

ON OPTIMALITY GAPS IN THE HALFIN–WHITT REGIME

BY BARIS ATA AND ITAI GURVICH

Northwestern University

We consider optimal control of a multi-class queue in the Halfin–Whitt regime, and revisit the notion of asymptotic optimality and the associated optimality gaps. The existing results in the literature for such systems provide asymptotically optimal controls with optimality gaps of $o(\sqrt{n})$ where n is the system size, for example, the number of servers. We construct a sequence of asymptotically optimal controls where the optimality gap grows logarithmically with the system size. Our analysis relies on a sequence of Brownian control problems, whose refined structure helps us achieve the improved optimality gaps.

1. Introduction. Queueing models with many-servers are prevalent in modeling call centers and other large-scale service systems. They are used for optimizing staffing and making dynamic control decisions. The complexity of the underlying queueing model renders such optimization problems intractable for exact analysis, and one needs to resort to approximations. A prominent mode of approximate analysis is to study such systems in the so-called Halfin–Whitt (HW) heavy-traffic regime; cf. [8]. Roughly speaking, the analysis of a queueing system in the HW regime proceeds by scaling up the number of servers and the arrival rate of customers in such a way that the system load approaches one asymptotically. To be more specific, instead of considering a single system, one considers a sequence of (closely related) queueing systems indexed by a parameter n along which the arrival rates and the number of servers scale up so that the system traffic intensity ρ^n satisfies

$$(1) \quad \sqrt{n}(1 - \rho^n) \rightarrow \beta \quad \text{as } n \rightarrow \infty.$$

In the context of dynamic control, passing to a formal limit of the (properly scaled) system dynamics equations as $n \rightarrow \infty$ gives rise to a *limit* diffusion control problem, which is often more tractable than the original dynamic control problem it approximates. The approximating diffusion control problem typically provides useful structural insights and guides the design of good policies for the original system. Once a candidate policy is proposed for the original problem of interest, its asymptotic performance can be studied in the HW regime. The ultimate goal is

Received June 2010; revised December 2010.

MSC2010 subject classifications. 60K25, 90B20, 90B36, 49L20, 60F17.

Key words and phrases. Multiclass queues, many servers, heavy-traffic, Halfin–Whitt regime, optimal control, optimality gaps, asymptotic optimality, strong approximations for queues.

to establish that the proposed policy performs well. To this end, a useful criterion is the notion of asymptotic optimality, which provides assurance that the optimality gap associated with the proposed policy vanishes asymptotically *under diffusion scaling* as $n \rightarrow \infty$. Hence, asymptotic optimality in this context is equivalent to showing that the optimality gap is $o(\sqrt{n})$.

A central reference for our purposes is the recent paper by Atar, Mandelbaum and Reiman [2], where the authors apply all steps of the above scheme to the important class of problems of dynamically scheduling a multiclass queue with many identical servers in the HW regime. Specifically, [2] considers a sequence of systems indexed by the number of servers n , where the number of servers and the arrival rates of the various customer classes increase with n such that the heavy-traffic condition holds; cf. equation (1). Following the scheme described above, the authors derive an approximate diffusion control problem through a formal limiting argument. They then show that the diffusion control problem admits an optimal Markov policy, and that the corresponding HJB equation (a semilinear elliptic PDE) has a unique classical solution. Using the Markov control policy and the HJB equation, the authors propose scheduling control policies for the original (sequence of) queueing systems of interest. Finally, they prove that the proposed sequence of policies is asymptotically optimal under diffusion scaling. Namely, the optimality gap of the proposed policy for the n th system is $o(\sqrt{n})$. A similar approach is applied to more general networks in [1]. In this paper, we study a similar queueing system (see Section 2). Our goal, however, is to provide an improved optimality gap which, in turn, requires a substantially different scheme than the one alluded to above.

Approximations in the HW regime for performance analysis have been used extensively for the study of fixed policies. Given a particular policy, it may often be difficult to calculate various performance measures in the original queueing system. Fortunately, the corresponding approximations in the HW regime are often more tractable. The machinery of strong approximations (cf. Csörgo and Horváth [4]) often plays a central role in such analysis. In the context of many-server heavy-traffic analysis, with strong approximations, the arrival and service processes (under suitable assumptions on the inter-arrival and service times) can be approximated by a diffusion process so that the approximation error on finite intervals is $O(\log n)$ (where n is the number of servers as before). Therefore, it is natural to expect that, under a given policy, the error in the diffusion approximations of the various performance metrics is $O(\log n)$, which is indeed verified for various settings in the literature (see, e.g., [11]).

A natural question is then whether one can go beyond the analysis of fixed policies and achieve an optimality gap that is logarithmic in n also under dynamic control, improving upon the usual optimality gap of $o(\sqrt{n})$. More specifically, can one propose a sequence of policies (one for each system in the sequence) where the optimality gap for the policy (associated with the n th system) is logarithmic in n ?

While one hopes to get logarithmic optimality gaps as suggested by strong approximations, it is not a priori clear if this can be achieved under dynamic control. The purpose of this paper is to provide a resolution to this question. Namely, we study whether one can establish such a strong notion of asymptotic optimality and if so, then how should one go about constructing policies which are asymptotically optimal in this stronger sense.

Our results show that such strengthened bounds on optimality gaps can be attained. Specifically, we construct a sequence of asymptotically optimal policies, where the optimality gap is logarithmic in n . Our analysis reveals that identifying (a sequence of) candidate policies requires a new approach. To be specific, we advance a sequence of diffusion control problems (as opposed to just one) where the diffusion coefficient in each system depends on the state and the control. This is contrary to the existing work on the asymptotic analysis of queueing systems in the HW regime. In that stream of literature, the diffusion coefficient is typically a (deterministic) constant. Indeed, Borkar [3] views the constant diffusion coefficient as a characterizing feature of the problems stemming from the heavy-traffic approximations in the HW regime. Interestingly, it is essential in our work to have the diffusion coefficient depend on the state and the control for achieving the logarithmic optimality gap. In essence, incorporating the impact of control on the diffusion coefficient allows us to track the policy performance in a more refined manner.

While the novelty of having the diffusion coefficient depend on the control facilitates better system performance, it also leads to a more complex diffusion control problem. In particular, the associated HJB equation is fully nonlinear; it is also nonsmooth under a linear holding cost structure. In what follows, we show that each of the HJB equations in the sequence has a unique smooth solution on bounded domains and that each of the diffusion control problems (when considered up to a stopping time) admits an optimal Markov control policy. Interpreting this solution appropriately in the context of the original problem gives rise to a policy under which the optimality gap is logarithmic in n . As in the performance analysis of fixed policies, strong approximations will be used in the last step, where we propose a sequence of controls for the original queueing systems, and show that we achieve the desired performance. However, it is important to note that strong approximation results alone are not sufficient for our results. Rather, for the improved optimality gaps we need the refined properties of the solutions to the HJB equations. Specifically, gradient estimates for the sequence of solutions to the HJB equations (cf. Theorem 4.1) play a central role in our proofs.

Our analysis restricts attention to a linear holding cost structure. However, we expect the analysis to go through for some other cost structures including convex holding costs. Indeed, the analysis of the convex holding cost case will probably be simpler as one tends to get “interior” solutions in that case as opposed to the corner solutions in the linear cost case, which causes nonsmoothness. One could also enrich the model by allowing abandonment. We expect the analysis to go through

with no major changes in these cases as well; see the discussion of possible extensions in Section 7. For purposes of clarity, however, we chose not to incorporate these additional/alternative features because we feel that the current set-up enables us to focus on and clearly communicate the main idea: the use of a novel Brownian model with state/control dependent diffusion coefficient to obtain improved optimality gaps.

Organization of the paper. Section 2 formulates the model and states the main result. Section 3 introduces a (sequence of) Brownian control problem(s), which are then analyzed in Section 4. A performance analysis of our proposed policy appears in Section 5. The major building blocks of the proof are combined to establish the main result in Section 6 and some concluding remarks appear in Section 7.

2. Problem formulation. We consider a queueing system with a single server-pool consisting of n identical servers (indexed from 1 to n) and a set $\mathcal{I} = \{1, \dots, I\}$ of job classes as depicted in Figure 1. Jobs of class- i arrive according to a Poisson process with rate λ_i and wait in their designated queue until their service begins. Once admitted to service, the service time of a class- i job is distributed as an exponential random variable with rate $\mu_i > 0$. All service and interarrival times are mutually independent.

Heavy-traffic scaling. We consider a sequence of systems indexed by the number of servers n . The superscript n will be attached to various processes and parameters to make the dependence on n explicit. (It will be omitted from parameters and




FIG. 1. A multiclass queue with many servers.

other quantities that do not change with n .) We assume that $\lambda_i^n = a_i \lambda^n$ for all n , where λ^n is the total arrival rate and $a_i > 0$ for $i \in \mathcal{I}$ with $\sum_i a_i = 1$. This assumption is made for simplicity of notation and presentation. Nothing changes in our results if one assumes, instead, that $\lambda_i^n/n \rightarrow \lambda_i$ and $\sqrt{n}(\lambda_i^n/n - \lambda_i) \rightarrow \hat{\lambda}_i$ as $n \rightarrow \infty$ where $\lambda_i/\sum_{k \in \mathcal{I}} \lambda_k = a_i > 0$.

The nominal load in the n th system is then given by

$$R^n = \sum_i \frac{\lambda_i^n}{\mu_i} = \lambda^n \sum_i \frac{a_i}{\mu_i},$$

so that defining $\bar{\mu} = [\sum_i a_i/\mu_i]^{-1}$ we have that $R^n = \lambda^n/\bar{\mu}$, which corresponds to the nominal number of servers required to handle all the incoming jobs. The heavy-traffic regime is then imposed by requiring that the number of servers deviates from the nominal load by a term that is a square root of the nominal load. Formally, we impose this by assuming that λ^n is such that

$$(2) \quad n = R^n + \beta \sqrt{R^n}$$

for some $\beta \in (-\infty, \infty)$ that does not scale with n . Also, we define the relative load imposed on the system by class- i jobs, denoted by v_i , as follows:

$$(3) \quad v_i = \frac{a_i/\mu_i}{\sum_{k \in \mathcal{I}} a_k/\mu_k}.$$

Note that $\sum_{i \in \mathcal{I}} v_i = 1$, and $v_i n$ can be interpreted as a first-order (fluid) estimate for the number of servers busy serving class- i customers.

2.1. System dynamics. Let $Q_i^n(t)$ and $X_i^n(t)$ denote the number of class- i jobs in the queue and in the system, respectively, at time t in the n th system. Similarly, let $Z_i^n(t)$ denote the number of servers working on class- i jobs at time t . Clearly, for all i, n, t , the following holds:

$$X_i^n(t) = Z_i^n(t) + Q_i^n(t).$$

In our setting, a control corresponds to determining how many of the class- i jobs currently in the system are placed in queue and in service for $i \in \mathcal{I}$. We take the process Z^n as our control in the n th system. Note that one can equivalently take the queue length process Q^n as the control. (The knowledge of either process is sufficient to pin down the evolution of the system given the arrival, service processes and the initial conditions.) Clearly, the control process must satisfy certain requirements for admissibility, including the usual nonanticipativity requirement. We defer a precise mathematical definition of admissible controls for now (see Definition 2.2). However, it should be clear that, given the process Z^n , one can construct the other processes of interest.

To be specific, consider a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and $2I$ mutually independent *unit-rate* Poisson processes $(\mathcal{N}_i^a(\cdot), \mathcal{N}_i^d(\cdot), i \in \mathcal{I})$ on that space.

Given the *primitives* $(\mathcal{N}_i^d(\cdot), \mathcal{N}_i^a(\cdot), X_i^n(0), Z_i^n(0); i \in \mathcal{I})$ and the control process Z^n , we construct the processes X^n, Q^n as follows: for $t \geq 0$ and $i \in \mathcal{I}$

$$(4) \quad X_i^n(t) = X_i^n(0) + \mathcal{N}_i^a(\lambda_i^n t) - \mathcal{N}_i^d\left(\mu_i \int_0^t Z_i^n(s) ds\right),$$

$$(5) \quad Q_i^n(t) = X_i^n(t) - Z_i^n(t).$$

The processes Z^n, Q^n, X^n must jointly satisfy the constraints

$$(6) \quad (Q^n(t), X^n(t), Z^n(t)) \in \mathbb{Z}_+^{3I}, \quad e \cdot Z^n(t) \leq n,$$

where e is the I -dimensional vector of ones.

Controls can be preemptive or nonpreemptive. Under a nonpreemptive control, a job that is assigned to a server keeps the server busy until its service is completed. In particular, given a nonpreemptive control Z^n , the process Z_i^n can decrease only through service completions of class- i jobs. In contrast, the class of preemptive controls is broader. While it includes nonpreemptive policies, it also includes controls that (occasionally) may preempt a job’s service. The preempted job is put back in the queue and its service is resumed at a later time (possibly by a different server). Hence, the class of preemptive controls subsumes the class of nonpreemptive ones (which is also immediate from Definition 1 in [2]) and the cost of an optimal policy among preemptive ones gives a lower bound for that among the nonpreemptive ones.

In what follows, we will largely focus on preemptive controls, which are easier to work with, and derive a specific policy which is near optimal in that class. The specific policy we derive is, however, nonpreemptive, and therefore, is near optimal among the nonpreemptive policies as well. More specifically, the policy we propose belongs to a class which we refer to as *tracking policies*.

To facilitate the definition of tracking policies, define $\mathcal{U} \subset \mathbb{R}_+^I$ as

$$(7) \quad \mathcal{U} = \left\{ u \in \mathbb{R}_+^I : \sum_i u_i = 1 \right\}.$$

Also, for all i and $t \geq 0$, let

$$(8) \quad \check{X}_i^n(t) = X_i^n(t) - v_i n.$$

Hence, the process \check{X}_i^n captures the oscillations of the process X_i^n around its “fluid” approximation $v_i n$. Throughout our analysis, for $x \in \mathbb{R}$ we let $(x)^+ = \max\{0, x\}$ and $(x)^- = \max\{0, -x\}$.

DEFINITION 2.1. Given a function $h : \mathbb{R}^I \rightarrow \mathcal{U}$, an h -tracking policy makes resource allocation decisions in the n th system as follows:

- (i) It is nonpreemptive. That is, once a server starts working on a job, it continues without interruption until that job’s service is completed.

(ii) It is work conserving. That is, the number of busy servers satisfies $e \cdot Z^n(t) = (e \cdot X^n(t)) \wedge n$ for all $t > 0$. In particular, no server is idle as long as there are n or more jobs in the system.

(iii) When a class- i job arrives to the system it joins the queue of class i if all servers are busy processing other jobs. Otherwise, the lowest-indexed idle server starts working on that job.

(iv) A server that finishes processing a job at a time t , idles if all queues are empty. Otherwise, she starts working on a job of class $i \in \mathcal{K}(t-)$ with probability $\lambda_i^n / \sum_{k \in \mathcal{K}(t-)} \lambda_k^n$, where, for $t > 0$, the set $\mathcal{K}(t-)$ is defined by

$$(9) \quad \mathcal{K}(t-) = \{k \in \mathcal{I} : Q_k(t) - h_k(\check{X}^n(t-))(e \cdot \check{X}^n(t-))^+ > 0\}.$$

Finally, if $(e \cdot \check{X}^n(t-))^+ > 0$ and $\mathcal{K}(t-) = \emptyset$, she picks for service a customer from the lowest index nonempty queue.

REMARK 2.1. For our optimality-gap bounds and, in particular, for the proof of Theorem 5.1 it is important that the policy be such that each of the job classes in the set $\mathcal{K}(t)$ gets a sufficient share of the capacity. This prevents excessive oscillation of the queues that may compromise the optimality gaps. Such oscillations could arise if, for example, the policy chooses for service a job of class

$$i = \min \arg \max_{k \in \mathcal{I}} \{Q_k(t-) - h_k(\check{X}^n(t-))(e \cdot \check{X}^n(t-))^+ : Q_k^n(t-) > 0\}.$$

Randomization is just one way to overcome such oscillations and, as the proofs (specifically that of Theorem 5.1) reveal, any choice rule that guarantees a sufficient share of the capacity to a class in $\mathcal{K}(t-)$ will suffice.

Our main result shows that a (nonpreemptive) tracking policy can achieve a near optimal performance among preemptive policies. Note that in our setting under preemption, one can restrict attention to work-conserving policies, that is, policies under which the servers never idle as long as there are jobs to work on.¹ More precisely, a control is work conserving if the following holds for all $t > 0$:

$$(10) \quad e \cdot Q^n(t) = (e \cdot \check{X}^n(t))^+.$$

Hereafter, we focus on work-conserving controls. Each such control can be mapped into a ratio control, which specifies what fraction of the total number of jobs in queue belongs to each class. To that end, let

$$(11) \quad U_i^n(t) = \frac{Q_i^n(t)}{(e \cdot Q^n(t)) \vee 1}.$$

¹By a coupling argument, this can be shown to hold with general queueing costs provided that there are no abandonments and that the service times are exponential; see, for example, the coupling argument on page 1126 of [2].

Note that the original control Z^n can be recovered from the ratio control U^n as follows:

$$Z_i^n(t) = X_i^n(t) - U_i^n(t)(e \cdot \check{X}^n(t))^+.$$

Equations (4)–(6) can then be replaced by

$$(12) \quad X_i^n(t) = X_i^n(0) + \mathcal{N}_i^a(\lambda_i^n t) - \mathcal{N}_i^d \left(\mu_i \int_0^t (X_i^n(s) - U_i^n(s)(e \cdot \check{X}^n(t))^+) ds \right),$$

$$(13) \quad Q_i^n(t) = U_i^n(t)(e \cdot \check{X}^n(t))^+,$$

$$(14) \quad Z_i^n(t) = X_i^n(t) - Q_i^n(t),$$

$$(15) \quad \check{X}_i^n(t) = X_i^n(t) - v_i n,$$

$$(16) \quad U^n(t) \in \mathcal{U}, \quad Q^n(t) \in \mathbb{Z}_+^I, \quad X^n(t) \in \mathbb{Z}_+^I.$$

Define the filtration

$$\bar{\mathcal{F}}_t = \sigma \{ \mathcal{N}_i^a(s), \mathcal{N}_i^d(s); i \in \mathcal{I}, s \leq t \}$$

and the σ -field

$$(17) \quad \bar{\mathcal{F}}_\infty = \bigvee_{t \geq 0} \bar{\mathcal{F}}_t.$$

Informally, $\bar{\mathcal{F}}_\infty$ contains the information about the entire evolution of the processes $(\mathcal{N}_i^a, \mathcal{N}_i^d, i \in \mathcal{I})$. A natural notion of admissibility requires that the control is nonanticipative so that it only uses historical information about the process X^n and about the arrivals and service completions up to the decision epoch. To accommodate randomized policies (as the h -tracking policy) we allow the control to use other information too as long as this information is independent of $\bar{\mathcal{F}}_\infty$.

DEFINITION 2.2. A process $U = (U_i(t), t \geq 0, i \in \mathcal{I})$ is a ratio control for the n th system if there exists a process $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ such that, together with the primitives, (\mathbb{X}^n, U) satisfies (12)–(16). The process U is an admissible ratio control if, in addition, it is adapted to the filtration $\mathcal{G} \vee \mathcal{F}_t^n$ where

$$\mathcal{F}_t^n = \sigma \left\{ \mathcal{N}_i^a(\lambda_i^n s), X_i^n(s), \mu_i \int_0^s Z_i^n(u) du, \mathcal{N}_i^d \left(\mu_i \int_0^s Z_i^n(u) du \right); i \in \mathcal{I}, 0 \leq s \leq t \right\},$$

and \mathcal{G} is a σ -field that is independent of $\bar{\mathcal{F}}_\infty$. The process \mathbb{X}^n is then said to be the queueing process associated with the ratio control U . We let Π^n be the set of admissible ratio controls for the n th system.

Ratio controls are work conserving by definition, but they need not be non-preemptive in general. However, note that given a function $h: \mathbb{R}^I \rightarrow \mathcal{U}$, the (non-preemptive) h -tracking policy corresponds to a ratio control U_h , which is non-preemptive. To be specific, given the primitives and the h -tracking policy, one can construct the corresponding queueing process $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ (see the construction after Lemma A.1). Then the ratio control U_h is constructed using the relation (11) so that \mathbb{X}^n and U_h jointly satisfy (12)–(16). Hence, one can speak of the ratio control and the queueing process associated with an h -tracking policy. Note that since the tracking policy makes resource allocation decisions using only information on the state of the system at the decision epoch (together with a randomization that is independent of the history), the resulting ratio control is admissible in the sense of Definition 2.2. The terms ratio control and h -tracking policy appear in several places in the paper. It will be clear from the context whether we refer to an arbitrary ratio control or to one associated with an h -tracking policy.

We close this section by stating the main result of the paper. To that end, let

$$(18) \quad \mathcal{X}^n = \{(x, q) \in \mathbb{Z}_+^{2I} : q = u(e \cdot x - n)^+ \text{ for some } u \in \mathcal{U}\}.$$

That is, \mathcal{X}^n is the set on which (X^n, Q^n) can obtain values under work conservation. In this set $e \cdot q = (e \cdot x - n)^+$ so that positive queue and idleness do not co-exist. We let $\mathbb{E}_{x,q}^U[\cdot]$ denote the expectation with respect to the initial condition $(X^n(0), Q^n(0)) = (x, q)$ and an admissible ratio control U . Given a ratio control U and initial conditions (x, q) , the expected infinite horizon discounted cost in the n th system is given by

$$(19) \quad C^n(x, q, U) = \mathbb{E}_{x,q}^U \left[\int_0^\infty e^{-\gamma s} c \cdot Q^n(s) ds \right],$$

where $c = (c_1, \dots, c_I)'$ is the strictly positive vector of holding cost rates and $\gamma > 0$ is the discount rate. For $(x, q) \in \mathcal{X}^n$, the value function is given by

$$V^n(x, q) = \inf_{U \in \Pi^n} \mathbb{E}_{x,q}^U \left[\int_0^\infty e^{-\gamma s} c \cdot Q^n(s) ds \right].$$

We next state our main result.

THEOREM 2.1. *Fix a sequence $\{(x^n, q^n), n \in \mathbb{Z}_+\}$ such that $(x^n, q^n) \in \mathcal{X}^n$ and $|x^n - vn| \leq M\sqrt{n}$ for all n and some $M > 0$. Then, there exists a sequence of tracking functions $\{h^n, n \in \mathbb{Z}_+\}$ together with constants $C, k > 0$ (that do not depend on n) such that*

$$C^n(x^n, q^n, U_h^n) \leq V^n(x^n, q^n) + C \log^k n \quad \text{for all } n,$$

where U_h^n is the ratio control associated with the h^n -tracking policy.

The constant k in our bound may depend on all system and cost parameters but not on n . In particular, it may depend on $(\mu_i, c_i, a_i; i \in \mathcal{I})$ and β . Its value is explicitly defined after the statement of Theorem 4.1.

Theorem 2.1 implies, in particular, that the optimal performance for nonpreemptive policies is close to that among the larger family of preemptive policies. Indeed, we identify a nonpreemptive policy (a tracking policy) in the queueing model whose cost performance is close to the optimal value of the preemptive control problem.

The rest of the paper is devoted to the proof of Theorem 2.1, which proceeds by studying a sequence of auxiliary Brownian control problems. The next subsection offers a heuristic derivation and a justification for the relevance of the sequence of Brownian control problems to be considered in later sections.

2.2. Toward a Brownian control problem. We proceed by deriving a sequence of approximating Brownian control problems heuristically, which will be instrumental in deriving a near-optimal policy for our original control problem. It is important to note that we derive an approximating Brownian control problem for each n as opposed to deriving a single approximating problem (for the entire sequence of problems). This distinction is crucial for achieving an improved optimality gap for n large because it allows us to tailor the approximation to each element of the sequence of systems.

To this end, let

$$l_i^n = \lambda_i^n - \mu_i v_i n \quad \text{for } i \in \mathcal{I}.$$

Fixing an admissible control U^n for the n th system [and centering as in (8)], we can then write (12) as

$$(20) \quad \check{X}_i^n(t) = \check{X}_i^n(0) + l_i^n t - \mu_i \int_0^t (\check{X}_i^n(s) - U_i^n(s)(e \cdot \check{X}^n(s))^+) ds + \check{W}_i^n(t),$$

where

$$(21) \quad \begin{aligned} \check{W}_i^n(t) = & \mathcal{N}_i^a(\lambda_i^n t) - \lambda_i^n t + \mu_i \int_0^t (\check{X}_i^n(s) - U_i^n(s)(e \cdot \check{X}^n(s))^+) ds \\ & - \mathcal{N}_i^d\left(\mu_i \int_0^t (\check{X}_i^n(s) + v_i n - U_i^n(s)(e \cdot \check{X}^n(s))^+) ds\right). \end{aligned}$$

In words, $\check{W}_i^n(t)$ captures the deviations of the Poisson processes from their means. It is natural to expect that an approximation result of the following form will hold: $(\check{X}_i^n, \check{W}_i^n; i \in \mathcal{I})$ can be approximated by $(\hat{X}_i^n, \hat{W}_i^n; i \in \mathcal{I})$ where

$$\begin{aligned} \hat{X}_i^n(t) &= \hat{X}_i^n(0) + l_i^n t - \mu_i \int_0^t (\hat{X}_i^n(s) - U_i^n(s)(e \cdot \hat{X}^n(s))^+) ds + \hat{W}_i^n(t), \\ \hat{W}_i(t) &= \tilde{B}_i^a(\lambda_i^n t) + \tilde{B}_i^S\left(\mu_i \int_0^t (\hat{X}_i^n(s) + v_i n - U_i^n(s)(e \cdot \hat{X}^n(s))^+) ds\right) \end{aligned}$$

and \tilde{B}^a, \tilde{B}^s are I -dimensional independent standard Brownian motions. Moreover, by a time-change argument we can write (see, e.g., Theorem 4.6 in [9])

$$(22) \quad \begin{aligned} \hat{X}_i^n(t) = & \hat{X}_i^n(0) + l_i^n t - \mu_i \int_0^t \hat{X}_i^n(s) - U_i^n(s)(e \cdot \hat{X}_i^n(s))^+ ds \\ & + \int_0^t \sqrt{\lambda_i^n + \mu_i(\hat{X}_i^n(s) + v_i n - U_i^n(s)(e \cdot \hat{X}_i^n(s))^+)} dB_i(s), \end{aligned}$$

where B is an I -dimensional standard Brownian motion constructed by setting

$$(23) \quad \begin{aligned} B_i(t) = & \int_0^t \frac{d\tilde{B}_i^S(\mu_i \int_0^s (\hat{X}_i^n(u) + v_i n - U_i^n(u)(e \cdot \hat{X}_i^n(u))^+ du)}{\sqrt{\mu_i(\hat{X}_i^n(s) + v_i n - U_i^n(s)(e \cdot \hat{X}_i^n(s))^+}} \\ & + \frac{\tilde{B}_i^a(\lambda_i^n t)}{\lambda_i^n t}. \end{aligned}$$

Taking a leap of faith and arguing heuristically, we next consider a Brownian control problem with the system dynamics

$$(23) \quad \hat{X}^n(t) = x + \int_0^t b^n(\hat{X}^n(s), \hat{U}^n(s)) ds + \int_0^t \sigma^n(\hat{X}^n(s), \hat{U}^n(s)) dB(t),$$

where \hat{U}^n will be an admissible control for the Brownian system and

$$(24) \quad b_i^n(x, u) = l_i^n - \mu_i(x_i - u_i(e \cdot x)^+)$$

and

$$(25) \quad \sigma_i^n(x, u) = \sqrt{\lambda_i^n + \mu_i v_i n + \mu_i(x_i - u_i(e \cdot x)^+)}.$$

Note that the Brownian control problem will only be used to propose a candidate policy, whose near optimality will be verified from first principles without relying on the heuristic derivations of this section.

To repeat, the preceding definition is purely formal and provided only as a means of motivating our approach. In what follows, we will directly state and analyze an auxiliary Brownian control problem motivated by the above heuristic derivation. The analysis of the auxiliary Brownian control problem lends itself to constructing near optimal policies for our original control problem. To be more specific, the system dynamics equation (23), and in particular, the fact that its variance is state and control dependent, is crucial to our results. Indeed, it is this feature of the auxiliary Brownian control problems that yields an improved optimality gap.

Needless to say, one needs to take care in interpreting (23)–(25), which are meaningful only up to a suitably defined hitting time. In particular, to have σ^n well defined, we restrict attention to the process while it is within some bounded domain. Actually, it suffices for our purposes to fix $\kappa > 0$ and $m \geq 3$ and consider the Brownian control problem only up to the hitting time of a ball of the form

$$(26) \quad \mathcal{B}_\kappa^n = \{x \in \mathbb{R}^I : |x| < \kappa \sqrt{n} \log^m n\},$$

where $|\cdot|$ denotes the Euclidian norm. We will fix the constant m throughout and suppress the dependence on m from the notation. Setting

$$(27) \quad n(\kappa) = \inf\{n \in \mathbb{Z}_+ : \sigma^n(x, u) \geq 1 \text{ for all } x \in \mathcal{B}_\kappa^n, u \in \mathcal{U}\},$$

the diffusion coefficient is strictly positive for all $n \geq n(\kappa)$ and $x \in \mathcal{B}_\kappa^n$. Note that, for all $i \in \mathcal{I}$, $x \in \mathcal{B}_\kappa^n$ and $u \in \mathcal{U}$,

$$(\sigma_i^n(x, u))^2 \geq \lambda_i^n + \mu_i v_i n - 2\mu_i \kappa \sqrt{n} \log^m n,$$

so that $(\sigma_i^n(x, u))^2 \geq \mu_i v_i n / 2 \geq 1$ for all sufficiently large n and, consequently, $n(\kappa) < \infty$.

REMARK 2.2. In what follows, and, in particular, through the proof of Theorem 2.1, the reader should note that while choosing the size of the ball to be ϵn (with ϵ small enough) would suffice for the nondegeneracy of the diffusion coefficient, that choice would be too large for our optimality gap proofs.

3. An approximating diffusion control problem (ADCP). Motivated by the discussion in the preceding section, we define admissible systems as follows.

DEFINITION 3.1 (Admissible systems). Fix $\kappa > 0$, $n \in \mathbb{Z}_+$ and $x \in \mathbb{R}^I$. We refer to $\theta = (\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P}, \hat{U}, B)$ as an admissible (κ, n) -system if:

- (a) $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ is a complete filtered probability space.
- (b) $B(\cdot)$ is an I -dimensional standard Brownian motion adapted to (\mathcal{F}_t) .
- (c) \hat{U} is \mathcal{U} -valued, \mathcal{F} -measurable and (\mathcal{F}_t) progressively measurable.

The process \hat{U} is said to be the control associated with θ . We also say that \hat{X} is a controlled process associated with the initial data x and an admissible system θ if \hat{X} is a continuous (\mathcal{F}_t) -adapted process on $(\Omega, \mathcal{F}, \mathbb{P})$ such that, almost surely, for $t \leq \hat{\tau}_\kappa^n$,

$$\hat{X}(t) = x + \int_0^t b^n(\hat{X}(s), \hat{U}(s)) ds + \int_0^t \sigma^n(\hat{X}(s), \hat{U}(s)) d\tilde{B}(t),$$

where $b^n(\cdot, \cdot)$ and $\sigma^n(\cdot, \cdot)$ are as defined in (24) and (25), respectively, and $\hat{\tau}_\kappa^n = \inf\{t \geq 0 : \hat{X}(t) \notin \mathcal{B}_\kappa^n\}$. Given $\kappa > 0$ and $n \in \mathbb{Z}_+$, we let $\Theta(\kappa, n)$ be the set of admissible (κ, n) -systems.

The Brownian control problem then corresponds to optimally choosing an admissible (κ, n) -system with associated control $(\hat{U}(t), t \geq 0)$ that achieves the minimal cost in the optimization problem

$$(28) \quad \hat{V}^n(x, \kappa) = \inf_{\theta \in \Theta(\kappa, n)} \mathbb{E}_x^\theta \left[\int_0^{\hat{\tau}_\kappa^n} e^{-\gamma s} \sum_{i \in \mathcal{I}} c_i \hat{U}_i(s) (e \cdot \hat{X}(s))^+ ds \right],$$

where $\mathbb{E}_x^\theta[\cdot]$ is the expectation operator when the initial state is $x \in \mathbb{R}^I$ and the admissible system θ . Hereafter, we refer to (28) as the ADCP on \mathcal{B}_κ^n . The following lemma shows that the Definition 3.1 is not vacuous. The proof appears in the Appendix.

LEMMA 3.1. *Fix the initial state $x \in \mathbb{R}^I$, $\kappa > 0$, $n \geq n(\kappa)$ and an admissible (κ, n) -system θ . Then, there exists a unique controlled process \hat{X} associated with x and θ .*

To facilitate future analysis, note from the definition of $\hat{\tau}_\kappa^n$ and (28) that

$$(29) \quad \hat{V}^n(x, \kappa) \leq \frac{1}{\gamma} (e \cdot c) \kappa \sqrt{n} \log^m n.$$

DEFINITION 3.2 (Markov controls). We say that an admissible (κ, n) -system $\theta = (\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P}, \hat{U}, B)$ with the associated controlled process \hat{X}^n induces a Markov control if there exists a function $g^n(\cdot) : \mathcal{B}_\kappa^n \rightarrow \mathcal{U}$ such that $\hat{U}(t) = g^n(\hat{X}^n(t))$ for $t \leq \hat{\tau}_\kappa^n$. We extend the function g^n to \mathbb{R}^I as follows:

$$(30) \quad h^n(x) = \begin{cases} g^n(x), & x \in \mathcal{B}_\kappa^n, \\ e_1, & \text{otherwise,} \end{cases}$$

where e_1 is the I -dimensional vector whose first component is 1 while the others are 0. We refer to $h^n(\cdot)$ as the tracking function associated with the admissible system θ .

In what follows, a policy \hat{U} will be called optimal for the approximating diffusion control problem (ADCP) on \mathcal{B}_κ^n if there exists an admissible (κ, n) -system $\theta = (\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P}, \hat{U}, B)$ such that

$$\hat{V}^n(x, \kappa) = \mathbb{E}_x^\theta \left[\int_0^{\hat{\tau}_\kappa^n} e^{-\gamma s} \sum_{i \in \mathcal{I}} c_i \hat{U}_i(s) (e \cdot \hat{X}(s))^+ ds \right].$$

Recall that X and U are used to denote performance relevant stochastic processes in both the Brownian model and the original queueing model, and that we add a hat, that is, we use \hat{X} and \hat{U} in the context of the Brownian model. To avoid confusion, the reader should keep in mind that hat-processes correspond to the ADCP while the ones with no hats correspond to the original queueing model.

Roadmap for the remainder of the paper. The main result in Theorem 2.1 builds on the following steps:

1. In Section 4 we show that for each n , the HJB equation associated with the ADCP has a unique and sufficiently smooth solution. Using that solution we

advance an optimal Markov control for the ADCP together with the corresponding tracking function. We also identify useful gradient bounds on the solutions to the sequence of HJB equations; cf. Theorem 4.1.

2. In Section 5 we conduct a performance analysis of h -tracking policies in the queueing system; cf. Theorem 5.1.
3. The result of Theorem 5.1 together with the gradient estimates in Theorem 4.1 are combined in a Taylor expansion-type argument in Section 6 to complete the proof of Theorem 2.1.

As a convention, throughout the paper we use the capital letter C to denote a constant that does not depend on n . The value of C may change from line to line within the proofs but it will be clear from the context.

4. Solution to the ADCP. This section provides a solution for the ADCP on \mathcal{B}_κ^n for each $n \in \mathbb{Z}$ and $\kappa > 0$. The HJB equation is a fully nonlinear and nonsmooth PDE. As such, it requires extra care when compared with the usual semilinear PDEs that arise in the analysis of *asymptotically* optimal controls in the Halfin–Whitt regime. We will build on existing results in the theory of PDEs and proceed through the following steps: (a) establish the existence and uniqueness of classical solutions; (b) relate this unique solution to the value function of the ADCP and (c) establish useful gradient estimates on the solution for the HJB equation. The last step is not necessary for existence and uniqueness but is important for the analysis of optimality gaps.

In what follows, we fix $\kappa > 0$ and $n \geq n(\kappa)$ and suppress the dependence of the solution to the HJB equation on n and κ . The following notation is needed to introduce the HJB equation. Given a twice continuously differentiable function ϕ , define

$$\phi_i = \frac{\partial \phi}{\partial x_i} \quad \text{and} \quad \phi_{ii} = \frac{\partial^2 \phi}{\partial x_i^2}.$$

Also, define the operator A_u^n for $u \in \mathcal{U}$ as follows:

$$(31) \quad A_u^n \phi = \sum_{i \in \mathcal{I}} b_i^n(\cdot, u) \phi_i + \frac{1}{2} \sum_{i \in \mathcal{I}} (\sigma_i^n(\cdot, u))^2 \phi_{ii}.$$

Defining

$$L(x, u) = \sum_{i \in \mathcal{I}} c_i u_i (e \cdot x)^+$$

for $x \in \mathbb{R}_+^I$ and $u \in \mathcal{U}$, the HJB equation is given by

$$(32) \quad 0 = \inf_{u \in \mathcal{U}} \{L(x, u) + A_u^n \phi(x) - \gamma \phi(x)\}.$$

Substituting $b^n(\cdot, \cdot)$ and $\sigma^n(\cdot, \cdot)$ into (32) gives

$$(33) \quad \begin{aligned} 0 = & -\gamma\phi(x) + (e \cdot x)^+ \cdot \min_{i \in \mathcal{I}} \left\{ c_i + \mu_i \phi_i(x) - \frac{1}{2} \mu_i \phi_{ii}(x) \right\} \\ & + \sum_{i \in \mathcal{I}} (l_i^n - \mu_i x_i) \phi_i(x) + \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i (v_i n + x_i)) \phi_{ii}(x). \end{aligned}$$

Our analysis of the HJB equation (33) draws on existing results on fully non-linear PDEs, and, in particular, the results on Bellman–Pucci type equations; cf. Chapter 17 of [6].

In what follows, fixing a set $\mathcal{B} \subseteq \mathbb{R}_+^I$, $\mathcal{C}^2(\mathcal{B})$ denotes the space of twice continuously differentiable functions from \mathcal{B} to \mathbb{R} . For $u \in \mathcal{C}^2(\mathcal{B})$, we let Du and D^2u denote the gradient and the Hessian of u , respectively. The space $\mathcal{C}^{2,\alpha}(\mathcal{B})$ is then the subspace of $\mathcal{C}^2(\mathcal{B})$ members of which also have second derivatives that are Hölder continuous of order α . That is, a twice continuously differentiable function $u : \mathbb{R}^I \rightarrow \mathbb{R}$ is in $\mathcal{C}^{2,\alpha}(\mathcal{B})$ if

$$\sup_{x,y \in \mathcal{B}, x \neq y} \frac{|D^2u(x) - D^2u(y)|}{|x - y|^\alpha} < \infty,$$

where $|\cdot|$ denotes the Euclidian norm. We define $d_x = \text{dist}(x, \partial\mathcal{B}) = \inf\{|x - y|, y \in \partial\mathcal{B}\}$ where $\partial\mathcal{B}$ stands for the boundary of \mathcal{B} and we let $d_{x,z} = \min\{d_x, d_z\}$. Also, we define

$$(34) \quad |u|_{2,\alpha,\mathcal{B}}^* = \sum_{j=0}^2 [u]_{j,\mathcal{B}}^* + \sup_{x,y \in \mathcal{B}, x \neq y} d_{x,y}^{2+\alpha} \frac{|D^2u(x) - D^2u(y)|}{|x - y|^\alpha},$$

where $[u]_{j,\mathcal{B}}^* = \sup_{x \in \mathcal{B}} d_x^j |D^j u(x)|$ for $j = 0, 1, 2$. Note that d_x^j denote the j th power of d_x and, similarly, $d_{x,y}^{2+\alpha}$ is the $(2 + \alpha)$ th power of $d_{x,y}$. Finally, we let $|u|_{0,\mathcal{B}}^* = [u]_{0,\mathcal{B}}^* = \sup_{x \in \mathcal{B}} |u(x)|$.

In the statement of the following theorem, e_j is the I -dimensional vector with 1 in the j th place and zeros elsewhere. Also, \mathcal{B}_κ^n , m and $n(\kappa)$ are as defined in (26) and (27), respectively.

THEOREM 4.1. *Fix $\kappa > 0$ and $n \geq n(\kappa)$. Then, there exists $0 < \alpha \leq 1$ (that does not depend on n) and a unique classical solution $\phi_\kappa^n \in \mathcal{C}^{0,1}(\bar{\mathcal{B}}_\kappa^n) \cap \mathcal{C}^{2,\alpha}(\mathcal{B}_\kappa^n)$ to the HJB equation (33) on \mathcal{B}_κ^n with the boundary condition $\phi_\kappa^n = 0$ on $\partial\mathcal{B}_\kappa^n$. Furthermore, there exists a constant $C > 0$ (that does not depend on n) such that*

$$(35) \quad |\phi_\kappa^n|_{2,\alpha,\mathcal{B}_\kappa^n}^* \leq C \sqrt{n} \log^{k_0} n,$$

where $k_0 = 4m(1 + 1/\alpha)$. In turn, for any $\vartheta < 1$,

$$(36) \quad \sup_{x \in \mathcal{B}_{\vartheta\kappa}^n} |D\phi_\kappa^n(x)| \leq \frac{C}{1 - \vartheta} \log^{k_1} n \quad \text{and} \quad \sup_{x \in \mathcal{B}_{\vartheta\kappa}^n} |D^2\phi_\kappa^n(x)| \leq \frac{C}{1 - \vartheta} \frac{\log^{k_2} n}{\sqrt{n}}$$

with $k_1 = k_0 - m$ and $k_2 = k_0 - 2m$. Also,

$$(37) \quad \begin{aligned} & \sup_{u \in \mathcal{U}} \left| \sum_{i \in \mathcal{I}} ((\phi_\kappa^n)_{ii}(y) - (\phi_\kappa^n)_{ii}(x)) (\sigma_i^n(x, u))^2 \right| \\ & \leq \frac{C}{1 - \vartheta} \log^{k_1} n \end{aligned}$$

for all $x, y \in \mathcal{B}_{\vartheta\kappa}^n$ with $|x - y| \leq 1$.

Note that (36) follows immediately from (35) through the definition of the operation $|\cdot|_{2,\alpha,\mathcal{B}_\kappa^n}^*$ in (34). Henceforth, we will use $k_i, i = 0, 1, 2$ for the values given in the statement of Theorem 4.1. Moreover, the constant k appearing in the statement of Theorem 2.1 is equal to $k_0 + 3$.

Theorem 4.1 facilitates a verification result, which we state next followed by the proof of Theorem 4.1. Below, $\hat{V}^n(x, \kappa)$ is the value function of the ADCP; cf. equation (28).

THEOREM 4.2. *Fix $\kappa > 0$ and $n \geq n(\kappa)$. Let ϕ_κ^n be the unique solution to the HJB equation (33) on \mathcal{B}_κ^n with the boundary condition $\phi_\kappa^n = 0$ on $\partial\mathcal{B}_\kappa^n$. Then, $\phi_\kappa^n(x) = \hat{V}^n(x, \kappa)$ for all $x \in \mathcal{B}_\kappa^n$. Moreover, there exists a Markov control which is optimal for the ADCP on \mathcal{B}_κ^n . The tracking function $h_\kappa^{*,n}$ associated with this optimal Markov control is defined by $h_\kappa^{*,n}(x) = e_{i^n(x)}$, where*

$$(38) \quad i^n(x) = \min \arg \min_{i \in \mathcal{I}} \left\{ c_i + \mu_i(\phi_\kappa^n)_i(x) - \frac{1}{2} \mu_i(\phi_\kappa^n)_{ii}(x) \right\} (e \cdot x)^+.$$

The HJB equation (33) has two sources of nondifferentiability. The first source is the minimum operation and the second is the nondifferentiability of the term $(e \cdot x)^+$. The first source of nondifferentiability is covered almost entirely by the results in [6]. To deal with the nondifferentiability of the function $(e \cdot x)^+$, we use a construction by approximations. The proof of existence and uniqueness in Theorem 4.1 follows an approximation scheme where one replaces the nonsmooth function $(e \cdot x)^+$ by a smooth (parameterized by a) function $f_a(e \cdot x)$. We show that the resulting ‘‘perturbed’’ PDE has a unique classical solution and that as $a \rightarrow \infty$ the corresponding sequence of solutions converges, in an appropriate sense, to a solution to (33) which will be shown to be unique. Note that this argument is repeated for each fixed n and κ .

To that end, given $a > 0$, define

$$(39) \quad f_a(y) = \begin{cases} y, & y \geq \frac{1}{4a}, \\ ay^2 + \frac{1}{2}y + \frac{1}{16a}, & -\frac{1}{4a} \leq y \leq \frac{1}{4a}, \\ 0, & \text{otherwise.} \end{cases}$$

Replacing $(e \cdot x)^+$ with $f_a(e \cdot x)$ in (33) gives the following equation:

$$(40) \quad 0 = -\gamma\phi(x) + f_a(e \cdot x) \cdot \min_{i \in \mathcal{I}} \left\{ c_i + \mu_i\phi_i(x) - \frac{1}{2}\mu_i\phi_{ii}(x) \right\} + \sum_{i \in \mathcal{I}} (l_i^n - \mu_i x_i)\phi_i(x) + \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + x_i))\phi_{ii}(x).$$

To simplify this further, let $\Gamma = \mathcal{B}_\kappa^n \times \mathbb{R}_+ \times \mathbb{R}^I \times \mathbb{R}^{I \times I}$ and for all $y \in \Gamma$, define the function

$$(41) \quad F_a^k[y] = \min\{F_a^1[y], \dots, F_a^I[y]\},$$

where for $k \in \mathcal{I}$ and $y = (x, z, p, r) \in \Gamma$,

$$(42) \quad F_a^k[y] = f_a(e \cdot x) \left[c_k + \mu_k p_k - \frac{1}{2}\mu_k r_{kk} \right] + \sum_{i \in \mathcal{I}} (l_i^n - \mu_i x_i) p_i + \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + x_i)) r_{ii} - \gamma z.$$

Then, (40) can be rewritten as

$$(43) \quad F_a[x, u(x), Du(x), D^2u(x)] = 0.$$

In the following statement we use the gradient notation introduced at the beginning of this section.

PROPOSITION 4.1. *Fix $\kappa > 0$, $n \geq n(\kappa)$ and $a > 0$. A unique classical solution $\phi_{\kappa,a}^n \in C^{0,1}(\bar{\mathcal{B}}_\kappa^n) \cap C^{2,\alpha}(\mathcal{B}_\kappa^n)$ exists for the PDE (40) on \mathcal{B}_κ^n with the boundary condition $\phi_{\kappa,a}^n = 0$ on $\partial\mathcal{B}_\kappa^n$. Moreover,*

$$(44) \quad |\phi_{\kappa,a}^n|_{2,\alpha,\mathcal{B}_\kappa^n}^* \leq C |\phi_{\kappa,a}^n|_{0,\mathcal{B}_\kappa^n}^* \log^{k_0} n \leq \tilde{C}$$

for $k_0 = 4m(1 + 1/\alpha)$ where $0 < \alpha \leq 1$ and $C > 0$ do not depend on a and n and \tilde{C} does not depend on a . Also, $\phi_{\kappa,a}^n$ is Lipschitz continuous on the closure $\bar{\mathcal{B}}_\kappa^n$ with a Lipschitz constant that does not depend on a (but can depend on κ and n).

We postpone the proof of Proposition 4.1 to the [Appendix](#) and use it to complete the proof of Theorem 4.1, followed by the proof of Theorem 4.2.

Proof of Theorem 4.1. Since we fix n and κ , they will be suppressed below. We proceed to show the existence by an approximation argument. To that end, fix a sequence $\{a^k; k \in \mathbb{Z}\}$ with $a^k \rightarrow \infty$ as $k \rightarrow \infty$ and let ϕ_{a^k} be the unique solution to (40) as given by Proposition 4.1. The next step is to show that ϕ_{a^k} has a subsequence that converges in an appropriate sense to a function ϕ , which is, in fact, a solution to the HJB equation (33). To that end, let

$$(45) \quad \mathcal{C}_*^{2,\alpha}(\mathcal{B}) = \{u \in C^{2,\alpha}(\mathcal{B}) : |u|_{2,\alpha,\mathcal{B}}^* < \infty\}.$$

Then, $\mathcal{C}_*^{2,\alpha}(\mathcal{B})$ is a Banach space (see, e.g., Exercise 5.2 in [6]). Since the bound in (44) is independent of a , we have that $\{\phi_{a^k}\}$ is a bounded sequence in $\mathcal{C}_*^{2,\alpha}(\mathcal{B})$ and hence, contains a convergent subsequence. Let u be a limit point of the sequence $\{\phi_{a^k}\}$. Since the gradient estimates in Proposition 4.1 are independent of a , they hold also for the limit function u , that is,

$$(46) \quad |u|_{2,\alpha,\mathcal{B}}^* \leq C|u|_{0,\mathcal{B}}^* \log^{k_0} n \leq \tilde{C}$$

for constants α and C that are independent of n . Proposition 4.1 also guarantees that the global Lipschitz constant is independent of a so that we may conclude that $u \in \mathcal{C}^{0,1}(\bar{\mathcal{B}})$ and that $u = 0$ on $\partial\mathcal{B}$.

We will now show that u solves (33) uniquely. To show that u solves (33), we need to show that $F[u] = 0$ (where $F[\cdot]$ is defined similar to $F_a[\cdot]$ with $(e \cdot x)^+$ replacing $f_a(e \cdot x)$). To that end, let $\{a^k, k \in \mathbb{Z}\}$ be the corresponding convergent subsequence [i.e., such that $\phi_{a^k} \rightarrow u$ in $\mathcal{C}_*^{2,\alpha}(\mathcal{B})$]. Henceforth, to simplify notation, we write

$$F_{a^{k_l}}[\phi_{a^{k_l}}(x)] = F_{a^{k_l}}[x, \phi_{a^{k_l}}(x), D\phi_{a^{k_l}}(x), D^2\phi_{a^{k_l}}(x)]$$

(and similarly for $F[\cdot]$). Fix $\delta > 0$ and let $\mathcal{B}(\delta) = \{x \in \mathbb{R}^I : |x| < \kappa \sqrt{n} \log^m n - \delta\}$. Note that since $\phi_{a^{k_l}} \rightarrow u$ in $\mathcal{C}_*^{2,\alpha}(\mathcal{B})$ we have, in particular, the convergence of $(\phi_{a^{k_l}}(x), D\phi_{a^{k_l}}(x), D^2\phi_{a^{k_l}}(x)) \rightarrow (u(x), Du(x), D^2u(x))$ uniformly in $x \in \mathcal{B}(\delta)$. The equicontinuity of the function $F^a[\cdot]$ on Γ guarantees then that

$$(47) \quad |F_{a^{k_l}}[\phi_{a^{k_l}}(x)] - F_{a^{k_l}}[u(x)]| \leq \epsilon$$

for all l large enough and $x \in \mathcal{B}(\delta)$. Note that $\sup_{x \in \mathbb{R}^I} |f_a(e \cdot x) - (e \cdot x)^+| \leq \epsilon$ for all a large enough so that,

$$(48) \quad \sup_{x \in \mathcal{B}} |F_{a^{k_l}}[u(x)] - F[u(x)]| \leq \epsilon$$

for all l large enough. Combining (47) and (48), we then have

$$\sup_{x \in \mathcal{B}} |F_{a^{k_l}}[\phi_{a^{k_l}}(x)] - F[u(x)]| \leq 2\epsilon$$

for all l large enough and $x \in \mathcal{B}(\delta)$. By definition $F^{a^k}[\phi_{a^{k_l}}(x)] = 0$ for all $x \in \mathcal{B}$ and since ϵ was arbitrary we have that $F[u(x)] = 0$ for all $x \in \mathcal{B}(\delta)$. Finally, since δ was arbitrary we have that $F[u(x)] = 0$ for all $x \in \mathcal{B}$. We already argued that $u = 0$ on $\partial\mathcal{B}$, so that u solves (33) on \mathcal{B} with $u = 0$ on $\partial\mathcal{B}$. This concludes the proof of existence of a solution to (33) that satisfies the gradient estimates (35).

Finally, the uniqueness of the solution to (33) follows from Corollary 17.2 in [6] noting that the function $F[x, z, p, r]$ is indeed continuously differentiable in the (z, p, r) arguments and it is decreasing in z for all (x, p, r) .

Using Theorem 4.2 [which only uses the existence and uniqueness of the solution $\phi_\kappa^n(x)$ that we already established] together with (29) we have that

$$|\phi_\kappa^n|_{0,\mathcal{B}_\kappa^n} = \sup_{x \in \mathcal{B}_\kappa^n} \hat{V}^n(x, \kappa) \leq \frac{1}{\gamma} \kappa \sqrt{n} \log^m n.$$

The bounds (35) and (36) now follow from (46) and we turn to prove (37).

To that end, since ϕ_κ^n solves (33), fixing $x, y \in \mathcal{B}_\kappa^n$ we have

$$(49) \quad \left| \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + x_i)) (\phi_\kappa^n)_{ii}(x) - \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + y_i)) (\phi_\kappa^n)_{ii}(y) \right| \\ \leq \gamma |\phi_\kappa^n(x) - \phi_\kappa^n(y)| \\ + \left| (e \cdot x)^+ \cdot \min_{i \in \mathcal{I}} \left\{ c_i + \mu_i(\phi_\kappa^n)_i(x) - \frac{1}{2} \mu_i(\phi_\kappa^n)_{ii}(x) \right\} \right. \\ \left. - (e \cdot y)^+ \cdot \min_{i \in \mathcal{I}} \left\{ c_i + \mu_i(\phi_\kappa^n)_i(y) - \frac{1}{2} \mu_i(\phi_\kappa^n)_{ii}(y) \right\} \right|.$$

We will now bound each of the elements on the right-hand side. To that end, let $i(x)$ be as defined in (38) and for each $x, z \in \mathcal{B}_{\vartheta\kappa}^n$ define

$$M_{i(x)}^n(z) = c_{i(x)} + \mu_{i(x)}(\phi_\kappa^n)_{i(x)}(z) - \frac{1}{2} \mu_{i(x)}(\phi_\kappa^n)_{i(x)i(x)}(z).$$

Using (36), we have by the mean value theorem that

$$(50) \quad |\phi_\kappa^n(x) - \phi_\kappa^n(y)| \leq |x - y| \max_{i \in \mathcal{I}} \sup_{z \in \mathcal{B}_{\vartheta\kappa}^n} |(\phi_\kappa^n)_i(z)| \leq C \log^{k_1} n$$

for all $x, y \in \mathcal{B}_{\vartheta\kappa}^n$ with $|x - y| \leq 1$, and we turn to bound the second element on the right-hand side of (49). Here, there are two cases to consider. Suppose first that $i(x) = i(y) = i$. Then, using (36) and the mean value theorem we have

$$|(\phi_\kappa^n)_i(x) - (\phi_\kappa^n)_i(y)| \leq |x - y| \max_{i \in \mathcal{I}} \sup_{z \in \mathcal{B}_{\vartheta\kappa}^n} |(\phi_\kappa^n)_{ii}(z)| \leq C \frac{\log^{k_2} n}{\sqrt{n}}$$

and, in turn, that

$$(51) \quad |M_i^n(x) - M_i^n(y)| \leq C \frac{\log^{k_2} n}{\sqrt{n}}$$

for all $x, y \in \mathcal{B}_{\vartheta\kappa}^n$ with $|x - y| \leq 1$. Now, $|x| \vee |y| \leq \kappa \sqrt{n} \log^m n$ for all $x, y \in \mathcal{B}_{\vartheta\kappa}^n$ and, by (36), $\sup_{z \in \mathcal{B}_{\vartheta\kappa}^n} |(\phi_\kappa^n)_{ii}(z)| \vee |(\phi_\kappa^n)_i(z)| \leq C \log^{k_1} n$ so that

$$(52) \quad |(e \cdot x)^+ M_i^n(x) - (e \cdot y)^+ M_i^n(y)| \\ \leq \kappa \sqrt{n} \log^m n |M_i^n(x) - M_i^n(y)| + \sup_{z \in \mathcal{B}_{\vartheta\kappa}^n} |M_i^n(z)| \\ \leq C \log^{k_1} n.$$

If, on the other hand, $i(x) \neq i(y)$ then by the definition of $i(\cdot)$,

$$c_{i(x)} + \mu_{i(x)}(\phi_\kappa^n)_{i(x)}(x) - \frac{1}{2} \mu_{i(x)}(\phi_\kappa^n)_{i(x)i(x)}(x) \\ \leq c_{i(y)} + \mu_{i(y)}(\phi_\kappa^n)_{i(y)}(x) - \frac{1}{2} \mu_{i(y)}(\phi_\kappa^n)_{i(y)i(y)}(x)$$

and

$$\begin{aligned} c_{i(y)} + \mu_{i(y)}(\phi_\kappa^n)_{i(y)}(y) - \frac{1}{2}\mu_{i(y)}(\phi_\kappa^n)_{i(y)i(y)}(y) \\ \leq c_{i(x)} + \mu_{i(x)}(\phi_\kappa^n)_{i(x)}(y) - \frac{1}{2}\mu_{i(x)}(\phi_\kappa^n)_{i(x)i(x)}(y). \end{aligned}$$

That is,

$$(53) \quad M_{i(x)}^n(x) \leq M_{i(y)}^n(x) \quad \text{and} \quad M_{i(y)}^n(y) \leq M_{i(x)}^n(y).$$

Using (36) as before we have for $x, y \in \mathcal{B}_{\vartheta\kappa}^n$ with $|x - y| \leq 1$ and $i(x) \neq i(y)$ that

$$|M_{i(x)}^n(x) - M_{i(x)}^n(y)| + |M_{i(y)}^n(x) - M_{i(y)}^n(y)| \leq C \frac{\log^{k_2} n}{\sqrt{n}}.$$

By (53) we then have that

$$\begin{aligned} |M_{i(x)}^n(x) - M_{i(y)}^n(y)| &\leq |M_{i(x)}^n(x) - M_{i(x)}^n(y)| \\ &\quad + |M_{i(y)}^n(x) - M_{i(y)}^n(y)| \\ &\leq C \frac{\log^{k_2} n}{\sqrt{n}} \end{aligned}$$

for all such x and y . In turn, since $|x| \vee |y| \leq \kappa\sqrt{n} \log^m n$,

$$(54) \quad |(e \cdot x)^+ M_{i(x)}^n(x) - (e \cdot y)^+ M_{i(y)}^n(y)| \leq C \log^{k_1} n$$

for $x, y \in \mathcal{B}_{\vartheta\kappa}^n$ with $|x - y| \leq 1$ and $i(x) \neq i(y)$. Plugging (50), (52) and (54) into the right-hand side of (49) we get

$$\begin{aligned} (55) \quad &\left| \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + x_i)) (\phi_\kappa^n)_{ii}(x) - \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + y_i)) (\phi_\kappa^n)_{ii}(y) \right| \\ &\leq C \log^{k_1} n \end{aligned}$$

for all $x, y \in \mathcal{B}_{\vartheta\kappa}^n$ with $|x - y| \leq 1$. Finally, recall that

$$\sigma_i^n(x, u) = \sqrt{\lambda_i^n + \mu_i v_i n + \mu_i(x_i - u_i(e \cdot x)^+)}$$

so that for all $u \in \mathcal{U}$,

$$\begin{aligned} &\left| \sum_{i \in \mathcal{I}} ((\phi_\kappa^n)_{ii}(y) - (\phi_\kappa^n)_{ii}(x)) (\sigma_i^n(x, u))^2 \right| \\ &= \left| \sum_{i \in \mathcal{I}} (\phi_\kappa^n)_{ii}(y) (\lambda_i^n + \mu_i v_i n + \mu_i(x_i - u_i(e \cdot x)^+)) \right. \\ &\quad \left. - (\phi_\kappa^n)_{ii}(x) (\lambda_i^n + \mu_i v_i n + \mu_i(x_i - u_i(e \cdot x)^+)) \right| \end{aligned}$$

$$\begin{aligned}
&\leq \left| \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + x_i)) (\phi_\kappa^n)_{ii}(x) \right. \\
&\quad \left. - \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + y_i)) (\phi_\kappa^n)_{ii}(y) \right| \\
&\quad + \left| \frac{1}{2} \sum_{i \in \mathcal{I}} (\phi_\kappa^n)_{ii}(x) \mu_i u_i(e \cdot x)^+ - (\phi_\kappa^n)_{ii}(y) \mu_i u_i(e \cdot y)^+ \right| \\
&\quad + \left| \frac{1}{2} \sum_{i \in \mathcal{I}} (\phi_\kappa^n)_{ii}(y) \mu_i (x_i - y_i) \right|.
\end{aligned}$$

The last two terms above are bounded by $C \log^{k_1} n$ by (36) and using $|x| \vee |y| \leq \kappa \sqrt{n} \log^m n$. Together with (55) this establishes (37) and concludes the proof of the theorem.

Proof of Theorem 4.2. Fix an initial condition $x \in \mathcal{B}_\kappa^n$ and an admissible (κ, n) -system $\theta = (\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P}, \hat{U}, B)$ and let \hat{X}^n be the associated controlled process. Using Itô's lemma for the function $\varphi(t, x) = e^{-\gamma t} \phi_\kappa^n(x)$ in conjunction with the inequality

$$L(x, u) + A_u \phi_\kappa^n(x) - \gamma \phi_\kappa^n(x) \geq 0 \quad \text{for all } x \in \mathcal{B}_\kappa^n, u \in \mathcal{U}$$

[recall that ϕ_κ^n solves (33)] we have that

$$\begin{aligned}
(56) \quad \phi_\kappa^n(x) &\leq \mathbb{E}_x^\theta \int_0^{t \wedge \hat{\tau}_\kappa^n} e^{-\gamma s} L(\hat{X}^n(s), \hat{U}(s)) ds + \mathbb{E}_x^\theta e^{-\gamma(t \wedge \hat{\tau}_\kappa^n)} \phi_\kappa^n(\hat{X}^n(t \wedge \hat{\tau}_\kappa^n)) \\
&\quad - \mathbb{E}_x^\theta \sum_{i \in \mathcal{I}} \int_0^{t \wedge \hat{\tau}_\kappa^n} e^{-\gamma s} (\phi_\kappa^n)_i(\hat{X}^n(s)) \sigma_i^n(\hat{X}^n(s), \hat{U}(s)) dB(s).
\end{aligned}$$

Here, $\hat{\tau}_\kappa^n$ is as defined in Definition 3.1 and it is a stopping time with respect to (\mathcal{F}_t) because of the continuity of \hat{X}^n . We now claim that

$$\mathbb{E}_x^\theta [e^{-\gamma t \wedge \hat{\tau}_\kappa^n} \phi_\kappa^n(\hat{X}^n(t \wedge \hat{\tau}_\kappa^n))] \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

Indeed, as ϕ_κ^n is bounded on \mathcal{B}_κ^n , on the event $\{\hat{\tau}_\kappa^n = \infty\}$ we have that

$$e^{-\gamma(t \wedge \hat{\tau}_\kappa^n)} \phi_\kappa^n(\hat{X}^n(t \wedge \hat{\tau}_\kappa^n)) \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

On the event $\{\hat{\tau}_\kappa^n < \infty\}$ we have $\hat{X}^n(\hat{\tau}_\kappa^n) \in \partial \mathcal{B}$ and, by the definition of $\hat{\tau}_\kappa^n$, that $\phi_\kappa^n(\hat{X}^n(\hat{\tau}_\kappa^n)) = 0$. The convergence in expectation then follows from the bounded convergence theorem (using again the boundedness of ϕ_κ^n on \mathcal{B}_κ^n). The last term in (56) equals zero by the optional stopping theorem.

Letting $t \rightarrow \infty$ in (56) and applying the monotone convergence theorem, we then have

$$\phi_\kappa^n(x) \leq \mathbb{E}_x^\theta \left[\int_0^{\hat{\tau}_\kappa^n} e^{-\gamma s} L(\hat{X}^n(s), \hat{U}(s)) ds \right].$$

Since the admissible system θ was arbitrary, we have that $\phi_\kappa^n(x) \leq \hat{V}^n(x, \kappa)$. To show that this inequality is actually an equality, let

$$(57) \quad h_\kappa^n(x) = e_{i^n(x)},$$

where $e_{i^n(x)}$ is as defined in the statement of the theorem.

The continuity of ϕ_κ^n guarantees that the function $i^n(x)$ is Lebesgue measurable, and so is, in turn, $h_\kappa^n(\cdot)$. Consider now the autonomous SDE:

$$(58) \quad \hat{X}^n(t) = x + \int_0^t \hat{b}^n(\hat{X}^n(s)) ds + \int_0^t \hat{\sigma}^n(\hat{X}^n(s)) dB(s),$$

where $\hat{b}^n(y) = b^n(y, h_\kappa^n(y))$ and $\hat{\sigma}^n(y) = \sigma^n(y, h_\kappa^n(y))$ on \mathcal{B}_κ^n . Then, \hat{b}^n and $\hat{\sigma}^n$ are bounded and measurable on the bounded domain \mathcal{B}_κ^n . Also, as the matrix $\hat{\sigma}^n$ is diagonal and the elements on the diagonal are strictly positive on \mathcal{B}_κ^n , it is positive definite there. Hence, a weak solution exists for the autonomous SDE (see, e.g., Theorem 6.1 of [10]). In particular, there exists a probability space $(\tilde{\Omega}, \mathcal{G}, \tilde{\mathbb{P}})$, a filtration (\mathcal{G}_t) that satisfies the usual conditions, a Brownian motion $B(t)$ and a continuous process \hat{X}^n —both adapted to (\mathcal{G}_t) , so that \hat{X}^n satisfies the autonomous SDE (58). Finally, since \hat{X}^n has continuous sample paths and it is adapted, it is also progressively measurable (see, e.g., Proposition 1.13 in [9]) and, by measurability of $h_\kappa^n(\cdot)$, so is the process $\hat{U}(t) = h_\kappa^n(\hat{X}^n(t))$. Consequently, $\theta = (\tilde{\Omega}, \mathcal{G}, \mathcal{G}_t, \tilde{\mathbb{P}}, \hat{U}, B)$ is an admissible system in the sense of Definition 3.1 and \hat{X}^n is the corresponding controlled process.

To see that θ is optimal for the ADCP on \mathcal{B}_κ^n , note that for $s < \hat{\tau}_\kappa^n$, we have by the HJB equation (32) that

$$L(\hat{X}^n(s), \hat{U}(s)) + A_{\hat{U}(s)} \phi_\kappa^n(\hat{X}^n(s)) - \gamma \phi_\kappa^n(\hat{X}^n(s)) = 0.$$

Applying Itô's rule as before, together with the bounded and dominated convergence theorems, we then have that

$$\phi_\kappa^n(x) = \mathbb{E}_x^\theta \left[\int_0^{\hat{\tau}_\kappa^n} e^{-\gamma s} L(\hat{X}^n(s), \hat{U}(s)) ds \right]$$

and the proof is complete.

5. The performance analysis of tracking policies. This section shows that given an optimal Markov control policy for the ADCP together with its associated tracking function $h_\kappa^{*,n}$, the nonpreemptive tracking policy imitates, in a particular sense, the performance of the Brownian system.

THEOREM 5.1. Fix κ and $\kappa' < \kappa$ as well as a sequence $\{(x^n, q^n), n \in \mathbb{Z}_+\}$ such that $(x^n, q^n) \in \mathcal{X}^n$, and $|x^n - \nu n| \leq M\sqrt{n}$ for all n and some $M > 0$. Let ϕ_κ^n and $h_{\kappa}^{*,n}$ be as in Theorem 4.2 and define

$$\psi^n(x, u) = L(x, u) + A_u^n \phi_\kappa^n(x) - \gamma \phi_\kappa^n(x) \quad \text{for } x \in \mathcal{B}_\kappa^n, u \in \mathcal{U}.$$

Let U_h^n be the ratio control associated with the $h_{\kappa}^{*,n}$ -tracking policy and let $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ be the associated queueing process with the initial conditions $Q^n(0) = q^n$ and $\check{X}^n(0) = x^n - \nu n$ and define

$$\tau_{\kappa', T}^n = \inf\{t \geq 0: \check{X}^n(t) \notin \mathcal{B}_{\kappa'}^n\} \wedge T \log n.$$

Then,

$$\mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} |\psi^n(\check{X}^n(s), U_h^n(s)) - \psi^n(\check{X}^n(s), h_{\kappa}^{*,n}(\check{X}^n(s)))| ds \right] \leq C \log^{k_0+3} n$$

for a constant C that does not depend on n .

Theorem 5.1 is proved in the Appendix. The proof builds on the gradient estimates in Theorem 4.1 and on a state-space collapse-type result for certain sub-intervals of $[0, \tau_{\kappa', T}^n]$.

REMARK 5.1. Typically one establishes a stronger state-space collapse result showing that the actual queue and the desired queue values are close in supremum norm. The difficulty with the former approach is that the tracking functions here are nonsmooth. While it is plausible that one can smooth these functions appropriately (as is done, e.g., in [2]), such smoothing might compromise the optimality gap. Fortunately, the weaker integral criterion implied by Theorem 5.1 suffices for our purposes.

6. Proof of the main result. Fix $\kappa > 0$ and let ϕ_κ^n be the solution to (33) on \mathcal{B}_κ^n (see Theorem 4.1). We start with the following lemma where $b_i^n(\cdot, \cdot)$ and $\sigma_i^n(\cdot, \cdot)$ are as in (24) and (25), respectively.

LEMMA 6.1. Let U^n be an admissible ratio control and let $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ be the queueing process associated with U^n . Fix $\kappa' < \kappa$ and $T > 0$ and let

$$\tau_{\kappa', T}^n = \inf\{t \geq 0: \check{X}^n(t) \notin \mathcal{B}_{\kappa'}^n\} \wedge T \log n.$$

Then, there exists a constant C that does not depend on n (but may depend on T , κ and κ') such that

$$\begin{aligned} \mathbb{E}[e^{-\gamma \tau_{\kappa', T}^n} \phi_\kappa^n(\check{X}^n(\tau_{\kappa', T}^n))] &\leq \phi_\kappa^n(\check{X}^n(0)) + \mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} A_{U^n(s)}^n \phi_\kappa^n(\check{X}^n(s)) ds \right] \\ &\quad - \gamma \mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} \phi_\kappa^n(\check{X}^n(s)) ds \right] + C \log^{k_1+1} n \\ &\leq \mathbb{E}[e^{-\gamma \tau_{\kappa', T}^n} \phi_\kappa^n(\check{X}^n(\tau_{\kappa', T}^n))] + 2C \log^{k_1+1} n. \end{aligned}$$

We will also use the following lemma where $c = (c_1, \dots, c_l)$ are the cost coefficients (see Section 2).

LEMMA 6.2. *Let (x^n, q^n) be as in the conditions of Theorem 2.1. Then, there exists a constant C that does not depend on n such that*

$$(59) \quad \mathbb{E}_{x^n, q^n}^U \left[\int_{\tau_{\kappa', T}^n}^{\infty} e^{-\gamma s} (e \cdot c) (e \cdot \check{X}^n(s))^+ ds \right] \leq C \log^2 n$$

and

$$(60) \quad \mathbb{E}_{x^n, q^n}^U [e^{-\gamma \tau_{\kappa', T}^n} \phi_{\kappa}^n(\check{X}^n(\tau_{\kappa', T}^n))] \leq C \log^2 n$$

for all n and any admissible ratio control U .

We postpone the proof of Lemma 6.1 to the end of the section and that of Lemma 6.2 to the Appendix and proceed now to prove the main result of the paper.

Proof of Theorem 2.1. Let $h_{\kappa}^{*,n}$ be the ratio function associated with the optimal Markov control for the ADCP (as in Theorem 4.1). Since κ is fixed we omit the subscript κ and use $h^n = h_{\kappa}^{*,n}$. Let U_h^n be the ratio associated with the h^n -tracking policy.

The proof will proceed in three main steps. First, building on Theorem 5.1 we will show that

$$(61) \quad \mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} L(\check{X}^n(s), U_h^n(s)) ds \right] \leq \phi_{\kappa}^n(\check{X}^n(0)) + C \log^{k_0+3} n.$$

Using Lemma 6.2, this implies

$$(62) \quad \begin{aligned} C^n(x^n, q^n, U_h^n) &= \mathbb{E} \left[\int_0^{\infty} e^{-\gamma s} L(\check{X}^n(s), U_h^n(s)) ds \right] \\ &\leq \phi_{\kappa}^n(\check{X}^n(0)) + C \log^{k_0+3} n. \end{aligned}$$

Finally, we will show that for any ratio control U^n ,

$$(63) \quad \phi_{\kappa}^n(\check{X}^n(0)) \leq \mathbb{E} \left[\int_0^{\infty} e^{-\gamma s} L(\check{X}^n(s), U^n(s)) ds \right] + C \log^{k_1+1} n,$$

where we recall that $k_1 = k_0 - m$. In turn,

$$V^n(x^n, q^n) \geq \phi_{\kappa}^n(x^n - vn) - C \log^{k_1+1} n \geq C^n(x^n, q^n, U_h^n) - 2C \log^{k_1+1} n,$$

which establishes the statement of the theorem.

We now turn to prove each of (61) and (63).

Proof of (61). To simplify notation we fix $\kappa > 0$ throughout and let $h^n(\cdot) = h_\kappa^{*,n}$. Using Lemma 6.1 we have

$$(64) \quad \begin{aligned} & \mathbb{E}[e^{-\gamma\tau_{\kappa',T}^n} \phi_\kappa^n(\check{X}^n(\tau_{\kappa',T}^n))] \\ & \leq \phi_\kappa^n(\check{X}^n(0)) + \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} A_{U_h^n(s)}^n \phi_\kappa^n(\check{X}^n(s)) ds\right] \\ & \quad - \gamma \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} \phi_\kappa^n(\check{X}^n(s)) ds\right] + C \log^{k_1+1} n. \end{aligned}$$

From the definition of h^n as a minimizer in the HJB equation we have that

$$\begin{aligned} 0 &= \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} A_{h^n(\check{X}^n(s))}^n \phi_\kappa^n(\check{X}^n(s)) ds\right] \\ & \quad - \gamma \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} \phi_\kappa^n(\check{X}^n(s)) ds\right] \\ & \quad + \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} L(\check{X}^n(s), h^n(\check{X}^n(s))) ds\right]. \end{aligned}$$

By Theorem 5.1 we then have that

$$(65) \quad \begin{aligned} C \log^{k_0+3} n &\geq \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} A_{U_h^n(s)}^n \phi_\kappa^n(\check{X}^n(s)) ds\right] \\ & \quad - \gamma \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} \phi_\kappa^n(\check{X}^n(s)) ds\right] \\ & \quad + \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} L(\check{X}^n(s), U_h^n(s)) ds\right] \\ & \geq 0. \end{aligned}$$

Since ϕ_κ^n is nonnegative, combining (64) and (65) we have that

$$\mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} L(\check{X}^n(s), U_h^n(s)) ds\right] \leq \phi_\kappa^n(\check{X}^n(0)) + C \log^{k_0+3} n,$$

which concludes the proof of (61).

Proof of (63). We now show that $V^n(x, q) \geq \phi_\kappa^n(\check{X}^n(0)) - C \log^{k_1+1} n$. To that end, fix an arbitrary ratio control U^n and recall that by the HJB equation,

$$A_u^n \phi_\kappa^n(x) - \gamma \phi_\kappa^n(x) + L(x, u) \geq 0$$

for all $u \in \mathcal{U}$ and $x \in \mathcal{B}_\kappa^n$. In turn, using the second inequality in Lemma 6.1 we have that

$$\begin{aligned} & \mathbb{E}[e^{-\gamma\tau_{\kappa',T}^n} \phi_\kappa^n(\check{X}^n(\tau_{\kappa',T}^n))] \\ & \geq \phi_\kappa^n(\check{X}^n(0)) - \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} L(\check{X}^n(s), U^n(s)) ds\right] \\ & \quad - 2C \log^{k_1+1} n. \end{aligned}$$

Using Lemma 6.2, we have, however, that

$$\mathbb{E}[e^{-\gamma\tau_{\kappa',T}^n} \phi_\kappa^n(\check{X}^n(\tau_{\kappa',T}^n))] \leq C \log^2 n$$

for a redefined constant C so that

$$\begin{aligned} C \log^2 n & \geq \phi_\kappa^n(\check{X}^n(0)) - \mathbb{E}\left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} L(\check{X}^n(s), U^n(s)) ds\right] \\ & \quad - 2C \log^{k_1+1} n \\ & \geq \phi_\kappa^n(\check{X}^n(0)) - \mathbb{E}\left[\int_0^\infty e^{-\gamma s} L(\check{X}^n(s), U^n(s)) ds\right] \\ & \quad - 2C \log^{k_1+1} n \end{aligned}$$

and, finally,

$$\phi_\kappa^n(\check{X}^n(0)) \leq \mathbb{E}\left[\int_0^\infty e^{-\gamma s} L(\check{X}^n(s), U^n(s)) ds\right] + C \log^{k_1+1} n$$

for a redefined constant $C > 0$. This concludes the proof of (63) and of the theorem.

We end this section with the proof of Lemma 6.1 in which the following auxiliary lemma will be of use.

LEMMA 6.3. *Fix $\kappa > 0$ and an admissible ratio control U^n and let $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ be the corresponding queueing process. Let*

$$\tau_{\kappa,T}^n = \inf\{t \geq 0 : \check{X}^n(t) \notin \mathcal{B}_\kappa^n\} \wedge T \log n,$$

and $(\check{W}_i^n, i \in \mathcal{I})$ be as defined in (21). Then, for each $i \in \mathcal{I}$, the process $\check{W}_i^n(\cdot \wedge \tau_{\kappa,n}^n)$ is a square integrable martingale w.r.t to the filtration $(\mathcal{F}_{t \wedge \tau_{\kappa,T}^n}^n)$ as are the processes

$$\mathcal{M}_i^n(\cdot) = (\check{W}_i^n(\cdot \wedge \tau_{\kappa,T}^n))^2 - \int_0^{\cdot \wedge \tau_{\kappa,T}^n} (\sigma_i^n(\check{X}^n(s), U^n(s)))^2 ds$$

and

$$\mathcal{V}_i^n(\cdot) = (\check{W}_i^n(\cdot \wedge \tau_{\kappa,T}^n))^2 - \sum_{s \leq \cdot \wedge \tau_{\kappa,T}^n} (\Delta \check{W}_i^n(s))^2.$$

Lemma 6.3 follows from basic results on martingales associated with time-changes of Poisson processes. The detailed proof appears in the [Appendix](#).

Proof of Lemma 6.1. Note that, as in (20), \check{X}^n satisfies

$$\check{X}_i^n(t) = \check{X}_i^n(0) + \int_0^t b_i^n(\check{X}^n(s), U^n(s)) ds + \check{W}_i^n(t),$$

and is a semi martingale. Applying Itô’s formula for semimartingales (see, e.g., Theorem 5.92 in [14]) we have for all $t \leq \tau_{\kappa', T}^n$, that

$$\begin{aligned} e^{-\gamma t} \phi_\kappa^n(\check{X}^n(t)) &= \phi_\kappa^n(\check{X}^n(0)) \\ &+ \sum_{s \leq t: |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} [\phi_\kappa^n(\check{X}^n(s)) - \phi_\kappa^n(\check{X}^n(s-))] \\ &- \sum_{i \in \mathcal{I}} \sum_{s \leq t: |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} (\phi_\kappa^n)_i(\check{X}^n(s)) \Delta \check{X}_i^n(s) \\ &+ \sum_{i \in \mathcal{I}} \int_0^t e^{-\gamma s} (\phi_\kappa^n)_i(\check{X}^n(s-)) b_i^n(\check{X}^n(s), U^n(s)) ds \\ &- \gamma \int_0^t e^{-\gamma s} \phi_\kappa^n(\check{X}^n(s)) ds \end{aligned}$$

and, after rearranging terms, that

$$\begin{aligned} e^{-\gamma t} \phi_\kappa^n(\check{X}^n(t)) &= \phi_\kappa^n(\check{X}^n(0)) + \frac{1}{2} \sum_{i \in \mathcal{I}} \sum_{s \leq t: |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} (\phi_\kappa^n)_{ii}(\check{X}^n(s-)) (\Delta \check{X}_i^n(s))^2 \\ &+ \sum_{i \in \mathcal{I}} \int_0^t e^{-\gamma s} (\phi_\kappa^n)_i(\check{X}^n(s-)) b_i(\check{X}^n(s), U^n(s)) ds \\ &+ C^n(t) - \gamma \int_0^t e^{-\gamma s} \phi_\kappa^n(\check{X}^n(s)) ds, \end{aligned}$$

where

$$\begin{aligned} C^n(t) &= \sum_{s \leq t: |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} \left[\phi_\kappa^n(\check{X}^n(s)) - \phi_\kappa^n(\check{X}^n(s-)) \right. \\ &- \sum_{i \in \mathcal{I}} (\phi_\kappa^n)_i(\check{X}^n(s-)) \Delta \check{X}_i^n(s) \\ &\left. - \frac{1}{2} \sum_{i \in \mathcal{I}} (\phi_\kappa^n)_{ii}(\check{X}^n(s-)) (\Delta \check{X}_i^n(s))^2 \right]. \end{aligned}$$

Setting $t = \tau_{\kappa', T}^n$ as defined in the statement of the lemma and taking expectations on both sides we have

$$\begin{aligned}
 & \mathbb{E}[e^{-\gamma \tau_{\kappa', T}^n} \phi_{\kappa}^n(\check{X}^n(t))] \\
 &= \phi_{\kappa}^n(\check{X}^n(0)) + \sum_{i \in \mathcal{I}} \mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} (\phi_{\kappa}^n)_i(\check{X}^n(s-)) b_i^n(\check{X}^n(s), U^n(s)) ds \right] \\
 (66) \quad &+ \frac{1}{2} \sum_{i \in \mathcal{I}} \mathbb{E} \left[\sum_{s \leq t : |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) (\Delta \check{X}_i^n(s))^2 \right] \\
 &+ \mathbb{E}[C^n(\tau_{\kappa', T}^n)] - \gamma \mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} \phi_{\kappa}^n(\check{X}^n(s)) ds \right].
 \end{aligned}$$

We will now examine each of the elements on the right-hand side of (66). First, note that $\Delta \check{X}_i^n(s) = \Delta \check{W}_i^n(s)$ and, in particular,

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{s \leq \tau_{\kappa', T}^n : |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) (\Delta \check{X}_i^n(s))^2 \right] \\
 &= \mathbb{E} \left[\sum_{s \leq \tau_{\kappa', T}^n : |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) (\Delta \check{W}_i^n(s))^2 \right].
 \end{aligned}$$

Using the fact that \mathcal{V}_i^n , as defined in Lemma 6.3, is a martingale as well as the fact that $\phi_{\kappa}^n(\check{X}^n(s))$ and its derivative processes are bounded up to $\tau_{\kappa'}^n$, we have that the processes

$$(67) \quad \bar{\mathcal{V}}_i^n(\cdot) := \int_0^{\cdot \wedge \tau_{\kappa', T}^n} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) d\mathcal{V}_i^n(s)$$

and

$$(68) \quad \bar{\mathcal{M}}_i^n(\cdot) := \int_0^{\cdot \wedge \tau_{\kappa', T}^n} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) d\mathcal{M}_i^n(s)$$

are themselves martingales with $\bar{\mathcal{V}}_i^n(0) = \bar{\mathcal{M}}_i^n(0) = 0$ and in turn, by optional stopping, that $\mathbb{E}[\bar{\mathcal{V}}_i^n(\tau_{\kappa', T}^n)] = \mathbb{E}[\bar{\mathcal{M}}_i^n(\tau_{\kappa', T}^n)]$ (see, e.g., Lemma 5.45 in [14]). In turn, by the definition of $\mathcal{M}_i^n(\cdot)$ and $\mathcal{V}_i^n(\cdot)$ we have

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{s \leq \tau_{\kappa', T}^n : |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) (\Delta \check{W}_i^n(s))^2 \right] \\
 &= \mathbb{E} \left[\int_0^t (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) d(\check{W}_i^n(s))^2 \right] \\
 &= \mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) (\sigma_i^n(\check{X}^n(s), U^n(s)))^2 ds \right].
 \end{aligned}$$

Plugging this back into (66) we have that

$$\begin{aligned} & \mathbb{E}[e^{-\gamma\tau_{\kappa',T}^n} \phi_{\kappa}^n(\check{X}^n(t))] \\ &= \phi_{\kappa}^n(\check{X}^n(0)) + \sum_{i \in \mathcal{I}} \mathbb{E} \left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} (\phi_{\kappa}^n)_i(\check{X}^n(s-)) b_i^n(\check{X}^n(s), U^n(s)) ds \right] \\ &+ \frac{1}{2} \sum_{i \in \mathcal{I}} \mathbb{E} \left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) (\sigma_i^n(\check{X}^n(s), U^n(s)))^2 ds \right] \\ &- \gamma \mathbb{E} \left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} \phi_{\kappa}^n(\check{X}^n(s)) ds \right] + \mathbb{E}[C^n(\tau_{\kappa',T}^n)], \end{aligned}$$

which, using the definition of A_u^n in (31), yields

$$\begin{aligned} & \mathbb{E}[e^{-\gamma\tau_{\kappa',T}^n} \phi_{\kappa}^n(\check{X}^n(t))] \\ &= \phi_{\kappa}^n(\check{X}^n(0)) + \mathbb{E} \left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} A_{U^n(s)}^n \phi_{\kappa}^n(\check{X}^n(s)) ds \right] \\ &- \gamma \mathbb{E} \left[\int_0^{\tau_{\kappa',T}^n} e^{-\gamma s} \phi_{\kappa}^n(\check{X}^n(s)) ds \right] \\ &+ \mathbb{E}[C^n(\tau_{\kappa',T}^n)]. \end{aligned}$$

To complete the proof it then remains only to show that there exists a constant C such that

$$|\mathbb{E}[C^n(\tau_{\kappa',T}^n)]| \leq C \log^{k_1+1} n.$$

To that end, note that by Taylor's expansion,

$$\begin{aligned} \phi_{\kappa}^n(\check{X}^n(s)) &= \phi_{\kappa}^n(\check{X}^n(s-)) + \sum_{i \in \mathcal{I}} (\phi_{\kappa}^n)_i(\check{X}^n(s-)) \Delta \check{X}_i^n(s) \\ &+ \frac{1}{2} \sum_{i \in \mathcal{I}} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-) + \eta_{\check{X}^n(s-)}^n) \Delta \check{X}_i^n(s), \end{aligned}$$

where $\eta_{\check{X}^n(s-)}^n$ is such that $\check{X}^n(s-) + \eta_{\check{X}^n(s-)}^n$ is between $\check{X}^n(s-)$ and $\check{X}^n(s-) + \Delta \check{X}^n(s)$. In turn, adding and subtracting a term, we have that

$$\begin{aligned} & \phi_{\kappa}^n(\check{X}^n(s)) - \phi_{\kappa}^n(\check{X}^n(s-)) - \sum_{i \in \mathcal{I}} (\phi_{\kappa}^n)_i(\check{X}^n(s-)) \Delta \check{X}_i^n(s) \\ (69) \quad & - \frac{1}{2} \sum_{i \in \mathcal{I}} (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-)) (\Delta \check{X}_i^n(s))^2 \\ & = \sum_{i \in \mathcal{I}} \frac{1}{2} ((\phi_{\kappa}^n)_{ii}(\check{X}^n(s-) + \eta_{\check{X}^n(s-)}^n) - (\phi_{\kappa}^n)_{ii}(\check{X}^n(s-))) (\Delta \check{X}_i^n(s))^2. \end{aligned}$$

Since the jumps are of size 1 and, with probability 1, there are no simultaneous jumps, we have that $|\eta_{\check{X}^n(s-)}| \leq 1$. Adding the discounting, summing and taking expectations we have

$$(70) \quad \mathbb{E}[C^n(t)] \leq \mathbb{E} \left[\sum_{s \leq t : |\Delta \check{X}^n(s)| > 0} e^{-\gamma s} \sum_{i \in \mathcal{I}} \frac{1}{2} \max_{y: |y| \leq 1} ((\phi_\kappa^n)_{ii}(\check{X}^n(s-) + y) - (\phi_\kappa^n)_{ii}(\check{X}^n(s-))) (\Delta \check{X}_i^n(s))^2 \right],$$

and a lower bound can be created by minimizing over y instead of maximizing. Using again the fact that $\Delta \check{X}_i^n(t) = \Delta \check{W}_i^n(t)$ and that $\check{\mathcal{M}}_i^n$ and $\check{\mathcal{V}}_i^n$ as defined in (68) and (67) are martingales, we have that

$$(71) \quad \mathbb{E}[C^n(t)] \leq \mathbb{E} \left[\int_0^t \sum_{i \in \mathcal{I}} \frac{1}{2} \max_{y: |y| \leq 1} ((\phi_\kappa^n)_{ii}(\check{X}^n(s-) + y) - (\phi_\kappa^n)_{ii}(\check{X}^n(s-))) \times (\sigma_i^n(\check{X}^n(s), U^n(s)))^2 ds \right].$$

From (37) we have that

$$(72) \quad \frac{1}{2} \left| \sum_{i \in \mathcal{I}} ((\phi_\kappa^n)_{ii}(y) - (\phi_\kappa^n)_{ii}(x)) (\sigma_i^n(x, u))^2 \right| \leq C \log^{k_1} n$$

for all $u \in \mathcal{U}$ and $x, y \in \mathcal{B}_{\kappa'}^n$, with $|x - y| \leq 1$. The proof is then concluded by plugging (72) into (71), setting $t = \tau_{\kappa', T}^n$ and recalling that we can repeat all the above steps to obtain a lower bound in (71) by replacing $\max_{y: |y| \leq 1}$ with $\min_{y: |y| \leq 1}$ in (70).

7. Concluding remarks. This paper proposes a novel approach for solving problems of dynamic control of queueing systems in the Halfin–Whitt many-server heavy-traffic regime. Its main contribution is the use of Brownian approximations to construct controls that achieve optimality gaps that are logarithmic in the system size. This should be contrasted with the optimality gaps of size $o(\sqrt{n})$ that are common in the literature on asymptotic optimality. A distinguishing feature of our approach is the use of a *sequence* of Brownian control problems rather than a single (limit) problem. Having an entire sequence of approximating problems allows us to perform a more refined analysis, resulting in the improved optimality gap.

In further contrast with the earlier literature, in each of these Brownian problems the diffusion coefficient depends on both the system state and the control. Incorporating the impact of control on diffusion coefficients allows us to track the performance of the policy better but, at the same time, it leads to a more complex

diffusion control problem in which the associated HJB equation is fully nonlinear and nonsmooth. For *each* Brownian problem, we show that the HJB equation has a sufficiently smooth solution that coincides with the value function and that admits an optimal Markov policy. Most importantly, we derive useful gradient estimates that apply to the whole sequence and bound the growth rate of the gradients with the system size. These bounds are crucial for controlling the approximation errors when analyzing the original queueing system under the proposed tracking control.

The motivating intuition behind our approximation scheme is that the value functions of each queueing system and its corresponding Brownian control problem ought to be close. In particular, the optimal control for the Brownian problem should perform well for the queueing system. Moreover, the optimal Markov control of the Brownian problem can be approximated by a ratio (or tracking) control for the queueing system. While these observations are “correct” at a high level, they need to be qualified further. Our analysis underscores two sources of approximation errors that need to be addressed in order to obtain the refined optimality gaps. First, the value function of the Brownian control problem may be substantially different than that of the (preemptive) optimal control problem for the queueing system. This difference must be quantified relative to the system size, which we do indirectly through the gradient estimates for the value function of the Brownian control problem; this is manifested, for example, in the proof of Lemma 6.1.

The second source of error is in trying to imitate the optimal ratio control of the approximating Brownian control by a tracking control in the corresponding queueing system. The error arises because we insist on having a nonpreemptive control for the queueing system. Whereas under a preemptive control, one may be able to rearrange the queues instantaneously to match the tracking function of the Brownian system, this is not possible with nonpreemptive controls. Instead, we carefully construct and analyze the performance of the proposed nonpreemptive tracking policy. In doing so, we prove that the tracking control imitates closely the Brownian system with respect to a specific integrated functional of the queueing dynamics (see Theorem 5.1 and Remark 5.1). Here too, the gradient estimates for the value function of the Brownian system play a key role.

While the focus of this paper has been a relatively simple model to illustrate the key ideas behind our approach and the important steps in the analysis, we expect that similar results can be established in the cases of impatient customers, more general cost structures as well as more general network structures.

As suggested by the preceding analysis, the viability of these extensions and others will depend on whether it is possible to (a) solve the sequence of Brownian control problems and establish the necessary gradient estimates and (b) establish the corresponding approximation result for the nonpreemptive tracking control.

While we expect that the results of [6] on fully nonlinear elliptic PDEs can be invoked for the more general settings, extending our analysis which builds on those results may not be always straightforward. In particular, it is not immediately

obvious how to generalize the proof of the tracking result in Theorem 5.1 to more general settings.

Nevertheless, we can make some observations about the extensions mentioned above:

- *General convex costs.* As discussed in the [Introduction](#), the analysis of the convex holding cost case will probably be simpler as one tends to get “interior” solutions in that case as opposed to the corner solutions in the linear cost case, which causes nonsmoothness. We expect that the enhanced smoothness (relative to the linear holding cost case) will simplify the analysis of the HJB equations as well as that of the tracking performance.
- *Abandonment.* Our starting point in the analysis is that, among preemptive policies, work conserving policies are optimal. This is not, in general, true when customers are impatient and may abandon while waiting (see the discussion in Section 5.1 of [2]). As is the case in [2], our analysis will go through also for the case of impatient customers provided that the cost structure is such that work conservation is optimal among preemptive policies.
- *General networks.* Inspired by the generalization of [2], by Atar [1], to tree-like networks, we expect, for example, that such a generalization is viable in our setting as well. Indeed, we expect that the analysis of the (sequence of) HJB equations and the sequence of ADCPs be fairly similar for the tree-like network setting. We expect that, in that more general setting, it would be more complicated to bound the performance of the tracking policies as in Theorem 5.1.

APPENDIX

Proof of Lemma 3.1. Up to τ_κ^n , both functions $b^n(\cdot, u)$ and $\sigma^n(\cdot, u)$ are bounded and Lipschitz continuous (uniformly in u). With these conditions satisfied, strong existence and uniqueness follow as in Appendix D of [5]. Specifically, strong existence follows by successive approximations as in the proof of Theorem 2.9 of [9] and uniqueness follows as in Theorem 2.5 there.

Proof of Proposition 4.1. Fix $\kappa > 0$, $n \in \mathbb{Z}_+$ and $a > 0$. Recall that (40) corresponds to finding $\phi_{\kappa,a}^n \in \mathcal{C}^2(\mathcal{B})$ such that

$$(73) \quad 0 = F_a[x, \phi_{\kappa,a}^n(x), D\phi_{\kappa,a}^n(x), D^2\phi_{\kappa,a}^n(x)], \quad x \in \mathcal{B},$$

and so that $\phi_{\kappa,a}^n = 0$ on $\partial\mathcal{B}$ where $F_a[\cdot]$ is as defined in (41). Then, Proposition 4.1 will follow from Theorem 17.18 in [6] upon verifying certain conditions. The gradient estimates will also follow from [6] by carefully tracing some constants to identify their dependence on κ , n and a .

To that end, note that the function $F_a^i(x, z, p, r)$ [as defined in (42)] is linear in the (z, p, r) arguments for all $k \in \mathcal{I}$ and $x \in \mathcal{B}$. In turn, this function is concave in these arguments. Hence, to apply Theorem 17.18 of [6] it remains to establish that

condition (17.53) of [6] is satisfied for each of these functions. In the following we suppress the constant $a > 0$ from the notation. It suffices to show that there exist constants $\underline{\Delta} \leq \bar{\Delta}$ and η such that uniformly in $k \in \mathcal{I}$, $y = (x, z, p, r) \in \Gamma$, and $\xi \in \mathbb{R}^I$

$$(74) \quad 0 < \underline{\Delta} |\xi^2| \leq \sum_{i,j} F_{i,j}^k[y] \xi_i \xi_j \leq \bar{\Delta} |\xi|^2,$$

$$(75) \quad \max\{|F_p^k[y]|, |F_z^k[y]|, |F_{rx}^k[y]|, |F_{px}^k[y]|, |F_{zx}^k[y]|\} \leq \eta \underline{\Delta},$$

$$(76) \quad \max\{|F_x^k[y]|, |F_{xx}^k[y]|\} \leq \eta \underline{\Delta} (1 + |p| + |r|),$$

where

$$F_{i,j}^k(x, z, p, r) = \frac{\partial}{\partial r_{ij}} F^k(x, z, p, r), \quad F_{x_l}^k(x, z, p, r) = \frac{\partial}{\partial x_l} F^k(x, z, p, r)$$

and

$$(F_{rx}^k(x, z, p, r))_{ilj} = \frac{\partial^2}{\partial r_{il} \partial x_j} F^k(x, z, p, r).$$

The other cross-derivatives are defined similarly. We will show that we can choose $\underline{\Delta} = \varepsilon_0 n$, $\bar{\Delta} = \varepsilon_1 n$, $\eta = \varepsilon_2$ for constants $\varepsilon_0, \varepsilon_1$ and ε_2 that do not depend on n and a —this will be important in establishing the aforementioned gradient estimates. To establish (74) note that, given $\xi \in \mathbb{R}^I$,

$$(77) \quad F_{ij}^k \xi_i \xi_j = \begin{cases} \frac{1}{2}(\lambda_i^n + \mu_i(v_i n + x_i)) \xi_i^2, & \text{for } i = j, i \neq k, \\ \frac{1}{2}(\lambda_i^n + \mu_i(v_i n + x_i)) \xi_i^2 - \frac{1}{2} f(e \cdot x), & \text{for } i = j = k, \\ 0, & \text{otherwise.} \end{cases}$$

Hence,

$$\sum_{i,j} F_{ij}^k \xi_i \xi_j = \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i(v_i n + x_i)) \xi_i^2 - \frac{1}{2} f(e \cdot x) \xi_k^2.$$

Consequently, for $(x, z, r, p) \in \Gamma$ we have that

$$\sum_{i,j} F_{ij}^k \xi_i \xi_j \leq I(\lambda + \mu_{\max} n + \mu_{\max} \kappa \sqrt{n} \log^m n) \sum_{i \in \mathcal{I}} \xi_i^2 + \frac{1}{2} \kappa \sqrt{n} \log^m n \xi_k^2,$$

where $\mu_{\max} = \max_k \mu_k$. In particular, we can choose $\varepsilon_1 > 0$ so that for all $n \in \mathbb{Z}$,

$$\sum_{i,j} F_{ij}^k \xi_i \xi_j \leq \varepsilon_1 n.$$

To obtain the lower bound note that, for $y \in \Gamma$,

$$\sum_{i,j} F_{ij}^k \xi_i \xi_j \geq \frac{1}{2} \left(\min_{i \in \mathcal{I}} \lambda_i^n + \min_{i \in \mathcal{I}} \mu_i \kappa \sqrt{n} \log^m n \right) \sum_{i \in \mathcal{I}} \xi_i^2 - \frac{1}{2} \xi_k^2 \kappa \sqrt{n} \log^m n.$$

Hence, we can find $\varepsilon_0 > 0$ such that for all n ,

$$\sum_{i,j} F_{ij}^k \xi_i \xi_j \geq \varepsilon_0 n.$$

Note that above ε_0 and ε_1 can depend on κ but they do not depend on n and a . Hence, we have established (74) and we turn to (75). To that end, note that

$$(78) \quad \begin{aligned} F_{pk}^k(x, z, p, r) &= f(e \cdot x) + l_k^n - \mu_k x_k \quad \text{and} \\ F_{pi}^k(x, z, p, r) &= l_i^n - \mu_i x_i \quad \text{for } i \neq k. \end{aligned}$$

Therefore,

$$\begin{aligned} |F_p^k| &\leq (e \cdot x)^+ + 1 + \sum_i |l_i^n + \mu_i x_i| \\ &\leq I\kappa \sqrt{n} \log^m n + 1 + I \max_i (|l_i^n| + \mu_i \kappa \sqrt{n} \log^m n), \end{aligned}$$

where we used the simple observation that $f(e \cdot x) \leq (e \cdot x)^+ + 1$. Clearly, we can choose ε_2 so that $|F_p^k| \leq \varepsilon_2 \varepsilon_0 \sqrt{n} \log^m n$. Also $F_z^k = -\gamma$ and $F_{zx} = 0$ so that by re-choosing ε_2 large enough we have $\max\{|F_z^k[y]|\}, |F_{zx}^k[y]|\} \leq \varepsilon_2 \varepsilon_0 \sqrt{n} \log n$. Finally, by (77) we have that

$$\begin{aligned} F_{rijx_i}^k &= 0 \quad \text{for } i \neq j, \\ F_{riix_j}^k &= 0 \quad \text{for } i \neq k, i \neq j, \\ F_{riix_i}^k &= \frac{1}{2} \mu_i \quad \text{for } i \neq k, \\ F_{riix_i}^k &= \frac{1}{2} \mu_i \quad \text{for } i \neq k, \\ F_{rkkx_k}^k &= \frac{1}{2} \mu_k - \frac{1}{2} \frac{\partial}{\partial x_k} f(e \cdot x), \\ F_{rkkx_j}^k &= \frac{1}{2} \frac{\partial}{\partial x_k} f(e \cdot x) \quad \text{for } j \neq k. \end{aligned}$$

Thus,

$$|F_{rx}^k|^2 \leq \sum_{l \in \mathcal{I}} \frac{1}{2} \left| \frac{\partial}{\partial x_l} f(e \cdot x) \right|^2 + \frac{1}{2} \mu_{\max} \leq \frac{1}{2} (1 + \mu_{\max}),$$

where we used the fact that $f(\cdot)$ is continuously differentiable with Lipschitz constant 1 (independently of a). Finally,

$$F_{x_i}^k = \frac{\partial}{\partial x_i} f(e \cdot x) \left(c_k + \mu_k p_k - \frac{1}{2} \mu_k r_{kk} \right) - \mu_i p_i + \frac{1}{2} \mu_i r_{ii},$$

so that

$$(79) \quad |F_{x_i}^k| \leq |c_k| + \mu_k |p| + \frac{1}{2} \mu_k |r| + \mu_i |p| + \frac{1}{2} \mu_i |r|.$$

Also, note that

$$F_{x_i x_j}^k = \frac{\partial}{\partial x_i \partial x_j} f(e \cdot x) \left(c_k + p_k - \frac{1}{2} r_{kk} \right),$$

so that

$$F_{x_i x_j}^k = \begin{cases} 2[c_k + \mu_k p_k - \frac{1}{2} \mu_k r_{kk}], & \text{if } |e \cdot x| \leq \frac{1}{4}, \\ 0, & \text{otherwise.} \end{cases}$$

Combining the above gives

$$|F_{xx}^k| \leq \varepsilon_2 \varepsilon_0 (1 + |p| + |r|)$$

for suitably redefined ε_2 which concludes the proof that the conditions (74)–(76) hold with $\bar{\Lambda} = \varepsilon_1 n$, $\underline{\Lambda} = \varepsilon_0 n$ and $\eta = \varepsilon_2$. Having verified these conditions, the existence and uniqueness of the solution $\phi_{\kappa, a}^n$ to (40) now follows from Theorem 17.18 in [6].

To obtain the gradient estimates in (44) we first outline how the solution $\phi_{\kappa, a}^n$ is obtained in [6] as a limit of solutions to smoothed equations (we refer the reader to [6], page 466, for the more elaborate description). To that end, let F_a^i be as defined in (42) and for $y \in \Gamma$ define

$$(80) \quad F^h[y] = G_h(F_a^1[y], \dots, F_a^I[y]),$$

where

$$G_h(y) = h^{-I} \int_{\bar{y} \in \mathbb{R}^I} \rho\left(\frac{y - \bar{y}}{h}\right) G_0(\bar{y}) d\bar{y}$$

and $G_0(x) = \min_{i \in \mathcal{I}} x_i$ and $\rho(\cdot)$ is a mollifier on \mathbb{R}^I (see [6], page 466). F^h satisfies all the bounds in (74)–(76) uniformly in h ; cf. [6], page 466. Then, there exists a unique solution u^h for the equations

$$(81) \quad F^h[u^h] = 0$$

on \mathcal{B}_κ^n with $u^h = 0$ on $\partial \mathcal{B}_\kappa^n$.

The solution $\phi_{\kappa, a}^n$ is now obtained as a limit of $\{u^h\}$ in the space $C_*^{2, \alpha}(\mathcal{B})$ as defined in (45). Moreover, since the gradient bounds are shown in [6] to be independent of h , it suffices for our purposes to fix h and focus on the construction of the gradient bounds.

Our starting point is the bound at the bottom of page 461 of [6] by which

$$(82) \quad |u^h|_{2, \alpha, \mathcal{B}_\kappa^n}^* \leq \check{C}(a, n)(1 + |u^h|_{2, \mathcal{B}_\kappa^n}^*),$$

where $|u^h|_{2, \mathcal{B}_\kappa^n}^* = \sum_{j=0}^2 |u^h|_{j, \mathcal{B}}^*$ and $[\cdot]_{j, \mathcal{B}}^*$, $j = 0, 1, 2$, are as defined in Section 4. The constant $\alpha(a, n)$ depends only on the number of classes I and on $\bar{\Lambda}/\underline{\Lambda}$ (see [6],

top of page 461) and this fraction equals, in our context, to $\varepsilon_1/\varepsilon_0$ and is thus constant and independent of n and a .

We will address the constant $\check{C}(a, n)$ shortly. We first argue how one proceeds from (82). Fix $0 < \delta < 1$, let $\epsilon = \delta/\check{C}(a, n)$ and $C(\epsilon) = 2/(\epsilon/8)^{1/\alpha}$ (see [6], top of page 132). Then, applying an interpolation inequality (see [6], bottom of page 461 and Lemma 6.32 on page 130), it is obtained that

$$|u^h|_{2,0,\mathcal{B}_k^n}^* \leq C(\epsilon)|u^h|_{0,\Omega}^* + \epsilon|u^h|_{2,\alpha,\mathcal{B}_k^n}^*.$$

Plugging this back into (82) one then has

$$|u^h|_{2,\alpha,\mathcal{B}_k^n}^* \leq \check{C}(a, n) \left(1 + \bar{C}\check{C}(a, n)^{1/\alpha}|u^h|_{0,\mathcal{B}_k^n}^* + \frac{\delta}{\check{C}(a, n)}|u^h|_{2,\alpha,\mathcal{B}_k^n}^* \right)$$

for a constant \bar{C} that depends only on δ and α . In turn,

$$|u^h|_{2,\alpha,\mathcal{B}_k^n}^* \leq \bar{C}(\check{C}(a, n))^{1+1/\alpha}|u^h|_{0,\mathcal{B}_k^n}^*$$

for a constant \bar{C} that does not depend on a or n .

Hence, to obtain the required bound in (44) it remains only to bound $\check{C}(a, n)$. Following [6], building on equation (17.51) of [6], $\check{C}(a, n)$ is the (minimal) constant that satisfies

$$(83) \quad C(1 + M_2)(1 + \tilde{\mu}R_0 + \bar{\mu}R_0^2) \leq \check{C}(a, n)(1 + |u^h|_{2,\mathcal{B}}^*),$$

where (as stated in [6], bottom of page 460) the (redefined) constant C depends only on the number of class I and on $\bar{\Delta}/\underline{\Delta} = \varepsilon_1/\varepsilon_0$. The constants $\tilde{\mu}$ and $\bar{\mu}$ are defined in [6] and we will explicitly define them shortly. Here one should not confuse $\bar{\mu}$ with the average service rate in our system. In what follows $\bar{\mu}$ will only be used as the constant in [6]. We now bound constants $\tilde{\mu}$ and $\bar{\mu}$. These are defined by

$$\begin{aligned} \tilde{\mu} &= \frac{D_0}{\underline{\Delta}(1 + M_2)}, & \bar{\mu} &= \frac{C(I)}{\underline{\Delta}} \left(\frac{A_0^2}{\underline{\Delta}\epsilon} + \frac{B_0}{1 + M_2} \right), \\ D_0 &= \sup_{x,y \in \mathcal{B}} \{ |F_x^h(y, u^h(y), Du^h(y), D^2u^h(x))| \\ &\quad + |F_z^h(y, u^h(y), Du^h(y), D^2u^h(x))||Du^h(y)| \\ &\quad + |F_p^h(y, u^h(y), Du^h(y), D^2u^h(x))||D^2u^h(y)| \}, \\ A_0 &= \sup_{\mathcal{B}} \{ |F_{rx}^h| + |F_p^h| \}, \\ B_0 &= \sup_{\mathcal{B}} \{ |F_{px}| |D^2u^h| + |F_z| |D^2u^h| + |F_{zx}| |Du^h| + |F_{xx}| \}, \end{aligned}$$

where $C(I)$ is a constant that depends only on the number of classes I , $\epsilon \in (0, 1)$ is arbitrary and fixed (independent of n and a) and $M_2 = \sup_{\mathcal{B}} |D^2u^h|$. The constants

$\bar{\mu}$, $\tilde{\mu}$ and M_2 are defined in [6], pages 456–460, and A_0 and B_0 are as on page 461 there.

We note that F_z^h is a constant, F_p^h is bounded by $\bar{C}\sqrt{n}\log^m n$ for some constant \bar{C} [see (44)] that depends only on κ and, by (79), $|F_x^h| \leq \varepsilon_2\varepsilon_0(1 + |p| + |r|)$. In turn, $D_0 \leq 4\varepsilon_2\varepsilon_0\sqrt{n}\log^m n \sup_{\mathcal{B}}(1 + |Du^h| + |D^2u^h|)$. Arguing similarly for A_0 and B_0 we find that there exists a constant \bar{C} (that does not depend on n and a) such that

$$A_0 \leq \bar{C}\sqrt{n}\log^m n \quad \text{and} \quad B^0 \leq \bar{C} \sup_{\mathcal{B}}(1 + |Du^h| + |D^2u^h|),$$

which in turn implies the existence of a redefined constant \bar{C} such that

$$\tilde{\mu} \leq \frac{\bar{C}\log^m n}{\sqrt{n}(1 + M_2)} \sup_{\mathcal{B}}(1 + |Du^h| + |D^2u^h|)$$

and

$$\bar{\mu} \leq \frac{\bar{C}\log^{2m} n}{n} + \frac{\bar{C}}{n(1 + M_2)} \sup_{\mathcal{B}}(1 + |Du^h| + |D^2u^h|).$$

The proof of the bound is concluded by plugging these back into (83) and setting $R_0 = \kappa\sqrt{n}\log^m n$ there to get that

$$\check{C}(a, n) \leq C \log^{4m(1+1/\alpha)} n$$

for some C that does not depend on a and n .

The constant \check{C} on the right-hand side of (44) (which can depend on n but does not depend on a) is argued as in the proof of Theorem 17.17 in [6] and we conclude the proof by noting that the global Lipschitz constant (that we allow to depend on n) follows from Theorem 7.2 in [13].

We next turn to proof of Theorem 5.1. First, we will explicitly construct the queueing process under the h -tracking policy and state a lemma that will be of use in the proof of the theorem. Define $A_i^n(t) = \mathcal{N}_i^a(\lambda_i^n t)$ so that A_i^n is the arrival process of class- i customers. Given a ratio control U^n and the associated queueing process $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n, \check{W}^n)$ is as defined in (21). Also, we define

$$D^n(t) = \sum_{i \in \mathcal{I}} \mathcal{N}_i^d \left(\mu_i \int_0^t Z_i^n(s) ds \right).$$

That is, $D^n(t)$ is the total number of service completions by time t in the n th system.

For the construction of the queueing process under the tracking policy we define a family of processes $\{\mathcal{A}_{i,\mathcal{H}}^n, i \in \mathcal{I}, \mathcal{H} \subset \mathcal{I}\}$ as follows: let $\{\xi_{\mathcal{K}}^l; l \in \mathbb{Z}_+, \mathcal{K} \subset \mathcal{I}\}$ be a family of i.i.d uniform $[0, 1]$ random variables independent of $\bar{\mathcal{F}}_\infty$ as defined in (17). For each $\mathcal{K} \subset \mathcal{I}$, define the processes $(\mathcal{A}_{i,\mathcal{H}}^n, i \in \mathcal{I})$ by

$$(84) \quad \mathcal{A}_{i,\mathcal{H}}^n(t) = \sum_{l=1}^{D^n(t)} 1 \left\{ \frac{\sum_{k < i, k \in \mathcal{H}} \lambda_k}{1 \vee \sum_{k \in \mathcal{H}} \lambda_k} < \xi_{\mathcal{K}}^l \leq \frac{\sum_{k \leq i, k \in \mathcal{H}} \lambda_k}{1 \vee \sum_{k \in \mathcal{H}} \lambda_k} \right\}.$$

We note that for any strict subset $\mathcal{H} \subset \mathcal{I}$ and $i \in \mathcal{K}$, the probability that a jump of $D^n(t)$ results in a jump of $\mathcal{A}_{i,\mathcal{H}}^n$ is equal to $\lambda_i^n / \sum_{k \in \mathcal{H}} \lambda_k^n = a_i / \sum_{k \in \mathcal{H}} a_k$ and is strictly greater than $\lambda_i^n / \sum_{k \in \mathcal{I}} \lambda_k^n = a_i$. We define

$$(85) \quad \epsilon_i = \min_{\mathcal{H} \subset \mathcal{I}} a_i - \frac{a_i}{\sum_{k \in \mathcal{H}} a_k},$$

and note that $\epsilon_i > 0$ by our assumption that $a_i > 0$ for all $i \in \mathcal{I}$ (see Section 2). Let $\bar{\epsilon} = \min_i \epsilon_i / 4$.

Note that at time intervals in which $i \in \mathcal{K}(\cdot) = \mathcal{H}$ (see Definition 2.1) for some $\emptyset \neq \mathcal{H} \subset \mathcal{I}$, the process $\mathcal{A}_{i,\mathcal{H}}^n$ jumps with probability $\lambda_i^n / \sum_{k \in \mathcal{H}} \lambda_k^n$ whenever a server becomes available (i.e., upon a jump of D^n). In turn, we will use the processes $\{\mathcal{A}_{i,\mathcal{H}}^n, i \in \mathcal{I}, \mathcal{H} \subset \mathcal{I}\}$ to generate (randomized) admissions to service of class- i customers under the h -tracking policy.

More specifically, under the h -tracking policy (see Definition 2.1) a customer from the class- i queue enters service in the following events:

- (i) A class- i customer that arrives at time t enters service immediately if there are idle servers, that is, if $(e \cdot \check{X}^n(t-))^- > 0$.
- (ii) If a server becomes available at time t (corresponding to a jump of D^n) and t is such that $i \in \mathcal{K}(t-) = \mathcal{H} \subset \mathcal{I}$, then a customer from the class- i queue is admitted to service at time t with probability $\lambda_i^n / \sum_{k \in \mathcal{H}} \lambda_k^n$. This admission to service corresponds to a jump of the process $\mathcal{A}_{i,\mathcal{H}}^n$ as defined in (84).
- (iii) If a server becomes available at time t (corresponding to a jump of D^n) and t is such that $\mathcal{K}(t-) = \emptyset$ and $i = \min\{k \in \mathcal{I} : Q_k^n(t) > 0\}$, then a class- i customer is admitted to service.

Formally, the queueing process $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ satisfies

$$\begin{aligned} Z_i^n(t) &= Z_i^n(0) + \int_0^t 1\{(e \cdot \check{X}^n(s))^- > 0\} dA_i^n(s) \\ &\quad + \sum_{\mathcal{H} \subset \mathcal{I}} \int_0^t 1\{i \in \mathcal{K}(s-), \mathcal{K}(s-) = \mathcal{H}\} d\mathcal{A}_{i,\mathcal{H}}^n(s) \\ &\quad + \int_0^t 1\{\mathcal{K}(s-) = \emptyset, i = \min\{k \in \mathcal{I} : Q_k^n(s-) > 0\}\} dD^n(s) \\ &\quad - \mathcal{N}_i^d\left(\mu_i \int_0^t Z_i^n(s) ds\right), \quad i \in \mathcal{I}, \\ X_i^n(t) &= X_i^n(0) + A_i^n(t) - \mathcal{N}_i^d\left(\mu_i \int_0^t Z_i^n(s) ds\right), \quad i \in \mathcal{I}, \\ Q_i^n(t) &= X_i^n(t) - Z_i^n(t), \quad i \in \mathcal{I}. \end{aligned}$$

The second, third and fourth terms on the right-hand side of the equation for Z_i^n correspond, respectively, to the events described by items (i)–(iii) above. Finally,

\check{X}^n is defined from X^n as in (8). The fact that the above system of equations has a unique solution is proved by induction on arrival and service completions times (see, e.g., the proof of Theorem 9.2 of [11]). Clearly, \mathbb{X}^n satisfies (12)–(16) with U_i^n there constructed from Q^n using (11).

We note that, with this construction, the tracking policy is admissible in the sense of Definition 2.2. Also, it will be useful for the proof of Theorem 5.1 to note that with this construction, if $[s, t]$ is an interval such that $i \in \mathcal{K}(u) \subset \mathcal{I}$ for all $u \in [s, t]$ then

$$(86) \quad Q_i^n(t) - Q_i^n(s) = A_i^n(t) - A_i^n(s) - \sum_{\mathcal{H} \subset \mathcal{I}} \int_s^t 1\{\mathcal{K}(u-) = \mathcal{H}\} dA_{i,\mathcal{H}}^n(u).$$

Before proceeding to the proof of Theorem 5.1 the following lemma provides preliminary bounds for arbitrary ratio controls.

LEMMA A.1. *Fix $\kappa, T > 0$ and a ratio control U^n , let $\mathbb{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ be the associated queueing process and define*

$$\tau_{\kappa,T}^n = \inf\{t \geq 0 : \check{X}^n(t) \notin \mathcal{B}_\kappa^n\} \wedge T \log n.$$

Then, there exist constants $C_1, C_2, K_0 > 0$ (that depend on T and κ but that do not depend on n or on the ratio control U^n) such that for all $K > K_0$ and all n large enough,

$$(87) \quad \mathbb{P}\left\{ \sup_{0 \leq t \leq 2T \log n} |\check{W}^n(t)| > K \sqrt{n} \log n \right\} \leq C_1 e^{-C_2 K \log n},$$

$$(88) \quad \mathbb{P}\{ |\check{X}^n(t) - \check{X}^n(s)| > ((t-s) + (t-s)^2) K \sqrt{n} \log n + K \log n, \text{ for some } s < t \leq 2T \log n \} \leq C_1 e^{-C_2 K \log n},$$

$$(89) \quad \mathbb{P}\{ |A_i^n(t) - A_i^n(s) - \lambda_i^n(t-s)| > \bar{\epsilon}n(t-s) + K \log n \text{ for some } s < t \leq \tau_{\kappa,T}^n \} \leq C_1 e^{-C_2 K \log n}, \quad i \in \mathcal{I},$$

$$(90) \quad \mathbb{P}\left\{ \left| D^n(t) - D^n(s) - \sum_i \mu_i v_i n(t-s) \right| > \bar{\epsilon}n(t-s) + K \log n \text{ for some } s < t \leq \tau_{\kappa,T}^n \right\} \leq C_1 e^{-C_2 K \log n},$$

$$(91) \quad \mathbb{P} \left\{ \mathcal{A}_{i, \mathcal{H}}^n(t) - \mathcal{A}_{i, \mathcal{H}}^n(s) - \lambda_i^n(t-s) \leq \frac{\epsilon_i}{2} n(t-s) - K \log n \right. \\ \left. \text{for some } s < t \leq \tau_{\kappa, T}^n \right\} \\ \leq C_1 e^{-C_2 K \log n}, \quad i \in \mathcal{I}, \mathcal{K} \subset \mathcal{I}.$$

PROOF. Equation (87) follows from strong approximations (see, e.g., Lemma 2.2. in [4]) and known bounds on the supremum of Brownian motion (see, e.g., equation 2.1.53 in [4]). Equation (88) then follows using this bound together with (52) in [2] but with $X^n(t) - X^n(s)$ instead of $X^n(t)$ (in the notation of [2] \check{W}^n is \hat{W}^n). Equations (89)–(91) follow by carefully constructing and bounding the increments. We outline the proof of (89) and the others follow similarly. To that end, note that given K and for all n large enough

$$\{|A_i^n(t) - A_i^n(s) - \lambda_i^n(t-s)| \leq \bar{\epsilon} n(t-s) + K \log n, \\ \text{for all } 0 \leq s \leq t \leq 2T \log n\} \\ \supseteq \left\{ \max_{l \leq N_i^n} \max_{j \geq 0: j \log^l n \leq 3T \log n} \frac{|A_i^{j,l,n} - \lambda_i^n / \log^l n|}{\sqrt{\lambda_i^n / \log^l n}} \leq K \sqrt{\log n} \right\},$$

where $A_i^{j,l,n} = A_i^n((j+1)/\log^l n) - A_i^n(j/\log^l n)$ and $N = \max\{l: \log^l n \leq \lambda_i^n / \log n\}$. Indeed, given an interval $[s, t)$ we can construct it from smaller intervals. Starting with $l = 0$, we fit as many intervals of size 1 into $[s, t)$, we then continue to fit as many intervals of size $1/\log n$ to the uncovered part of the interval and continue sequentially in l . We omit the simple and detailed construction. Note that with such construction, given an interval $[s, t)$, its covering uses at most $\log n$ intervals of size $\log^l n$ for each $l \geq 0$. Also, note that $N_i^n \leq C \log n$ for all n and some constant C . From here, using strong approximations (or bounds for Poisson random variables as in [7]) we have, for each j and l , that

$$\mathbb{P} \left\{ \frac{|A_i^{j,l,n} - \lambda_i^n / \log^l n|}{\sqrt{\lambda_i^n / \log^l n}} > K \sqrt{\log n} \right\} \leq C_1 e^{-C_2 K \log n}.$$

Since the number of intervals considered is of the order of $n \log n$, the bound follows with redefined constants C_1 and C_2 . \square

Proof of Theorem 5.1. Since κ is fixed throughout we use $h^n(\cdot) = h_{\kappa}^{*,n}(\cdot)$. As in the statement of the theorem, let

$$\psi^n(x, u) = L(x, u) + A_u^n \phi_{\kappa}^n(x) - \gamma \phi_{\kappa}^n(x) \quad \text{for } x \in \mathcal{B}_{\kappa}^n, u \in \mathcal{U},$$

so that by the definition of $A_u^n(x)$ we have

$$(92) \quad \begin{aligned} \psi^n(x, u) = & -\gamma \phi_\kappa^n(x) \\ & + (e \cdot x)^+ \cdot \sum_{i \in \mathcal{I}} u_i \left\{ c_i + \mu_i (\phi_\kappa^n)_i(x) - \mu_i \frac{1}{2} (\phi_\kappa^n)_{ii}(x) \right\} \\ & + \sum_{i \in \mathcal{I}} (l_i^n - \mu_i x_i) (\phi_\kappa^n)_i(x) + \frac{1}{2} \sum_{i \in \mathcal{I}} (\lambda_i^n + \mu_i (v_i n + x_i)) (\phi_\kappa^n)_{ii}(x). \end{aligned}$$

Defining, as before,

$$M_i^n(z) = c_i + \mu_i (\phi_\kappa^n)_i(z) - \frac{1}{2} \mu_i (\phi_\kappa^n)_{ii}(z),$$

we have that

$$\psi^n(x, u) - \psi^n(x, v) = (e \cdot x)^+ \left(\sum_{i \in \mathcal{I}} v_i M_i^n(x) - \sum_{i \in \mathcal{I}} u_i M_i^n(x) \right).$$

Let U^n be the ratio control associated with the h^n -tracking policy, let $\check{X}^n = (X^n, Q^n, Z^n, \check{X}^n)$ be the associated queueing process and define

$$(93) \quad \begin{aligned} \check{\psi}^n(s) &= \psi^n(\check{X}^n(s), U^n(s)) - \psi^n(\check{X}^n(s), h^n(\check{X}^n(s))) \\ &= (e \cdot \check{X}^n(s))^+ \sum_{i \in \mathcal{I}} h_i^n(\check{X}^n(s)) M_i^n(\check{X}^n(s)) \\ &\quad - (e \cdot \check{X}^n(s))^+ \sum_{i \in \mathcal{I}} U_i^n(s) M_i^n(\check{X}^n(s)). \end{aligned}$$

Recall that, by construction, $Q_i^n(s) = (e \cdot \check{X}^n(s))^+ U_i^n(s)$ so that (93) can be rewritten as

$$\begin{aligned} \check{\psi}^n(s) &= \psi^n(\check{X}^n(s), U^n(s)) - \psi^n(\check{X}^n(s), h^n(\check{X}^n(s))) \\ &= (e \cdot \check{X}^n(s))^+ \sum_{i \in \mathcal{I}} h_i^n(\check{X}^n(s)) M_i^n(\check{X}^n(s)) \\ &\quad - \sum_{i \in \mathcal{I}} Q_i^n(s) M_i^n(\check{X}^n(s)). \end{aligned}$$

The theorem will be proved if we show that

$$(94) \quad \mathbb{E} \left[\int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} |\check{\psi}^n(s)| ds \right] \leq C \log^{k_0+3} n.$$

To that end, define a sequence of times $\{\tau_l^n\}$ as follows:

$$\tau_{l+1}^n = \inf\{t > \tau_l^n : h^n(\check{X}^n(t)) \neq h^n(\check{X}^n(\tau_l^n))\} \wedge \tau_{\kappa', T}^n \quad \text{for } l \geq 0,$$

where $\tau_0^n = \eta^n \wedge \tau_{\kappa', T}^n$ and

$$(95) \quad \eta^n = t_0 \frac{\log^m n}{\sqrt{n}}$$

for $t_0 = 4\kappa/\epsilon_i$ with $\epsilon_i = \min_{\mathcal{H} \subset \mathcal{I}} a_i - \frac{a_i}{\sum_{k \in \mathcal{H}} a_k}$ as in (85). Finally, we define $r^n = \sup\{l \in \mathbb{Z}_+ : \tau_l^n \leq \tau_{\kappa', T}^n\}$ and set $\tau_{r^n+1}^n = \tau_{\kappa', T}^n$. We then have

$$\begin{aligned} & \int_0^{\tau_{\kappa', T}^n} e^{-\gamma s} |\check{\psi}^n(s)| ds \\ &= \sum_{l=1}^{r^n+1} \int_{\tau_{l-1}^n}^{\tau_l^n} e^{-\gamma s} |\check{\psi}^n(s)| ds \\ &= \sum_{l=1}^{r^n+1} \left(\int_{\tau_{l-1}^n}^{(\tau_{l-1}^n + \eta^n) \wedge \tau_l^n} e^{-\gamma s} |\check{\psi}^n(s)| ds + \int_{\tau_{l-1}^n + \eta^n}^{\tau_l^n \vee (\tau_{l-1}^n + \eta^n)} e^{-\gamma s} |\check{\psi}^n(s)| ds \right). \end{aligned}$$

The proof is now divided into three parts. We will show that, under the conditions of the theorem,

$$(96) \quad \mathbb{E} \left[\sup_{1 \leq l \leq r^n+1} \sup_{\tau_{l-1}^n \leq s < (\tau_{l-1}^n + \eta^n) \wedge \tau_l^n} |\check{\psi}^n(s)| \right] \leq C \log^{k_0+2} n,$$

$$(97) \quad \mathbb{E} \left[\sup_{1 \leq l \leq r^n+1} \sup_{(\tau_{l-1}^n + \eta^n) \leq s < \tau_l^n \vee (\tau_{l-1}^n + \eta^n)} |\check{\psi}^n(s)| \right] \leq C \log^{k_0+2} n,$$

where we define $\sup_{(\tau_{l-1}^n + \eta^n) \leq s < \tau_l^n \vee (\tau_{l-1}^n + \eta^n)} |\check{\psi}^n(s)| = 0$ if $\tau_l^n \leq \tau_{l-1}^n + \eta^n$. Finally, we will show that

$$(98) \quad \mathbb{E} \left[\int_0^{\eta^n \wedge \tau_{\kappa', T}^n} |\check{\psi}^n(s)| ds \right] \leq C \log^{k_0} n.$$

The proof of (96) hinges on the fact that, sufficiently close to a change point τ_l^n , all the customer classes, i , for which $h_i^n(\check{X}^n(s)) = 1$ for some s in a neighborhood of τ_l^n , will have similar values of $M_i^n(\check{X}^n(s))$. This will follow from our gradient estimates for ϕ_κ^n . The proof of (97) hinges on the fact that, η^n time units after a change point τ_l^n the queues of all the classes for which $h^n(\check{X}^n(\tau_l^n)) = 0$ are small because, under the tracking policy, these classes receive a significant share of the capacity.

Toward formalizing this intuition, define the following event on the underlying probability space:

$$\tilde{\Omega}(K) = \{ |\check{X}^n(t) - \check{X}^n(s)| \leq K \sqrt{n} \log^2 n(t-s) + K \log n, \text{ for all } s < t \leq \tau_{\kappa', T}^n \}$$

$$\cap_{\mathcal{H} \subset \mathcal{I}} \left\{ \mathcal{A}_{i, \mathcal{H}}^n(t) - \mathcal{A}_{i, \mathcal{H}}^n(s) - \lambda_i^n(t-s) \geq \frac{\epsilon_i}{2} n(t-s) - K \log n \right. \\ \left. \text{for all } s < t \leq \tau_{\kappa', T}^n \right\}$$

$$\cap_{i \in \mathcal{I}} \{ |A_i^n(t) - A_i^n(s) - \lambda_i^n(t-s)| \leq \bar{\epsilon} n(t-s) + K \log n \text{ for all } s < t \leq \tau_{\kappa', T}^n \}.$$

For each $0 \leq t \leq \tau_{k',T}^n$ and $i \in \mathcal{I}$ let

$$(99) \quad \zeta_i^n(t) = \sup\{s \leq t : h_i^n(\check{X}^n(s)) = 1\},$$

$$(100) \quad \tilde{\zeta}_i^n(t) = \inf\{s \geq t : h_i^n(\check{X}^n(s)) = 1\} \wedge \tau_{k',T}^n$$

and

$$(101) \quad \hat{\zeta}_i^n(t) = \begin{cases} \zeta_i^n(t) + \eta^n, & \text{if } Q_i^n(t) > 4K \log n, \\ t, & \text{otherwise.} \end{cases}$$

Then, we claim that on $\tilde{\Omega}(K)$ and for all t with $\tilde{\zeta}_i^n(t) > \hat{\zeta}_i^n(t)$,

$$(102) \quad \sup_{\hat{\zeta}_i^n(t) \leq s < \tilde{\zeta}_i^n(t)} |(e \cdot \check{X}^n(s))^+ U_i^n(s) - (e \cdot \check{X}^n(s))^+ h_i^n(\check{X}^n(s))| \leq 12K \log n.$$

Note that since $h_i^n(\cdot) \in \{0, 1\}$, the above is equivalently written as

$$(103) \quad \sup_{\hat{\zeta}_i^n(t) \leq s < \tilde{\zeta}_i^n(t)} Q_i^n(s) \leq 12K \log n.$$

In words, when the process $\check{X}^n(t)$ enters a region in which $h_i^n(\check{X}^n(\cdot)) = 0$ the queue of class i will be drained up to $12K \log n$ within at most η^n time units and it will remain there up to $\tilde{\zeta}_i^n(t)$. We postpone the proof of (102) and use it in proceeding with the proof of the theorem.

To that end, fix $l \geq 0$ and let

$$j_l^* = \min \arg \min_{i \in \mathcal{I}} M_i^n(\check{X}^n(\tau_l^n)).$$

Then, by the definition of the function h^n in (38) we have that $h_{j_l^*}^n(\check{X}^n(\tau_l^n)) = 1$ and $h_i(\check{X}^n(\tau_l^n)) = 0$ for all $i \neq j_l^*$. In particular,

$$\begin{aligned} \check{\psi}^n(s) &= (e \cdot \check{X}^n(s))^+ h_{j_l^*}^n(\check{X}^n(s)) M_{j_l^*}^n(\check{X}^n(s)) \\ &\quad - \sum_{i \in \mathcal{I}} Q_i^n(s) M_i^n(\check{X}^n(s)) \end{aligned}$$

for all $s \in [\tau_l^n, (\tau_l^n + \eta^n) \wedge \tau_{l+1}^n)$. Let

$$\mathcal{J}(\tau_l^n) = \{i \in \mathcal{I} : Q_i^n(\tau_l^n -) > 4K \log n\}.$$

Then, simple manipulations yield

$$(104) \quad \begin{aligned} |\check{\psi}^n(s)| &\leq \sum_{i \notin \mathcal{J}(\tau_l^n) \cup \{j_l^*\}} Q_i^n(s) |M_i^n(\check{X}^n(s))| \\ &\quad + |M_{j_l^*}^n(\check{X}^n(s))| \left| (e \cdot \check{X}^n(s))^+ - \sum_{i \in \mathcal{J}(\tau_l^n) \cup \{j_l^*\}} Q_i^n(s) \right| \\ &\quad + \sum_{i \in \mathcal{J}(\tau_l^n)} Q_i^n(s) |M_i^n(\check{X}^n(s)) - M_{j_l^*}^n(\check{X}^n(s))|. \end{aligned}$$

We turn to bound each of the elements on the right-hand side of (104). First, note that for all $i \notin \mathcal{J}(\tau_i^n) \cup \{j_i^*\}$ it follows from (103) that

$$\sup_{\tau_i^n \leq s < (\tau_i^n + \eta^n) \wedge \tau_{i+1}^n} Q_i^n(s) \leq 12K \log n.$$

Also, by (36) we have for all $i \in \mathcal{I}$ that

$$(105) \quad \sup_{0 \leq s \leq \tau_{\kappa', T}^n} |M_i^n(\check{X}^n(s))| \leq C \log^{k_1} n,$$

so that

$$(106) \quad \sum_{i \notin \mathcal{J}(\tau_i^n) \cup \{j_i^*\}} Q_i^n(s) |M_i^n(\check{X}^n(s))| \leq 12IKC \log^{k_1+1} n$$

for all $s \in [\tau_i^n, (\tau_i^n + \eta^n) \wedge \tau_{i+1}^n]$ and a constant C that does not depend on n . From (103) and from the fact that $\sum_{i \in \mathcal{I}} Q_i^n(s) = (e \cdot \check{X}^n(s))^+$ we similarly have that

$$(107) \quad |M_{j_i^*}^n(\check{X}^n(s))| \left| (e \cdot \check{X}^n(s))^+ - \sum_{i \in \mathcal{J}(\tau_i^n) \cup \{j_i^*\}} Q_i^n(s) \right| \leq 12ICK \log^{k_1+1} n.$$

To bound the last element on the right-hand side of (104) note that for each $i \in \mathcal{J}(\tau_i^n)$ there exists $\tau_i^n - \eta^n \leq t \leq \tau_i^n$ such that $h_j^n(\check{X}^n(t)) = 1$. Otherwise, we would have a contradiction to (102). We now claim that for each $i \in \mathcal{J}(\tau_i^n)$,

$$(108) \quad |M_i^n(\check{X}^n(s)) - M_{j_i^*}^n(\check{X}^n(s))| \leq \frac{C \log^{k_1+2} n}{\sqrt{n}}$$

for all s in $[\tau_i^n - \eta^n, \tau_i^n + \eta^n]$. Indeed, by the definition of $\tilde{\Omega}(K)$, we have that $|\check{X}^n(t) - \check{X}^n(s)| \leq C \log^{m+2} n$ for all s, t in $[\tau_i^n - \eta^n, \tau_i^n + \eta^n]$. As in the proof of (37) [see, e.g., (51)] we have that

$$(109) \quad |M_i^n(x) - M_i^n(y)| \leq \frac{C \log^{k_2+m+2} n}{\sqrt{n}}, \quad i \in \mathcal{I},$$

for $x, y \in \mathcal{B}_{\kappa'}^n$, with $|x - y| \leq C \log^{m+2} n$. In turn,

$$(110) \quad |M_i^n(\check{X}^n(t)) - M_i^n(\check{X}^n(s))| \leq \frac{C \log^{k_2+m+2} n}{\sqrt{n}} = \frac{C \log^{k_1+2} n}{\sqrt{n}}$$

for all $i \in \mathcal{I}$ and all $s, t \in [\tau_i^n - \eta^n, \tau_i^n + \eta^n]$ where we used the fact that $k_1 = k_2 + m$. Since, for each $j \in \mathcal{J}(\tau_i^n)$, there exists $\tau_i^n - \eta^n \leq t \leq \tau_i^n$ such that $h_j^n(\check{X}^n(t)) = 1$ we have, by the definition of h^n that $j \in \arg \min_{i \in \mathcal{I}} M_i^n(\check{X}^n(t))$ for such t so that (108) now follows from (110). Finally, recall that $\sum_{i \in \mathcal{I}} Q_i^n(t) = (e \cdot \check{X}^n(s))^+ \leq \kappa \sqrt{n} \log^m n$ for all $s \leq \tau_{\kappa', T}^n$ and that $k_0 = k_1 + m$ so that by (108)

$$\sum_{i \in \mathcal{J}(\tau_i^n)} Q_i^n(s) |M_i^n(\check{X}^n(s)) - M_{j_i^*}^n(\check{X}^n(s))| \leq C \log^{k_0+2} n.$$

Plugging this into (104) together with (106) and (107) we then have that, on $\tilde{\Omega}(K)$,

$$\sup_{\tau_{l-1}^n \leq s < (\tau_{l-1}^n + \eta^n) \wedge \tau_l^n} |\check{\psi}^n(s)| \leq C \log^{k_1+m+2} n = CK \log^{k_0+2} n.$$

This argument is repeated for each l . To complete the proof of (96) note that, using (105) together with $\sup_{0 \leq s \leq \tau_{k',T}^n} |e \cdot \check{X}^n(s)| \leq \kappa \sqrt{n} \log^m n$, we have that $\sup_{0 \leq s \leq \tau_{k',T}^n} |\check{\psi}^n(s)| \leq C \sqrt{n} \log^{k_1+m} n$. Applying Hölder's inequality we have that

$$\begin{aligned} & \mathbb{E} \left[\sup_{1 \leq l \leq r^n + 1} \sup_{\tau_{l-1}^n \leq s < (\tau_{l-1}^n + \eta^n)} |\check{\psi}^n(s)| \right] \\ & \leq \mathbb{E} \left[\sup_{1 \leq l \leq r^n + 1} \sup_{\tau_{l-1}^n \leq s < (\tau_{l-1}^n + \eta^n)} |\check{\psi}^n(s)| 1_{\{\tilde{\Omega}(K)\}} \right] \\ & \quad + \mathbb{E} \left[\max_{1 \leq l \leq r^n + 1} \sup_{\tau_{l-1}^n \leq s < (\tau_{l-1}^n + \eta^n)} |\check{\psi}^n(s)| 1_{\{\tilde{\Omega}(K)^c\}} \right] \\ & \leq C \log^{k_0+2} n + C \sqrt{n} \log^{k_1+m} n C_1 e^{-(C_2 K/2) \log n} \end{aligned}$$

for redefined constants C_1, C_2 and (96) now follows by choosing K large enough.

We turn to prove (97). Rearranging terms in (93) we write

$$\check{\psi}^n(s) = \sum_{i \in \mathcal{I}} M_i^n(\check{X}^n(s)) ((e \cdot \check{X}^n(s))^+ h_i^n(\check{X}^n(s)) - (e \cdot \check{X}^n(s))^+ U_i^n(s)),$$

so that equation (97) now follows directly from (102) and (105) through an application of Hölder's inequality.

Finally, to establish (98), note that from the definition of $\tau_{k',T}^n$,

$$\begin{aligned} \sup_{0 \leq t \leq \eta^n \wedge \tau_{k',T}^n} |\check{\psi}^n(s)| & \leq I \sup_{0 \leq t \leq \eta^n \wedge \tau_{k',T}^n} |\check{X}^n(t)| \sum_{i \in \mathcal{I}} M_i^n(\check{X}^n(t)) \\ & \leq I \sup_{0 \leq t \leq \eta^n \wedge \tau_{k',T}^n} C \log^{k_1} n |\check{X}^n(t)| \leq C \kappa \sqrt{n} \log^{k_1+m} n. \end{aligned}$$

In turn,

$$(111) \quad \mathbb{E} \left[\int_0^{\tau_0^n} e^{-\gamma t} |\check{\psi}^n(t)| dt \right] \leq C \log^{k_1+m} n = C \log^{k_0} n.$$

We have thus proved (96)–(98) and to conclude the proof of the theorem it remains only to establish (102). To that end, let $\check{\zeta}_i^n(t)$, $\tilde{\zeta}_i^n(t)$ and $\hat{\zeta}_i^n(t)$ be as in (99)–(101). Fix an interval $[l, s] \in (\hat{\zeta}_i^n(t), \tilde{\zeta}_i^n(t))$ such that $Q_i^n(u) > 2K \log n$ for all $u \in [l, s]$. By the definition of the tracking policy, (86) holds on this interval so that, on $\omega \in \tilde{\Omega}(K)$,

$$\begin{aligned} (112) \quad Q_i^n(l) - Q_i^n(s) & \leq \bar{\epsilon} n(t-s)n - \frac{\epsilon_i}{2} n(t-s) + 2K \log n \\ & \leq -\frac{\epsilon_i}{4} n(t-s) + 2K \log n. \end{aligned}$$

Equation (102) now follows directly from (112). Indeed, note for all $t \leq \tau_{\kappa', T}$, $Q_i^n(t) \leq (e \cdot \check{X}^n(t))^+ \leq |\check{X}^n(t)| \leq \kappa \sqrt{n} \log^m n$. Hence, $Q_i^n(\check{\zeta}_i^n(t)) \leq \kappa \sqrt{n} \log^m n$. In turn, using (112) and assuming that $\check{\zeta}_i^n(t) \geq \zeta_i^n(t) + \eta^n$ we have that $Q_i^n(\zeta_{0,i}^n(t)) \leq 4K \log n$ for some time $\zeta_{0,i}^n(t) \leq \check{\zeta}_i^n(t) + \eta^n$ with η^n as defined in (95). Also, let

$$\zeta_{2,i}^n(t) = \inf\{t \geq \zeta_{0,i}^n(t) : Q_i^n(t) \geq 12K \log n\}$$

and

$$\zeta_{1,i}^n(t) = \sup\{t \leq \zeta_{2,i}^n(t) : Q_i^n(t) \leq 8K \log n\}.$$

Note that (112) applies to any subinterval $[l, s]$ of $[\zeta_{1,i}^n(t), \zeta_{2,i}^n(t)]$. In turn, $\zeta_{2,i}^n(t) \leq \check{\zeta}_i^n(t)$ would constitute a contradiction to (112) so that we must have that $Q_i^n(s) \leq 12K \log n$ for all $s \in [\zeta_{0,i}^n(t), \check{\zeta}_i^n(t)]$ with $\zeta_{0,i}^n(t) \leq \check{\zeta}_i^n(t) + \eta^n$. Finally, note that $\zeta_{0,i}^n(t)$ can be taken to be t if $Q_i^n(t) \leq 4K \log n$.

This concludes the proof of (102) and, in turn, the proof of the theorem.

Proof of Lemma 6.2. Let $T, \tau_{\kappa', T}^n$ and (x^n, q^n) be as in the statement of the lemma. We first prove (59). To that end, we claim that, for all T large enough,

$$(113) \quad \mathbb{E}_{x^n, q^n} \left[\int_{T \log n}^{\infty} e^{-\gamma s} (e \cdot c)(e \cdot \check{X}^n(s))^+ ds \right] \leq C \log^2 n$$

for some $C > 0$ and all $n \in \mathbb{Z}$. This is a direct consequence of Lemma 3 in [2] that, in our notation, guarantees that

$$\mathbb{E}_{x^n, q^n} [|\check{X}^n(t)|] \leq C(1 + |x^n| + \sqrt{n}(t + t^2))$$

for all $t \geq 0$ and some constant $C > 0$. We use (113) to prove Lemma 6.2. The assertion of the lemma will be established by showing that

$$\mathbb{E}_{x^n, q^n} \left[\int_{\tau_{\kappa', T}^n}^{2T \log n} e^{-\gamma s} (e \cdot c)(e \cdot \check{X}^n(s))^+ ds \right] \leq C \log^2 n.$$

To that end, applying Hölder's inequality, we have

$$(114) \quad \begin{aligned} & \mathbb{E}_{x^n, q^n} \left[\int_{\tau_{\kappa', T}^n}^{2T \log n} e^{-\gamma s} (e \cdot c)(e \cdot \check{X}^n(s))^+ ds \right] \\ & \leq \mathbb{E}_{x^n, q^n} \left[(2T \log n - \tau_{\kappa', T}^n)^+ \sup_{0 \leq t \leq 2T \log n} (e \cdot c)(e \cdot \check{X}^n(t))^+ ds \right] \\ & \leq \sqrt{\mathbb{E}_{x^n, q^n} [(2T \log n - \tau_{\kappa', T}^n)^+]} \\ & \quad \times \sqrt{\mathbb{E}_{x^n, q^n} \left[\left(\sup_{0 \leq t \leq 2T \log n} (e \cdot c)(e \cdot \check{X}^n(t))^+ ds \right)^2 \right]}. \end{aligned}$$

Using Lemma A.1 we have that

$$(115) \quad \mathbb{E}_{x^n, q^n} \left[\left(\sup_{0 \leq t \leq 2T \log n} (e \cdot c)(e \cdot \check{X}^n(t))^+ ds \right)^2 \right] \leq Cn \log^6 n$$

for some $C > 0$ (that can depend on T). Also, since $m \geq 3$,

$$\mathbb{P}\{\tau_{\kappa', T}^n < 2T \log n\} \leq \mathbb{P}\left\{ \sup_{0 \leq t \leq 2T \log n} |\check{X}^n(t)| > \kappa' \sqrt{n} \log^3 n - M \sqrt{n} \right\}.$$

Choosing κ' (and in turn κ large enough) we then have, using Lemma A.1, that

$$(116) \quad \mathbb{P}\{\tau_{\kappa', T}^n < 2T \log n\} \leq \frac{C}{n^2}$$

and hence, that

$$(117) \quad \mathbb{E}_{x^n, q^n} \left[((2T \log n - \tau_{\kappa', T}^n)^+)^2 \right] \leq C.$$

Plugging (115) and (117) into (114) we then have that

$$(118) \quad \mathbb{E}_{x^n, q^n} \left[\int_{\tau_{\kappa', T}^n}^{2T \log n} e^{-\gamma s} (e \cdot c)(e \cdot \check{X}^n(s))^+ ds \right] \leq C \log^2 n.$$

To conclude the proof we will show that (60) follows from our analysis thus far. Indeed,

$$\begin{aligned} & \mathbb{E}[e^{-\gamma \tau_{\kappa', T}^n} \phi_{\kappa}^n(\check{X}^n(\tau_{\kappa', T}^n))] \\ & \leq \mathbb{E}_{x^n, q^n}^U \left[\int_{\tau_{\kappa', T}^n}^{2T \log n} e^{-\gamma s} \sup_{0 \leq s \leq 2T \log n} (e \cdot c)(e \cdot \check{X}^n(s))^+ ds \right]. \end{aligned}$$

The right-hand side here is bounded by $C \log^2 n$ by the same argument that leads to (118).

Proof of Lemma 6.3. Recall that \check{W}^n is defined by $\check{W}_i^n(t) = M_{i,1}^n(t) - M_{i,2}^n(t)$, where

$$\begin{aligned} M_{i,1}^n(t) &= \mathcal{N}_i^a(\lambda_i^n t) - \lambda_i^n t, \\ M_{i,2}^n(t) &= \mathcal{N}_i^d \left(\mu_i \int_0^t (\check{X}_i^n(s) + v_i n - U_i^n(s)(e \cdot \check{X}^n(s))^+) ds \right) \\ &\quad - \mu_i \int_0^t (\check{X}_i^n(s) - U_i^n(s)(e \cdot \check{X}^n(s))^+) ds. \end{aligned}$$

The fact that each of the processes $M_{i,1}^n(t)$ and $M_{i,2}^n(t)$ are square integrable martingales with respect to the filtration (\mathcal{F}_t^n) follows as in Section 3 of [12] and specifically as in Lemma 3.2 there.

Since, with probability 1, there are no simultaneous jumps of \mathcal{N}_i^a and \mathcal{N}_i^d , the quadratic variation process satisfies

$$\begin{aligned} [\check{W}_i^n]_t &= [M_{i,1}^n]_t + [M_{i,2}^n]_t \\ &= \sum_{s \leq t} (\Delta M_{i,1}^n(s))^2 + \sum_{s \leq t} (\Delta M_{i,2}^n(s))^2, \end{aligned}$$

where the last equality follows again from Lemma 3.1 in [12] (see also Example 5.65 in [14]). Finally, the predictable quadratic variation process satisfies

$$\begin{aligned} \langle \check{W}_i^n \rangle_t &= \langle M_{i,1}^n \rangle_t + \langle M_{i,2}^n \rangle_t \\ &= \lambda_i^n t + \mu_i \int_0^t (\check{X}_i^n(s) + v_i n - U_i^n(s)(e \cdot \check{X}_i^n(s))^+) ds \\ &= \int_0^t (\sigma_i^n(\check{X}_i^n(s), U_i^n(s)))^2 ds, \end{aligned}$$

where the second equality follow again follows from Lemma 3.1 in [12] and the last equality from the definition of $\sigma_i^n(\cdot, \cdot)$ [see (25)]. By Theorem 3.2 in [12] ($(\check{W}_i^n(t))^2 - [\check{W}_i^n]_t, t \geq 0$) and $(\check{W}_i^n(t))^2 - [M_{i,1}^n]_t, t \geq 0$) are both martingales with respect to (\mathcal{F}_i^n) . In turn, by the optional stopping theorem so are the processes $\mathcal{M}_i^n(\cdot)$ and $\mathcal{V}_i^n(\cdot)$ as defined in the statement of the lemma. Finally, it is easy to verify that these are square integrable martingales using the fact the time changes are bounded for all finite t .

REFERENCES

- [1] ATAR, R. (2005). Scheduling control for queueing systems with many servers: Asymptotic optimality in heavy traffic. *Ann. Appl. Probab.* **15** 2606–2650. [MR2187306](#)
- [2] ATAR, R., MANDELBAUM, A. and REIMAN, M. I. (2004). Scheduling a multi class queue with many exponential servers: asymptotic optimality in heavy traffic. *Ann. Appl. Probab.* **14** 1084–1134. [MR2071417](#)
- [3] BORKAR, V. S. (2005). Controlled diffusion processes. *Probab. Surv.* **2** 213–244. [MR2178045](#)
- [4] CSÖRGO, M. and HORVÁTH, L. (1993). *Weighted Approximations in Probability and Statistics*. Wiley, Chichester. [MR1215046](#)
- [5] FLEMING, W. H. and SONER, H. M. (2006). *Controlled Markov Processes and Viscosity Solutions*. Springer, New York. [MR2179357](#)
- [6] GILBARG, D. and TRUDINGER, N. (2001). *Elliptic Partial Differential Equations of Second Order*. Springer, Berlin. [MR1814364](#)
- [7] GLYNN, P. W. (1987). Upper bounds on Poisson tail probabilities. *Oper. Res. Lett.* **6** 9–14. [MR0891601](#)
- [8] HALFIN, S. and WHITT, W. (1981). Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* **29** 567–588. [MR0629195](#)
- [9] KARATZAS, I. and SHREVE, S. (1991). *Brownian Motion and Stochastic Calculus*, 2nd ed. Springer, New York. [MR1121940](#)
- [10] KRYLOV, N. (2008). *Controlled Diffusion Processes*. Springer, New York. [MR2723141](#)
- [11] MANDELBAUM, A., MASSEY, W. and REIMAN, M. (1998). Strong approximations for Markovian service networks. *Queueing Syst.* **30** 149–201. [MR1663767](#)

- [12] PANG, G., TALREJA, R. and WHITT, W. (2007). Martingale proofs of many-server heavy-traffic limits for Markovian queues. *Probab. Surv.* **4** 193–267. [MR2368951](#)
- [13] TRUDINGER, N. S. (1983). Fully nonlinear, uniformly elliptic equations under natural structure conditions. *Trans. Amer. Math. Soc.* **278** 751–769. [MR0701522](#)
- [14] VAN DER VAART, A. W. (2006). Martingales, diffusions and financial mathematics. Lecture notes. Available at <http://www.math.vu.nl/sto/onderwijs/mdfm/>.

KELLOGG SCHOOL OF MANAGEMENT
NORTHWESTERN UNIVERSITY
EVANSTON, ILLINOIS 60208
USA
E-MAIL: b-ata@kellogg.northwestern.edu
i-gurvich@kellogg.northwestern.edu