

Uncertainty in the Demand for Service: The Case of Emergency Departments and Call Centers

Shimrit Maman, Rafael
Avishai Mandelbaum, Faculty of Industrial Engineering and
Management, Technion
Sergey Zeltyn, IBM Research Lab, Haifa

Technion, MSOM Annual Conference

June 29, 2010

Talk Contents

- 1 Motivation
- 2 Research Outline
- 3 Related Work
- 4 Model Definition
- 5 Case Studies: Financial Call Center and Emergency Department
- 6 Theoretical Results
- 7 Numerical Experiments
- 8 Practical Guidelines
- 9 Future Research Challenges

Motivation

Standard assumption in service system modeling: arrival process is **Poisson** with **known parameters**.

Emergency departments and call centers: known arrival rates for each basic interval (say, one hour in EDs, 15 min in CCs).

Motivation

Standard assumption in service system modeling: arrival process is **Poisson** with **known parameters**.

Emergency departments and call centers: known arrival rates for each basic interval (say, one hour in EDs, 15 min in CCs).

Application of standard approach to basic interval (say, next Tuesday, 9am-10am):

- Derive Poisson parameters from historical data and some forecasting procedure.

Motivation

Standard assumption in service system modeling: arrival process is **Poisson** with **known parameters**.

Emergency departments and call centers: known arrival rates for each basic interval (say, one hour in EDs, 15 min in CCs).

Application of standard approach to basic interval (say, next Tuesday, 9am-10am):

- Derive Poisson parameters from historical data and some forecasting procedure.
- Plug parameters into a queueing model (Erlang-C, Erlang-A, Queueing Network, Skills-Based Routing models, ...).

Motivation

Standard assumption in service system modeling: arrival process is **Poisson** with **known parameters**.

Emergency departments and call centers: known arrival rates for each basic interval (say, one hour in EDs, 15 min in CCs).

Application of standard approach to basic interval (say, next Tuesday, 9am-10am):

- Derive Poisson parameters from historical data and some forecasting procedure.
- Plug parameters into a queueing model (Erlang-C, Erlang-A, Queueing Network, Skills-Based Routing models, ...).
- Set staffing levels according to model and service constraint (e.g., 80% of CC customers answered within 30 sec).

Motivation

Standard assumption in service system modeling: arrival process is **Poisson** with **known parameters**.

Emergency departments and call centers: known arrival rates for each basic interval (say, one hour in EDs, 15 min in CCs).

Application of standard approach to basic interval (say, next Tuesday, 9am-10am):

- Derive Poisson parameters from historical data and some forecasting procedure.
- Plug parameters into a queueing model (Erlang-C, Erlang-A, Queueing Network, Skills-Based Routing models, ...).
- Set staffing levels according to model and service constraint (e.g., 80% of CC customers answered within 30 sec).

Is standard Poisson assumption valid? As a rule it is not, one observes larger variability of the arrival process than the one expected from the Poisson hypothesis.

Research Outline

- Design model for overdispersed arrival rate.
- Validate arrival model via data analysis.
- Plug arrival model into relevant queueing models.
- Derive asymptotic results relevant for real-life staffing problems and provide practical guidelines.
- Validate queueing model via numerical experiments.

Related Work



Steckley, Henderson and Mehrotra

Forecast Errors in Service Systems. 2007.



Koole and Jongbloed

Managing uncertainty in call centers using poisson mixtures. 2001.



Halfin and Whitt

Heavy-traffic Limits for Queues with many Exponential Servers. 1981.



Zeltyn and Mandelbaum

Call centers with impatient customers: Exact analysis and many-server asymptotics of the $M/M/n+G$ queue. 2005.



Feldman, Mandelbaum, Massey and Whitt

Staffing of Time-Varying Queues to Achieve Time-Stable Performance. 2008



Whitt

Staffing a call center with uncertain arrival rate and absenteeism. 2006.

Queue with Overdispersed Arrival Rate: Model Definition

The $M^?|M|n + G$ Queue:

- λ - **Expected** arrival rate of a Poisson arrival process.
- μ - Exponential service rate.
- n service agents.
- G - Patience distribution: time that a customer is willing to wait in queue.

Queue with Overdispersed Arrival Rate: Model Definition

The $M^?|M|n + G$ Queue:

- λ - **Expected** arrival rate of a Poisson arrival process.
- μ - Exponential service rate.
- n service agents.
- G - Patience distribution: time that a customer is willing to wait in queue.

Random Poisson Arrival Rate:

$$M = \lambda + \lambda^c X, \quad c \leq 1,$$


where X is a random variable with zero mean and finite variance.

- $c \leq 1/2$: Conventional variability $\sim$ QED staffing regime.
- $1/2 < c < 1$: Moderate variability $\sim$ QED-c regime (**new**).
- $c = 1$: Extreme variability $\sim$ Efficiency-Driven regime.

Financial Call Center: Data Description

- Israeli bank.
- Arrival counts to the Retail queue are studied.
- 263 regular weekdays ranging between April 2007 and April 2008.
- Holidays with different daily patterns are excluded.
- Each day is divided into 48 half-hour intervals.


Financial Call Center: Arrival Rates



(1) Sundays;
(2) Mondays;

(3) Tuesdays and Wednesdays;
(4) Thursdays;

Financial Call Center: Overdispersion Phenomenon



Poisson CV = $1/\sqrt{\text{mean arrival rate}}$.

Sampled CV's $\gg$ Poisson CV's $\Rightarrow$ **Over-Dispersion**

Moderate and Extreme Variability: Relation between Mean and Standard Deviation of Arrival Rate

Number of arrivals during a basic interval (say, Tue, 9-10am):
Poisson Y with random rate $M = \lambda + \lambda^c \cdot X$, where $E(X) = 0$,
standard deviation $\sigma(X) > 0$ and $1/2 < c \leq 1$. Then,

$$\ln(\sigma(Y)) \approx c \cdot \ln(\lambda) + \ln(\sigma(X)).$$

Moderate and Extreme Variability: Relation between Mean and Standard Deviation of Arrival Rate

Number of arrivals during a basic interval (say, Tue, 9-10am):
 Poisson Y with random rate $M = \lambda + \lambda^c \cdot X$, where $E(X) = 0$,
 standard deviation $\sigma(X) > 0$ and $1/2 < c \leq 1$. Then,

$$\ln(\sigma(Y)) \approx c \cdot \ln(\lambda) + \ln(\sigma(X)).$$

Proof:

$$\text{Var}(Y) = \lambda^{2c} \cdot \text{Var}(X) + \lambda + \lambda^c E(X),$$

and

$$\sigma(Y)/\lambda^c = \sqrt{\sigma^2(X) + \lambda^{1-2c}},$$


$$\lim_{\lambda \rightarrow \infty} (\ln(\sigma(Y)) - c \ln(\lambda)) = \ln(\sigma(X)).$$

Therefore, for large λ ,


$$\ln(\sigma(Y)) \approx c \cdot \ln(\lambda) + \ln(\sigma(X)).$$

Financial Call Center: Fitting Regression Model

Tue-Wed, 30 min resolution




Tue-Wed, 5 min resolution




Financial Call Center: Fitting Regression Model

Tue-Wed, 30 min resolution



Tue-Wed, 5 min resolution



Results:

- Two clusters exist: midnight-10:30am and 10:30am-midnight.
- Very good fit ($R^2 > 0.97$).
- Significant linear relations for different weekdays and time-resolution (5-30 min):

$$\ln(\sigma(Y)) = c \cdot \ln(\lambda) + \ln(\sigma(X)).$$

Financial Call Center: Outline of Additional Results


- Good fit of a well-known Gamma mixture model (Jongbloed and Koole ['01]) to data of Financial Call Center.
- Relation between our main model and Gamma Poisson mixture model is established.
- Distribution of X is derived under Gamma assumption: it is **asymptotically normal** given a large arrival rate.

Emergency Department: Data Description

- Rambam hospital Emergency Department.
- 194 weeks between from January 2004 till October 2007 (five war weeks are excluded from data).
- The analysis is performed using two resolutions: hourly arrival rates (168 intervals in a week) and three-hour arrival rates (56 intervals in a week).


Emergency Department: Arrival Rates

Average Hourly Number of Arrivals




Emergency Department: Over-Dispersion Phenomenon

One-hour resolution



Three-hour resolution



- Moderate over-dispersion.
- Conventional variability ($c = 1/2$) seems to be a reasonable assumption for hourly resolution.

QED- c Regime: Fixed Arrival Rate

QED- c staffing rule:

$$n = \frac{\lambda}{\mu} + \beta \left(\frac{\lambda}{\mu} \right)^c + o(\sqrt{\lambda}), \quad \beta \in \mathbb{R}, \quad c \in (1/2, 1).$$

QED- c Regime: Fixed Arrival Rate

QED- c staffing rule:

$$n = \frac{\lambda}{\mu} + \beta \left(\frac{\lambda}{\mu} \right)^c + o(\sqrt{\lambda}), \quad \beta \in \mathbb{R}, \quad c \in (1/2, 1).$$

Assume an $M|M|n + G$ queue with **fixed arrival rate** λ .

Take λ to ∞ :

- $\beta > 0$: Over-staffing.
- $\beta < 0$: Under-staffing.

For both cases we provide asymptotically equivalent expressions (or bounds) for $P\{W_q > 0\}$, $P\{Ab\}$ and $E[W_q]$, where W_q - waiting time.

Proofs: Based on $M/M/n+G$ building blocks from Zeltyn and Mandelbaum['05], carried out via the Laplace Method for asymptotic calculation of integrals.

QED-c Regime: Random Arrival Rate

Theorem

Assume random arrival rate $M = \lambda + \lambda^c \mu^{1-c} X$, $c \in (1/2, 1)$, $E[X] = 0$, finite $\sigma(X) > 0$, and staffing according to the QED-c staffing rule with the corresponding c . Then, as $\lambda \rightarrow \infty$,

- a. Delay probability: $P_{M,n}\{W_q > 0\} \sim 1 - F(\beta).$
- b. Abandonment probability: $P_{M,n}\{Ab\} \sim \frac{E[X - \beta]_+}{n^{1-c}}.$
- c. Average waiting time: $E_{M,n}[W_q] \sim \frac{E[X - \beta]_+}{n^{1-c} \cdot g_0},$
where g_0 =patience density at the origin.

Proofs: based on conditioning on values of X and results for QED-c staffing rule.


QED-c Regime: Numerical Experiments

Examples: Consider two distributions of X


- Uniform distribution on $[-1,1]$,
- Standard Normal distribution.

$$(1) \beta = -0.5, c = 0.7$$

Delay Probability



Abandonment Probability



Practical Guidelines

- Determine “uncertainty coefficient” c via regression analysis.
- If $1/2 < c < 1$, assume that X is asymptotically normal, calculate standard deviation from regression model.
- Apply QED- c (or QED, Efficiency-Driven) asymptotic results in order to determine appropriate staffing.

Outline of Additional Results

- **Queueing Theory.** Asymptotic performance measures derived and constraint satisfaction problems solved for:
 - QED regime ($c = 1/2$).
 - Efficiency-Driven regime ($c = 1$), discrete and continuous distribution of X .
- **Numerical Experiments.** Very good fit between asymptotic results and the exact ones (simulation).
- **Time-varying Expected Arrival Rate.** Enhance ISA algorithm developed by Feldman with the features of random arrival rate in the $M_t/M/n + G$ queue.
Goal: determine time-dependent staffing levels aiming to achieve a time-stable delay probability.

Future Research Challenges

- Incorporating **forecasting errors** into our model (in the spirit of Steckley et al., 2007).
- **Scaling problem:** dependence of c on the basic interval duration.
- **Time-varying queueing models:** achieving time-stable performance measures (probability to abandon, average wait).
- Validation of $M^?/M/n+M$ (or $M^?/M/n+G$) model in call center environment (and probably other service systems).