# Finite-Time Minimax Bounds and an Optimal Lyapunov Policy in Queueing Control

Yujie Liu   Vincent Y. F. Tan   Yunbei Xu | National University of Singapore

Paper: https://arxiv.org/pdf/2506.18278 | Code: github.com/LYJ-aa/lyapopt-queueing-code

**Poster message**
Short horizons need their own theory: LyapOpt can match minimax finite-time rates where MaxWeight can be too coarse.

finite-time guarantees | minimax lower bounds | full Lyapunov drift | MaxWeight gap | nonstationary queues

## What problem is the paper solving?

**Classical queueing theory often optimizes stability or diffusion limits.** This paper asks a more operational question:

**At a fixed horizon $T$, how small can the worst-case expected backlog be?**

The answer is a **minimax benchmark** over capacity-constrained systems, plus a simple policy that targets it directly.

- Lower bound for *any* scheduling policy.
- New policy LyapOpt: minimizes a one-step quadratic Lyapunov lookahead.
- Separation: MaxWeight optimizes only first-order drift and can be suboptimal in finite time.

## Queueing model and minimax class

**State dynamics.** For $n$ queues, schedule $d(t) \in D_t$, realized departures $D(t)$, arrivals $A(t)$:

$$Q(t+1) = \max\{Q(t) - D(t), 0\} + A(t+1).$$

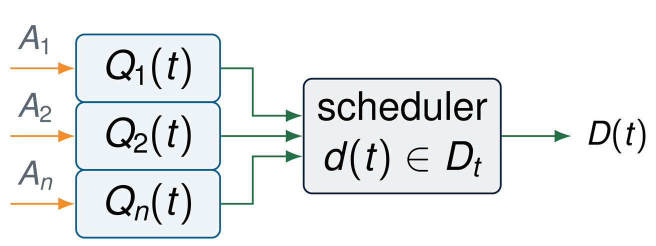


**Capacity region.** For $U \subset \mathbb{R}^n_+$, $\Pi(U) = \{\gamma \geq 0 : \exists u \in \mathrm{conv}(U),\ \gamma \leq u\}$. Arrivals are feasible when $\lambda(t) \in \rho\Pi(D_t)$.
The model class $\mathcal{M}^\rho(B, C_a, C_d)$ limits capacity and variability:

$$\frac{1}{n}\sum_i d_i^2 \leq B^2, \quad \frac{1}{n}\sum_i \mathrm{Var}(A_i(t)) \leq C_a^2, \quad \frac{1}{n}\sum_i \mathrm{Var}(S_d(t)_i) \leq C_d^2.$$

**Minimax risk:**

$$\Gamma^\rho(B, C_a, C_d, T) = \inf_{\Phi} \sup_{(D,A,S)\in\mathcal{M}^\rho} \mathbb{E}^{\Phi}_{A,S}\left[\sum_i Q_i(T)\right].$$

## Main lower bound: what no policy can beat

For universal $c > 0$ and sufficiently large $T$,

$$\Gamma^\rho(B, C_a, C_d, T) \geq c \min\left\{ n\sqrt{(C_a^2 + C_d^2)(T-1)},\ \frac{n(C_a^2 + C_d^2)}{B(1-\rho)} \right\} + n\rho B.$$

- **Heavy traffic** $\rho \uparrow 1$: fluctuations force $\Omega(n\sqrt{(C_a^2 + C_d^2)T})$ backlog.
- **Interior regime** $\rho < 1$: congestion plateaus at order $(1-\rho)^{-1}$.
- Proof reduces total backlog to the maximum of a drifted random walk, then couples it to Brownian motion.

## Regime summary at a glance

| regime | scaling | interpretation |
|---|---|---|
| **heavy traffic** | $n\sqrt{(C_a^2 + C_d^2)T}$ | horizon-length fluctuations dominate backlog |
| **interior** | $\frac{n(C_a^2 + C_d^2)}{B(1-\rho)}$ | negative drift eventually caps congestion |

Main message: finite-horizon queueing performance depends on *both* stochastic variability and distance to the boundary, and the active term changes with $(T, \rho)$.

## LyapOpt: optimize the full one-step drift

**Policy definition.** At time $t$, choose the schedule that minimizes the next pre-arrival squared queue length:

$$d^{\text{LyapOpt}}(t) \in \arg\min_{d\in D_t} \sum_{i=1}^{n} \left(\max\{Q_i(t) - d_i, 0\}\right)^2.$$

$Q(t)$ → solve convex / finite optimization over $D_t$ → $d^{\text{LyapOpt}}(t)$ → $Q(t+1)$

minimize quadratic lookahead, not just a linear score

The derandomized quadratic drift separates into

$$f(Q(t), d) = \underbrace{2\sum_i Q_i(t)(\lambda_i(t) - d_i)}_{\text{first-order drift}} + \underbrace{\sum_i (d_i^2 - \lambda_i(t)^2)}_{\text{second-order geometry}}.$$

MaxWeight: linear first-order term | LyapOpt: full quadratic drift

## Finite-time guarantee for LyapOpt

For $\rho \in (0, 1]$, the paper proves

$$\mathbb{E}^{\text{LyapOpt}}\left[\sum_i Q_i(T)\right] \leq n\sqrt{\sum_{t=1}^{T-1} \frac{1}{n}\mathbb{E}\left[\min_{d\in D_t} f(Q(t), d)\right] + (C_a^2 + C_d^2)(T-1)} + \sum_i \mathbb{E}[A_i(T)].$$

When $\lambda(t) \in \mathrm{star}(D_t)$,

$$\mathbb{E}^{\text{LyapOpt}}\left[\sum_i Q_i(T)\right] \leq n\sqrt{(C_a^2 + C_d^2)(T-1)} + \sum_i \mathbb{E}[A_i(T)].$$

which matches the heavy-traffic minimax rate up to constants.

## Proof architecture

1. Minimax hard queueing instance → 2. Total backlog $\geq$ net-input max → 3. Drifted random walk lower bound
4. Brownian coupling transfers $\sqrt{T}$ / $(1-\rho)^{-1}$ scaling

The key inequality is a one-dimensional reduction:

$$\sum_i Q_i(T) \geq \max_{1\leq k\leq T-1} \sum_{t=k}^{T-1} \left(\sum_i A_i(t) - \sum_i D_i(t)\right) + \sum_i A_i(T).$$

Then the maximum of a negatively drifted process behaves like

$$\mathbb{E}\left[\sup_{0\leq s\leq t} (\sigma W_s - \theta s)\right] \asymp \min\{\sigma\sqrt{t}, \sigma^2/\theta\}.$$

## Stability is retained

For $\rho < 1$, LyapOpt also inherits classical Lyapunov-drift stability:

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}^{\text{LyapOpt}}\left[\sum_i Q_i(t)\right] \leq \frac{n^2(B^2 + C_a^2 + C_d^2)}{2(1-\rho)\alpha}.$$

Headline: **short-term minimax optimality** plus **stability**.

## Why MaxWeight can lose in finite time

MaxWeight chooses an extreme schedule by maximizing a linear backlog score:

$$d^{\text{MaxWeight}}(t) \in \arg\max_{d\in D_t} \langle Q(t), d\rangle.$$

Its finite-time upper bound contains the capacity term $B$:

$$\mathbb{E}^{\text{MaxWeight}}\left[\sum_i Q_i(T)\right] \leq n\sqrt{(B^2 + C_a^2 + C_d^2)(T-1)} + \sum_i \mathbb{E}[A_i(T)].$$

For a two-queue hard family, the paper proves a matching obstruction in heavy traffic:

$$\liminf_{\rho\uparrow 1} \sup_{\mathcal{M}^\rho} \mathbb{E}^{\text{MaxWeight}}\left[\sum_i Q_i(T)\right] \geq c\sqrt{(C_a^2 + C_d^2)(T-1)} + \frac{BT^{1/3}}{4\sqrt{3}}.$$

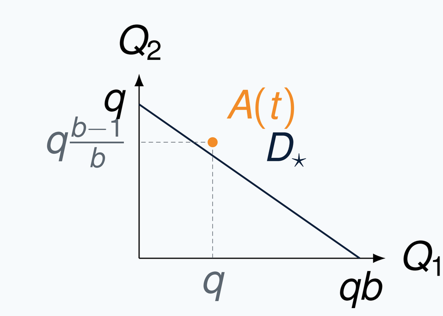


Asymmetric service line segment

- Fast option for $Q_1$: service $(b, 0)$.
- Slow option for $Q_2$: service $(0, 1)$.
- MaxWeight over-prioritizes the fast queue by chasing extreme points.
- LyapOpt tracks the arrival-rate geometry and avoids deterministic backlog.

## Empirical picture

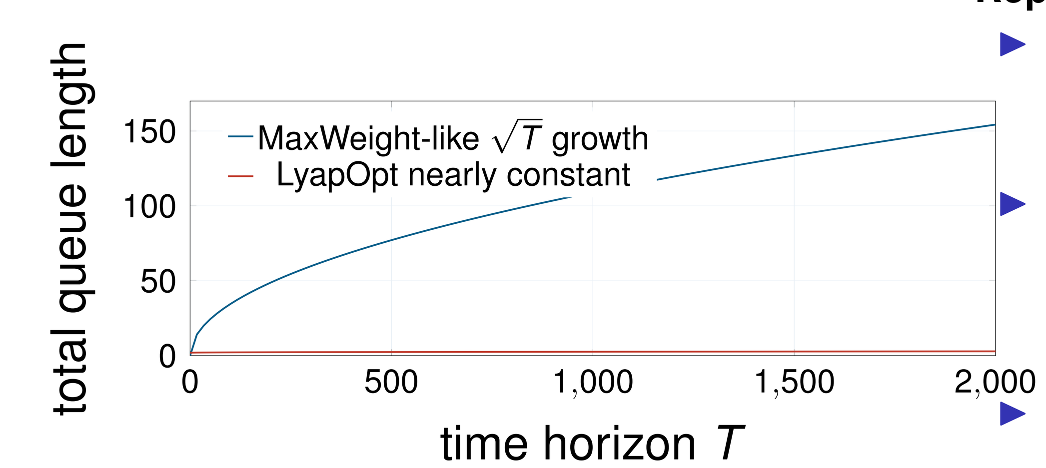


Schematic of the deterministic gap in the hard instance.

**Reported simulation trends.**

- Hard instance: LyapOpt stays far below MaxWeight as $T$ or $B$ grows.
- General random instances with $n = 8$: LyapOpt has lower total backlog and lower queue imbalance.
- In the interior regime $\rho = 0.9$, both stabilize quickly, but LyapOpt stabilizes at a lower backlog.

**Random-instance study at $t = 1000$: ratio = LyapOpt expected backlog / MaxWeight expected backlog.**

| queues $n$ | ratio $\leq 1$ | ratio $\leq 0.9$ | ratio $\leq 0.8$ |
|---|---|---|---|
| 2 | 63.30% | 9.90% | 0.10% |
| 3 | 97.30% | 78.75% | 38.30% |
| 4 | 99.70% | 95.60% | 85.85% |
| 8 | 100.00% | 99.50% | 92.00% |

## Takeaways for conference discussion

1. **Finite time is not just steady state before convergence.** It exposes parameter dependences hidden by diffusion limits.
2. **The second-order drift term matters.** Geometry of $D_t$ can determine backlog before asymptotics wash it out.
3. **LyapOpt is simple:** choose the schedule that minimizes the post-service squared queue vector.
4. **Open directions:** computational oracles, sharper constants, and learning variants with unknown statistics.

Source: Liu, Tan, Xu, *Finite-Time Minimax Bounds and an Optimal Lyapunov Policy in Queueing Control*, arXiv:2506.18278v3.