Lyapunov Optimization: An Introduction

Hai-Liang Zhao
hliangzhao97@gmail.com

January 25, 2019

Outline

Outline

Outline

Outline

Outline

A Brief Introduction

Lyapunov optimization refers to the use of a Lyapunov function to optimally control a dynamical system.

Lyapunov function

Lyapunov functions are used extensively in control theory to ensure different forms of **system stability**. The state of a system at a particular time is often described by a multi-dimensional vector. A Lyapunov function is **a nonnegative scalar measure** of this multi-dimensional state. Typically, the function is defined to grow large when the system moves towards undesirable states. System stability is achieved by taking control actions that make **the Lyapunov function drift** in the negative direction towards zero.

Lyapunov drift (the Lyapunov function drift)

Adding a weighted penalty term to the Lyapunov drift and **minimizing the sum** leads to the drift-plus-penalty algorithm for joint network stability and penalty minimization.

Outline

(Continuous) Optimization Problems

Optimization Problem

The standard form of a continuous optimization problem is

$$\begin{aligned} \mathcal{P}_1 : \quad & \min_{\mathbf{x}\in\mathbb{R}^n} f(\mathbf{x}) \\ s.t. \quad & c_i(\mathbf{x}) \leq 0, i = 1, 2, ..., k \\ & h_j(\mathbf{x}) = 0, j = 1, 2, ..., l, \end{aligned}$$

where

1. $f : \mathbb{R}^n \to \mathbb{R}$ is the objective function to be minimized over the n-variable vector $\mathbf{x}$;
2. $c_i(\mathbf{x}) \leq 0$ are called inequality constraints;
3. $h_j(\mathbf{x}) = 0$ are called equality constraints, and
4. $k, l \geq 0$.

Stochastic Optimization Problems

Random events happen in every time slot t for $t \in \mathcal{T}$.
In every time slot,

1. $\mathbf{w}(t) \triangleq [w_1(t), w_2(t), ..., w_n(t)] \in \Omega^n$ are the **i.i.d.** random events;
2. $\boldsymbol{\alpha}(t) \triangleq [\alpha_1(t), \alpha_2(t), ..., \alpha_m(t)] \in \mathcal{A}^m$ are the control actions.

According to the random events $\mathbf{w}(t)$ already happened, the system (decision maker) takes control actions $\boldsymbol{\alpha}(t)$ in a certain way. Therefore, for the optimization goal $p(t)$, we have

$$p(t) = P(\mathbf{w}(t), \boldsymbol{\alpha}(t)), \quad (1)$$

where $P(\cdot)$ is a **certain** function (e.g., the way we calculate latency).
Besides the optimization goal $p(t)$, other variables in the system who can affect the optimality $y_k(t), k \in \mathcal{K}$ (power consumption, avaliable bandwidth, etc.) are also impacted by the taken control actions. Thus we have

$$y_k(t) = Y_k(\mathbf{w}(t), \boldsymbol{\alpha}(t)), k \in \{1, ..., K\}, \quad (2)$$

where $\forall k \in \{1, ..., K\}, Y_k(\cdot)$ are all **certain** functions.

Stochastic Optimization Problems

Stochastic Optimization Problem

Minimize a **time average** optimization goal under serveral constraints in a time horizon which is slotted. The independent variables are the control actions to cope with the random events.

$$\mathcal{P}_2 : \min_{\forall t, \boldsymbol{\alpha}(t) \in \mathcal{A}^m} \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[p(t)] \tag{3}$$

$$s.t. \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[y_k(t)] \leq 0, k \in \{1, ..., K\}, \tag{4}$$

where $p(t)$ is obtained by (1), $y_k(t)$ is obtained by (2).

1. Why **Expectation** $\mathbb{E}[\cdot]$?
2. Why time average?

Virtual Queues

For every constraints $y_k(t), k \in \mathcal{K}$ define a virtual queue with initial backlog 0:

$$Q_k(t+1) = \max\{Q_k(t) + y_k(t), 0\}, k \in \{1, ..., K\}. \quad (5)$$

How to control queues to ensure (2) always stand up?

According to (5), we have $y_k(t) \leq Q_k(t+1) - Q_k(t)$. Then we calculate the sum on time slots and take the Expectation:

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}[y_k(t)] \leq \frac{\mathbb{E}[Q_k(T)]}{T}, k \in \{1, ..., K\}. \quad (6)$$

If the following constraint set up:

$$\lim_{T\to\infty} \frac{\mathbb{E}[Q_k(T)]}{T} = 0, k \in \{1, ..., K\}, \quad (7)$$

The constraint (2) always set up.

The Derivate Stochastic Optimization Problem

The derivate problem

The new problem derivated from $\mathcal{P}_2$ and (7) is defined as

$$\mathcal{P}_3: \min_{\forall t, \boldsymbol{\alpha}(t) \in \mathcal{A}^m} \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[p(t)] \tag{8}$$

$$s.t. \quad (7).$$

Actually, a queue $Q_k(t)$ that satisfies the above limit equation is said to be **mean rate stable**.

Other degree of Stability:

$$\exists \delta \leq 0, \lim_{T \to \infty} \frac{\mathbb{E}[Q_k(T)]}{T} \leq \delta, k \in \{1, ..., K\}. \tag{9}$$

How to solve $\mathcal{P}_3$ with asymtotically optimality?

Lyapunov function

A Lyapunov function is **a nonnegative scalar measure** of this multi-dimensional state. Denote the queue vector as $\boldsymbol{\Theta}(t) \triangleq [Q_1(t), ..., Q_K(t)]$, the Lyapunov function is defined by

$$L(\boldsymbol{\Theta}(t)) \triangleq \frac{1}{2}\sum_{k=1}^{K} Q_k(t)^2. \quad (10)$$

Lyapunov drift is defined as $\Delta(\boldsymbol{\Theta}(t)) \triangleq L(\boldsymbol{\Theta}(t+1)) - L(\boldsymbol{\Theta}(t))$.

Because of (5), $Q_k(t+1)^2 \leq (Q_k(t) + y_k(t))^2, k \in \{1, ..., K\}$. Calculate the sum of all queues' backlog, we can obtain that

$$\begin{aligned}\Delta(\boldsymbol{\Theta}(t)) &= \frac{1}{2}\sum_{k=1}^{K} Q_k(t+1)^2 - \frac{1}{2}\sum_{k=1}^{K} Q_k(t)^2 \\ &\leq \frac{1}{2}\sum_{k=1}^{K} y_k(t)^2 + \sum_{k=1}^{K} Q_k(t)y_k(t) \leq B + \sum_{k=1}^{K} Q_k(t)y_k(t). \quad (11)\end{aligned}$$

Drift-Plus-Penalty Expression

In order to keep the queue stable, we have to minimize the optimization goal $p(t)$ and Lyapunov drift simultaneously, tuned by parameter V.

Drift-plus-penalty problem

$$\mathcal{P}_4 : \min_{\forall t, \boldsymbol{\alpha}(t) \in \mathcal{A}^m} \mathbb{E}[\Delta(\boldsymbol{\Theta}(t)) + V \cdot p(t) | \boldsymbol{\Theta}(t)] \tag{12}$$

$$s.t. \quad (7).$$

We cannot solve $\mathcal{P}_4$ in every independent time slot! Thus we have

The upper bound of drift-plus-penalty

$$\mathcal{P}_5 : \min_{\forall t, \boldsymbol{\alpha}(t) \in \mathcal{A}^m} \mathbb{E}[B + V \cdot p(t) + \sum_{k=1}^{K} Q_k(t) y_k(t) | \boldsymbol{\Theta}(t)] \tag{13}$$

$$s.t. \quad (7).$$

Drift-plus-penalty algorithm

In every time slot take the following procedure:

1. At the beginning of the tth time slot, obtain random events and queues' backlog: $\mathbf{w}(t), \mathbf{\Theta}(t)$.
2. Solve the following problem to obatin the optimal control action $\boldsymbol{\alpha}^*(t)$:

$$\boldsymbol{\alpha}^*(t) = \underset{\boldsymbol{\alpha}(t)\in\mathcal{A}^m}{\text{argmin}}\ \mathbb{E}[B + V \cdot p(t) + \sum_{k=1}^{K} Q_k(t) y_k(t) | \mathbf{\Theta}(t)]. \tag{14}$$

3. $\forall k \in \{1, ..., K\}$, according to (5) to update $Q_k(t)$.
4. $t \leftarrow t + 1$.

▷ Why constraint (7) is not considered in the algorithm?
▷ How to solve the problem in **Step. 2**?

Outline

Average Penalty Analysis

$\mathbf{w}$-only policy is a stationary and random method to obtain $\boldsymbol{\alpha}^{\dagger}(t)$ depend only on the observed $\mathbf{w}(t)$. i.e., for every possible random event $\mathbf{w}(t) \in \Omega^n$, $\mathbf{w}$-only policy decide $\boldsymbol{\alpha}^{\dagger}(t)$ according to a conditional probability distribution:

$$\boldsymbol{\alpha}^{\dagger}(t) = \underset{\boldsymbol{\alpha}(t) \in \mathcal{A}^m}{\operatorname{argmax}} \Pr(\boldsymbol{\alpha}(t) | \mathbf{w}(t)). \tag{15}$$

Optimal $\mathbf{w}$-only policy

Optimal $\mathbf{w}$-only policy is an $\mathbf{w}$-only policy under the following conditions:

$$P(\mathbf{w}(t), \boldsymbol{\alpha}^{\star}(t)) = p^{\star}, \tag{16}$$

$$Y_k(\mathbf{w}(t), \boldsymbol{\alpha}^{\star}(t)) \leq 0, k \in \{1, ..., K\}, \tag{17}$$

where

$$p^{\star} = \min \Big(\lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[p(t)] \Big). \tag{18}$$

Performance Analysis

Analysis on optimality and queue size (proof on link)

▷ Optimality gap: $O(\frac{1}{V})$

$$\sum_{t=0}^{T-1} \mathbb{E}[p(t)|\boldsymbol{\Theta}(t)] \le p^{\star} + \frac{B}{V}. \tag{19}$$

▷ Average queue size: $O(V)$

$$\begin{aligned} \frac{1}{T}\sum_{t=0}^{T-1}\sum_{k=1}^{K} \mathbb{E}[Q_k(t)|\boldsymbol{\Theta}(t)] &\le \frac{B'}{\epsilon} + \mathbb{E}[L(\boldsymbol{\Theta}(0))] - \mathbb{E}[L(\boldsymbol{\Theta}(T))] \\ &\le \frac{B + V(p_{max} - p_{min})}{\epsilon}. \end{aligned} \tag{20}$$

▷ Constraint (7) can always be satisfied:

$$\forall k \in \{1, ..., K\}, \lim_{T \to \infty} \frac{\mathbb{E}[Q_k(T)]}{T} = 0.$$

Outline

Conclusions

1. Lyapunov Optimization can only obtain asymtotically optimality, i.e. $O(\frac{1}{V})$.
2. A longer time horizon is needed to obtain a better solution.
3. Relation between the length of time horizon and the time average queue size?

Time for Case Study :-)