Submitted to *Management Science*
manuscript MS-0xxx-2007.1

Maximum Happiness: Amusement Park Rides as Closed Queueing Networks

Ronald S. Tibben-Lembke
Managerial Sciences Department, University of Nevada, Reno, NV 89503, rtl@unr.edu

In 2006, 28% of U.S. residents visited a theme park, and the number one reason given for visiting was "the rides." The rides are central to their attraction, and new ones can cost over $100 million to develop, and yet and they have never been studied in the OR/MS literature. We study amusement park rides as closed queueing networks, dividing them into "carousels," "roller coasters," and "tightly coupled" rides. We estimate throughput as a function of the number of cars, the time to load and unload riders, the ride time, and the number of parallel loading and unloading zones, and develop formulas that can easily be built into a spreadsheet tool to help ride designers see how potential changes to the ride could increase rider throughput. We begin with the deterministic case, continue to the stochastic case, and consider a case study of California Screamin', the marquee ride at Disney's California Adventure, before concluding with a computational study.

Key words: Service Operations, Queueing Theory, Amusement Parks, Roller Coasters
History: This paper was first submitted on July 30, 2007.

1. Introduction

Amusement parks are a huge business. In the 2002 U.S. Economic Census, amusement and theme parks in the U.S. had estimated revenues of $8.3 billion (U.S. Census Bureau 2004). In 2006, the Walt Disney theme parks alone had global revenues of nearly $10b (Disney 2006). According to the International Association of Amusement Parks & Attractions (IAAPA) there were 335 million visitors to U.S. amusement parks in 2006, in which 1.5 billion rides were taken. Twenty-eight percent of the U.S. population visited an amusement park in 2006, and 50% planned to visit one in the next year (IAAPA 2007).

Amusement parks have been little studied in the OR/MS literature, despite their size and importance in popular culture, and despite the fact that the operation of the rides and other attractions is clearly central to the parks' ability to draw visitors. IAAPA surveys show that rides are the number one reason people visit amusement parks, and 46% say roller coasters are their favorite ride (IAAPA 2007).

New roller coasters can cost as much as $10-12 million (O'Brien 1996), and major new rides like Universal Studio's "Spiderman" ride cost $100 million to build, and Disney World's "Mission: Space" cost $150 million (Krisner 2002). Disneyland recently spent $100 million just to change the theme of a submarine ride from *Captain* Nemo to *Finding* Nemo (Jefferson 2007). For major investments like this, the parks must carefully plan the capacity of the rides to gain as much benefit for the park as possible. The formulas derived in this paper can easily be used in a spreadsheet, to allow ride designers to see how, for example, a modification that reduces average loading time, or separating unloading and loading into separate zones will affect ride throughput, even taking into account variability in loading and unloading time. When combined with estimated construction costs of the various changes, the engineers can then easily determine the most cost-effective ways to increase rider throughput.

All major American amusement parks charge a flat daily rate for admission, or sell multi-day or annual passes. The park gets paid the same, regardless of how many rides each guest takes. However, amusement parks clearly have an incentive to maximize the number of rides a guest can take per day: time spent queueing cannot be spent buying souvenirs or eating in restaurants, and more rides taken should lead to better customer satisfaction and greater future profits. Disneyland believes that if patrons cannot experience at least 9 attractions, spending on souvenirs goes down, because people "feel cheated out of experiencing a full range of attractions," and "are walking out unhappy and not buying souvenirs" (Reckard 2001).

In general, the capacity of a ride is not a fixed number. As demand varies during the day, more workers are placed on a ride as needed to help riders load and unload faster and increase throughput (Rajaram and Ahmadi 2003), and more cars may be brought online (Ahmadi 1997). We assume that each ride is moving as many people as possible, given its current configuration, and address the issue of improving the ride's maximum throughput.

If we define the loading time for a ride as T_L, the riding time as T_R, and the unloading time as T_U, then we may define a total changeover time for a car as $T_C = T_U + T_L$. If a ride has multiple cars, for safety reasons, there is likely a minimum allowable time between successive cars, which we will call T_S. If $T_C < T_S$, cars will be ready to leave in less time than T_S, and if the previous car left recently, the second car will have to wait until T_S has passed since it departed. During that time, it cannot leave the loading zone, and the unloading/loading zone will be blocked. As we will see, this safety condition has a profound impact on the way the rides are modeled, because the inter-departure times do not have the memoryless property.

In a roller coaster, T_S is likely determined based on the time required to safely stop cars, should the car before it unexpectedly stop on the tracks. T_S could not be reduced unless the ride were made less exciting or reconfigured in some other way. Also, ASTM standard F-2291, Section 6 (ASTM 2006, p. 1348) requires the cars be far enough apart that a rider cannot be injured by hitting the ride or another rider, say by sticking an arm outside the car.

Spacing may also be affected by concerns about the quality of the ride experience. In many rides, not being too close to the cars before you adds to the enjoyment (so when the monster jumps out at the person in front of you, you don't see it, for example), so the minimum spacing of cars would be depend on the ride layout. We can express the minimum car spacing either in terms of time, T_S, or distance, d_S. If the ride travels at speed s, the two are equivalent: $T_S = d_S/s$.

Below, we calculate the cycle time between subsequent cars for most types of rides. Once we estimate the cycle time per car, CT, the total throughput (measured in the number of cars per time period) TH = 1/CT. If $\overline{p}$ is the average number of passengers per car, the throughput measured in the number of passengers per time period. can easily be determined: $\mathrm{TH}_p = \overline{p} * \mathrm{TH}$.

The results presented apply of course not only to amusement park rides, but to any closed queueing network in a service or production environment with a similar operating structure. In a heat-treat, drying, baking, or painting tunnel operation, a minimum time between departing pieces could be required, for quality control. However, as amusement park rides were the impetus for the research, we will use them as motivating examples throughout.

In §2, we describe the literature and in §3 discuss rider behavior and attempts to increase TH. In §4, we consider rides as deterministic closed queueing networks, and in §5 we include stochastic changeover times. In §6 we present a case study of California Screamin' from Disney's California Adventure, and in §7 we present a computational study. Finally, in §8 we summarize and present suggestions for future research.

2. Literature

We study a tandem network consisting of a small number of stations, some of which may have multiple parallel stations. With unloading, loading, and the riding portion, we have only three stations. Routings are deterministic, with the only possibility for variation coming if there are multiple unloading and/or loading stations. We begin with deterministic models, and move to stochastic models, but as Haxholdt et al. (2003) show, even deterministic models can be difficult to study.

Because our networks are very simple in comparison to the models studied in the stochastics literature, our review will be brief. Study of closed queueing networks began with Jackson (1957, 1963), and Koenigsberg (1958), who studies it in the context of mining. Gordon and Newell (1967) show the equivalence between closed and open networks. Buzein (1973) presents computational methods. Koenigsberg (1982) presents an overview of the previous 25 years' worth of results on closed queueing networks.

Although amusement parks have been studied in the leisure and travel literature (e.g. Formica and Olson 1998), very little has appeared in the OR/MS literature. Ahmadi (1997) presents a model for determining how much capacity each ride should have, and how to influence customers' routing through the park. Although Ahmadi seeks to maximize the minimum number of rides a customer would be able to take, he notes that a survey done by the theme park he was working with showed that "beyond a threshold level of average number of rides, further rides provided little improvement in customer satisfaction" (Ahmadi 1997, p.1), an interesting finding he suggests worthy of further research. Despite that survey result, we will continue with the assumption that parks would like to get as much throughput as possible out of their multi-million-dollar rides. Rajaram and Ahmadi (2003) determine which rides to increase capacity of, at different times of day, based on the rates at which customers go from one area of the park to another, to drive customers to the areas of the park where they spend more money on souvenirs and food.

Ahmadi (1997) notes that ride capacity is typically linear in the number of cars (p. 2), but his model does not make this assumption, assuming that the cost for a given level of capacity is known. Rajaram and Ahmadi (2003) make a similar assumption, that the number of employees needed to achieve a given level of capacity is known. As we will show, for many rides, ride throughput will not in general be linear in the number of cars used.

It is very difficult to determine exactly how capacity decisions are currently made, because many rides are designed by in-house engineers (called "Imagineers" at Disney), or by independent ride engineering firms, and both are highly secretive (Kirsner 2002).

3. Rider Behavior and Throughput

In the author's experience, what constitutes an acceptable wait seems to be relative to the number of visitors in the park. Unfortunately, riders cannot accurately estimate the wait time because the queue area is often not directly observable. An estimated wait time is often posted at the entrance, but guests often doubt their reliability. As a result, although the number of people who would like to take a ride may vary throughout the day with park attendance, the number of people who ride the ride will likely vary quite differently. However, it is said that nature abhors a vacuum, and an amusement park equivalent may be "every ride will have a line."

Assumption 1. *Every ride will always have a queue of riders waiting for it: in queueing terminology, no server will ever starve.*

If a park doubled the throughput of a marquee ride, the shorter line would attract more riders. Perhaps people consciously or unconsciously conclude that a ride is worth waiting in line a certain amount of time for, and the line grows to reflect that valuation. If there typically is no line, a rational park would modify or replace the ride to meet guests' tastes. Given Assumption 1, we assume an inexhaustible supply of riders and do not consider their queueing behavior.

3.1. Loading and Unloading

As we will see, T_L and T_U are crucial determinants of CT. The loading process can be broken into several steps. In the language of just-in-time, any steps that can be done “offline” (while the ride is going) instead of “inline” (while the ride is stopped), will reduce the changeover time, and get the ride moving again more quickly.

First, exactly which customers will be riding on the next ride need to be identified, including which customers will be riding together. At the Disneyland Dumbo ride, while the ride is operating, an employee walks along those waiting, asking who is riding together, handing one plastic feather to each group. Before the ride has stopped, everyone who will be loading knows they will be on the next ride. By contrast, at Disneyland’s Astro Orbiter, a ride nearly identical in structure, (but with rockets instead of flying elephants), once the ride stops, the operator lets people into the ride, and as each person comes up to the gate, asks whom they will be riding with. The operator attempts to keep track mentally of the number of groups that have been allowed in, but often, once all of the passengers allowed in have been seated, it is discovered one more group is needed. These differences in procedure are a large part of the reason the author observed an average T_C of 3.75 minutes for the Astro Orbiter, compared with 2.4 for Dumbo. As we will see below, differences in T_C make a huge difference in carousel TH.

The Astro Orbiter example also shows how attempts to increase $\overline{p}$ by finding more riders may increase T_L so much that TH_p declines, overall. For some rides, single rider lines can increase p with no impact on T_L.

After determining who will be riding, each customer must be assigned to a car. In most roller coasters, customers stand in front of the gate for their assigned car, before the cars arrive. For most carousels, customers run for the car they want, and then look around for a different one, if someone else arrives first. Once riders select a car, they climb in, stow any belongings, fasten the safety restraint, and the operator needs to make sure each person is safely restrained. In Dumbo, a simple lap belt is used, which is easily verified by the operator, but the notched leather belt on the

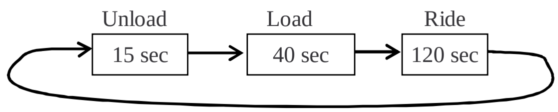

Figure 1 Ride from Cars' Perspective

Astro Orbiter is more difficult to snug, and requires the operator to lean inside the car to verify safety. These differences also contribute to the T_C differences of the rides.

Other operational differences between employees can also have significant impact on total throughput. For example, the author observed one carousel operating with a cycle time of 4 minutes. When a different operator began, the cycle time went down to 3.5 minutes, with no observable decrease in $\overline{p}$, because the second crew was able to safely load the riders in less time.

3.2. Rides from the Cars' Perspective

From the perspective of the riders, they first load, then ride, then unload. For example, if $T_L = 40$ sec., $T_R = 120$ sec., and $T_U = 15$ sec., then clearly the run time for the ride is 175 seconds. If we look at the ride from the perspective of a car, as in Figure 1, it returns from a ride, unloads, then loads and goes out on another ride. Clearly, the run time for a one-car ride remains the same at 175 seconds, regardless of the perspective we look at it from. Since we focus on the throughput of the cars, not the riders' queueing experience, we will use the cars' perspective.

We will consider ride capacity in terms of the number of cars it can process per hour, defining a "car" as the set of riders who board at the same time. For a roller coaster, one car is all of the people who get on and depart at the same time, which might actually look like a short train. For a carousel, a car is all of the people who get on the carousel and ride it at the same time.

4. Rides as Deterministic Closed Queueing Networks

In this section, we consider rides as deterministic closed queueing networks. In the next section, we consider stochastic times.

4.1. Carousels: Single-Car Rides

The simplest ride to analyze consists of a single car: on a carousel, there is only one carousel. Everybody gets on and rides at the same time, and when the ride is over, everyone gets off. As

Figure 2 Multiple Carousel Rides

listed in the Appendix, many other rides operate in the same way: bumper cars, stage shows, or any other attraction where everyone enters together, rides together, and exits at the same time.

If a ride has several cars that operate independently, with independent personnel running them, then we may consider them as several independent carousels. For example, a "moon shot" type ride (of which the Maliboomer is an example at Disney's California Adventure) may have several different towers, each with its own cars. If each tower has its own personnel to help with loading and unloading and to run it, they run independently, and we should treat them as independent carousels. However, on a roller coaster, the cars are not independent, because they interact through their use of a common track and loading and unloading area.

Figure 2 shows a timeline for a carousel. A wavy line represents the loading period, a straight line represents the ride time, and a zig-zag line represents unloading. In the figure, we see two complete rides, and a third ride load and begin but not complete.

Clearly the run time for each load of passengers is $\mathrm{RT} = T_R + T_C$. Because there is only one car $(n = 1)$, the cycle time is equal to the run time, and $\mathrm{CT} = \mathrm{RT}$.

$$\mathrm{CT}(n=1) = T_R + T_C \tag{1}$$

Because there is only one car, $\mathrm{TH} = 1/(T_R + T_C)$, and $\mathrm{TH}_p = \overline{p}/(T_R + T_C)$. Assuming the ride length is not reduced, ride throughput is maximized by minimizing $T_C = T_U + T_L$.

4.2. Roller Coasters: Multiple-Car Rides

If multiple cars can load or operate simultaneously, the analysis is more complicated. Any ride with multiple separate cars that share some resources we will call a roller coaster. As the Appendix shows, many kinds of rides may be included in this category.

In this section, we assume the time required to ride the course is greater than the changeover time $T_R > T_C$, which is probably typical. (We consider the other case, which we call "short-ride"

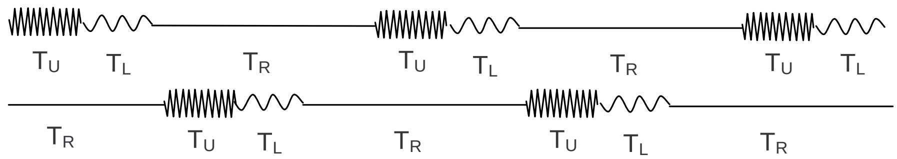

Figure 3 Two-Car Roller Coaster

roller coasters, in § 4.4.) Given this relationship, there are three possible ranges of values for T_S. If $T_R > T_C \geq T_S$, the safety time required between cars is shorter than or equal to the time it takes the next cars to get ready, so the safety factor is not significant. We will consider this case first, below. The second case is when $T_R > T_S > T_C$, which we consider in § 4.3. Finally, if $T_S \geq T_R > T_C$, the safety time is greater than or equal to the time it takes the car to take the ride, so there can only be one car on the ride at any time, and we consider this with the short-ride roller coaster.

Lemma 1. *If ride and changeover times are deterministic, and $T_C \geq T_S$, that is, safety is not a factor, the throughput of the ride is given by*

$$CT(T_R > T_C \geq T_S) = \begin{cases} (T_R + T_C)/n & \text{if} \quad n < (T_R + T_C)/T_C \\ T_C & \text{if} \quad n \geq (T_R + T_C)/T_C \end{cases} \tag{2}$$

Proof: Consider $n = 2$, with $T_R > T_C \geq T_S$. In Figure 3, the top timeline represents car 1, and the bottom is car 2. While the first car is out on the course, the second car reloads and leaves before the first car returns. The run time for each car is $\text{RT} = T_R + T_C$. In any time period of length RT, two cars unload, so the average time between rides is $\text{CT} = \text{RT}/2 = (T_R + T_C)/2$. For n cars, if $T_R \geq (n-1) * T_C$, the cars will still not have to wait, and $\text{CT} = T_R/n$.

If more cars added, cars will have to wait for each other at some point. Because $T_C \geq T_S$, cars will accumulate waiting to unload. With n cars on the ride, cars will accumulate when the ride time, T_R, is less than the time for the other $n-1$ cars to unload and load, when $T_R \leq (n-1) * T_C$. Rearranging, cars will accumulate when $n \geq (T_R + T_C)/T_C$.

For the sake of example, assume that this happens for a three cars as illustrated in Figure 4, where $T_C < T_R < 2T_C$. As soon as the first car leaves, the second car reloads and goes out. Finally, the third car comes in, reloads, and goes out again. Because $T_R < 2T_C$, when the first car returns,

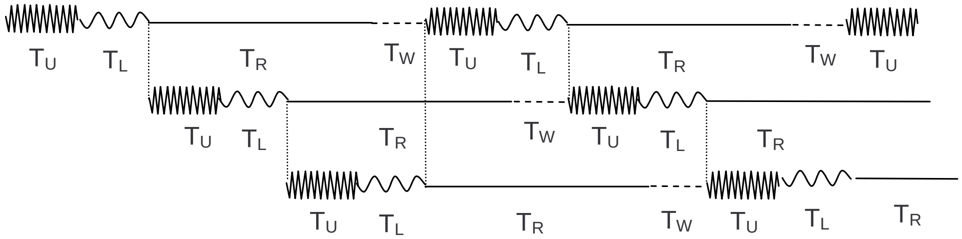

Figure 4 **Three-Car Roller Coaster with Queueing**

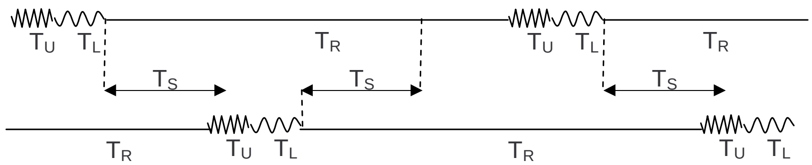

Figure 5 **Two-Car Roller Coaster with Safety Time and no Safety Delays**

it will face a waiting time $T_W = 2T_C - T_R$. Total run time is now given by the sum of loading, unloading, riding, and waiting time: $\mathrm{RT} = T_C + T_R + T_W = 3T_C$. In general, the run time for a car is its changeover time T_C, plus the changeover time of the other $n-1$ cars, for a total of $\mathrm{RT} = nT_C$. In any period RT, n cars return, so $\mathrm{CT} = \mathrm{RT}/n = T_C$. $\square$

Buying additional cars is probably easier and cheaper than making other substantial changes to the ride. Thus, when safety is not an issue, ride throughput is maximized by making sure $n \geq (T_R + T_C)/T_C$. Assuming the park wants to avoid the substantial cost and service implications of reducing T_R, further increases in TH will depend upon reductions in T_C. If T_C is reduced enough, it could happen that $T_C < T_S$, and the results of the next section will need to be used.

4.3. Safety is a Factor

If safety is a factor, (that is $T_R > T_S > T_C$), cars may have to wait before departing, and we assume any waiting cars block the reloading zone until the safety interval has passed. However, if the ride time is long, relative to the safety time, and there are only a few cars, it is possible that the cars may not have to wait at all because of T_S.

In Figure 5, there are two cars with a safety factor that is considerably shorter than T_R, but still longer than T_C. At the start, car 2 is on the ride, and car 1 reloads and leaves. When the safety

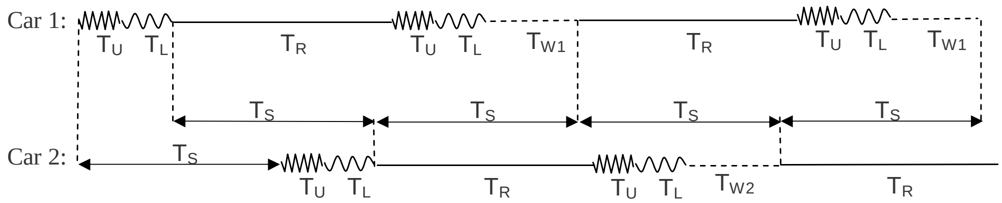

Figure 6 Two-Car Roller Coaster with Safety Time

interval after car 1's departure has ended, car 2 is just unloading, so T_S does not delay car 2. Car 2 leaves, and the safety interval after its departure expires before car 1 has even returned. In any interval $\mathrm{RT} = T_C + T_R$, two cars unload, so $\mathrm{CT} = \mathrm{RT}/2$. If there are n cars on the ride, and the departure of one car is never delayed because of the safety interval after another car's departure, in any interval RT, n cars will unload, so $\mathrm{CT} = \mathrm{RT}/n$.

LEMMA 2. *If ride and changeover times are deterministic, and $T_R > T_S > T_C$, that is, safety concerns are a factor, the throughput of the ride is given by*

$$CT(T_R > T_S > T_C) = \begin{cases} (T_R + T_C)/n & \text{if} \quad n < (T_R + T_C)/T_S \\ T_S & \text{if} \quad n \geq (T_R + T_C)/T_S \end{cases} \tag{3}$$

Proof: If there is only one car on the ride, clearly, every $T_R + T_C$, one car finishes. As cars are added, eventually there will be enough cars that the bottleneck never starves, because cars will always queue before reloading, and TH will equal the bottleneck rate, in this case, $\mathrm{CT} = T_S$. If the other $n-1$ cars can reload during the time it takes any one car to go on the ride, T_R, there will be no queueing because $(n-1) * T_C < T_R$. For large n, $n-1$ cars cannot reload during T_R, and the server never starves. This occurs when $(n-1) * T_C \geq T_R$, or $n \geq (T_R + T_C)/T_C$. For smaller n, in any interval $T_C + T_R$, all n cars complete the ride, giving average $\mathrm{CT} = (T_C + T_R)/n$. □

Figure 5 illustrates how TH is determined by $(T_R + T_C)/n$ when $n < (T_R + T_C)/T_S$. Figure 6 shows two cars where the safety factor is larger, relative to T_R, and $n = 2 > (T_C + T_R)/T_S$. The first car re-loads and leaves. Car 2 arrives as quickly as possible, T_S after car 1, and departs T_S after car 1. Car 1 returns and is ready to leave $T_R + T_C$ after its first departure, but it cannot re-depart until T_S has passed since car 2 left. Car 2 left T_S after car 1, so car 1 cannot leave until $2T_S$ after its last departure, so it will wait $T_{W1} = 2T_S - (T_R + T_C)$. Time $2T_S$ passed between car 2's departures, and it spent T_R on the ride and T_C changing, so it also waited $2T_S - (T_R + T_C)$. For either car, $\mathrm{RT} = T_C + T_R + T_W = 2T_S$, and there are two arrivals in that period, so $\mathrm{CT} = 2T_S/2 = T_S$. Because $2T_S > T_C + T_R$, we know $n = 2 > (T_C + T_R)/T_S$. As predicted by eq. (3), $\mathrm{CT} = T_S$.

Again we see that throughput is maximized by ensuring that the number of cars is large enough. Beyond that, increases in TH will also come from reductions in T_C. Reducing T_S would also increase TH, but that may require expensive to accomplish.

Although it seemed unlikely at the outset, by determining the CT separately, depending on the relative size of T_S and T_C, the two CT equations (2) and (3) may be combined.

Theorem 1. *If ride and changeover times are deterministic, TH is given by*

$$CT(T_R > T_C) = \begin{cases} (T_R + T_C)/n & \text{if} \quad n < (T_R + T_C)/\max\{T_C, T_S\} \\ \max\{T_C, T_S\} & \text{if} \quad n \geq (T_R + T_C)/\max\{T_C, T_S\}. \end{cases} \tag{4}$$

Proof: Consider Equation (4). If $T_C \geq T_S$, (as is the case in Lemma 1), then the break point $(T_R + T_C)/\max\{T_C, T_S\}$ evaluates to $(T_R + T_C)/T_C$, the same as in Lemma 1. For such values of n, $\mathrm{CT} = (T_C + T_R)/n$, also as Lemma 1. For larger n, $\mathrm{CT} = \max\{T_C, T_S\} = T_C$, also as in Lemma 1. Both CT expressions and the breakpoint are the same as in Lemma 1. Thus Equation (4) returns the correct values for $T_C \geq T_S$.

If $T_C < T_S$, $(T_R + T_C)/\max\{T_C, T_S\} = (T_R + T_C)/T_S$, the same as in Lemma 2. For $n < (T_R + T_C)/T_S$, $\mathrm{CT} = (T_C + T_R)/n$, the same as in Lemma 2. For larger n, $\mathrm{CT} = \max\{T_C, T_S\} = T_S$, also as Lemma 2. Both CT expressions and the breakpoint are the same as in Lemma 2. Thus Equation (4) returns the correct values for $T_S > T_C$. □

4.4. Short-Ride Roller Coasters

Above, we assumed that the ride time is longer than the changeover time, which seems reasonable, because most roller coasters take only a matter of seconds to load and unload, and the rides last

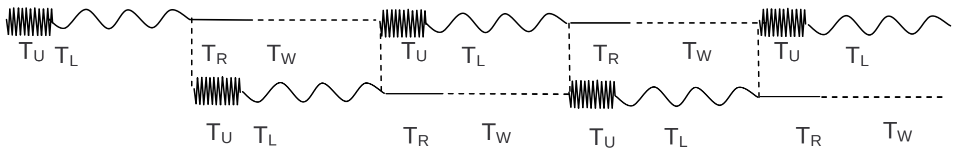

Figure 7 **Two-Car Short-Ride Roller Coaster ($T_C > T_S$)**

longer than that. If $T_R \leq T_C$, the ride time is very short. Such a ride may seem unlikely. However, it is possible to imagine a brief but intense ride which requires a longer loading time for the staff to make sure the riders are secure, for example, something like a "flying" roller coaster, where riders lay down and "fly" like Superman.

If there is only one car, such a ride is clearly a carousel. A two-car roller coaster would look like Figure 7. After the first car leaves, the second car reloads and leaves. While the second car is loading, the first car returns and waits (because $T_R \leq T_C$) a time of $T_W = T_C - T_R$. The ride time for $n = 2$ becomes $\text{RT} = T_C + T_R + T_W = 2T_C$, and $\text{CT} = \text{RT}/2 = T_C$. Since the first car returns before the second car is ready to leave, safety time cannot be an issue.

THEOREM 2. *For short ride roller coasters ($T_R < T_C$),*

$$CT(T_R \leq T_C) = \begin{cases} T_R + T_C & \text{if} \quad n = 1 \\ T_C & \text{if} \quad n \geq 2. \end{cases} \tag{5}$$

Proof: Follows from the above. □

Because TH is maximized by $n = 2$, it seems unlikely additional cars would be added. Additional TH would need to be gained from reductions in T_C.

For completeness, we briefly consider $T_S \geq T_R > T_C$. It seems hard to imagine a ride where say $T_R = 30$, and $T_S = 45$ seconds, so a car returns after 30 seconds, but for safety reasons, cars can only depart every 45 seconds. Perhaps if some equipment has to reset itself or cool off, such a case could arise. It is quite similar to the short-ride case, in that there can be at most one car on the ride at any time. If there is only one car, it is a carousel, but $\text{CT} = \max\{T_S, T_R + T_C\}$. In general,

$$\text{CT}(T_S \geq T_R > T_C) = \begin{cases} \max\{T_S, T_R + T_C\} & \text{if} \quad n = 1 \\ T_S & \text{if} \quad n \geq 2. \end{cases} \tag{6}$$

4.5. Separate Loading and Unloading Zones

Loading and unloading involves humans climbing into and out of ride cars. Many changes may be made to reduce these times as much as possible, but eventually each ride must reach a point where no further reductions in either time can be made cost-effectively. The next logical step (or perhaps an earlier logical step) to consider is to separate the loading and unloading zones, so one car can be unloading while the car before it is loading.

Thus far, we assumed loading and unloading happen in the same place, so the total time to change riders is the sum of the unloading and the loading time, $T_C = T_U + T_L$. If separate loading and unloading zones are created, the time between subsequent cars being ready to depart the unloading and loading zones is equal to the maximum of the unloading and loading times: $T_C = \max\{T_U, T_L\}$.

For a carousel, there is no change in CT, because there is only one car. By creating examples similar to Figures 3-7 above, but with separate unloading and loading zones, it can easily be shown that equations (2)-(6) continue to be valid, using $T_C = \max\{T_U, T_L\}$.

Because the definition of when safety is a factor or not depends whether or not $T_S > T_C$, reducing T_C by creating separate zones may mean that safety will become a factor where it previously was not a factor. For short-ride roller coasters, the reduction in T_C brought by separating loading and unloading may mean $T_R \not\leq T_C$, and the ride would no longer fall into the short-ride category. Finally, regardless of which equation is used, a different term in these equations may become relevant, as the reduction in T_C changes the values of the conditions for which term is to be used.

This process of separation, can, of course, be taken even further. On the Space Mountain ride at Disneyland, the loading process is broken into two separate loading zones: in the first, riders climb in and out and pull down the safety harnesses. In the second zone, an employee verifies that the safety harnesses are snugly down over the riders. As more separate loading and unloading zones are added, their process times must be added to the T_C calculation: $T_C = \max\{T_U, T_{L1}, T_{L2}, \ldots\}$.

Adding additional loading and unloading zones may require significant expense, but depending on the relative sizes of T_U, T_L and T_S, may lead to increases in TH sufficent to justify the expense.

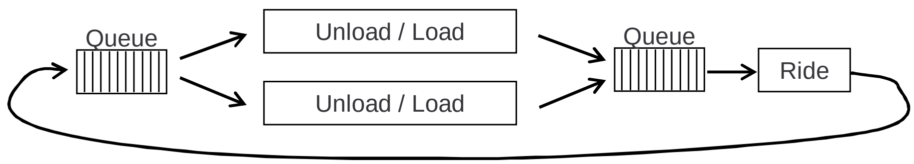

Figure 8 Parallel Unloading/Loading Zones for Roller Coaster

4.6. Parallel Unloading/Loading Zones

Instead of separate loading and unloading zones, another option is to create parallel loading and unloading zones, where more than one car can unload and/or load at the same time, as in Figure 8. Any incoming cars finding the unloading zones occupied wait for the first available unloading zone. If the cars are on separate tracks, each with its own dedicated unloading/loading zone, (as is the case with the Matterhorn at Disneyland, where they even have separate lines), the two are independent roller coasters. Assuming the cars share a common set of tracks, after each car loads, it may move to an outbound queue, to wait for T_S to pass so that it may depart.

Let m be the number of parallel (identical) loading/unloading zones, and assume only one car may be in each zone at one time (i.e., there are not separate loading and unloading zones) with changeover time of T_C seconds. The m zones together can turn out a maximum rate of one car every T_C/m seconds. If $T_C/m \geq T_S$, then the changeover process determines the throughput.

Lemma 3. *For n cars and m parallel unloading/loading zones, if $T_R > T_C/m \geq T_S$ the CT of the ride is given by*

$$CT(T_R > T_C/m \geq T_S) = \begin{cases} (T_R + T_C)/n & \text{if} \quad n < m(T_R + T_C)/T_C \\ T_C/m & \text{if} \quad n \geq m(T_R + T_C)/T_C. \end{cases} \tag{7}$$

Proof: Assume that safety is a not factor, that is, the time between subsequent cars being ready for departure is longer than the safety interval. With m parallel servers and deterministic processing time of T_C, the inter-departure time is T_C/m seconds, so this is equivalent to assuming $T_C/m > T_S$. Also assume sufficient cars are present to prevent the bottleneck re-loading processes from starving. From Little's Law, there must be $T_R/(T_C/m)$ cars out riding on the ride at any time. We know that m cars are in the loading zones, so a total of at least $m * T_R/T_C + m = m(T_R + T_C)/T_C$ cars must be in use to maximize throughput.

If $n < m(T_R+T_C)/T_C$, the bottleneck starves on occasion. In any time T_R, one car will travel the whole ride, and by the time it reloads and is ready to go, the other $n-1$ cars will have departed, and be out of its way. Thus, in any period T_R+T_C, n cars depart, so CT $= (T_R+T_C)/n$. As a continuity check, it can be easily verified both expressions give the same value when $n = m(T_R+T_C)/T_C$. □

Thus, we see that Lemma 1 is just a special case of Lemma 3, where $m = 1$.

LEMMA 4. *For n cars and m parallel unloading/loading zones, if $T_R > T_S > T_C/m$ the CT of the ride is given by*

$$CT(T_R > T_S > T_C/m) = \begin{cases} (T_R+T_C)/n & \text{if} \quad n < (T_R+T_C)/T_S \\ T_S & \text{if} \quad n \geq (T_R+T_C)/T_S. \end{cases} \tag{8}$$

Proof: Assume safety is a factor, that is $T_S > T_C/m$. When there are enough cars on the ride to ensure that the departure process never starves, and a car is always ready to depart every T_S seconds, CT $= T_S$. The time for one car to reload and take the ride is T_R+T_C. If cars depart every T_S seconds, a total of $(T_R+T_C)/T_S$ cars could depart during that time. If $n \geq (T_R+T_C)/T_S$, the departure process will never starve, and CT $= T_S$. When $n < (T_R+T_C)/T_S$, we know CT $> T_S$. In any period T_R+T_C, only n cars can be output, so CT $= (T_R+T_C)/n$. □

THEOREM 3. *If ride and changeover times are deterministic, n cars are used, and m parallel unloading/loading stations, the throughput of the ride is given by*

$$CT(T_R > T_C/m) = \begin{cases} (T_R+T_C)/n & \text{if} \quad n < (T_R+T_C)/\max\{T_C/m, T_S\} \\ \max\{T_C/m, T_S\} & \text{if} \quad n \geq (T_R+T_C)/\max\{T_C/m, T_S\}. \end{cases} \tag{9}$$

Proof: Similar to proof of Theorem 1. □

Thus we see that Theorem 1 is really a special case of Theorem 3, where $m = 1$.

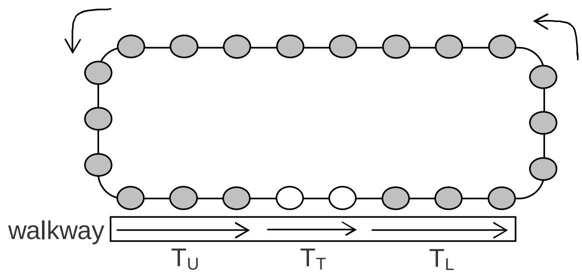

Figure 9 Tightly-Coupled Cars

4.7. Tightly Coupled Cars (“Peoplemovers”)

In some rides the cars are tightly coupled in a chain or moving beltway, and never stop moving: the only way to stop a car is to stop the whole ride. For example, at Disneyland, the Buzz Lightyear Astro Blasters ride is configured this way. See Figure 9 for a diagram. To load, riders step onto a moving walkway going the same speed as the cars, and climb in. At the end, they step out of the cars onto the other end of the same walkway. The car then travels to the loading zone. Let T_T be the cars' travel time between loading and unloading.

A moving sidewalk or “peoplemover” would seem necessary for cars to be this tightly coupled, because the cars must be stationary relative to the unloading/loading surface. Otherwise, the whole ride would need to stop every few seconds for loading and unloading. Although these rides generally move at a constant speed, it would seem possible to configure them to move slower in loading and unloading and faster during the ride, like a high-speed ski lift.

Multiple cars can be in the loading or unloading process at the same time. Because all of the cars are tightly coupled, there is no time spent waiting for the previous car to unload. Every $\mathrm{RT} = T_U + T_T + T_L + T_R$ seconds, each car returns to the unloading zone, and another carload of passengers climbs out, so $\mathrm{CT} = \mathrm{RT}/n = (T_U + T_T + T_L + T_R)/n$.

Given the length of the loading and unloading zones, the time can easily be calculated. For example, if $s = 5$ ft./sec., the cars would pass through a 50 ft. unloading zone in 10 seconds.

If a passenger is unable to board the car during the time the car is in the loading zone, the ride will have to be stopped until the rider has safely boarded. Stopping the ride obviously inconveniences all of the other riders, and loses throughput. The same is true of the unloading zone. Therefore, the park would reasonably want to construct the unloading and loading zones to be large enough to ensure that some percentage of passengers will be able to load in it, say $p = \Pr\{\text{rider safely exits}\}$.

If d_U is the distance (length) of the unloading zone, and d_L is that of loading zone, then the time the car is in the unloading zone, $T_U = d_u/s$, and $T_L = d_L/s$. If loading rate is exponentially distributed with mean λ_L, then the probability of a rider getting loaded, and the ride not needing to stop, is $1 - e^{-\lambda_L T_L}$. If $p_{\min}$ is the desired probability of not stopping the ride, it is easy to show that for a given value of s,

$$d_L \geq \frac{-\ln(1 - p_{\min})}{\lambda_L} s = \frac{p'}{\lambda_L} s$$

where we define $p' = -\ln(1 - p_{\min})$. An equivalent expression is derived for d_U.

The goal in designing the ride is to minimize $\text{CT} = \text{RT}/n$. Defining D as the total distance the car travels, $D = d_L + d_R + d_U + d_T$. Because $\text{RT} = D/s$, we will try to make d_L and d_U as short as possible, given s, and set $d_L = p' * s/\lambda_L$, and $d_U = p' * s/\lambda_U$. The values of λ_U and λ_L are a function of the ride design, so we assume them as small as feasible, and take them as given.

At some rides, the space d_T is used to start boarding riders who will are identified as needing additional time, (for example, riders arriving in a wheelchair, who bypass the regular line) who then have $T_T + T_L$ to load. The extra T_T can greatly reduce the odds that the entire ride will need to be stopped. This space can also be used for additional unloading time if needed. Because of this important purpose, the park may specify a minimum time T_T, or $T_{T\min}$. Minimizing d_T lowers CT, and there is no reason to extend it greater than the minimum, so $T_T = T_{T\min}$, and $d_T = T_{T\min} * s$.

For customer satisfaction purposes, there is likely a minimum time that the entertainment portion of the ride should last, $T_{R\min}$. Reducing T_R lowers CT, so it will be set at its lowest value, $T_R = T_{R\min}$, and $d_R = T_{R\min} * s$ Combining these results, we then have run time

$$\text{RT} = T_{R\min} + T_{T\min} + \frac{p'}{\lambda_L} + \frac{p'}{\lambda_U}.$$

At a speed of s, each car travels the total distance of the ride in RT, so the total distance is $\text{RT} * s$. We may write the total distance D of the ride as

$$D = \text{RT} * s = \left(T_{R\min} + T_{T\min} + \frac{p'}{\lambda_L} + \frac{p'}{\lambda_U}\right) * s \tag{10}$$

Given the minimum distance d_S, the distance between cars, D/n, must be $D/n \geq d_S$, or $n \leq D/d_S$, which we can write as $n \leq \mathrm{RT} * s/d_S$. Because n must be integer, it will not always be possible to set $n = \mathrm{RT} * s/d_S$. In general, $n = \lfloor \mathrm{RT} * s/d_S \rfloor$. Thus, we have

$$\mathrm{CT} = \mathrm{RT}/n = \frac{\mathrm{RT}}{\lfloor \mathrm{RT} * s/d_S \rfloor} \approx \frac{d_S}{s}, \tag{11}$$

$$\mathrm{TH} = 1/\mathrm{CT} = \frac{\lfloor \mathrm{RT} * s/d_S \rfloor}{\mathrm{RT}} \approx \frac{s}{d_S}. \tag{12}$$

Strictly speaking, $\mathrm{CT} \geq d_S/s$ and $\mathrm{TH} \leq s/d_S$, but the error in the approximation can be shown to be $(\mathrm{RT} * s/d_S) / \lfloor \mathrm{RT} * s/d_S \rfloor$. For $n = 80$ (as the case in the Buzz Lightyear ride), the error from the approximation is $\leq 81/80 = 1.25\%$, which may not be a major difference.

From equation (12), the park can increase TH as much as it wants to, by increasing the speed, s. As equation (10) shows, the required length of the ride is directly proportional to s. Doubling the speed of the ride therefore means doubling the length of the ride, which likely means a doubling of the cost of building the ride. Actually, doubling the speed would likely more than double the cost, because a faster-moving ride would likely require sturdier tracks and wheels, etc., which would increase the cost per linear foot to build.

Using equations (12) and (10), the park can see how increasing s increases TH, but increases the cost. Depending on how much the park is willing to pay for increases in TH, it can use the equations to find its optimal TH.

If the park wants to define a minimum time T_S between riders, so that, for example, the monster that pops out can get safely back into place in time to surprise the next car, it must set $n = \lfloor \mathrm{RT}/T_S \rfloor \leq R_T/T_S$, and equations (11) and (12) become

$$\mathrm{CT} = \mathrm{RT}/n = \frac{\mathrm{RT}}{\lfloor \mathrm{RT}/T_S \rfloor} \approx T_S, \tag{13}$$

$$\mathrm{TH} = 1/\mathrm{CT} = \frac{\lfloor \mathrm{RT}/T_S \rfloor}{\mathrm{RT}} \approx \frac{1}{T_S}. \tag{14}$$

If the park specifies a minimum time between cars T_S, CT cannot be reduced below T_S.

5. Stochastic Closed Queueing Networks

To consider rides as stochastic closed queueing networks, we assume that changeover times, T_C are exponentially distributed, with rate λ. We assume $m = 1$, and one zone for loading and loading.

In the case of loading and unloading of humans from an amusement park ride, the exponential distribution would seem easy to defend: one can never tell when one person will take longer to get in the car, get the safety harness fastened, etc., or to collect any personal items and get out of the car, coming back to look for something they think they left, etc. Once a car is in motion, the time to complete the ride T_R should be nearly deterministic, regulated by the laws of physics and the machinery of the ride, including computerized ride control systems.

If cars depart according to an exponential distribution, and return to unload a fixed amount of time later, clearly those inter-arrival times have an exponential distribution (Burke 1956). If there is only one car on the ride, there is no need to worry about a safety interval, T_S.

THEOREM 4. *If a ride consists of one car, and the changeover time, T_C, is exponentially distributed with mean $1/\lambda$, total ride $TH = \lambda$.*

Proof: The changeover process is a single-server Poisson process. With a single car, safety is not a concern, so T_S is not a factor in departure times, and throughput is that of a single-server Poisson process. □

However, if the number of cars, n, is greater than 1, T_S must be taken into account. We define τ as the time between the departure of one car and the arrival of the next car. For a given value of τ, the probability density function of the inter-departure times is a truncated exponential:

$$f(t|\tau) = \begin{cases} 0 & : \quad t < T_S - \tau \\ 1 - e^{-\lambda(T_S - \tau)} & : \quad t = T_S - \tau \\ e^{-\lambda t} & : \quad t > T_S - \tau. \end{cases} \tag{15}$$

It is not possible for a car depart less than $T_S - \tau$ after the previous car. Any car that is ready in less time than that is held for that time, so the probability of leaving after $T_S - \tau$ is the probability that the car is ready in $T_S - \tau$ or less. For times longer than that, the probability is just the exponential distribution.

We define T_D as the time to departure, the time a car occupies the server, performing the changeover and then waiting for $T_S - \tau$ to pass. For a given realization of the changeover time, T_C, and τ, we find $T_D = \min\{T_C, T_S - \tau\}$.

THEOREM 5. *If a car enters the server time τ after the previous car left, the expected time to departure (the time it occupies the server), $E[T_D|\tau]$ is given by:*

$$E[T_D|\tau] = \frac{1}{\lambda} e^{-\lambda(T_S - \tau)} + T_S - \tau. \tag{16}$$

Proof: Given the definitions of $f(t|\tau)$, and T_D, $E[T_D|\tau] = \int_{T_C=0}^{\infty} \min\{T_C, T_S - \tau\} f(T_C|\tau) dT_C$ becomes:

$$\begin{aligned} E[T_D|\tau] &= (T_S - \tau) \int_{T_C=0}^{T_S-\tau} f(T_C|\tau) dT_C + \int_{T_C=T_S-\tau}^{\infty} T_C f(T_C|\tau) dT_C \\ &= (T_S - \tau)\left(1 - e^{-\lambda(T_S-\tau)}\right) + \int_{T_C=T_S-\tau}^{\infty} T_C f(T_C|\tau) dT_C. \end{aligned}$$

Given the memoryless nature of exponential service times, this easily becomes

$$E[T_D|\tau] = (T_S - \tau)\left(1 - e^{-\lambda(T_S-\tau)}\right) + \left(\frac{1}{\lambda} + T_S - \tau\right) e^{-\lambda(T_S-\tau)}$$

which simplifies to the desired result. □

To compute $E[\text{CT}]$, we need to determine $E[T_D] = E\left[E[T_D|\tau]\right]$. Unfortunately, $E\left[E[T_D|\tau]\right] = \int_{\tau=0}^{\infty} E[T_D|\tau] f_\tau(\tau) d\tau$ cannot be computed unless we can express $f_\tau(\tau)$, the probability that a car arrives τ after the other left. In a closed queueing network with deterministic ride time, $f_\tau(\tau)$ is equivalent to the probability of the inter-departure time. Unfortunately, T_D, violates the memoryless property, because the time until the next car departs depends on, τ, how long ago the previous car departed. This makes the system difficult to study using traditional queueing tools. However, when τ is known, equation (15) can be used. If there are enough cars in the system to keep the ride at 100% capacity, a queue will always be present in front of unloading, and $\tau = 0$ for all cars.

LEMMA 5. *If the number of cars, n, is such that $n \geq n^* = (T_R + T_S)/T_S$, then the p.d.f. of inter-departure times is a truncated exponential distribution, as in equation (15), with $f(t|\tau = 0)$.*

Proof: Assume n is great enough to guarantee the server always has a queue of waiting cars. Cars will enter the server immediately upon the departure of the previous car, and $\tau = 0$ for all cars, so the probability of interdeparture times is given by $f(t|\tau = 0)$.

While a car goes on the ride, the other $n-1$ cars will occupy the server for at least $(n-1) * T_S$ seconds before the first car can re-enter the server. If $(n-1) * T_S \geq T_R$, the server will still be blocked when the first car returns, and there will still be a queue, and τ will still be 0. Thus, τ will always be 0 if $n * T_S \geq T_R + T_S$. □

THEOREM 6. *If the number of cars, n, is such that $n \geq n^* = (T_R + T_S)/T_S$, then*

$$CT = \frac{1}{\lambda} e^{-\lambda T_S} + T_S.$$

Proof: From Lemma 5, if $n \geq n^* = (T_R + T_S)/T_S$, the interdeparture time is given by $f(t|\tau = 0)$. Each car will block the server for at least T_S, and on average will block the server for $E[T_D|\tau = 0]$. On average, CT $= E[T_D|\tau = 0]$, so TH $= 1/E[T_D|\tau = 0]$. Evaluating $E[T_D|\tau = 0]$, we obtain the desired result. □

Setting $n = n^*$ from Theorem 6, we guarantee that the server will be busy when a car returns. If $\tau = 0$ for all cars, while a car goes on the ride, the other $n-1$ cars will actually occupy the server for an average time of $(n-1) * E[T_D|\tau = 0] \geq (n-1)T_S$. We could attempt to set n to ensure that, on average, the server is just finishing the last previous car as a car returns, and re-derive Lemma 5 using $(n-1) * E[T_D|\tau = 0] \geq T_R$ instead of $(n-1) * T_S \geq T_R$. To do so, it would be argued that, on average, the previous $n-1$ cars would just be finishing as any given car arrived, and, on average, no capacity would be lost. Unfortunately, because of the non-memoryless nature of the problem, that would not be valid. The dependent nature of the safety window means that fluctuations of faster and slower changeover times are not allowed to cancel each other out. If some cars leave more quickly than expected, the server can sit idle, and capacity is lost. In the computational results section, we will also see that this value of n does not produce maximum TH.

Thus, $n \geq n^* = (T_R + T_S)/T_S$ cars will guarantee maximum possible throughput. For such n, the expected time for $n-1$ cars to reload is $(n-1) * E[T_D|\tau = 0] = (n-1) * \left(\frac{1}{\lambda} e^{-\lambda T_S} + T_S\right)$. Once a car leaves, the other $n-1$ will occupy the server this long. Subtracting the run time from this, we find the expected queue time $E[T_Q]$:

$$E[T_Q] = (n-1) * \left(\frac{1}{\lambda} e^{-\lambda T_S} + T_S\right) - T_R. \tag{17}$$

For $n \geq n^*$, summing the expected queue time, ride time and departure time, we can find the expected run time RT for each car, the total time the passengers spend on a ride:

$$\begin{aligned} E[\mathrm{RT}] &= E[T_Q] + T_R + E[T_D|\tau = 0] \\ &= n * \left(\frac{1}{\lambda} e^{-\lambda T_S} + T_S\right). \end{aligned} \tag{18}$$

To summarize these results, when $n = 1$, the problem is trivial. When $n \geq n^*$, the process is definitely not memoryless, but we can still compute CT, $E[\mathrm{RT}]$, and $E[T_Q]$. If $1 < n < n^*$, the problem cannot be solved theoretically, and throughput can only be estimated through simulation.

5.1. Tightly-Coupled Rides

In section 4.7, the length of the loading and unloading zones was determined by the minimum probability $p_{\min}$ of not having to stop the ride. But the calculations of TH and CT were done assuming the ride never stopped. For a given target TH, there will be $\mathrm{TH} * (1 - p_{\min})$ riders who do not load successfully, and cause the ride to stop. The average delay for each of these riders will be $1/\lambda_L$, so expected TH lost due to riders not loading in time T_L will be $\mathrm{TH} * (1 - p_{\min})/\lambda_L$.

For unloading riders, T_U is set to guarantee $p_{\min}$ percent of riders will successfully unload in time T_U. However, the ride actually will only stop if the riders do not successfully unload in time $T_U + T_T$. The probability of not unloading in that time is $e^{-(T_U+T_T)\lambda_U}$. Each of these riders will cause a delay of $1/\lambda_U$, so the total throughput lost from riders not unloading is $\mathrm{TH} * e^{-(T_U+T_T)\lambda_U}/\lambda_U$. Combining these results with equation (12), the expected TH is

$$E[\mathrm{TH}] \approx \frac{s}{d_S} \left(1 - \frac{(1 - p_{\min})}{\lambda_L}\right) \left(1 - \frac{e^{-(T_U+T_T)\lambda_U}}{\lambda_U}\right). \tag{19}$$

6. Case Study: California Screamin'

California Screamin' was built at a cost of $50 million, and is the marquee ride of Disney's California Adventure, Disney's newest U.S. park (Muller 2003). The ride was selected because of its prominence, but also because as an outdoor roller coaster, the departures are easy to time, because after waiting for T_S, cars blast off from a standing start, right next to the boardwalk, unlike Disneyland's marquee roller coaster, Space Mountain, which is an indoor ride.

The author timed 96 departures: 41 runs with 4 cars running, and 55 with 5 cars running. The ride has 7 cars, up to 6 of which can be running at any one time. Five of the cars can hold 24 riders, and 2 can hold 23 (ulitimaterollercoaster.com 2007). We will use 24 riders per car in our calculations. Data were not collected on $\bar{p}$, but an unadvertised single rider line appears to help keep $\bar{p}$ quite high.

With 4 cars, the ride had an average departure time of 53.68 seconds, for $\mathrm{TH}_p = 1{,}609$ riders per hour. With 5 cars, the average departure time was 40.75 seconds, for $\mathrm{TH}_p = 2{,}120$, an increase of 32%. The stated ride capacity is 2,200 per hour (LAT 2001), which would require an average departure time of 39.27 seconds. With 4 cars, the ride appeared to be operating with $T_S = 40$ sec. With 5 cars, it appeared to be operating with $T_S = 37$, which helps to explain how a 25% increase in n could lead to a 32% in TH_p.

The ride has two $m = 2$ unloading/loading zones, but as the stochastic case for $m > 1$ is outside the scope of the current paper, we will estimate CT using the deterministic formula. The average time to unload was approximately 6 seconds, and roughly 45 seconds were required to load and perform the safety check, and we will assume $T_C = 51$ seconds. The ride is listed as lasting 2:36, or $T_R = 156$ seconds, which is in line with the author's experience.

From Theorem 3, if $T_S = 37$ seconds, $T_R = 156$, $T_C = 51$, and $m = 2$, we get $n^* = 6$, and it would seem Disneyland was correct in building the ride to accommodate 6 cars. For $n = 6$, $\mathrm{CT} = T_S = 37$. For $n = 5$, $\mathrm{CT} = 41.4$, which is quite close to the observed value of 40.75. For $n = 4$, assuming $T_S = 40$, $\mathrm{CT} = 51.2$ seconds, which is also close to the observed value of 53.68.

	Theoretical		Observed		Stated
n	CT	TH	CT	TH	TH
6	37.0	2,335			2,200
5	41.4	2,087	40.8	2,120	
4	51.2	1,688	53.7	1,609	

Table 1 Theoretical, Observed and Stated Throughput for California Screamin'

The theoretical and observed CT and TH_p values for California Screamin' are summarized in Table 1. Because we are using the deterministic formula, we would expect changeover variability to result in actual TH_p being lower than predicted. For $n = 4$, observed TH_p is 4.68% less than predicted. For $n = 5$, observed TH_p is 1.48% higher than predicted. Because T_C is based on a very small sample size, it may be an overestimate, which would explain this difference. The ride was never observed with $n = 6$, but for the deterministic formula, $\mathrm{CT} = T_S = 37$, and TH=2,335. Disney claims trains depart every 36 seconds (rcdb.com 2007, Marden 2007), which would give the ride capacity of 2,400 riders/hour.

For TH = 2,335, Disney underestimates TH_p by 5.8%. If $T_S = 36$, their estimate underestimates capacity by 9.1%. Attempts to learn how Disney creates its estimates were unsuccessful. Clearly, future work is needed to more accurately predict TH of these rides in a stochastic environment.

7. Computational Study

As shown, we cannot determine a closed-form expression for the throughput with stochastic changeover times, except for the two extreme points where $n = 1$ and $n \geq n^* = (T_R + T_S)/T_S$. The result for $n = 1$ is unlikely to be useful, although Gadget's Go Coaster in Disneyland was often observed to operate this way. At the other extreme, n^* is the minimum number of cars needed to guarantee maximum possible TH for the ride. Using n^* cars will guarantee maximum throughput, but it is possible that this number of cars will lead to a long queue time before unloading, which will hurt rider satisfaction. If a smaller number of cars would cause a very small reduction in throughput, but a significant reduction in queue time, the parks may decide to use $n < n^*$.

In order to study the ride throughput as a function of n, we must turn to simulation. For example, suppose $1/\lambda = E[T_C] = 0.375$ min., and $T_S = 0.25$ min. From Theorem 6, we obtain $n^* = 9$. If

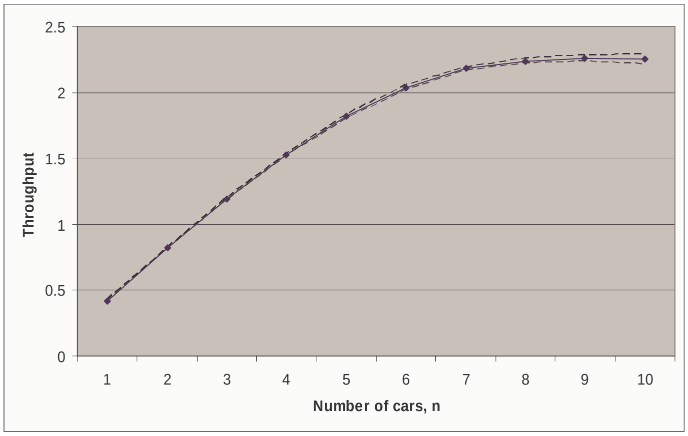

Figure 10 TH as n increases, $1/\lambda = E[T_C]$ = 0.375, T_S=0.25, $E[T_D]$=0.443

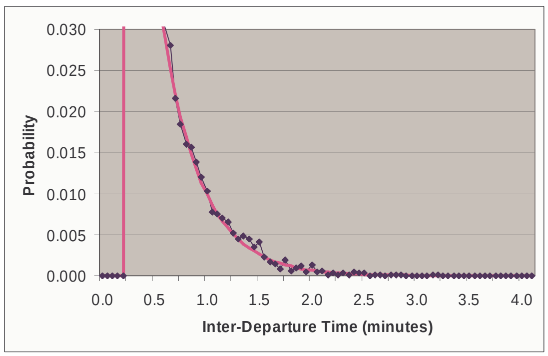

Figure 11 Exponential Inter-Departure Times when $n = n^* = 9$

$n \geq n^*$, then $E[T_D|\tau = 0] = 0.443$. If we set n using $n = (T_R + E[T_D|\tau = 0])/E[T_D|\tau = 0]$, we would obtain $n = \lceil 5.52 \rceil = 6$. We simulated 6,000 car rides, and ran each simulation 5 times, and the TH for different levels of n is shown in Figure 10. We calculated a 90% confidence interval for the average throughput values (shown by the dashed lines), based on the standard deviations of the throughput values. As the graph clearly shows, 6 cars has less throughput than 9. Maximum throughput is achieved by using $n \geq n^*$ from Theorem 6, and anything else is clearly suboptimal.

The need to simulate arises from the fact that we cannot characterize the arrival rate. Figure 11 shows the results from simulating 6,000 car rides, with $T_S = 0.25$, $E[T_C] = 0.25$, and $T_R = 2$. We used $n = n^* = 9$. As predicted by Lemma 5, the inter-departure times clearly approximate a truncated exponential distribution. In Figure 12, we see that inter-departure times may be very non-exponential. In this case, $T_S = 0.5$, $E[T_C] = 0.25$, and $T_R = 6$, which yields $n^* = 13$. In the

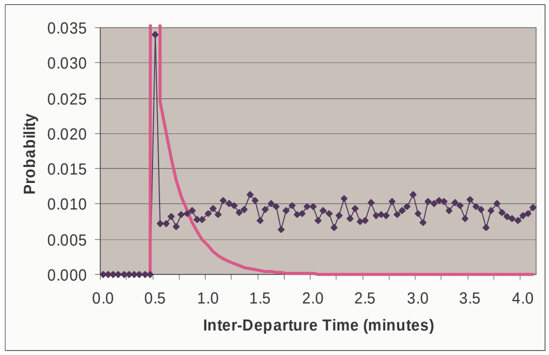

Figure 12 Non-Exponential Inter-Departure Times when $n < n^*$

simulation, $n = 2$, and 6,000 car rides were simulated. Clearly, the times are far from exponential, almost uniformly distributed. Thus, always assuming an exponential distribution would not appear to be a good approximation, further strengthening the argument for simulation.

Figure 13 shows car queue time from simulations of $6,000$ rides. For each, we used $T_S = 0.5$ min, and $T_R = 2.0$ min. In the top line we assumed $E[T_C] = 0.25$ min. The middle line has $E[T_C] = 0.5$ min, and the bottom line has $E[T_C] = 1$ min. For these parameters, $n^* = 5$. As n increases, the average queue time for re-loading increases. For each line, a heavy black line shows the queueing time predicted by equation (17). The equation is only valid for $n \geq n^* = 5$, but it clearly matches the simulation results perfectly.

For these values, if $E[T_C] = T_S = 0.5$, and we use $n = n^* = 5$, the queue time is 0.74 min, perhaps longer than riders might prefer for a 2 minute ride, but probably not unacceptable. If $E[T_C] = 0.25$ min, maximum TH is attained with an average queue time of 0.14 min. But if $E[T_C] = 1.0$ min, in order to achieve maximum throughput, the queue time balloons to 2.4 minutes, which is probably much longer than customers would be willing to tolerate for a 2 minute ride. Simulation is invaluable in this situation to consider the tradeoff between throughput and car queue time, and to see the potential benefits of reducing changeover time, T_C.

In Figure 14, the results of several sets of experiments are summarized. It shows the throughput per minute for different values of T_C. For all experiments, we assumed $T_S = 2$ and $T_R = 4$, which yields $n^* = (T_R + T_S)/T_S = 3$. For a given number of cars, n, the ratio of $E[T_C]/T_S$ was increased

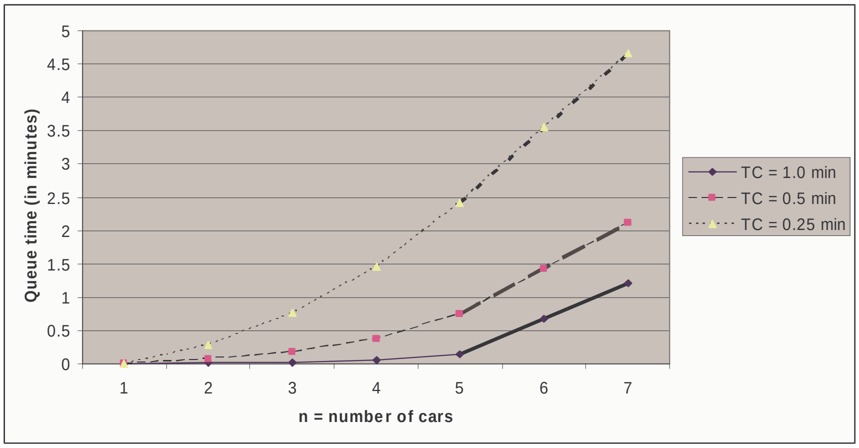

Figure 13 **Queue time vs.** n **for** $1/\lambda = 0.25, 0.5,$ **and** 1.0**, with** $T_S = 0.5$**,** $T_R = 2.0$

from 0.1 to 2.0, by increasing $E[T_C]$ from 0.2 to 4. The lines represent the throughput of different n. Not surprisingly, more cars means greater throughput, and $n > n^*$ leads to no additional increase. Also not surprisingly, faster changeover means greater throughput.

What is perhaps surprising is that the shape of the curve changes as n increases. For $n = 1$ or $n = 2$, the TH curve is convex. For $n \geq n^* = 3$, the lines are identical, and are convex for large values of T_C/T_S, and concave for small values. For larger values of n, a small reduction in T_C brings less TH increase than the same reduction would for fewer cars.

The most important lesson, perhaps, from the graph is that reducing T_C to below T_S can have significant benefits. Even for $n \geq n^*$, when $E[T_C] = T_S$, the throughput is only 60% of the throughput that could be achieved if $E[T_C]$ were reduced even to just $0.4T_S$.

Simulation is the only way to estimate the throughput for a $2 \leq n < n^*$, or to determine the benefits of reductions in changeover times. Using simulation, any of the many configuration possibilities for rides presented in this paper can be studied, to find the throughput to be gained by increasing the number of cars, separating the loading and unloading stations, adding parallel loading stations, or adding queueing areas after loading to prevent the station from becoming blocked by T_S.

8. Summary and Future Research

In this paper, we have derived formulas to estimate the throughput of many different amusement park rides. These formulas should be of use to any amusement park or ride engineering firm wanting

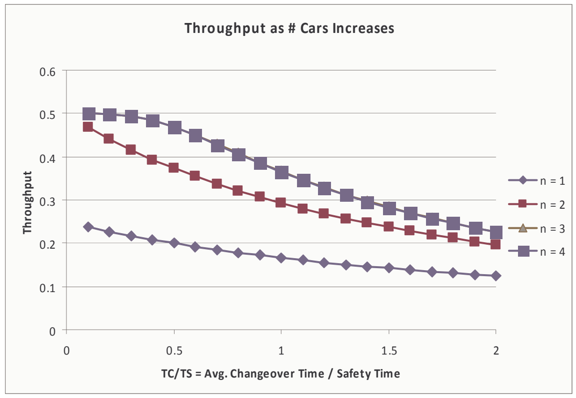

Figure 14 **Throughput for** $n = 1 - 4$**, when** $T_S = 2$**,** $T_R = 4$**,** $E[T_C] = 0.2 \ldots 4.0$

to estimate the throughput of a proposed ride, or to increase the throughput of an existing ride. The formulas can easily be incorporated into a spreadsheet, so that the throughput impact of T_S, T_L, or T_U reductions can easily be seen, as well as the impact of creating separate loading and unloading zones, or of adding parallel unloading/loading zones. If the firm can estimate the cost of various changes, it can use that information together with the formulas to easily determine the most cost-effective ways to increase rider throughput.

We have studied deterministic models of a number of structures of rides, but have only been able to study stochastic versions of a small number of those rides. Future research is needed to study these rides for $1 < n < n^*$ in a stochastic environment, and to estimate the impact of adding a queue before or after loading, observing the impact of separate unloading/loading zones, and/or parallel zones.

The biggest obstacle to further theoretical results is the fact that the memoryless property does not always hold for roller coasters, because of the safety interval. Theoretical results are needed to determine approximations for closed queueing networks with minimum intervals between jobs.

In the real world, differences in weight of passengers, the distribution of weight in the car etc., can lead to variations in ride time, and it seems likely that any actual observed differences in ride time would have a normal distribution. It would seem that this would be especially true on water-borne rides. For example, at Disneyland, the log flume ride (Splash Mountain) and the whitewater rafting ride (Grizzly's River Run) have queues in the middle of the ride, apparently to ensure that the necessary safety intervals are followed at key points in the ride, because of the greater variability in water-borne rides. Methods should be developed to study these types of rides, as well.

Acknowledgments

The author would like to thank his wife Jennie and children Grace and Noah for their patience when several family vacations at Disneyland degenerated into data-gathering expeditions.

9. Appendix

To show how all of the rides and attractions in an amusement park may be fit into the classification scheme we have used, we present a list of many of the rides and attractions in Disneyland and Disney's California Adventure.

Carousels	Roller Coasters
Astro Orbiter	Autopia
Dumbo the Flying Elephant	Buzz Lightyear Astro-Blasters*
The Enchanted Tiki Room	California Screamin'
Flik's Flyers	Casey Jr. Circus Train
Francis' Ladybug Boogie	Grizzly River Run
It's Tough to Be a Bug	Heimlich's Choo-Choo
Jumpin' Jellyfish	It's a Small World
King Arthur Carousel	Jungle Cruise
King Triton's Carousel	The Matterhorn
Mad Tea Party	Monsters, Inc.
Maliboomer	Mr. Toad's Wild Ride
Orange Stinger	Mulholland Madness
Soarin' Over California	Peter Pan's Flight
Tuck and Roll's Drive 'Em Buggies	Pinocchio's Daring Journey
Zephyr	Pirates of the Caribbean
	Space Mountain
	Splash Mountain
	Snow White's Scary Adventures
	Storybook Land Canal Boats
	Thunder Mountain

* = Tightly Coupled Ride

References

Ahmadi, R. 1997. Managing Capacity and Flow at Theme Parks. *Management Science*, **45** 1, 1-13.

ASTM. 2006. F 2291-06a Standard Practice for Design of Amusement Rides and Devices, in *Annual Book of ASTM Standards*, vol. 15.07. ASTM International: West Conshohocken, PA.

Burke, P. 1956. The Output of a Queuing System. *Operations Research*, **4**, 699-709.

Buzen, J. 1973. Computational Algorithms for Closed Queueing Networks with Exponential Servers. *Communications of the ACM*. **16** 9 527-531.

Disney Corp. 2006. *Annual Report.* Disney Corporation, Anaheim, CA.

Disney, R., and D. Konig. 1985. Queueing Networks: A Survey of Their Random Processes. *SIAM Review* **27** 3, 335-403.

Formica, S. and M. Olsen 1998. Trends in the amusement park industry. *International Journal of Contemporary Hospitality Management*. **10** 7, 297-308.

Gordon, W. and G. Newell. 1967. Closed Queueing Systems with Exponential Servers. *Operations Research* **15** 2 254-265.

Haxholdt, C., E. Larsen, A. van Ackere. 2003. Mode Locking and Chaos in a Deterministic Queueing Model with Feedback. *Management Science*, **49** 6, 816-830.

International Association of Amusement Parks and Attractions, 2007. Industry Information. Retrieved July 20, 2007 ⟨http://www.iaapa.org⟩.

Jackson, J. 1957. Networks of Waiting Lines. *Operations Research*, **5** 4, 518-521.

Jackson, J. 1963. Jobshop-like Queueing Systems. *Management Science*, **10** 1, 131-142.

Jefferson, D. 2007. Disney's New Magic. *Newsweek*, **149** 26, 14.

Kirsner, S. 2002. Rebuilding Tomorrowland. *Wired*, **10** 12.

Kleinrock, L. 1976. *Queueing Systems Volume II: Computer Applications.* John Wiley & Sons, New York.

Konigsberg, E. 1958. Cyclic Queues. *Operational Research Quarterly*. **9** 22-35.

Konigsberg, E. 1982. Twenty Five Years of Cyclic Queues and Closed Queue Networks: A Review. *Journal of the Operational Research Quarterly*. **33** 605-619.

Marden, D. 2007. email communication. July 15.

Muller, J. 2003. How to Build a Mountain. *Forbes*. **172**(9) 86-88.

O'Brien, T. 1996. Six Flags' Superman The Escape: 100 MPH, 415-Ft. Tower, 6.5 Seconds of Weightlessness. *Amusement News*, **108**(3), p. 18.

Rajaram, K. and R. Ahmadi. 2003. Flow Management to Optimize Retail Profits at Theme Parks. *Operations Research*, **51** 2, 175-184.

rcdb.com. 2007. California Screamin'. Retrieved July 30, 2007. ⟨ http://www.rcdb.com/id731.htm ⟩

Reckard, E. 2001. The Most Jam-Packed Theme Park on Earth? *Los Angeles Times*, Jan. 14, C1,C6.

ultimaterollercoaster.com. 2007. California Screamin'. Retrieved July 30, 2007. ⟨ http://www.ultimaterollercoaster.com/coasters/yellowpages/coasters/cascream_dca.shtml⟩

U.S. Department of Commerce. 2004. *Amusement, Gambling, and Recreation Industries: 2002 Economic Census, Arts, Entertainment and Recreation Industry Series.* U.S. Department of Commerce, Washington, D.C.