RRR for UUU: Exact Analysis of Pee Queue Systems with Perfect Urinal Etiquette

Kristen Gardner
Computer Science Department
Carnegie Mellon University
ksgardne@cs.cmu.edu

Ziv Scully
Computer Science Department
Carnegie Mellon University
zscully@cs.cmu.edu

ABSTRACT

Queueing systems with multiple servers operating in parallel, such as the M/M/k model, have been extensively studied. However, most prior literature examines the limited case in which that all servers may operate simultaneously. This is despite the fact that many practical queueing systems encounter constraints that make this an unrealistic assumption. In this work, we investigate one such setting: the men's lavatory, in which a strict etiquette requires that no two adjacent urinals be in use at the same time. We introduce and analyze a new queueing model, the Context-2 Unease Processing Network (C2UPN), which formalizes a row of urinals used with perfect etiquette. We derive exact results for a row of 3 urinals (UUU) using the Recursive Renewal Reward (RRR) technique. Remarkably, our method generalizes to many other urinal topologies, including longer rows and cyclic configurations.

ACM Reference format:
Kristen Gardner and Ziv Scully. 2017. RRR for UUU: Exact Analysis of Pee Queue Systems with Perfect Urinal Etiquette. In *Proceedings of SIGBOVIK 2017, Pittsburgh, PA, USA, March 31, 2017 (SIGBOVIK '17),* 5 pages.

1 INTRODUCTION

Imagine you are at a conference, and the second coffee break is approaching. Most attendees are filled with the free diuretics provided during the first coffee break and have only one question on their minds: how long will they have to wait in line to use the restroom once the break starts? It is of utmost importance to sustain only a short queueing time. Despite the universality of this problem, there is shockingly little theoretical work analyzing response time in bathrooms. Instead, most prior work on bathrooms focuses on anthropological studies of human behavior in bathrooms.

In this paper, we turn to queueing theory to present new exact analysis of response time in one specific category of bathroom queueing problem: the urinals-only setting. Here customers arrive to the system, wait in the queue for an available urinal, and depart the system upon completing their urination. While many bathrooms consist of both urinals and stalls, we choose to focus on the urinals-only setting for two reasons. First, often in bathrooms that offer both urinals and stalls, people opt to only use the stalls even if there are vacant urinals [2]. Second, the introduction of stalls necessitates the existence of both "type-1" and "type-2" customers that have different service times; this complicates the analysis.

Urinal queueing exhibits several unique properties not present in other bathroom queueing settings. Most notably, here we must consider the degree of *urinal etiquette* exhibited by urinators. Urinal etiquette has to do with the number of vacant urinals left between urinators. The three possible urinal etiquette degrees are:

1. No etiquette. Here an arriving urinator will use any vacant urinal, regardless of whether the adjacent urinals are vacant or occupied.
2. Partial etiquette. Here an arriving urinator may choose to use a vacant urinal that is adjacent to an occupied urinal (for example, if the urinator is experiencing a high degree of urgency). The urinator may also choose to join the queue if the only vacant urinals are adjacent to occupied urinals.
3. Perfect etiquette. Here an arriving urinator *never* uses a vacant urinal that is adjacent to an occupied urinal.

In all three settings, we assume that an arriving urinator will always use a vacant urinal that is not adjacent to an occupied urinal. We note that Justus argues that the "buffer zone" is mandatory, meaning that an arriving urinator may never use a vacant urinal that is next to an occupied urinal, except in cases of unusually high load [5]. We agree, hence in this paper we focus on the perfect etiquette setting.

Our main contribution in this paper is the first exact analysis of response time in urinal systems. We begin by studying 3- and 5-urinal systems configured in a standard topology in which the urinals are arranged in a straight line. Our approach involves modeling the system using a Markov chain and apply the Recursive Renewal Reward (RRR) technique to solve the chain exactly. In Section 4 we consider alternative urinal topologies and investigate the conditions under which our approach allows us to develop exact analysis. Finally in Section 5 we discuss directions for future work.

1.1 Related Work

As noted above, most of the related work focuses on human behavior. Cahill et al. conduct an extensive observational study of humans in bathrooms and find that while unacquainted individuals typically avoid conversation at urinals, people who already know each other often converse while urinating (though they avoid making eye contact) [2].[1] The authors also point out that closed stalls make excellent hiding places while conducting observational studies in bathrooms [2]. Empirical results also indicate that women spend a significantly longer time using bathrooms than men do [1]. Perhaps this is because 85% of women choose to "crouch" instead of sitting directly on public toilet seats, which can reduce the average flow rate by 21% [8]. We refer the reader to [9] for a survey of

[1] Our personal sense of urinal etiquette, finely tuned by more than two decades of combined urinal experience, discourages urinal conversation, particularly between students and their advisors.

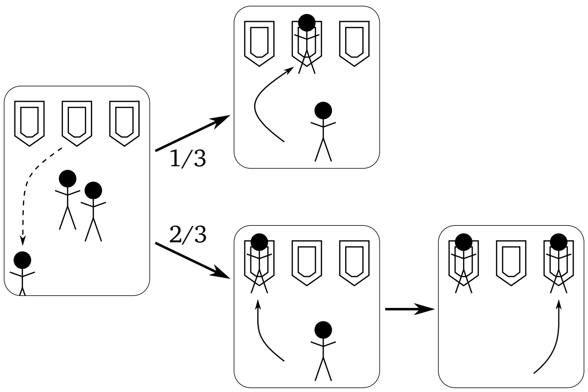

Figure 2.1: The urinal selection process with 3 urinals in a row (viz. the M/M/3/C2UPN queue). A single completion at the center urinal leaves the next urinator with a choice between 3 available urinals. If they choose an edge urinal, then a second urinator can also enter service.

other related behavioral results, perhaps to read while using the bathroom (69% of people use their phones while on the toilet [3]).

All of the above work, though interesting, is orthogonal to our mathematical approach to the urinal problem. To our knowledge, the only existing theoretical work on urinal usage is [6]. The paper considers a setting in which a urinator enters the system and needs to choose which urinal to occupy so as to maximize his privacy, i.e., the time until an adjacent urinal becomes occupied. Unfortunately the model makes several uncomfortable assumptions, including that the urination duration is infinite, so urinators never leave the bathroom. Our work allows for finite urination duration and assumes perfect etiquette, so privacy is always maintained.

2 SYSTEM MODEL

We model *urinals* as servers of fixed service rate 1. Lavatory users, or *urinators*, have i.i.d. service requirements drawn from a specified distribution and arrive according to a specified stochastic process. There is a queue of infinite capacity holding waiting urinators, who enter service when possible in first-come, first-served order.

In traditional queueing, all servers are available at all times. Unfortunately, life is not so simple in the men's lavatory. In practice, adjacent urinals cannot both serve urinators at the same time. We capture this as an *unease graph*: vertices represent urinals, and edges represent pairs of urinals that cannot be occupied simultaneously. While a urinal is serving a urinator, its neighbors in the unease graph become unavailable until service at the urinal completes.

The urinator waiting at the head of the queue enters service at the first urinal to become available. When multiple urinals to become available simultaneously, the urinator occupies a urinal chosen uniformly at random from the set of available urinals. Note that a single completion can enable more than one waiting urinator to enter service. Figure 2.1 illustrates an example of the urinal selection process.

In the vast majority of men's lavatories, the urinals are assembled in a small number of rows. In this case, the unease graph is a union of paths, with edges between adjacent urinals in the same row. We call a single row of urinals a *Context-2 Unease Processing Network* (C2UPN), as each occupied urinal makes up to 2 urinals unavailable.

For the remainder of this paper, we consider the Markovian case, which has exponential service and interarrival times. Urinator service requirements are distributed exponentially with rate μ, and the lavatory experiences a Poisson arrival stream of urinators with rate λ. In Kendall notation, this is the M/M/k/C2UPN queue.

A natural generalization of the C2UPN is to consider other urinal topologies arising from unease graphs beyond paths. Because the number of urinals made uneasy by an may be an arbitrary number N, we refer to this as the *Context-N Unease Processing Network* (CNUPN). We discuss some CNUPN systems in Section 4.

3 ANALYZING MARKOVIAN C2UPN QUEUES

Our chief weapon for analyzing M/M/k/C2UPN systems is *Recursive Renewal Reward* (RRR), a technique pioneered by Gandhi et al. [4]. Originally introduced to analyze the M/M/k queue with exponentially distributed setup costs, RRR enables exact analysis of many queueing systems with Markov chains that are infinite in just one dimension. Roughly speaking, we can think of a one-dimensionally infinite Markov chain as consisting of several "layers" of states. If the chain is eventually periodic and transitions between layers in the repeating region are all one-way, we can apply RRR to compute variety of metrics. Specifically, we find the z-transform of the number of jobs in the queue, from which various useful metrics, such as expected system response time, follow easily.

3.1 Summary of RRR

Here we give a high-level overview of RRR, referring the reader to Gandhi et al. [4] for a more detailed exposition. RRR solves the following problem: given a Markov chain with certain nice properties, find the average value of some time-varying *reward rate*, which depends only on the current state. For example, to find the average number of jobs in the queue, $\mathbf{E}[N]$, we set the reward rate in each state to the number of urinators in the queue. (We write N instead of the usual N_Q for the number in queue to reduce clutter.)

We designate some set of states as *home states*. For queueing systems, there is usually a single home state, namely the empty system. Given home states, a *renewal cycle* is the time interval between transitions into the set home states. A classic result of renewal-reward theory tells us that

$$\mathbf{E}[\text{reward rate}] = \frac{\mathbf{E}[\text{reward accumulated in a renewal cycle}]}{\mathbf{E}[\text{length of a renewal cycle}]}.$$

For certain Markov chains, such as that of the M/M/3/C2UPN, the two quantities right-hand side are very easy to compute!

It is clear how to use RRR to compute expectations. To compute more complicated metrics, such as $\mathbf{Var}(N)$, we just use a more complicated reward rate function. In particular, all moments of N can be computed from its z-transform, $\hat{N}(z)$. Because $\hat{N}(z) = \mathbf{E}[z^N]$, we can compute it by setting the reward rate of each state to be z^n, where n is the number of jobs in the queue when the chain is in that state.

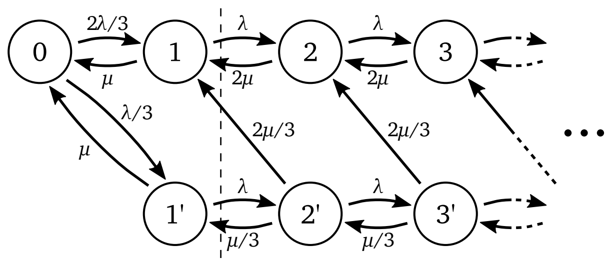

Figure 3.1: The M/M/3/C2UPN CTMC. The repeating portion is to the right of the dotted line, starting at 2 and $2'$.

3.2 Exact Analysis of the M/M/3/C2UPN Queue

The continuous time Markov chain (CTMC) for the M/M/3/C2UPN queue is shown in Figure 3.1. The chain has two layers: a "sad" layer, in which the middle urinal is occupied, and a "happy" layer, in which the middle urinal is idle. We label the happy and sad states with n total urinators in the system with n and n', respectively. (There is no state $0'$ because the middle urinal cannot be occupied if there are no urinators.)

Because our model is in continuous time, the probability of two urinators completing at the exact same time is 0. This means that in the repeating portion of the chain, it is possible to go from sad to happy (see Figure 2.1) but not vice versa, because a transition from happy to sad would require simultaneous departures to make the center urinal available. The only transition into the sad layer is from state 0, when a urinator arriving at an empty system occupies the center urinal.

We now apply RRR to find $\hat{N}(z)$. State 0 serves as our single home state. The first step is to compute the expected renewal cycle length, which we write as L. We will need the following definitions.

- Let L_i for $i \geq 1$ be the expected amount of time it takes to go left from i, possibly visiting states to the right in the meantime. Going left from i always ends up at $i-1$.
- Let L_i' for $i \geq 1$ be the expected amount of time it takes to go left from i'. Going left from i' may end up at either $i-1$ or $(i-1)'$.
- Let q be the probability that going left from i' lands at $i-1$, and let $q' = 1-q$ be the probability that going left from i' lands at $(i-1)'$.

Examining the Markov chain, we see that

$$L = \frac{1}{\lambda} + \frac{2}{3}L_1 + \frac{1}{3}L_1'$$
$$L_1 = \frac{1}{\mu+\lambda} + \frac{\lambda}{\mu+\lambda}(L_2 + L_1)$$
$$L_1' = \frac{1}{\mu+\lambda} + \frac{\lambda}{\mu+\lambda}(L_2' + qL_1 + q'L_1').$$

Because the Markov chain repeats after 2 and $2'$, we know $L_3 = L_2$ and $L_3' = L_2'$, giving us

$$L_2 = \frac{1}{2\mu+\lambda} + \frac{\lambda}{2\mu+\lambda}(2L_2)$$
$$L_2' = \frac{1}{\mu+\lambda} + \frac{\lambda}{\mu+\lambda}((1+q')L_2' + (1-q')L_2).$$

Given q', these linear equations are easily solved. Inspecting the Markov chain, we see that q' satisfies

$$q' = \frac{\mu}{\mu+\lambda}\left(\frac{1}{3}\right) + \frac{\lambda}{\mu+\lambda}(q')^2.$$

Let $\rho = \lambda/\mu$. Knowing $q' \in (0,1)$, we can pick the correct solution to the quadratic,

$$q' = \frac{1}{2\rho}\left(1 + \rho - \sqrt{(1+\rho)^2 - \frac{4}{3}\rho}\right).$$

As expected, this is decreasing in ρ: the more likely arrivals are compared to departures, the more likely it is that we transition from sad to happy at some point before going left.

We now compute the expected reward accumulated per renewal cycle, which we right as R. Let R_i be the total reward accumulated while going left from $i \geq 0$, and similarly for R_i'. This includes both the reward from state i or i' and reward from states to the right that are visited before going left. The reward rate for state i is $z^{(i-2)^+}$ (that is, $z^{\max\{i-2,0\}}$), and the reward rate for state i' is $z^{(i-1)^+}$.

Examining the Markov chain, we see that

$$R = \frac{1}{\lambda} + \frac{2}{3}R_1 + \frac{1}{3}R_1'$$
$$R_1 = \frac{1}{\mu+\lambda} + \frac{\lambda}{\mu+\lambda}(R_2 + R_1)$$
$$R_1' = \frac{1}{\mu+\lambda} + \frac{\lambda}{\mu+\lambda}(R_2' + qR_1 + q'R_1').$$

Because the Markov chain repeats after 2 and $2'$, we know $R_3 = zR_2$ and $R_3' = zR_2'$, giving us

$$R_2 = \frac{1}{2\mu+\lambda} + \frac{\lambda}{2\mu+\lambda}(1+z)R_2$$
$$R_2' = \frac{z}{\mu+\lambda} + \frac{\lambda}{\mu+\lambda}((q'+z)R_2' + (1-q')R_2).$$

This is the same q' as in the computation for L, so this is just a system of linear equations.

Sparing the reader the remaining details, we skip to the simple closed-form solution,

$$\hat{N}(z) = \frac{(\rho-2)(-3\rho+X+3)(\rho(z-1)-1)(\rho^2(X(z-1)+4z-7)+\rho(X(z-2)+3z-8)-2(X+3)+3\rho^3(z-1))}{(3\rho^2+5\rho+(\rho+2)X+6)(\rho z-2)(X+\rho(3-6z)+3)},$$

where $X = \sqrt{9\rho^2+6\rho+9}$. We leave using this result to compute $\mathbf{E}[N]$ and other metrics as a simple exercise for the reader.

3.3 Generalization to M/M/5/C2UPN

The same general approach works for longer rows of urinals, but the layer structure becomes more complex. Instead of solving a quadratic for q', we must solve a system of quadratics to find several probabilities. Remarkably, the next-largest interesting case, the M/M/5/C2UPN, admits a closed-form solution. One probability is

$$\frac{1}{2\rho}\left(1 + \rho - \frac{1}{\sqrt{219}}\sqrt{111\rho^2 + 36\sqrt{9\rho^4+3\rho^3+3\rho+9} + 55\rho + 111}\right),$$

and the rest are algebraic functions of it. From the probabilities, solving for L and R is routine. Whether the M/M/k/C2UPN has such an elegant solution for general k is a rich area for future research.

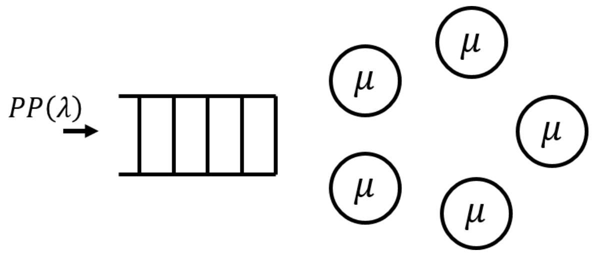

Figure 4.1: A 5-urinal CUP system.

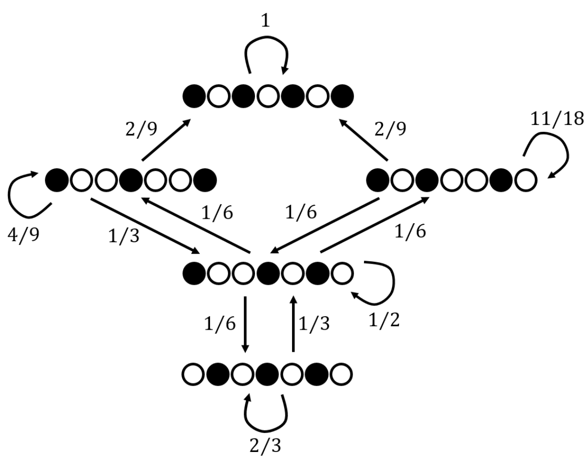

Figure 4.2: State transitions in the 7-state linear urinal configuration when the queue is non-empty.

4 ALTERNATIVE URINAL TOPOLOGIES

Thus far we have considered only the "standard" bathroom topology in which the urinals are arranged in a row, making the unease graph a path. But a multitude of urinal arrangements are possible. In this section we extend our results to an important class of alternative topologies: *Circular Urinal Positioning* (CUP).

In a CUP system, the urinals are arranged in a circle around a central pillar (see Figure 4.1), making the unease graph a cycle. Urinators arrive to the system as a Poisson process with rate λ and the urination duration is exponentially distributed with rate μ. Unlike in the line topology studied in the previous section, in a CUP system all of the urinals are symmetric in that there is not an endpoint with only a single neighboring urinal. While the CUP configuration makes it more challenging for urinators to find private urinals, we find that surprisingly, the CUP configuration makes analysis much more tractable in large bathrooms.

Consider the 7-urinal system. In a linear configuration, there are five possible states when the queue is non-empty and assuming perfect etiquette. The transitions between these states are shown in Figure 4.2. The complicated state transitions, and in particular the non-DAG structure, make it difficult to solve the resulting Markov chain using RRR. This is because it is possible to transition back and forth between pairs of states, meaning that finding the "leftward" probabilities will require solving a high-degree polynomial.

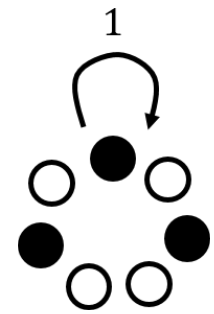

Figure 4.3: State transitions in the 7-urinal CUP configuration when the queue is non-empty.

Now suppose that instead our 7 urinals are arranged in a CUP structure. Now there is only one possible state when the queue is non-empty and all urinators observe perfect etiquette. The new state transition diagram is shown in Figure 4.3. Indeed, the system reduces to an M/M/3, which is easy to solve.

In general, the state space for any size urinal system is simpler in the CUP configuration than in the linear configuration. Due to the increased analytical tractability of large CUP systems, we suggest that all bathrooms be reconfigured so that the urinals satisfy a CUP structure.

5 CONCLUSIONS AND FUTURE WORK

In this paper we derived the first exact analysis of mean response time in urinal systems. Our approach, which uses the Recursive Renewal Reward technique, applies in the 3- and 5-urinal linear C2UPN systems, as well as in larger CUP systems.

There are several interesting and important directions for future work. Here we only consider the perfect etiquette setting, in which urinators never occupy adjacent urinals. When load is high, it may be necessary to move instead to the partial etiquette setting, in which urinators may occupy adjacent urinals if their urgency is sufficiently high. This complicates our Markov chain analysis because it introduces many new possible states for the urinal system. Furthermore, in the "probabilistic urgency" (p-urgency) setting, the exponential urination duration assumption may not be realistic. Empirical work has shown that a more realistic distribution is the sum of a delay before the start of urination, and a urination duration; both of these components depend on the proximity of the urinator to other urinators [7]. Hence in the partial etiquette/p-urgency setting, we also need to extend our results to general urination duration distributions.

An alternative direction for future work involves the strategic decision of which urinal to choose in a p-urgency system. We have assumed that a urinator will choose uniformly at random from among the permissible urinals, but this need not be the case. A common strategy is to choose the urinal that maximizes the distance between urinators. While this strategy is beneficial for ensuring one's individual privacy, it may reduce the overall system efficiency. An interesting direction for future work would be to investigate the Privacy-Efficiency Envelope.

We hope that this paper will serve as the start of a steady stream of future work on analyzing the performance of urinal and other crucial lavatory-related queueing systems.

REFERENCES

[1] Michelle A Baillie, Shawndel Fraser, and Michael J Brown. 2009. Do women spend more time in the restroom than men? *Psychological reports* 105, 3 (2009), 789–790.

[2] Spencer E Cahill, William Distler, Cynthia Lachowetz, Andrea Meaney, Robyn Tarallo, and Teena Willard. 1985. Meanwhile Backstage Public Bathrooms and the Interaction Order. *Journal of Contemporary Ethnography* 14, 1 (1985), 33–58.

[3] Michelle Castillo. 2016. Study reveals what people do with their phones in the bathroom. *CNBC* (July 2016).

[4] Anshul Gandhi, Sherwin Doroudi, Mor Harchol-Balter, and Alan Scheller-Wolf. 2013. Exact Analysis of the M/M/K/Setup Class of Markov Chains via Recursive Renewal Reward. In *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '13)*. ACM, New York, NY, USA, 153–166. DOI:http://dx.doi.org/10.1145/2465529.2465760

[5] Jeremy C Justus. 2006. Piss Stance: Private Parts in Public Places: An Analysis of the Men's Room and Gender Control. *Studies in Popular Culture* 28, 3 (2006), 59–70.

[6] Evangelos Kranakis and Danny Krizanc. 2010. The urinal problem. In *International Conference on Fun with Algorithms*. Springer, 284–295.

[7] R Dennis Middlemist, Eric S Knowles, and Charles F Matter. 1976. Personal space invasions in the lavatory: Suggestive evidence for arousal. *Journal of personality and social psychology* 33, 5 (1976), 541.

[8] K. H. MOORE, D. H. RICHMOND, J. R. SUTHERST, A. H. IMRIE, and J. L. HUTTON. 1991. Crouching over the toilet seat: prevalence among British gynaecological outpatients and its effect upon micturition. *BJOG: An International Journal of Obstetrics and Gynaecology* 98, 6 (1991), 569–572. DOI:http://dx.doi.org/10.1111/j.1471-0528.1991.tb10372.x

[9] Clint Rainey. 2015. Everything We Know About Human Bathroom Behavior. *NYMag* (May 2015).