that arise from classical filters such as Wiener filters and Kalman filters (see [13]). In both cases the impulse response function decays as $f(s) \sim \eta^s$ for some $\eta < 1$.

These examples highlight that the form of the impulse response function captures the degree of short-term/long-term correlation in prediction errors. The form of the correlation structure plays a key role in the performance results we prove, and its impact can be captured through the following definition. For any $k > 0$, let $\|f_k\|^2$ be the two norm square of prediction error covariance over k steps of prediction, i.e.,

$$\|f_k\|^2 = \text{tr}(\mathbb{E}[\delta y_k \delta y_k^T]) = \text{tr}(R_e \sum_{s=0}^k f(s)^T f(s)), \quad (3)$$

where $\delta y_k^T = y_{t+k} - y_{t+k|t} = \sum_{s=t+1}^{t+k} f(t+k-s)e(s)$. Derivation of (3) can be found in Appendix B.1 Equation (19).

2.3 The competitive difference

For any algorithm ALG that comes up with feasible actions $x_{ALG,t} \in F$, $\forall t$, the cost of the algorithm over the horizon can be written as

$$\text{cost}(ALG) = \sum_{t=1}^T h(x_{ALG,t}, y_t) + \beta \|x_{ALG,t} - x_{ALG,t-1}\| \quad (4)$$

We compare the performance of our online algorithm against the optimal offline algorithm OPT , which makes the optimal decision with full knowledge of the trajectory of y_t .

$$\text{cost}(OPT) = \min_{x_t \in F} \sum_{t=1}^T h(x_t, y_t) + \beta \|x_t - x_{t-1}\| \quad (5)$$

The results in this paper bound the *competitive difference* of algorithms for OCO with switching costs and prediction noise. Informally, the competitive difference is the additive gap between the cost of the online algorithm and the cost of the offline optimal.

To define the competitive difference formally in our setting we need to first consider how to specify the instance. To do this, let us first return to the specification of the prediction model in (2) and expand the summation all the way to time zero. This expansion highlights that the process y_t can be viewed as a random deviation around the predictions made at time zero, $y_{t|0} := \hat{y}_t$, which are specified externally to the model:

$$y_t = \hat{y}_t + \sum_{s=1}^t f(t-s)e(s). \quad (6)$$

Thus, an instance can be specified either via the process y_t or via the initial predictions $\hat{y}_t$, and then the random noise from the model determines the other. The latter is preferable for analysis, and thus we state our definition of competitive difference (and our theorems) using this specification.

DEFINITION 2. We say an online algorithm ALG has (expected) **competitive difference** at most $\rho(T)$ if:

$$\sup_{\hat{y}} \mathbb{E}_e [\text{cost}(ALG) - \text{cost}(OPT)] \leq \rho(T). \quad (7)$$

Note that the expectation in the definition above is with respect to the prediction noise, $(e(t))_{t=1}^T$, and so both terms $\text{cost}(ALG)$ and $\text{cost}(OPT)$ are random. Unlike ALG, the offline optimal algorithm OPT knows each exact realization of e before making the decision.

Importantly, though we specify our results in terms of the competitive difference, it is straightforward to convert them into results about the competitive ratio and regret, which are more commonly studied in the OCO literature. Recall that the **competitive ratio** bounds the ratio of the algorithm's cost to that of OPT, and the **regret** bounds the difference between the algorithm's cost and the offline *static* optimal.

Converting a result on the competitive difference into a result on the competitive ratio requires lower bounding the offline optimal cost, and such a bound can be found in Theorem 6 of [33]. Similarly, converting a result on the competitive difference into a result on the regret requires lower bounding the offline static optimal cost, and such a bound can be found in Theorem 2 of [13].

3. ALGORITHM DESIGN

There is a large literature studying algorithms for online convex optimization (OCO), both with the goal of designing algorithms with small regret and algorithms with small competitive ratio.

These algorithms use a wide variety of techniques. For example, there are numerous algorithms that maintain sub-linear regret, e.g., *online gradient descent* (OGD) based algorithms [54, 21] and *Online Newton Step and Follow the Approximate Leader* algorithms [21]. (Note that the classical setting does not consider switching costs; however, [4] shows that similar regret bounds can be obtained when switching costs are considered.) By contrast, there only exist algorithms that achieve constant competitive ratio in limited

settings, e.g., [34] shows that, when F is a one-dimensional normed space, there exists a deterministic online algorithm that is 3-competitive. This is because, in general, obtaining a constant competitive ratio is impossible in the worst-case: [10] has shown that any deterministic algorithm must be $\Omega(n)$ -competitive given metric decision space of size n and [9] has shown that any randomized algorithm must be $\Omega(\sqrt{\log n / \log \log n})$ -competitive.

However, all of the algorithms and results described above are in the worst-case setting and do not consider algorithms that have noisy predictions available. Given noisy predictions, the most natural family of algorithms to consider come from the family of Model Predictive Control (MPC) algorithms, which is a powerful, prominent class of algorithms from the control community. In fact, the only analytic results for OCO problems with predictions to this point have come from algorithms inspired by MPC, e.g., [8, 48, 31, 12]. (Note that there is a large literature on such algorithms in control theory, e.g., [18, 38] and the references therein, but the analysis needed for OCO is different than from the stability analysis provided by the control literature.)

To this point, two promising candidate algorithms have emerged in the context of OCO: Receding Horizon Control (RHC) [32] and Averaging Fixed Horizon Control (AFHC) [33]. We discuss these two algorithms in Section 3.1 below and then introduce our novel class of Committed Horizon Control (CHC) algorithms, which includes both RHC and AFHC as special cases, in Section 3.2. The class of CHC algorithms is the focus of this paper.

3.1 Two promising algorithms

At this point the two most promising algorithms for integrating noisy predictions into solutions to OCO problems are RHC and AFHC.

Receding Horizon Control (RHC): RHC operates by determining, at each timestep t , the optimal actions over the window $(t+1, t+w)$, given the starting state x_t and a prediction window (horizon) of length w .

To state this more formally, let $y_{\cdot|\tau}$ denote the vector $(y_{\tau+1|\tau}, \dots, y_{\tau+w|\tau})$, the prediction of y in a w timestep prediction window at time τ . Define $X^{\tau+1}(x_\tau, y_{\cdot|\tau})$ as the vector in F^w indexed by $t \in \{\tau+1, \dots, \tau+w\}$, which is the solution to

$$\begin{aligned} \min_{x_{\tau+1}, \dots, x_{\tau+w}} \quad & \sum_{t=\tau+1}^{\tau+w} h(x_t, y_{t|\tau}) + \sum_{t=\tau+1}^{\tau+w} \beta \|x_t - x_{t-1}\|, \quad (8) \\ \text{subject to} \quad & x_t \in F. \end{aligned}$$

ALGORITHM 1 (RECEDING HORIZON CONTROL). For all $t \leq 0$, set $x_{RHC,t} = 0$. Then, at each timestep $\tau \geq 0$, set

$$x_{RHC,\tau+1} = X_{\tau+1}^{\tau+1}(x_{RHC,\tau}, y_{\cdot|\tau}) \quad (9)$$

RHC has a long history in the control theory literature, e.g., [8, 18, 38, 12]. However, there are few results known in the OCO literature, and most such results are negative. In particular, the competitive ratio of RHC with perfect lookahead window w is $1 + O(1/w)$ in the one-dimensional setting. The performance is not so good in the general case. In particular, outside of the one-dimensional case the competitive ratio of RHC is $1 + \Omega(1)$, i.e., the competitive ratio does not decrease to 1 as the prediction window w increases in the worst case [34].

Averaging Fixed Horizon Control (AFHC): AFHC provides an interesting contrast to RHC. RHC ignores all history – the decisions and predictions that led it to be in the current state – while AFHC constantly looks both backwards and forwards. Specifically, AFHC averages the choices made

by Fixed Horizon Control (FHC) algorithms. In particular, AFHC with prediction window size w averages the actions of w FHC algorithms, each with different predictions available to it. At time t , a FHC algorithm determines the optimal actions $x_{t+1}, \dots, x_{t+w}$ given a prediction window (horizon) of length w as done in RHC. But, then FHC implements all actions in the trajectory $x_{t+1}, \dots, x_{t+w}$ instead of just the first action x_t . Fixed Horizon Control algorithms are individually more naive than RHC, but by averaging them AFHC can provide improved worst-case performance compared to RHC. To define the algorithm formally, let

$$\Omega_k = \{i : i \equiv k \pmod{w}\} \cap [-w+1, T] \text{ for } k = 0, \dots, w-1.$$

ALGORITHM 2 (FIXED HORIZON CONTROL, VERSION k). $FHC^{(k)}(w)$, is defined in the following manner. For all $t \leq 0$, set $x_{FHC,t}^{(k)} = 0$. At timeslot $\tau \in \Omega_k$ (i.e., before $y_{\tau+1}$ is revealed), for all $t \in \{\tau+1, \dots, \tau+w\}$, use (8) to set

$$x_{FHC,t}^{(k)} = X_t^{\tau+1} \left(x_{FHC,\tau}^{(k)}, y_{\cdot|\tau} \right). \quad (10)$$

Note that, for $k \geq 1$, the algorithm starts from $\tau = k - w$ rather than $\tau = k$ in order to calculate $x_{FHC,t}^{(k)}$ for $t < k$.

While individual FHC can have poor performance, surprisingly, by averaging different versions of FHC we can obtain an algorithm with good performance guarantee. Specifically, $AFHC$ is defined as follows.

ALGORITHM 3 (AVERAGING FIXED HORIZON CONTROL). For all k , at each timeslot $\tau \in \Omega_k$, use $FHC^{(k)}$ to determine $x_{FHC,\tau+1}^{(k)}, \dots, x_{FHC,\tau+w}^{(k)}$, and for $t = 1, \dots, T$, set

$$x_{AFHC,t} = \frac{1}{w} \sum_{k=0}^{w-1} x_{FHC,t}^{(k)}. \quad (11)$$

In contrast to RHC, AFHC has a competitive ratio of $1 + O(1/w)$ regardless of the dimension of the action space in the perfect lookahead model [33]¹. This improvement of AFHC over RHC is illustrated in Figure 1(a), which shows for a specific setting with perfect lookahead, AFHC approaches the offline optimal with increasing prediction window size while RHC is relatively constant. (The setting used for the figure uses a simple model of a data center with a multi-dimensional action space, and is described in Appendix A.)

Comparing RHC and AFHC: Despite the fact that the worst-case performance of AFHC is dramatically better than RHC, RHC provides better performance than AFHC in realistic settings when prediction can be inaccurate in the lookahead window. For example, Figure 1(b) highlights that RHC can outperform AFHC by an arbitrary amount if the predictions are noisy. Specifically, if we make predictions accurate for a small window γ and then inaccurate for the remaining $(w - \gamma)$ steps of the lookahead window, AFHC is affected by the inaccurate predictions whereas RHC only acts on the correct ones. The tradeoff between the worst-case bounds and average-case performance across AFHC and RHC is also evident in the results shown in Figure 3 of [33].

The contrast between Figure 1(a) and 1(b) highlights that, at this point, it is unclear when one should use AFHC/RHC. In particular, AFHC is more robust but RHC may be better in many specific settings. Further, the bounds we have described so far say nothing about the impact of noise on the performance (and comparison) of these algorithms.

¹Note that this result assumes that there exists $\epsilon_0 > 0$, s.t. $h(x, y) \geq \epsilon_0 \cdot x, \forall x, y$, and the switching cost is $\beta \cdot (x_t - x_{t-1})^+$ where $(x)^+ = \max(x, 0)$.

Figure 1: Total cost of RHC and AFHC, normalized by the cost of the offline optimal, versus: (a) prediction window size, (b) number of steps of perfect prediction with $w = 10$. Note (a) and (b) were produced under different cost settings, see Appendix A.

3.2 A general class of algorithms

The contrast between the performance of receding horizon control (RHC) and averaging fixed horizon control (AFHC) in worst-case and practical settings is a consequence of the fact that RHC is entirely “forward looking” while AFHC keeps an “eye on the past”. However, both algorithms are extreme cases in that RHC does not consider any information that led it to its current state, while AFHC looks back at w FHC algorithms – every set of predictions that led to the current state.

One way to view this difference between RHC and AFHC is in terms of *commitment*. In particular, AFHC has FHC algorithms that commit to the w decisions at each timestep and then the final choice of the algorithm balances these commitments by averaging across them. In contrast, RHC commits only one step at a time.

Building on this observation, we introduce the class of Committed Horizon Control (CHC) algorithms. The idea behind the class is to allow commitment of a fixed number, say v , of steps. The minimal level of commitment, $v = 1$, corresponds to RHC and the maximal level of commitment, $v = w$, corresponds to AFHC. Thus, the class of CHC algorithms allows variation between these extremes.

Formally, to define the class of CHC algorithms we start by generalizing the class of FHC algorithms to allow limited commitment. An FHC algorithm with commitment level v uses a prediction window of size w but then executes (commits to) only the first $v \in [1, w]$ actions which can be visualized by Figure 2. To define this formally, let

$$\Psi_k = \{i : i \equiv k \pmod{v}\} \cap [-v + 1, T] \text{ for } k = 0, \dots, v - 1.$$

Fixed horizon control with lookahead window w and commitment level v , $FHC^{(k)}(v, w)$, is defined in the following manner. For notational convenience, we write $x^{(k)} \equiv x_{FHC^{(k)}(v, w)}^{(k)}$.

ALGORITHM 4 (FHC WITH LIMITED COMMITMENT). *For all $t \leq 0$, set $x_{FHC^{(k)}(v, w), t}^{(k)} = 0$. At timeslot $\tau \in \Psi_k$ (i.e., before $y_{\tau+1}$ is revealed), for all $t \in \{\tau + 1, \dots, \tau + v\}$, use (8) to set*

$$x_t^{(k)} = X_t^{\tau+1} \left(x_\tau^{(k)}, y_{|\tau} \right). \quad (12)$$

Note that, for $k \geq 1$, the algorithm starts from $\tau = k - v$ rather than $\tau = k$ in order to calculate $x_t^{(k)}$. We can see that FHC with limited commitment is very similar to FHC as both use (8) to plan w timesteps ahead, but here only the first v steps are committed to action.

$CHC(v, w)$, the CHC algorithm with prediction window w and commitment level v , averages over v FHC algorithms

Figure 2: Fixed Horizon Control with commitment level v : optimizes once every v timesteps for the next w timesteps and commits to use the first v of them.

with prediction window w and commitment level v . Figure 3 provides an overview of CHC. For conciseness in the rest of the paper, we will use $x_t^{(k)}$ to denote the action decided by $FHC^{(k)}(v, w)$ at time t .

Figure 3: Committed Horizon Control: at each timestep, it averages over all v actions defined by the v FHC algorithms with limited commitment.

ALGORITHM 5 (COMMITTED HORIZON CONTROL). *At each timeslot $\tau \in \Psi_k$, use $FHC^{(k)}(v, w)$ to determine $x_{\tau+1}^{(k)}, \dots, x_{\tau+v}^{(k)}$, and at timeslot $t \in 1, \dots, T$, $CHC(v, w)$ sets*

$$x_{CHC, t} = \frac{1}{v} \sum_{k=0}^{v-1} x_t^{(k)} \quad (13)$$

RHC and AFHC are the extreme levels of commitment in CHC policies and, as we see in the analysis that follows, it is often beneficial to use intermediate levels of commitment depending on the structure of prediction noise.

4. AVERAGE-CASE ANALYSIS

We now present the main technical results of this paper, which analyze the performance of committed horizon control (CHC) algorithms and address several open challenges relating to the analysis of receding horizon control (RHC) and averaging fixed horizon control (AFHC). In this section we characterize the average case performance of CHC as a function of the commitment level v of the policy and properties of the prediction noise, i.e., the variance of prediction noise $e(s)$ and the form of the correlation structure, $f(s)$. Concentration bounds are discussed in Section 5. All proofs are presented in Appendix B.

Our main result establishes *bounds on the competitive difference of CHC under noisy predictions*. Since CHC generalizes RHC and AFHC, our result also provides the *first* analysis of RHC with noisy predictions and further enables a comparison between RHC and AFHC based on the properties of the prediction noise.

Prior to this paper, only AFHC has been analyzed in the case of OCO with noisy predictions [13]. Further, the analysis of AFHC in [13] depends delicately on the structure of

the algorithm and thus cannot be generalized to other policies, such as RHC. Our results here are made possible by a *novel* analytic technique that transforms the control strategy employed by OPT, one commitment length at a time, to the control strategy employed by $FHC^{(k)}(v, w)$. At each intermediate step, we exploit the optimality of $FHC^{(k)}(v, w)$ within the commitment length to bound the difference in costs; the sum of these costs over the entire transformation gives a bound on the difference in costs between OPT and $FHC^{(k)}(v, w)$. We then exploit Jensen's inequality to extend this bound on competitive difference to CHC.

Theorem 1 below presents our main result characterizing the performance of CHC algorithms under noisy predictions for functions that are α -Hölder continuous in the second argument; in particular, $\alpha = 1$ corresponds to the class of function that is Lipschitz continuous in the second argument.

THEOREM 1. *Assuming that the prediction error follows (2), then for h that is α -Hölder continuous in the second argument, we have*

$$\mathbb{E}\text{cost}(\text{CHC}) \leq \mathbb{E}\text{cost}(\text{OPT}) + \frac{2T\beta D}{v} + \frac{2GT}{v} \sum_{k=0}^{v-1} \|f_k\|^\alpha. \quad (14)$$

Note that, while Theorem 1 is stated in terms of the competitive difference, it can easily be converted into results about the competitive ratio and regret as explained in Section 2.

There are two terms in the bound on the competitive difference of CHC: (i) The first term $\frac{2T\beta D}{v}$ can be interpreted as the price of switching costs due to limited commitment; this term decreases as the commitment level v increases. (ii) The second term $\frac{2GT}{v} \sum_{k=0}^{v-1} \|f_k\|^\alpha$ represents the impact of prediction noise on the competitive difference and can be characterized by $\|f_k\|$ (defined in (3)), which is impacted by both the variance of $e(s)$ and the structural form of the prediction noise correlation, $f(s)$.

Theorem 1 allows us to immediately analyze the performance of RHC and AFHC as they are special cases of CHC. We present our results comparing the performance of RHC and AFHC by analyzing how the optimal level of commitment, v , depends on properties of the prediction noise.

In order to make concrete comparisons, it is useful to consider specific forms of prediction noise. Here, we consider four cases: (i) i.i.d. prediction noise, (ii) prediction noise with long range correlation, (iii) prediction noise with short range correlation, and (iv) prediction noise with exponentially decaying correlation. All four cases can be directly translated to assumptions on the correlation structure, $f(\cdot)$. Recall that many common predictors, e.g., Wiener and Kalman filters, yield f that is exponentially decaying.

i.i.d. prediction noise. The assumption of i.i.d. prediction noise is idealistic since it only happens when the forecast for y_t is optimal based on the information prior to time t for all $t = 1, \dots, T$ [23]. However, analysis of the i.i.d. noise is instructive and provides a baseline for comparison with more realistic models. In this case, Theorem 1 can be specialized as follows. Recall that $\mathbb{E}[e(s)e(s)^T] = R_e$, and $\text{tr}(R_e) = \sigma^2$.

COROLLARY 2. *Consider i.i.d. prediction error, i.e.,*

$$f(s) = \begin{cases} I, & s = 0 \\ 0, & \text{otherwise.} \end{cases}$$

If h satisfies is α -Hölder continuous in the second argument, then the expected competitive difference of CHC is upper bounded by

$$\mathbb{E}\text{cost}(\text{CHC}) \leq \mathbb{E}\text{cost}(\text{OPT}) + \frac{2T\beta D}{v} + 2GT\sigma^\alpha,$$

which is minimized when $v^ = w$.*

This can be proved by simply applying the form of $f(s)$ to (14). Corollary 2 highlights that, in the i.i.d. case, the level of commitment that minimizes the competitive difference always coincides with the lookahead window w , independent of all other parameters. This is intuitive since, when prediction noise is i.i.d., increasing commitment level does not increase the cost due to prediction errors. Combined with the fact that increasing the commitment level decreases the costs incurred by switching, we can conclude that AFHC is optimal in the i.i.d. setting.

Long range correlation. In contrast to i.i.d. prediction noise, another extreme case is when prediction noise has strong correlation over a long period of time. This is pessimistic and happens when past prediction noise has far-reaching effects on the prediction errors in the future, i.e., the current prediction error is sensitive to errors in the distant past. In this case, prediction only offers limited value since prediction errors accumulate. For long range correlation, we can apply Theorem 1 as follows.

COROLLARY 3. *Consider prediction errors with long range correlation such that*

$$\|f(s)\|_F = \begin{cases} c, & s \leq L \\ 0, & s > L, \end{cases}$$

where $L > w$. If h is α -Hölder continuous in the second argument, the expected competitive difference of CHC is upper bounded by

$$\begin{aligned} \mathbb{E}\text{cost}(\text{CHC}) - \mathbb{E}\text{cost}(\text{OPT}) &\leq \frac{2T\beta D(\alpha + 2) - 4GTc^\alpha\sigma^\alpha}{\alpha + 2} v^{-1} \\ &\quad + \frac{2^{3+\frac{\alpha}{2}} GTc^\alpha\sigma^\alpha}{\alpha + 2} v^{\alpha/2}. \end{aligned}$$

If $\frac{\beta D}{Gc^\alpha\sigma^\alpha} > \alpha(2w)^{1+\frac{\alpha}{2}} + 2$, then $v^ = w$; if $\frac{\beta D}{Gc^\alpha\sigma^\alpha} < \frac{2}{\alpha+2}$, then $v^* = 1$, otherwise v^* is in between 1 and w .*

Corollary 3 highlights that, in the case of long range correlation, the level of commitment that minimizes the competitive difference depends on the variance σ^2 , the switching cost β , the smoothness G , α , and diameter of the action space D .

The term $\frac{\beta D}{Gc^\alpha\sigma^\alpha}$ can be interpreted as a measure of the relative importance of the switching cost and the prediction loss. If $\frac{\beta D}{Gc^\alpha\sigma^\alpha} = \frac{2}{\alpha+2} \in O(1)$, i.e., the one step loss due to prediction error is on the order of the switching cost, then $v^* = 1$ and RHC optimizes the performance bound; if $\frac{\beta D}{Gc^\alpha\sigma^\alpha} = \alpha(2w)^{1+\frac{\alpha}{2}} + 2 \in \Omega(w)$, then $v^* = w$ and AFHC optimizes the performance bound. Otherwise, $v^* \in (1, w)$.

We illustrate these results in Figure 4(a) which plots the competitive difference as a function of the commitment level for various parameter values. The case for the dashed line satisfies $\frac{\beta D}{Gc^\alpha\sigma^\alpha} > \alpha(2w)^{1+\frac{\alpha}{2}} + 2$ and shows competitive difference decreases with increasing levels of commitment. Here, the window size is 100, and thus AFHC minimizes the competitive difference, validating Corollary 3. The dot-dashed line satisfies $\frac{\beta D}{Gc^\alpha\sigma^\alpha} < \frac{2}{\alpha+2}$ and shows the increase in competitive difference with commitment, highlighting that RHC is optimal. The solid line does not satisfy either of these conditions and depicts the minimization of competitive difference at intermediate levels of commitment (marked with a circle). Figure 4(b) illustrates the relationship between α and the optimal commitment level v^* (marked with a circle that corresponds to the same v^* as in Figure 4(a)). As α increases, the prediction loss increases, and thus the optimal commitment level decreases to allow for updated predictions.

Short range correlation. Long range correlation is clearly pessimistic as it assumes that the prediction noise is always

Figure 4: Illustration of Corollary 3, for long range dependencies. (a) shows the time averaged expected competitive difference as a function of the commitment level, and (b) shows the optimal commitment level as a function of α .

correlated within the lookahead window. Here, we study another case where prediction noise can be correlated, but only within a small interval that is less than the lookahead window w . This is representative of scenarios where only limited past prediction noises affect the current prediction. For such short range correlation, Theorem 1 gives us:

COROLLARY 4. Consider prediction errors with short range correlation such that

$$\|f(s)\|_F = \begin{cases} c, & s \leq L \\ 0, & s > L, \end{cases}$$

where $L \leq w$. If h is α -Hölder continuous in the second argument, the expected competitive difference of CHC is upper bounded by:

if $v > L$

$$\mathbb{E}\text{cost}(\text{CHC}) - \mathbb{E}\text{cost}(\text{OPT}) \leq \frac{2T\beta D}{v} + 2GT(c\sigma)^\alpha (L+1)^{\alpha/2} - \frac{2GT(c\sigma)^\alpha}{v} \frac{\alpha}{\alpha+2} ((L+1)^{\alpha/2}(\alpha L - 2) + 1);$$

if $v \leq L$

$$\mathbb{E}\text{cost}(\text{CHC}) - \mathbb{E}\text{cost}(\text{OPT}) \leq \frac{2T\beta D}{v} + \frac{4GTc^\alpha \sigma^\alpha}{v(\alpha+2)} ((v+1)^{(\alpha+2)/2} - 1).$$

If $\frac{\beta D}{Gc^\alpha \sigma^\alpha} > H(L)$, where $H(L) = \frac{1}{\alpha+2}((L+1)^{\alpha/2}(\alpha L - 2) + 1)$, then $v^* = w$; if $\frac{\beta D}{Gc^\alpha \sigma^\alpha} < \min(H(L), \frac{2}{\alpha+2})$, then $v^* = 1$, otherwise v^* is in between 1 and w .

Corollary 4 shows that the structure of the bound on the competitive difference itself depends on the relative values of v and L . In terms of the optimal commitment level, Corollary 4 shows that, similar to Corollary 3, the term $\frac{\beta D}{Gc^\alpha \sigma^\alpha}$ comes into play; however, unlike Corollary 3 (where $L > w$), the optimal commitment level now also depends on the length of the interval, L , within which prediction errors are correlated. Note that $H(L)$ is increasing in L . If $\frac{\beta D}{Gc^\alpha \sigma^\alpha} > H(L)$, i.e., the prediction loss and L are small compared to the switching cost, then $v^* = w$ and thus AFHC optimizes the performance bound. On the other hand, if the prediction loss and L are large compared to the switching cost, then $v^* = 1$, and thus RHC optimizes the bound; otherwise, v^* lies in between 1 and w , and thus intermediate levels of commitment under CHC perform better than AFHC and RHC.

Figure 5: Illustration of Corollary 4, for short range correlations. (a) shows the time averaged expected competitive difference as a function of the commitment level, and (b) shows the optimal commitment level as a function of α .

Note that when prediction noise is i.i.d., we have $L = 0$ and $H(L) < 0$; hence we have $\frac{\beta D}{Gc^\alpha \sigma^\alpha} > H(L)$ and thus $v^* = w$, which corresponds to the conclusion of Corollary 2.

We illustrate these results in Figure 5(a), which plots the competitive difference as a function of the commitment for various parameter values. The dashed line satisfies $\frac{\beta D}{Gc^\alpha \sigma^\alpha} > H(L)$ and shows the drop in competitive difference with increasing levels of commitment. The competitive difference is lowest when the commitment level is 100, which is also the window size, thus validating the optimality of AFHC as per Corollary 4. The dot-dashed line satisfies $\frac{\beta D}{Gc^\alpha \sigma^\alpha} < \min(H(L), \frac{2}{\alpha+2})$ and shows the increase in competitive difference with commitment, highlighting that RHC is optimal. The solid line does not satisfy either of these conditions and depicts the minimization of competitive difference at intermediate levels of commitment. Figure 5(b) illustrates the relationship between α and the optimal commitment level v^* . As α increases, loss due to prediction noise increases; as a result, v^* decreases.

Exponentially decaying correlation. Exponentially decaying correlation is perhaps the most commonly observed model in practice and is representative of predictions made via Wiener [49] or Kalman [28] filters. For clarity of illustration we consider the case of $\alpha = 1$ here. In this case, Theorem 1 results in the following corollary.

COROLLARY 5. Consider prediction errors with exponentially decaying correlation, i.e., there exists a $a < 1$, such that

$$\|f(s)\|_F = \begin{cases} ca^s, & s \geq 0 \\ 0, & s < 0. \end{cases}$$

If h is 1-Hölder continuous, then the expected competitive difference of CHC is upper bounded by

$$\mathbb{E}\text{cost}(\text{CHC}) - \mathbb{E}\text{cost}(\text{OPT}) \leq \frac{2T\beta D}{v} + \frac{2GTc\sigma}{1-a^2} - \frac{a^2(1-a^{2v})GTc\sigma}{v(1-a^2)^2}.$$

When $\frac{\beta D}{Gc\sigma} \geq \frac{a^2}{2(1-a^2)}$ the commitment that minimizes the performance bound is $v^* = w$, i.e., AFHC minimizes the performance bound. When $\frac{\beta D}{Gc\sigma} < \frac{a^2}{2(1+a)}$, $v^* = 1$, i.e., RHC minimizes the performance bound.

Corollary 5 shows that when the prediction noise σ and the correlation decay a are small, the loss due to switching costs is dominant, and thus commitment is valuable; on the other hand, when σ and a are large, then the loss due to inaccurate

Figure 6: Illustration of Corollary 5, for exponentially decaying correlations. (a) shows the time averaged expected competitive difference as a function of the commitment level, and (b) shows the optimal commitment level as a function of the decay parameter, a .

predictions is dominant, and thus a smaller commitment is preferable to exploit more updated predictions.

We illustrate these results in Figure 6(a), which plots the competitive difference as a function of the commitment for various parameter values. The dashed line satisfies $\frac{\beta D}{Gc\sigma} > \frac{a^2}{2(1-a^2)}$ and shows the drop in competitive difference with increasing levels of commitment. The competitive difference is lowest when the commitment level is 100, which is also the window size, thus validating the optimality of AFHC as per Corollary 5. The dot-dashed line satisfies $\frac{\beta D}{Gc\sigma} > \frac{a^2}{2(1+a)}$ and shows the increase in competitive difference with commitment, highlighting that RHC is optimal. The solid line does not satisfy either of these conditions and depicts the minimization of competitive difference at intermediate levels of commitment. Figure 6(b) illustrates the relationship between a and the optimal commitment level v^* . As a increases, correlation decays more slowly, thus the loss due to prediction noise becomes dominant; as a result, v^* decreases.

Strong convexity. All of our results to this point depend on the diameter of the action space D . While this is common in OCO problems, e.g., [54, 27], it is not desirable.

Our last result in this section highlights that it is possible to eliminate the dependence on D by making a stronger structural assumption on h – strong convexity. In particular, we say that $h(\cdot)$ is m –strongly convex in the first argument w.r.t. the norm of the switching cost $\|\cdot\|$ if $\forall x_1, x_2, y$,

$$h(x_1, y) - h(x_2, y) \geq \langle \partial_x h(x_2, y), (x_1 - x_2) \rangle + \frac{m}{2} \|x_1 - x_2\|^2.$$

Strong convexity is used in the online learning literature to obtain performance bounds that are independent of the diameter of action space, see, e.g., [22, 46]. Under the assumption of strong convexity, we obtain the following.

THEOREM 6. *If h is m -strongly convex in the first argument with respect to $\|\cdot\|$ and α -Hölder continuous in the second argument, we have*

$$\mathbb{E}\text{cost}(CHC) - \mathbb{E}\text{cost}(OPT) \leq \frac{2\beta^2 T}{mv} + 2GT \sum_{k=0}^{v-1} \|f_k\|^\alpha.$$

Theorem 6 is useful when the diameter of the feasible set D is large or unbounded; when D is small, we can apply Theorem 1 instead. As above, it is straightforward to apply the techniques in Corollaries 2 – 5 to compute v^* for strongly convex h under different types of prediction noise².

²We only need to change βD with β^2/m in the bounds of the corollaries to draw parallel conclusions.

5. CONCENTRATION BOUNDS

Our results to this point have focused on the performance of CHC algorithms *in expectation*. In this section, we establish bounds on the *distribution* of costs under CHC algorithms. In particular, we prove that, under a mild additional assumption, the likelihood of cost exceeding the average case bounds proven in Section 4 decays exponentially.

For simplicity of presentation, we state and prove the concentration result for CHC when the online parameter y is one-dimensional. In this case, $R_e = \sigma^2$, and the correlation function $f: \mathbb{N} \rightarrow \mathbb{R}$ is a scalar valued function. The results can be generalized to the multi-dimensional setting at the expense of considerable notational complexity in the proofs.

Additionally, for simplicity of presentation we assume (for this section only) that $\{e(t)\}_{t=1}^T$ are uniformly bounded, i.e., $\exists \epsilon > 0$, s.t. $\forall t, |e(t)| < \epsilon$. Note that, with additional effort, the boundedness assumption can be relaxed to $e(t)$ being subgaussian, i.e., $\mathbb{E}[\exp(e(t)^2/\epsilon^2)] \leq 2$, for some $\epsilon > 0$.³

Given $\{\hat{y}_t\}_{t=1}^T$, the competitive difference of CHC is a random variable that is a function of the prediction error $e(t)$. To state our concentration results formally, let $V_1 T$ be the upper bound of the expected competitive difference of CHC in (14), i.e., $V_1 T = \frac{2T\beta D}{v} + \frac{2GT}{v} \sum_{k=1}^v \|f_k\|^\alpha$.

THEOREM 7. *Assuming that the prediction error follows (2), and h is α -Hölder continuous in the second argument, we have*

$$\begin{aligned} &\mathbb{P}(\text{cost}(CHC) - \text{cost}(OPT) > V_1 T + u) \\ &\leq \exp\left(\frac{-u^2 \alpha^2}{2^{1+2\alpha} G^2 \epsilon^2 T F(v)}\right), \end{aligned}$$

for any $u > 0$, where $F(v) = \left(\frac{1}{v} \sum_{k=0}^{v-1} (v-k)^\alpha |f(k)|^\alpha\right)^2$.

This result shows that the competitive difference has a sub-Gaussian tail, which decays much faster than the normal large deviation bounds obtained by bounding moments, i.e., Markov Inequality, the rate of decay is dependent on the sensitivity of h to disturbance in the second argument (G, α), the size of variation (ϵ), and the correlation structure ($F(v)$). This is illustrated in Figure 7, where we show the distribution of the competitive difference of CHC under different prediction noise correlation assumptions. We can see that, for prediction noise that decays fast (i.i.d. and exponentially decaying noise with small a) in Figure 7(a), the distribution is tightly concentrated around the mean, whereas for prediction noise that are fully correlated (short range correlation and long range correlation) in Figure 7(b), the distribution is more spread out.

If we consider the time-averaged competitive difference, or the *regret* against the offline optimal, we can equivalently state Theorem 7 as follows.

COROLLARY 8. *Assuming that the prediction error follows (2), and h is α -Hölder continuous, then probability that the competitive difference of CHC exceeds V_1 can be bounded by*

$$\begin{aligned} &\mathbb{P}\left(\frac{1}{T} [\text{cost}(CHC) - \text{cost}(OPT)] > V_1 + u\right) \\ &\leq \exp\left(\frac{-u^2}{2^{1+2\alpha} G^2 \epsilon^{2\alpha} F(v)/T}\right), \end{aligned}$$

where $F(v) = \left(\frac{1}{v} \sum_{k=0}^{v-1} (v-k) |f(k)|^\alpha\right)^2$. Assuming $f(s) \leq C$ for $s = 0, \dots, v$, then $\lim_{T \rightarrow \infty} F(v)/T = 0$ if either $v \in O(1)$, or $f(s) \leq c\eta^s$ for some $\eta < 1$.

³This involves more computation and worse constants in the concentration bounds. Interested readers are referred to Theorem 12 and the following remark of [11] for a way to generalize the concentration bound.

Figure 7: The cumulative distribution function of average-case bounds under different correlation structures: (a) i.i.d prediction noise; exponentially decaying, $a = 2/3$; (b) long range; short range, $L = 4$. Competitive differences simulated with random re-activation of standard normal $e(t)$ 1000 times under the following parameter values: $T = 100, v = 10, \beta D = 1, G = 0.1, \alpha = 1, c = 1$.

Corollary 8 shows that, when either the commitment level v is constant, or the correlation $f(s)$ is exponentially decaying, the parameter of concentration $F(v)/T$ for the regret of CHC tends to 0. The full proof is given in Appendix B.6. To prove this result on the concentration of the competitive difference, we make heavy use of the fact that h is α -Hölder continuous in the second argument, which implies that the competitive difference is α -Hölder continuous in e . This allows application of the method of bounded difference, i.e., we bound the difference of $V(e)$ where one component of e is replaced by an identically-distributed copy. More specifically, we use the following lemma, the one-sided version of one due to McDiarmid:

LEMMA 9 ([39], LEMMA 1.2). *Let $X = (X_1, \dots, X_n)$ be independent random variables and Y be the random variable $f(X_1, \dots, X_n)$, where function f satisfies $|f(x) - f(x'_k)| \leq c_k$ whenever x and x'_k differ in the k th coordinate. Then for any $t > 0$, $\mathbb{P}(Y - \mathbb{E}Y > t) \leq \exp\left(\frac{-2t^2}{\sum_{k=1}^n c_k^2}\right)$.*

6. CONCLUDING REMARKS

Online convex optimization (OCO) problems with switching costs and noisy predictions are widely applicable in networking and distributed systems. Prior efforts in this area have resulted in two promising algorithms – Receding Horizon Control (RHC) and Averaging Fixed Horizon Control (AFHC). Unfortunately, it is not obvious when each algorithm should be used. Further, thus far, only AFHC has been analyzed in the presence of noisy predictions, despite the fact that RHC is seemingly more resistant to prediction noise in many settings.

In this paper, we provide the first analysis of RHC with noisy predictions. This novel analysis is made possible by the introduction of our new class of online algorithms, Committed Horizon Control (CHC), that allows for arbitrary levels of commitment, thus generalizing RHC and AFHC. Our analysis of CHC provides explicit results characterizing the optimal commitment level as a function of the variance and correlation structure of the prediction noise. In doing so, we characterize when RHC/AFHC is better depending on the properties of the prediction noise, thus addressing an important open challenge in OCO.

Our focus in this paper has been on the theoretical analysis of CHC and its implications for RHC and AFHC. The superiority of CHC suggests that it is a promising approach for integrating predictions into the design of systems, especially those that operate in uncertain environments. Going

forward, it will be important to evaluate the performance of CHC algorithms in settings where RHC and AFHC have been employed, such as dynamic capacity provisioning, geographical load balancing, and video streaming.

7. REFERENCES

- [1] M. A. Adnan, R. Sugihara, and R. K. Gupta. Energy efficient geographical load balancing via dynamic deferral of workload. In *IEEE Int. Conf. Cloud Computing (CLOUD)*, pages 188–195, 2012.
- [2] Hölder condition. https://en.wikipedia.org/wiki/Holder_condition.
- [3] G. Ananthanarayanan, A. Ghodsi, S. Shenker, and I. Stoica. Effective straggler mitigation: Attack of the clones. In *Proc. NSDI*, volume 13, pages 185–198, 2013.
- [4] L. Andrew, S. Barman, K. Ligett, M. Lin, A. Meyerson, A. Roytman, and A. Wierman. A tale of two metrics: Simultaneous bounds on competitiveness and regret. In *Conf. on Learning Theory (COLT)*, pages 741–763, 2013.
- [5] M. Asawa and D. Tenekeztzis. Multi-armed bandits with switching penalties. *IEEE Transactions on Automatic Control*, 41(3):328–348, 1996.
- [6] N. Bansal, A. Gupta, R. Krishnaswamy, K. Pruhs, K. Schewior, and C. Stein. A 2-competitive algorithm for online convex optimization with switching costs. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 40. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- [7] A. Bartolini, M. Cacciari, A. Tilli, and L. Benini. A distributed and self-calibrating model-predictive controller for energy and thermal management of high-performance multicores. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2011*, pages 1–6. IEEE, 2011.
- [8] A. Bemporad and M. Morari. Robust model predictive control: A survey. In *Robustness in identification and control*, pages 207–226. Springer, 1999.
- [9] A. Blum, H. Karloff, Y. Rabani, and M. Saks. A decomposition theorem and bounds for randomized server problems. In *Proc. Symp. Found. Comp. Sci. (FOCS)*, pages 197–207, Oct 1992.
- [10] A. Borodin, N. Linial, and M. E. Saks. An optimal on-line algorithm for metrical task system. *J. ACM*, 39(4):745–763, 1992.
- [11] S. Boucheron, G. Lugosi, and O. Bousquet. Concentration inequalities. In *Advanced Lectures on Machine Learning*, pages 208–240. Springer, 2004.
- [12] E. F. Camacho and C. B. Alba. *Model predictive control*. Springer, 2013.
- [13] N. Chen, A. Agarwal, A. Wierman, S. Barman, and L. L. H. Andrew. Online convex optimization using predictions. In *Proc. ACM SIGMETRICS*, pages 191–204. ACM, 2015.
- [14] N. Chen, L. Gan, S. H. Low, and A. Wierman. Distributional analysis for model predictive deferrable load control. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 6433–6438. IEEE, 2014.
- [15] L. Gan, A. Wierman, U. Topcu, N. Chen, and S. H. Low. Real-time deferrable load control: Handling the uncertainties of renewable generation. *SIGMETRICS Perform. Eval. Rev.*, 41(3):77–79, Jan. 2014.
- [16] A. Gandhi, Y. Chen, D. Gmach, M. Arlitt, and M. Marwah. Minimizing Data Center SLA Violations and Power Consumption via Hybrid Resource Provisioning. In *Proceedings of the 2011 International Green Computing Conference*, pages 49–56, Orlando, FL, USA, 2011.
- [17] A. Gandhi, M. Harchol-Balter, and I. Adan. Server farms with setup costs. *Performance Evaluation*, 67:1123–1138, 2010.
- [18] C. E. Garcia, D. M. Prett, and M. Morari. Model predictive control: theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.
- [19] D. Gmach, J. Rolia, L. Cherkasova, and A. Kemper. Workload analysis and demand prediction of enterprise data center applications. In *Proc. IEEE Int. Symp. Workload Characterization, IISWC '07*, pages 171–180. IEEE Computer Society, 2007.

- [20] S. Guha and K. Munagala. Multi-armed Bandits with Metric Switching Costs. In *Automata, Languages and Programming*, volume 5556 of *Lecture Notes in Computer Science*, pages 496–507. Springer Berlin Heidelberg, 2009.
- [21] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [22] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [23] Innovation process. [https://en.wikipedia.org/wiki/Innovation_\(signal_processing\)](https://en.wikipedia.org/wiki/Innovation_(signal_processing)).
- [24] V. Joseph and G. de Veciana. Variability aware network utility maximization. *arXiv preprint arXiv:1111.3728*, 2011.
- [25] V. Joseph and G. de Veciana. Jointly optimizing multi-user rate adaptation for video transport over wireless systems: Mean-fairness-variability tradeoffs. In *Proc. IEEE INFOCOM*, pages 567–575, 2012.
- [26] T. Kailath, A. H. Sayed, and B. Hassibi. *Linear estimation*. Prentice-Hall, Inc., 2000.
- [27] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [28] R. E. Kalman. A new approach to linear filtering and prediction problems. *J. Fluids Engineering*, 82(1):35–45, 1960.
- [29] S.-J. Kim and G. B. Giannakis. Real-time electricity pricing for demand response using online convex optimization. In *IEEE Innovative Smart Grid Tech. Conf.*, pages 1–5, 2014.
- [30] A. Krioukov, P. Mohan, S. Alspaugh, L. Keys, D. Culler, and R. Katz. NapSAC: Design and implementation of a power-proportional web cluster. In *Proceedings of the 1st ACM SIGCOMM Workshop on Green Networking*, Green Networking '10, pages 15–22, New Delhi, India, 2010.
- [31] D. Kusic, J. O. Kephart, J. E. Hanson, N. Kandasamy, and G. Jiang. Power and performance management of virtualized computing environments via lookahead control. *Cluster computing*, 12(1):1–15, 2009.
- [32] W. Kwon and A. Pearson. A modified quadratic cost problem and feedback stabilization of a linear system. *IEEE Transactions on Automatic Control*, 22(5):838–842, 1977.
- [33] M. Lin, Z. Liu, A. Wierman, and L. L. Andrew. Online algorithms for geographical load balancing. In *Int. Green Computing Conference (IGCC)*, pages 1–10. IEEE, 2012.
- [34] M. Lin, A. Wierman, L. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Trans. Networking*, 21(5):1378–1391, Oct 2013.
- [35] Z. Liu, I. Liu, S. Low, and A. Wierman. Pricing data center demand response. In *Proc. ACM Sigmetrics*, 2014.
- [36] Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen. Data center demand response: Avoiding the coincident peak via workload shifting and local generation. *Performance Evaluation*, 70(10):770–791, 2013.
- [37] L. Lu, J. Tu, C.-K. Chau, M. Chen, and X. Lin. Online energy generation scheduling for microgrids with intermittent energy sources and co-generation. In *Proc. ACM SIGMETRICS*, pages 53–66. ACM, 2013.
- [38] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- [39] C. McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.
- [40] B. Narayanaswamy, V. K. Garg, and T. Jayram. Online optimization for the smart (micro) grid. In *Proc. ACM Int. Conf. on Future Energy Systems*, page 19, 2012.
- [41] D. Niu, H. Xu, B. Li, and S. Zhao. Quality-assured cloud bandwidth auto-scaling for video-on-demand applications. In *Proc. IEEE INFOCOM*, pages 460–468, March 2012.
- [42] C. Rhemann, A. Hosni, M. Bleyer, C. Rother, and M. Gelautz. Fast Cost-volume Filtering for Visual Correspondence and Beyond. In *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, CVPR '11, pages 3017–3024, 2011.
- [43] A. P. Sage and J. L. Melsa. Estimation theory with applications to communications and control. Technical report, DTIC Document, 1971.
- [44] S. Sastry and M. Bodson. *Adaptive control: stability, convergence and robustness*. Courier Dover Publications, 2011.
- [45] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel. Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 33(9):1251–1270, 2014.
- [46] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [47] H. Wang, J. Huang, X. Lin, and H. Mohsenian-Rad. Exploring smart grid and data center interactions for electric power load balancing. *ACM SIGMETRICS Performance Evaluation Review*, 41(3):89–94, 2014.
- [48] X. Wang and M. Chen. Cluster-level feedback power control for performance optimization. In *High Performance Computer Architecture*, pages 101–110. IEEE, 2008.
- [49] N. Wiener. *Extrapolation, interpolation, and smoothing of stationary time series*, volume 2. MIT press, Cambridge, MA, 1949.
- [50] L. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *J. Machine Learning Research*, 9999:2543–2596, 2010.
- [51] F. Zanini, D. Atienza, L. Benini, and G. De Micheli. Multicore thermal management with model predictive control. In *Proc. IEEE. European Conf. Circuit Theory and Design (ECCTD)*, pages 711–714, 2009.
- [52] F. Zanini, D. Atienza, G. De Micheli, and S. P. Boyd. Online convex optimization-based algorithm for thermal management of MPSoCs. In *Proc. ACM Great Lakes Symp. VLSI*, pages 203–208, 2010.
- [53] X. Y. Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics & Optimization*, 42(1):19–33, 2000.
- [54] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proc. Int. Conf. Machine Learning (ICML)*, pages 928–936. AAAI Press, 2003.

APPENDIX

A. EXPERIMENTAL SETUP FOR FIG. 1

Setting for Figure 1(a): This example corresponds to a simple model of a data center. There are $(w + 1)$ types of jobs and $(w + 2)$ types of servers available to process these jobs. Each server has a different linear cost $\{a(t), b, c : 0 < a(t) < b < c\}$ (low, medium, high respectively) depending on the job type. The low cost is a monotonically increasing function of time that asymptotically approaches the constant medium cost (i.e. $a(t) = \alpha + (b - \alpha)\frac{t-1}{t}$, where $0 < \alpha < b$). The switching cost β only applies when a server is turned on (shut down costs can be included in the turning on cost) and has a magnitude greater than the difference between the medium and low costs (i.e. $\beta > b - \alpha$). The high cost is constant but greater than the difference between the medium and low costs multiplied by the prediction window size plus the switching cost. (i.e. $c > (b - \alpha)w + \beta$). One special server (server 0) can process all jobs with medium cost. Label all other servers 1 through $(w + 1)$ and all job types 1 through $(w + 1)$. Let server $s \in \{1, \dots, w + 1\}$ be able to process job type s with low cost, job type $s - 1$ with high cost, and all other job types with medium cost.

We assume perfect prediction within the prediction window, w . The trace that forms Figure 1(a), is one in which the whole work load is only with one job type at each timestep starting with job type 1 and sequentially cycles through all job types every $(w + 1)$ timesteps.

This forces RHC to switch every timestep and FHC to switch every w timesteps to avoid a future high cost but take advantage of a low cost at the current timestep.

The offline optimal puts all of the workload on server 0 that processes all jobs with medium cost and so never incurs a switching cost after the first timestep.

RHC and AFHC try to take advantage of the low cost but the trace tricks them with a high cost one timestep beyond the prediction window. Switching to server 0 is always slightly too expensive by $(b - \alpha)^{\frac{1}{t}}$ within the prediction window. The values used in Figure 1(a) are as follows: cycling workload of size 1 for 100 timesteps, $\alpha = 0.9$, $b = 1$, $\beta = 2$, $c = 0.1(w + 1) + 3$.

Setting for Figure 1(b): Similar to Figure 1(a), the setting in which this example was constructed corresponds to a simple model of a data center. The key difference is that predictions are noisy. There are $(w + 1)$ types of jobs and $(w + 1)$ types of servers available to process these jobs. Each server has a different linear cost $\{a, c : 0 < a < c\}$ (low, high respectively) depending on the job type. The switching cost β only applies when a server is turned on (shut down costs can be included in the turning on cost) and has a magnitude less than the difference between the high and low cost (i.e. $\beta < c - a$). Label all servers 1 through $(w + 1)$ and all job types 1 through $(w + 1)$. Let server $s \in \{1, \dots, w + 1\}$ be able to process job type s with low cost, and all other job types with high cost.

We assume perfect prediction within only the first γ timesteps of the prediction window, w . The trace that forms Figure 1(b), is one in which the whole work load is only with one job type at each timestep starting with job type 1 and sequentially cycles through all job types every $(w + 1)$ timesteps. Error in the last $w - \gamma$ timesteps of the prediction window is produced by making those predictions be equal to the prediction of the last perfect prediction (i.e. the γ -th timestep within the prediction window).

RHC equals the offline optimal solution in this setting which is to switch the whole workload at every timestep to the server with the unique low cost. AFHC on the other hand puts $(w - \gamma)/w$ of the workload on servers with high cost and only γ/w of the workload on the server with the unique low cost. The values used in Figure 1(b) are as follows: cycling workload of size 1 for 30 timesteps, $a = 1$, $c = 6$, $\beta = 0.1$.

B. PROOF OF ANALYTIC RESULTS

We first introduce some additional notation used in the proofs. For brevity, for any vector x we write $x_{i..j} = (x_i, \dots, x_j)$ for any $i \leq j$. Let x^* denote the offline optimal solution to (5), and let the cost of an online algorithm during time period $[t_1, t_2]$ with boundary conditions x_S, x_E and with online data $y_{t_1..t_2}$ be

$$g_{t_1, t_2}(x; x_S; x_E; y) = \sum_{t=t_1}^{t_2} h(x_t, y_t) + \beta \|x_S, x_{t_1}\| \\ + \sum_{t=t_1+1}^{t_2} \beta \|x_{t-1}, x_t\| + \beta \|x_{t_2}, x_E\|.$$

If x_E is omitted, then by convention $x_E = x_{t_2}$ (and thus $\beta \|x_{t_2} - x_E\| = 0$). If x_S is omitted, then by convention $x_S = x_{t_1-1}$. Note that $g_{t_1, t_2}(x)$ depends only on x_i for $t_1 - 1 \leq i \leq t_2$.

B.1 Proof of Theorem 1

To characterize the suboptimality of CHC in the stochastic case, we first analyze the competitive difference of fixed horizon control with commitment level v , $FHC^k(v)$. Without loss of generality, assume that $k = 0$. Subsequently we omit k and v in FHC for simplicity. Construct a sequence of T -tuples $(\xi^1, \xi^2, \dots, \xi^{M_1})$, where $M_1 = \#\{t \in [1, T] \mid t \bmod v = 1\} \leq \lceil T/v \rceil$, such that $\xi^1 = x^*$ is the offline optimal solution, and $\xi_t^\tau = x_{FHC, t}^\tau$ for all $t < \tau v + 1$ hence, $\xi^{M_1} = x_{FHC}$. At stage τ , to calculate $\xi^{\tau+1}$, apply

FHC to get $(\tilde{x}_{\tau v+1}, \dots, \tilde{x}_{\tau v+w}) = X^\tau(\xi_{\tau v}^\tau, y_{\cdot|\tau v})$, and replace $\xi_{\tau v+1:(\tau+1)v}^\tau$ with $\tilde{x}_{\tau v+1:(\tau+1)v}$ to get $\xi^{\tau+1}$, i.e.,

$$\xi^{\tau+1} = (\xi_1^\tau, \dots, \xi_{\tau v}^\tau, \tilde{x}_{\tau v+1}, \dots, \tilde{x}_{(\tau+1)v}, \xi_{(\tau+1)v+1}^\tau, \dots, \xi_T^\tau).$$

By examining the terms in ξ^τ and $\xi^{\tau+1}$, we have

$$g_{1, T}(\xi^{\tau+1}; y) - g_{1, T}(\xi^\tau; y) \\ = -\beta \|x_{(\tau+1)v+1}^* - x_{(\tau+1)v}^*\| + \beta \|x_{(\tau+1)v+1}^* - \tilde{x}_{(\tau+1)v}\| \\ - \beta \|x_{\tau v+1}^* - \xi_{\tau v}^\tau\| + \beta \|\tilde{x}_{\tau v+1} - \xi_{\tau v}^\tau\| \\ - \sum_{t=\tau v+1}^{(\tau+1)v} (h(x_t^*, y_t) + \beta \|x_t^* - x_{t-1}^*\|) \\ + \sum_{t=\tau v+1}^{(\tau+1)v} (h(\tilde{x}_t, y_t) + \beta \|\tilde{x}_t - \tilde{x}_{t-1}\|) \quad (15)$$

By construction of $(\tilde{x}_{\tau v+1}, \dots, \tilde{x}_{(\tau+1)v})$, it is the optimal solution for $g_{\tau v+1, (\tau+1)v}(x; \xi_{\tau v}^\tau; \tilde{x}_{(\tau+1)v+1}; y_{\cdot|\tau})$, hence

$$\sum_{t=\tau v+1}^{(\tau+1)v} (h(\tilde{x}_t, y_{t|\tau v}) + \beta \|\tilde{x}_t - \tilde{x}_{t-1}\|) \\ + \beta \|\tilde{x}_{\tau v+1} - \xi_{\tau v}^\tau\| + \beta \|\tilde{x}_{(\tau+1)v+1} - \tilde{x}_{(\tau+1)v}\| \\ \leq \sum_{t=\tau v+1}^{(\tau+1)v} (h(x_t^*, y_{t|\tau v}) + \beta \|x_t^* - x_{t-1}^*\|) \\ + \beta \|x_{\tau v+1}^* - \xi_{\tau v}^\tau\| + \beta \|\tilde{x}_{(\tau+1)v+1} - x_{(\tau+1)v}^*\|,$$

Substituting the above inequality into (15) and by triangle inequality, we have

$$g_{1, T}(\xi^{\tau+1}; y) - g_{1, T}(\xi^\tau; y) \\ \leq 2\beta \|x_{(\tau+1)v+1}^* - \tilde{x}_{(\tau+1)v+1}\| + \sum_{t=\tau v+1}^{(\tau+1)v} |h(x_t^*, y_{t|\tau v}) - h(x_t^*, y_t)| \\ + \sum_{t=\tau v+1}^{(\tau+1)v} |h(\tilde{x}_t, y_t) - h(\tilde{x}_t, y_{t|\tau v})| \\ \leq 2\beta D + 2G \sum_{t=\tau v+1}^{(\tau+1)v} \|y_t - y_{t|\tau v}\|_2^\alpha, \quad (16)$$

Summing these inequalities from $\tau = 0$ to $\tau = M_1$ and noting that $\xi^{M_1} = x_{FHC^1(v)}$ and $\xi^1 = x^*$, we have

$$\text{cost}(FHC^1(v)) \leq \text{cost}(OPT) + 2M_1\beta D + 2G \sum_{\tau=0}^{M_1} \sum_{t=\tau v+1}^{(\tau+1)v} \|y_t - y_{t|\tau v}\|_2^\alpha \\ = \text{cost}(OPT) + 2M_1\beta D + 2G \sum_{t=1}^T \|y_t - y_{t|\phi^1(t)}\|_2^\alpha. \quad (17)$$

where $\phi^k(t) = \arg \min_{u \in \Psi_k, u \leq t} |t - u|$. For $k = 1$, $\phi^1(t) = u$ whenever $u = \tau v$ and $t \in [u, u + v - 1]$ for some τ . We only have M_1 terms of the switching cost $\|x_{(\tau+1)v+1}^* - \tilde{x}_{(\tau+1)v+1}\|$ since $(M_1 + 1)v + 1 > T$. By the same argument, we have

$$\text{cost}(FHC^k(v)) \leq \text{cost}(OPT) + 2M_k\beta D + 2G \sum_{t=1}^T \|y_t - y_{t|\phi^k(t)}\|_2^\alpha$$

Recall that $x_{CHC,t} = \frac{1}{v} \sum_{k=1}^v x_{FHC,t}^{(k,v)}$, by convexity of the cost function and Jensen's inequality, we have

$$\begin{aligned} \text{cost}(CHC) &\leq \frac{1}{v} \sum_{k=0}^{v-1} \text{cost}(FHC^k(v)) \\ &\leq \text{cost}(OPT) + \frac{2 \sum_{k=0}^{v-1} M_k \beta D}{v} + \frac{2G}{v} \sum_{t=1}^T \sum_{k=0}^{v-1} \|y_t - y_{t|t-\phi^k(t)}\|_2^\alpha \\ &\leq \text{cost}(OPT) + \frac{2T\beta D}{v} + \frac{2G}{v} \sum_{t=1}^T \sum_{k=0}^{v-1} \|y_t - y_{t|t-\phi^k(t)}\|_2^\alpha \\ &\leq \text{cost}(OPT) + \frac{2T\beta D}{v} + \frac{2G}{v} \sum_{t=1}^T \sum_{k=0}^{v-1} \|y_t - y_{t|t-(k+1)}\|_2^\alpha. \end{aligned} \quad (18)$$

where the third inequality is because $\sum_{k=0}^{v-1} M_k = T$ since by definition M_k is the number of elements in $[1, T]$ that is congruent to k modulus v ; and the fourth inequality is because for all t , $t - \phi^k(t)$ always range from 1 to v when k goes from 0 to $v - 1$.

Finally, we show that $\mathbb{E} \|y_\tau - y_{\tau| \tau-(k+1)}\|_2^\alpha \leq \|f_k\|^\alpha$ to finish the proof. Note that for $\alpha = 2$, by (2), we have

$$\begin{aligned} \mathbb{E} \|y_\tau - y_{\tau| \tau-(k+1)}\|_2^2 &= \mathbb{E} \left\| \sum_{s=\tau-k}^{\tau} f(\tau-s)e(s) \right\|_2^2 \\ &= \mathbb{E} \text{tr} \left(\sum_{s_1, s_2=0}^k e(\tau-s_1)^T f(s_1)^T f(s_2) e(\tau-s_2) \right) \\ &= \text{tr} \left(\sum_{s_1, s_2=0}^k f(s_1)^T f(s_2) \mathbb{E} e(\tau-s_2) e(\tau-s_1)^T \right) \\ &= \text{tr} \left(R_e \sum_{s=0}^k f(s)^T f(s) \right) = \|f_k\|^2, \end{aligned} \quad (19)$$

where the second equality is due to cyclic invariance of trace and linearity of expectation, and third equality is due to the fact that $e(s)$ are uncorrelated. When $\alpha \leq 2$, $F(x) = x^{\alpha/2}$ is a concave function, hence by Jensen's inequality,

$$\begin{aligned} \mathbb{E} \|y_\tau - y_{\tau| \tau-(k+1)}\|_2^\alpha &= \mathbb{E} F(\|y_\tau - y_{\tau| \tau-(k+1)}\|_2^2) \\ &\leq F(\mathbb{E} \|y_\tau - y_{\tau| \tau-(k+1)}\|_2^2) = \|f_k\|^\alpha. \end{aligned}$$

B.2 Proof of Corollary 3

Taking expectation over the prediction error and assuming long range correlation, we have for all $k \leq v \leq w$

$$\begin{aligned} \|f_k\|^2 &= \sum_{s=0}^k \text{tr}(R_e f(s)^T f(s)) = \sum_{s=0}^k \langle R_e^{1/2}, f(s) \rangle^2 \\ &\leq \sum_{s=0}^k (\|R_e^{1/2}\|_F \|f(s)\|_F)^2 = (k+1)c^2\sigma^2, \end{aligned}$$

where the inequality is due to Cauchy-Schwarz and $\|f_k\| = \sqrt{k+1}c\sigma$. To compute competitive difference of CHC , note that

$$\begin{aligned} \sum_{k=0}^{v-1} \|f_k\|^\alpha &= \sum_{k=0}^{v-1} (\sqrt{k+1}c\sigma)^\alpha \leq c^\alpha \sigma^\alpha \int_1^{v+1} k^{\alpha/2} dk \\ &= \frac{2c^\alpha \sigma^\alpha}{\alpha+2} ((v+1)^{\frac{\alpha+2}{2}} - 1). \end{aligned} \quad (20)$$

Thus, by Theorem 1,

$$\begin{aligned} \mathbb{E} \text{cost}(CHC) - \mathbb{E} \text{cost}(OPT) &\leq \frac{2T\beta D}{v} + \frac{2GT}{v} \sum_{k=0}^{v-1} (\sqrt{k+1}c\sigma)^\alpha \\ &\leq \frac{2T\beta D}{v} + \frac{4GTc^\alpha \sigma^\alpha}{v(\alpha+2)} ((v+1)^{1+\alpha/2} - 1) \\ &= \frac{2T\beta D(\alpha+2) - 4GTc^\alpha \sigma^\alpha}{\alpha+2} v^{-1} + \frac{4GTc^\alpha \sigma^\alpha}{\alpha+2} v^{\alpha/2} \left(\frac{v+1}{v} \right)^{1+\alpha/2} \\ &\leq \frac{2T\beta D(\alpha+2) - 4GTc^\alpha \sigma^\alpha}{\alpha+2} v^{-1} + \frac{2^{3+\frac{\alpha}{2}} GTc^\alpha \sigma^\alpha}{\alpha+2} v^{\alpha/2}, \end{aligned}$$

where the last inequality is because $(v+1)/v \leq 2$ for $v \geq 1$. If $\beta D(\alpha+2) \leq 2Gc^\alpha \sigma^\alpha$, which implies $\frac{\beta D}{Gc^\alpha \sigma^\alpha} \leq \frac{2}{\alpha+2}$, then the right hand side is an increasing function of v , hence the commitment level that minimizes the performance guarantee is $v^* = 1$.

On the other hand, if $\beta D(\alpha+2) > 2Gc^\alpha \sigma^\alpha$, then let $A = \frac{\beta D(\alpha+2) - 2Gc^\alpha \sigma^\alpha}{\alpha+2}$, $B = \frac{2^{2+\alpha/2} Gc^\alpha \sigma^\alpha}{\alpha+2}$, then the right hand side is $F(v) = 2T(Av^{-1} + Bv^{\alpha/2})$, by examining the gradient $F'(v) = 2T(-Av^{-2} + \frac{\alpha}{2} Bv^{-(1-\alpha/2)})$, since $F'(v) > 0$ iff $v^2 F'(v) > 0$ and $v^2 F'(v)$ is an increasing function in v , we can see that when $v < (\frac{2A}{\alpha B})^{2/(\alpha+2)}$, $F'(v) < 0$ and $F(v)$ is a decreasing function, when $v > (\frac{2A}{\alpha B})^{2/(\alpha+2)}$, $F'(v) > 0$ and $F(v)$ is an increasing function, hence when $v = (\frac{2A}{\alpha B})^{2/(\alpha+2)}$, $F'(v) = 0$ is the global minimum point of $F(v)$.

Therefore, when $(\frac{2A}{\alpha B})^{2/(\alpha+2)} \geq w$, we have $v^* = w$, this happens when $\frac{\beta D}{Gc^\alpha \sigma^\alpha} > \alpha(2w)^{1+\frac{\alpha}{2}} + 2$. When $\frac{\beta D}{Gc^\alpha \sigma^\alpha} \in (\frac{2}{\alpha+2}, \alpha(2w)^{1+\frac{\alpha}{2}} + 2)$, v^* is between 1 to w , in this case, $v^* = \left(\frac{\beta D(\alpha+2) - 2Gc^\alpha \sigma^\alpha}{2^{2+\alpha/2} Gc^\alpha \sigma^\alpha} \right)^{2/(\alpha+2)}$.

B.3 Proof of Corollary 4

Taking expectation over the prediction error, when $k < L$, similar to the proof of Corollary 3, $\|f_k\|^2 \leq (k+1)c^2\sigma^2$. When $k \geq L$, $\|f_k\|^2 \leq (L+1)c^2\sigma^2$. Hence if $v > L$, we have

$$\begin{aligned} \sum_{k=0}^{v-1} \|f_k\|^\alpha &= \sum_{k=0}^{L-1} \|f_k\|^\alpha + \sum_{k=L}^{v-1} \|f_k\|^\alpha \\ &\leq \frac{2(c\sigma)^\alpha}{\alpha+2} ((L+1)^{\frac{\alpha+2}{2}} - 1) + (v-L)(c\sigma)^\alpha (L+1)^{\alpha/2} \\ &= v(c\sigma)^\alpha (L+1)^{\alpha/2} \\ &\quad + \frac{(c\sigma)^\alpha}{\alpha+2} ((L+1)^{\alpha/2} (2(L+1) - (\alpha+2)L) - 1) \\ &= v(c\sigma)^\alpha (L+1)^{\alpha/2} \\ &\quad - \frac{(c\sigma)^\alpha}{\alpha+2} ((L+1)^{\alpha/2} (\alpha L - 2) + 1), \end{aligned}$$

where the first inequality is by (20). Hence, by Theorem 1,

$$\begin{aligned} \mathbb{E} \text{cost}(CHC) - \mathbb{E} \text{cost}(OPT) &\leq \frac{2T\beta D}{v} + \frac{2GT}{v} \sum_{k=0}^{v-1} \|f_k\|^\alpha \\ &\leq \frac{2T\beta D}{v} + 2GT(c\sigma)^\alpha (L+1)^{\alpha/2} \\ &\quad - \frac{2GT}{v} \frac{(c\sigma)^\alpha}{\alpha+2} ((L+1)^{\alpha/2} (\alpha L - 2) + 1) \end{aligned}$$

The right hand side can be written as $2T(\frac{A-B}{v} + C)$, where $A = \beta D$, $B = G(c\sigma)^\alpha H(L)$, $C = G(c\sigma)^\alpha (L+1)^{\alpha/2}$ and

$H(L) = \frac{1}{\alpha+2} \left((L+1)^{\frac{\alpha}{2}} (\alpha L - 2) + 1 \right)$. When $A > B$, then the right hand side is a decreasing function in v , hence $v^* = w$; this happens when $\frac{\beta D}{G c^\alpha \sigma^\alpha} > H(L)$. When $A < B$, then the right hand side is an increasing function in v , hence we want v^* to be small, i.e., $v \leq L$, this happens when $\frac{\beta D}{G c^\alpha \sigma^\alpha} < H(L)$. When $v \leq L$, $\sum_{k=0}^{v-1} \|f_k\|^\alpha \leq \frac{2(c\sigma)^\alpha}{\alpha+2} ((v+1)^{1+\frac{\alpha}{2}} - 1)$, hence

$$\begin{aligned} \mathbb{E}\text{cost}(CHC) - \mathbb{E}\text{cost}(OPT) &\leq \frac{2T\beta D}{v} + \frac{2GT}{v} \sum_{k=0}^{v-1} \|f_k\|^\alpha \\ &\leq \frac{2T\beta D(\alpha+2) - 4GTc^\alpha \sigma^\alpha}{v(\alpha+2)} + \frac{4GT(c\sigma)^\alpha}{v(\alpha+2)} ((v+1)^{1+\alpha/2}) \end{aligned}$$

If $\frac{\beta D}{G c^\alpha \sigma^\alpha} \leq \frac{2}{\alpha+2}$, then the right hand side is an increasing function in v , hence $v^* = 1$.

B.4 Proof of Corollary 5

Taking expectation over the prediction error, assuming that there exists $a < 1$, such that for all s , $\|f(s)\|_F \leq ca^s$, we have

$$\begin{aligned} \|f_k\|^2 &= \sum_{s=0}^k \text{tr}(R_e f(s)^T f(s)) = \sum_{s=0}^k \langle R_e^{1/2}, f(s) \rangle^2 \\ &\leq \sum_{s=0}^k (\|R_e^{1/2}\|_F \|f(s)\|_F)^2 = \sum_{s=0}^k c^2 \sigma^2 a^{2s} = c^2 \sigma^2 \frac{1-a^{2k}}{1-a^2}, \end{aligned}$$

where the inequality is due to Cauchy-Schwarz, hence for h that is G -Lipschitz in the second argument, we have $\|f_k\| \leq \frac{c\sigma(1-a^{2(k+1)/2})}{1-a^2}$, where the inequality is because $\sqrt{1-a^2} \geq (1-a^2)$, and $1-a^{2(k+1)} \leq 1-a^{2(k+1)} + a^{4(k+1)}/4 = (1-a^{2(k+1)}/2)^2$, hence,

$$\begin{aligned} \mathbb{E}\text{cost}(CHC) - \mathbb{E}\text{cost}(OPT) &\leq \frac{2T\beta D}{v} + \frac{2GT}{v} \sum_{k=0}^{v-1} \frac{c\sigma(1-a^{2(k+1)/2})}{1-a^2} \\ &\leq \frac{2T\beta D}{v} + \frac{2GTc\sigma}{1-a^2} - \frac{GTc\sigma}{v(1-a^2)} \sum_{k=1}^v a^{2k} \\ &\leq \frac{2T\beta D}{v} + \frac{2GTc\sigma}{1-a^2} - \frac{GTc\sigma}{v(1-a^2)} \frac{a^2(1-a^{2v})}{1-a^2}. \end{aligned}$$

Let $A = 2T\beta D$, $B = \frac{a^2 GTc\sigma}{1-a^2}$ and $C = \frac{2GTc\sigma}{1-a^2}$, then

$$\begin{aligned} \mathbb{E}\text{cost}(CHC) - \mathbb{E}\text{cost}(OPT) &\leq \frac{A}{v} - \frac{B(1-a^{2v})}{v} + C \\ &= \frac{(A-B) + Ba^{2v}}{v} + C. \end{aligned}$$

Hence when $A \geq B$, $(A-B) + Ba^{2v} > 0$, and the RHS is a decreasing function in v , hence $v^* = w$, and this happens when $2T\beta D \geq \frac{a^2 GTc\sigma}{1-a^2}$ which implies $\frac{\beta D}{G c^\alpha \sigma^\alpha} \geq \frac{a^2}{2(1-a^2)}$.

On the other hand, if $A < B - Ba$, then $(A-B) + Ba^{2v} \leq A - B + Ba < 0$, and the right hand side is an increasing function in v , hence $v^* = 1$, and this happens when $2T\beta D < \frac{a^2 GTc\sigma}{1-a^2} (1-a)$, which implies $\frac{\beta D}{G c^\alpha \sigma^\alpha} \leq \frac{a^2}{2(1+a)}$.

B.5 Proof of Theorem 6

The proof follows in the same fashion as that of Theorem 1. Recall that we have

$$\begin{aligned} &g_{1,T}(\xi^{\tau+1}; y) - g_{1,T}(\xi^\tau; y) \\ &\leq 2\beta \|x_{(\tau+1)v+1}^* - \tilde{x}_{(\tau+1)v+1}\| + \sum_{t=\tau v+1}^{(\tau+1)v} |h(x_t^*, y_{t|\tau v}) - h(x_t^*, y_t)| \\ &\quad + \sum_{t=\tau v+1}^{(\tau+1)v} |h(\tilde{x}_t, y_t) - h(\tilde{x}_t, y_{t|\tau v})| \\ &\leq 2\beta \|x_{(\tau+1)v+1}^* - \tilde{x}_{(\tau+1)v+1}\| + 2G \sum_{t=\tau v+1}^{(\tau+1)v} \|y_t - y_{t|\tau v}\|_2^\alpha, \end{aligned} \tag{21}$$

Since h is m -strongly convex, $g_{\tau v+1, (\tau+1)v}$ is also m -strongly convex, hence

$$\begin{aligned} &g_{\tau v+1, (\tau+1)v}(\tilde{x}_{\tau v+1: (\tau+1)v}; \xi_{\tau v}^\tau; \tilde{x}_{(\tau+1)v+1}; y_{\cdot|\tau v}) \\ &\leq g_{\tau v+1, (\tau+1)v}(x_{\tau v+1: (\tau+1)v+1}^*; \xi_{\tau v}^\tau; \tilde{x}_{(\tau+1)v+1}; y_{\cdot|\tau v}) \\ &\quad - \partial g_{\tau v+1, (\tau+1)v}(\tilde{x}_{\tau v+1: (\tau+1)v}) \cdot (x_{\tau v: (\tau+1)v}^* - \tilde{x}_{\tau v: (\tau+1)v}) \\ &\quad - \frac{m}{2} \sum_{t=\tau v+1}^{(\tau+1)v} \|x_t^* - \tilde{x}_t\|^2 \end{aligned}$$

By the optimality of $\tilde{x}_{\tau v+1: (\tau+1)v}$ minimizes the cost function $g_{\tau v+1, (\tau+1)v}(x; \xi_{\tau v}^\tau; \tilde{x}_{(\tau+1)v+1}; y_{\cdot|\tau v})$, we have the first order condition

$$\partial g_{\tau v+1, (\tau+1)v}(\tilde{x}_{\tau v+1: (\tau+1)v}) \cdot (x_{\tau v+1: (\tau+1)v}^* - \tilde{x}_{\tau v+1: (\tau+1)v}) \geq 0,$$

hence

$$\begin{aligned} &\sum_{t=\tau v+1}^{(\tau+1)v} (h(\tilde{x}_t, y_{t|\tau v}) + \beta \|\tilde{x}_t - \tilde{x}_{t-1}\|) \\ &\quad + \beta \|\tilde{x}_{\tau v+1} - \xi_{\tau v}^\tau\| + \beta \|\tilde{x}_{(\tau+1)v+1} - \tilde{x}_{(\tau+1)v}\| \\ &\leq \sum_{t=\tau v+1}^{(\tau+1)v} (h(x_t^*, y_{t|\tau v}) + \beta \|x_t^* - x_{t-1}^*\| - \frac{m}{2} \|\tilde{x}_t - x_t^*\|^2) \\ &\quad + \beta \|x_{\tau v+1}^* - \xi_{\tau v}^\tau\| + \beta \|\tilde{x}_{(\tau+1)v+1} - x_{(\tau+1)v}^*\|, \end{aligned}$$

Substituting the above inequality into (21) and summing over τ , we have

$$\begin{aligned} \text{cost}(FHC^1(v)) - \text{cost}(OPT) &= \sum_{\tau=0}^M g_{1,T}(\xi^{\tau+1}; y) - g_{1,T}(\xi^\tau; y) \\ &\leq 2\beta \sum_{\tau=0}^{M-1} \|x_{(\tau+1)v+1}^* - \tilde{x}_{(\tau+1)v+1}\| - \sum_{\tau=0}^M \sum_{t=\tau v+1}^{(\tau+1)v} \frac{m}{2} \|\tilde{x}_t - x_t^*\|^2 \\ &\quad + 2G \sum_{\tau=0}^M \sum_{t=\tau v+1}^{(\tau+1)v} \|y_t - y_{t|\tau v}\|_2^\alpha \\ &\leq \sum_{\tau=1}^M \left(2\beta \|x_{\tau v+1}^* - \tilde{x}_{\tau v+1}\| - \frac{m}{2} \|x_{\tau v+1}^* - \tilde{x}_{\tau v+1}\|^2 \right) \\ &\quad + 2G \sum_{\tau=0}^M \sum_{t=\tau v+1}^{(\tau+1)v} \|y_t - y_{t|\tau v}\|_2^\alpha \\ &\stackrel{(a)}{\leq} \frac{2\beta^2 M}{m} + 2G \sum_{t=1}^T \|y_t - y_{t|\phi^{-1}(t)}\|_2^\alpha. \end{aligned}$$

where (a) is because for any t ,

$$\begin{aligned} & 2\beta \|x_t^* - \tilde{x}_t\| - \frac{m}{2} \|x_t^* - \tilde{x}_t\|^2 \\ &= -\frac{m}{2} (\|x_t^* - \tilde{x}_t\| - \frac{2\beta}{m})^2 + \frac{2\beta^2}{m} \leq \frac{2\beta^2}{m}. \end{aligned}$$

Summing over k from 0 to $v-1$ as in (18) and taking expectation on both sides as in (19) finishes the proof.

B.6 Proof of Theorem 7

By the proof of Theorem 1 and assuming one dimensional setting, we have

$$\begin{aligned} \text{cost}(CHC) - \text{cost}(OPT) &\leq \frac{2\beta TD}{v} + \frac{2G}{v} \sum_{\tau=1}^T \sum_{k=0}^{v-1} |y_\tau - y_{\tau-k}|^\alpha \\ &= \frac{2\beta TD}{v} + \frac{2G}{v} \sum_{\tau=1}^T \sum_{k=1}^v \left| \sum_{s=\tau-k+1}^{\tau} f(\tau-s)e(s) \right|^\alpha \end{aligned} \quad (22)$$

which is a function of the randomness of prediction noise $e = (e(1), \dots, e(T))$, let

$$l(e) := \frac{2\beta TD}{v} + \frac{2G}{v} \sum_{\tau=1}^T \sum_{k=1}^v \left| \sum_{s=\tau-k+1}^{\tau} f(\tau-s)e(s) \right|^\alpha$$

be the upper bound of the competitive difference of CHC in terms of the random variables $e(1), \dots, e(T)$. For every i , let $e(i)'$ be an independent and identical copy of $e(i)$, and let $e'_i = (e(1), \dots, e(i-1), e(i)', e(i+1), \dots, e(T))$ be the vector that differ from e by the replacing the i th coordinate with an identical copy of $e(i)$, then let δ_{is} be the kronecker delta, we have

$$\begin{aligned} |l(e) - l(e'_i)| &\leq \frac{2G}{v} \sum_{\tau=1}^T \sum_{k=1}^v \left| \sum_{s=\tau-k+1}^{\tau} \delta_{is} f(\tau-s)e(s) \right|^\alpha \\ &\quad - \frac{2G}{v} \sum_{\tau=1}^T \sum_{k=1}^v \left| \sum_{s=\tau-k+1}^{\tau} \delta_{is} f(\tau-s)e'_i(s) \right|^\alpha \\ &\leq \frac{2G}{v} \sum_{\tau=1}^T \sum_{k=1}^v \left| \sum_{s=\tau-k+1}^{\tau} \delta_{is} f(\tau-s)(e(s) - e'_i(s)) \right|^\alpha \end{aligned}$$

where the last inequality is due to the fact that $F(x) = x^\alpha$ is a α -Hölder continuous with constant 1, hence $\forall x, y, |x^\alpha - y^\alpha| \leq |x - y|^\alpha$. Therefore,

$$\begin{aligned} |l(e) - l(e'_i)| &\leq \frac{2G}{v} \sum_{\tau=1}^T \sum_{k=1}^v \left| \sum_{s=\tau-k+1}^{\tau} \delta_{is} f(\tau-s)(e(s) - e'_i(s)) \right|^\alpha \\ &\stackrel{(a)}{\leq} \frac{2G}{v} \sum_{k=1}^v \sum_{\tau=1}^T \sum_{s=\tau-k+1}^{\tau} \delta_{is} |f(\tau-s)|^\alpha |e(s) - e'_i(s)|^\alpha \\ &\stackrel{(b)}{=} \frac{2G}{v} |e(i) - e'_i(i)|^\alpha \sum_{k=1}^v \sum_{s=0}^{k-1} |f(s)|^\alpha \\ &\leq \frac{2^{1+\alpha} G \varepsilon^\alpha}{v} \sum_{s=0}^{v-1} (v-s) |f(s)|^\alpha \end{aligned} \quad (23)$$

where (a) is because for all $\alpha \leq 1$, $(\sum_{i=1}^k |a_i|)^\alpha \leq \sum_{i=1}^k |a_i|^\alpha$, to prove this, note that we only need to show that

$$\forall x, y \geq 0, (x+y)^\alpha \leq x^\alpha + y^\alpha \quad (24)$$

and iterate this inequality k times. To prove the above, note that when $x = 0$ or $y = 0$ or $\alpha = 1$, the inequality is trivially

true. Otherwise, wlog assume $x \geq y$, and let $t = y/x \leq 1$. (24) is equivalent to $(1+t)^\alpha \leq 1+t^\alpha$, this is true by examining the function $F(t) = 1+t^\alpha - (1+t)^\alpha$, note that $F(0) = 0$, and for $\alpha < 1$, $F(t)' = \alpha(t^{\alpha-1} - (1+t)^{\alpha-1}) > 0$, since a^x hence $F(t) \geq 0$ for all t , which proves (a). (b) is because let $s' = \tau - s$, then

$$\begin{aligned} & \sum_{k=1}^v \sum_{\tau=1}^T \sum_{s=\tau-k+1}^{\tau} \delta_{is} |f(\tau-s)|^\alpha |e(s) - e'_i(s)|^\alpha \\ &= \sum_{k=1}^v \sum_{\tau=1}^T \sum_{s'=0}^{k-1} \delta_{i(\tau-s')} |f(s')|^\alpha |e(\tau-s') - e'_i(\tau-s')|^\alpha \\ &= \sum_{k=1}^v \sum_{s'=0}^{k-1} |f(s')|^\alpha \left(\sum_{\tau=1}^T \delta_{i(\tau-s')} |e(\tau-s') - e'_i(\tau-s')|^\alpha \right) \\ &= |e(i) - e'_i(i)|^\alpha \sum_{k=1}^v \sum_{s'=0}^{k-1} |f(s')|^\alpha \end{aligned}$$

Let $c_i = 2^{1+\alpha} G \varepsilon^\alpha \sum_{k=0}^{v-1} \frac{(v-k)|f(k)|^\alpha}{v}$, then by (23), $|l(e) - l(e'_i)|^2 \leq c_i^2$, by Lemma 9, we have

$$\mathbb{P}(\text{cost}(CHC) - \text{cost}(OPT) \geq V_1 T + u)$$

$$\leq \exp \left(\frac{-2u^2}{\sum_{i=1}^T c_i^2} \right) = \exp \left(\frac{-u^2}{2^{1+2\alpha} G^2 \varepsilon^{2\alpha} T F(v)} \right)$$

where $F(v) = (\frac{1}{v} \sum_{k=0}^{v-1} (v-k)|f(k)|^\alpha)^2$.

B.7 Proof of Corollary 8

By Theorem 7,

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{T} [\text{cost}(CHC) - \text{cost}(OPT)] > V_1 + u \right) \\ &= \mathbb{P}(\text{cost}(CHC) - \text{cost}(OPT) > V_1 T + uT) \\ &\leq \exp \left(\frac{-2u^2}{\sum_{i=1}^T c_i^2} \right) = \exp \left(\frac{-u^2}{2^{1+2\alpha} G^2 \varepsilon^{2\alpha} F(v)/T} \right). \end{aligned}$$

If $v \in O(1)$, then $F(v)$ is bounded since $f(s)$ is bounded, hence $F(v)/T \rightarrow 0$ as $T \rightarrow \infty$.

Otherwise, if $f(s) \leq c\eta^s$ for $\eta < 1$, denote $a = \eta^\alpha$, then

$$\begin{aligned} \sum_{k=0}^{v-1} (v-k) |f(k)|^\alpha &\leq c^\alpha \left(v \sum_{k=0}^{v-1} a^k - a \sum_{k=0}^{v-1} \frac{d}{da} (a^k) \right) \\ &= c^\alpha \left(v \frac{1-a^v}{1-a} - a \frac{d}{da} \left(\frac{1-a^v}{1-a} \right) \right) = c^\alpha \frac{v-a(v+1)+a^{v+1}}{(1-a)^2} \end{aligned}$$

then

$$\begin{aligned} F(v) &= \left(\frac{1}{v} \sum_{k=0}^{v-1} (v-k) |f(k)|^\alpha \right)^2 \leq \left(\frac{c^\alpha (v - \eta^\alpha (v+1) + \eta^{\alpha(v+1)})}{(1-a)^2 v} \right)^2 \\ &\leq \left(\frac{c^\alpha}{(1-\eta^\alpha)^2} + \frac{c^\alpha \eta^\alpha}{(1-\eta^\alpha)^2 v} \right)^2 \in O(1). \end{aligned}$$

Hence in this case $F(v)/T \rightarrow 0$ as $T \rightarrow \infty$.