Estimating fluid queueing models using continuous-time hidden Markov models

Yi Zhern Yang

Supervised by Dr. Mark Fackrell

THE UNIVERSITY OF MELBOURNE

Introduction

Fluid queueing models are widely used to describe systems in which material, information or workload evolves continuously over time under stochastic control with applications in areas such as communication networks and cloud computing. In many practical settings, the underlying system dynamics are not directly observable, and only aggregated or indirect measurements are available. This project investigates the estimation of fluid queueing models by formulating them as continuous-time Markov models and developing an estimation framework based on discretisation and Expectation-Maximisation techniques.

Fluid queueing model (FQM)

Let $\{X(t), t \geq 0\}$ denote a *source process* modelled as a continuous-time Markov chain (CTMC) with finite state space

$$\mathcal{S} = \{1, 2, \cdots, N\}$$

and generator matrix $Q = (q_{ij})_{i,j \in \mathcal{S}}$.

Each state $i \in \mathcal{S}$ is associated with a net fluid flow rate r_i, forming the diagonal *flow matrix*

$$R = \text{diag}\,(r_1, r_2, \cdots, r_N)\,.$$

The resulting fluid dynamics are described by the *buffer content process* $\{C(t), t \geq 0\}$.

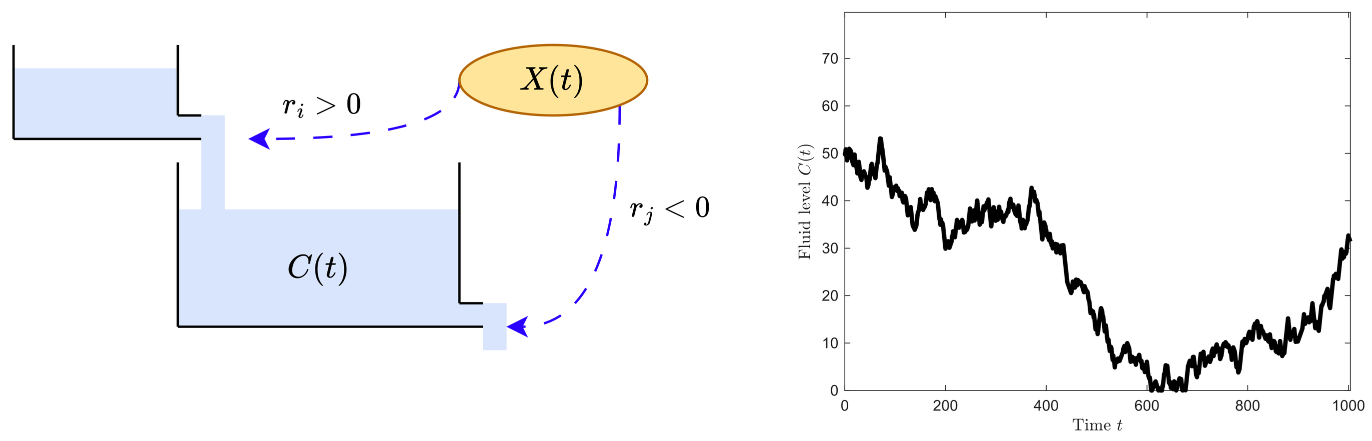

(a) Illustration of FQM (b) Simulation of FQM

Figure 1. Fluid Queueing Model

Research objective

To estimate the generator matrix Q of the underlying source process and the associated flow matrix R, given a finite set of fluid level observations $\{Y_t = C(t_k)\}_{k=1}^{M}$, recorded at discrete time points.

Model

Fluid queueing model as a continuous-time hidden Markov model

The source process evolves as a CTMC, while the fluid level depends on the current hidden state but does not itself form a Markov process. Observations are conditionally independent given the state of the source process, and the system evolves in continuous time. Consequently, the fluid queueing model can be formulated as a *continuous-time hidden Markov model (CT-HMM)*, where the source process is treated as the hidden state process and the fluid level as the observed process.

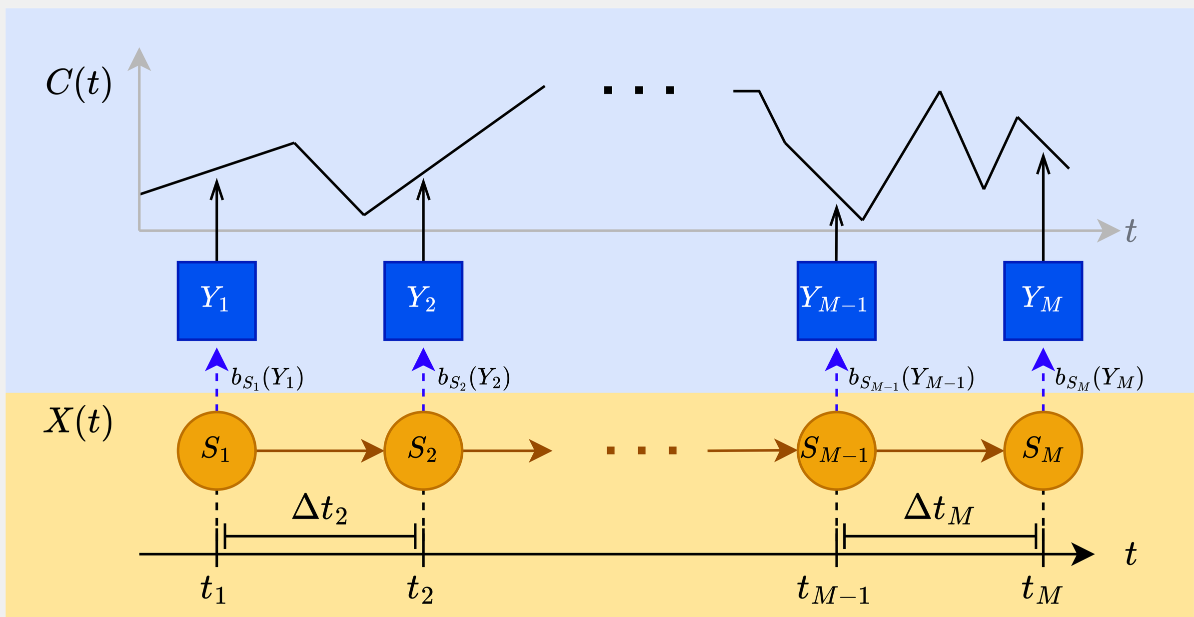

Figure 2. Illustration of a CT-HMM

Baum–Welch algorithm

When direct likelihood-based estimation is intractable due to incomplete observations, the *Expectation-Maximisation (EM) algorithm* provides a principled alternative. In Markovian settings, EM specialises to the *Baum-Welch algorithm*, which iteratively updates model parameters via an expectation step (E-step) and maximisation step (M-step) (Rabiner [2])

Forward-backward likelihood

The E-step relies on the forward-backward procedure, which computes posterior state and transition probabilities conditioned on the observations.

Initialisation

$$\alpha_1(i) = \pi_i b_i(Y_1), \quad i \in \mathcal{S} \qquad \beta_M(i) = 1, \quad i \in \mathcal{S}$$

Induction

$$\alpha_{t+1}(j) = \left[\sum_{i=1}^{N} \alpha_t(i) a_{ij}\right] b_j(Y_{t+1}) \qquad \beta_t(i) = \sum_{j=1}^{N} a_{ij} b_j(O_{t+1}) \beta_{t+1}(j)$$

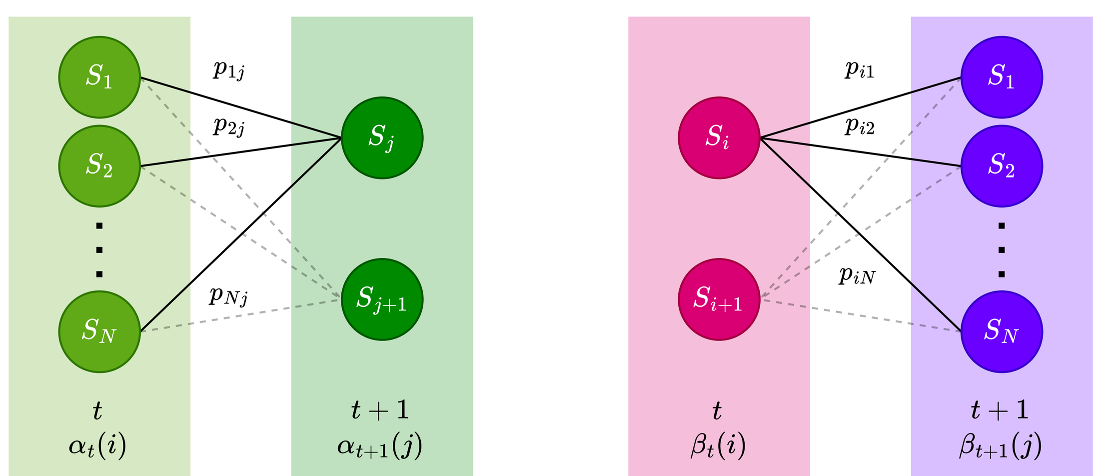

(a) Forward (b) Backward

Figure 3. Illustration of the forward-backward likelihood

E-step

Posterior state probabilities

$$\gamma_t(i) = \frac{\alpha_t(i)\beta_t(i)}{\sum_{k=1}^{N} \alpha_t(k)\beta_t(k)}$$

Posterior transition probabilities

$$\xi_t(i,j) = \frac{\alpha_t(i) a_{ij} b_j(y_{t+1}) \beta_{t+1}(j)}{\sum_{k=1}^{N}\sum_{\ell=1}^{N} \alpha_t(k) a_{k\ell} b_\ell(y_{t+1}) \beta_{t+1}(\ell)}$$

M-step

Updated transition probabilities

$$\hat{a}_{ij} = \frac{\sum_{t=1}^{M-1} \xi_t(i,j)}{\sum_{t=1}^{M-1}\sum_{k=1}^{N} \xi_t(i,k)}$$

Updated emission probabilities

$$\hat{b}_j(v_k) = \frac{\sum_{t:\, y_t = v_k} \gamma_t(j)}{\sum_{t=1}^{M} \gamma_t(j)}$$

Issues

Several challenges arise when applying standard hidden Markov model (HMM) techniques, preventing the direct use of the Baum-Welch algorithms

- **Issue 1:** Time evolves continuously, while classical HMM are discrete-time.
- **Issue 2:** Current fluid level is dependent on the previous fluid level.
- **Issue 3:** The rate of fluid flow is not directly observable.

Discretisation strategy

To address these challenges, the continuous-time process is *discretised* at fixed time intervals. By observing the fluid level at successive time points and computing finite differences, the continuous-time model is approximated by a discrete-time hidden Markov model. This approximation enables the application of the Baum-Welch algorithm for parameter estimation, with flow rates inferred from changes in successive fluid level observations.

Theorem

Given a sequence of times $\{a_k = kh\}_{k=0}^{n}$, we can then form $\hat{X}_k = X_{kh}$ and $\hat{Y}_k = Y_{kh}$ for $k \in \{0, 1, ..., n\}$. $(\hat{X}, \hat{Y})$ is a discrete-time hidden Markov model with hidden Markov chain $\hat{X}_k$ with transition probability matrix $\hat{Q}_h = e^{Qh}$; observations $\hat{Y}_k$ with emission functions $G(i, dy) = \mathbb{P}(\hat{Y}_k \in dy \mid \hat{X}_k = i)$ for $i \in E$ (Gàmiz, Limnios, Segovia-Garcìa [1])

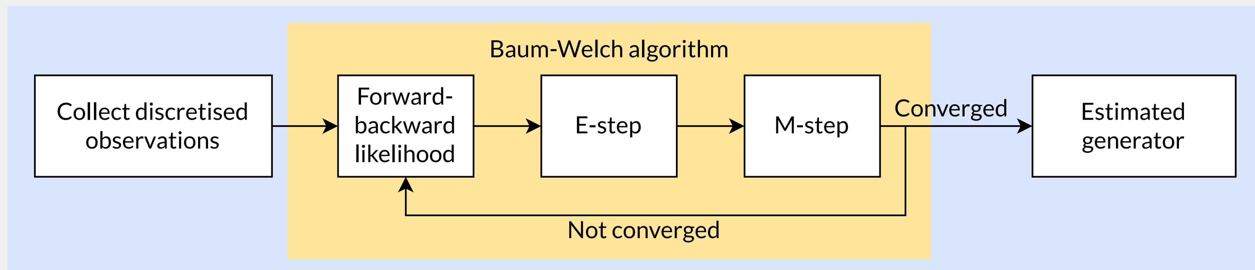

Figure 4. Process of estimation

Study methodology

Two numerical experiments were conducted. The first examined sensitivity to the initial parameter estimates. The second investigated estimation accuracy as the discretisation step size decreases, while the number of observations is held fixed. Both experiments follow the workflow illustrated below.

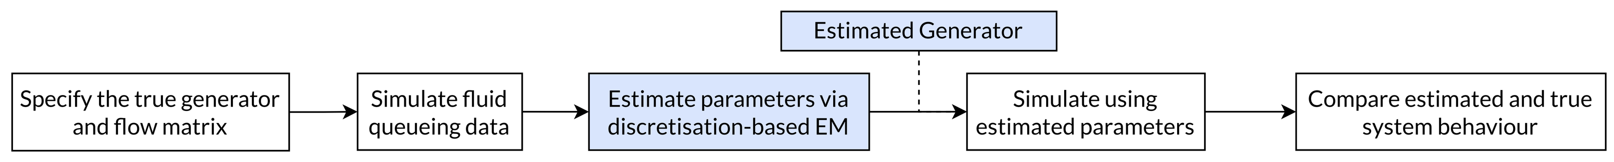

Figure 5. Experiment methodology

Results

Figure 6 summarises the performance of the discretisation-based EM estimator across repeated simulations, measured using the relative L_2-norm error.

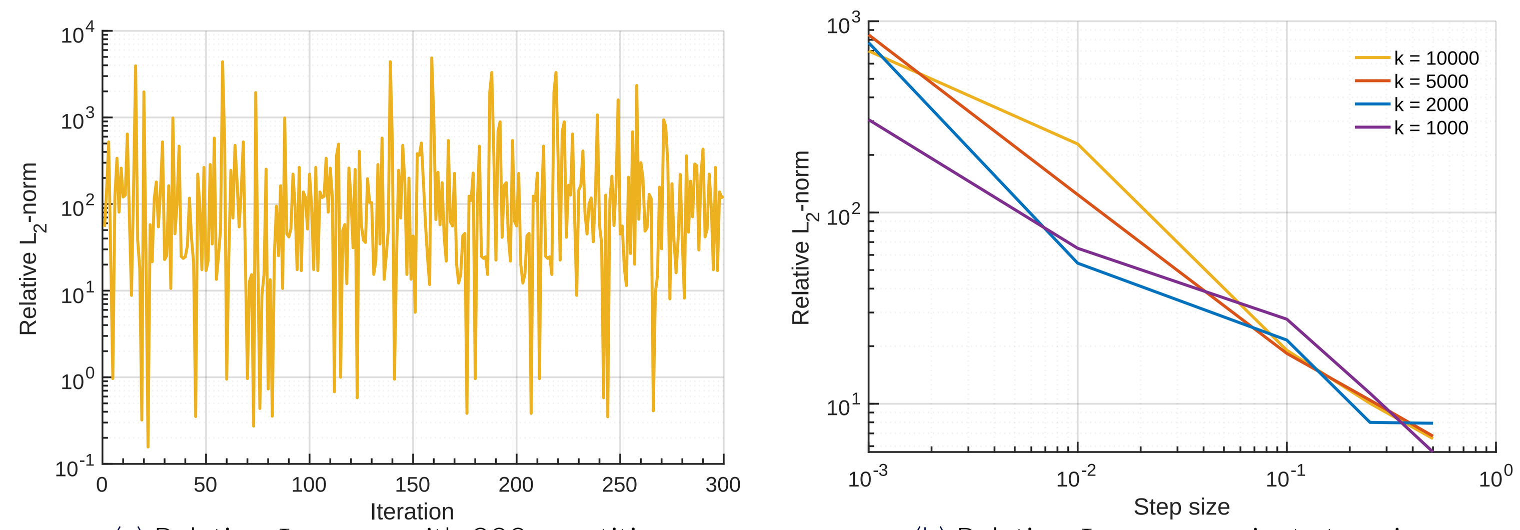

(a) Relative L_2-norm with 300 repetitions (b) Relative L_2-norm against step size

Figure 6. Results of experiments

Discussion

The results indicate that the EM algorithm is highly sensitive to the initial choice of the generator and emission parameters, reflecting convergence to local maxima. Furthermore, when the number of observations is fixed, reducing the discretisation step size leads to increased estimation error. This occurs because overly fine discretisation fails to capture the global dynamics of the fluid process over the observation horizon.

Conclusion and future work

Conclusion

This study demonstrates that the EM framework provides a viable approach for estimating the source process and flow rates in fluid queueing models. However, accurate estimation requires appropriate preprocessing of observations using discretisation. Due to the sensitivity to initialisation, repeated estimation and aggregation of results are necessary to obtain reliable parameter estimates.

Future work

- Develop estimation methods based directly on continuous-time Baum-Welch algorithms.
- Investigate theoretical consistency and identifiability.
- Apply the framework to real-world fluid queueing data.

Acknowledgement and references

The author gratefully acknowledges the supervision and guidance of Dr. Mark Fackrell throughout this project. This work was supported by the Vacation Scholarship program of the School of Mathematics and Statistics, The University of Melbourne.

[1] María Luz Gámiz, Nikolaos Limnios, and Mari Carmen Segovia-García. The continuous-time hidden Markov model based on discretization. Properties of estimators and applications. *Statistical Inference for Stochastic Processes*, 26(3):525–550, Oct 2023.

[2] Lawrence Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

yizhern@student.unimelb.edu.au

fackrell@unimelb.edu.au