been referred to as the "Theory of Bottlenecks." Since it deals with problems of congestion and delay I don't think I need to go into any long explanation justifying its importance. Let us consider one of the simplest problems to which the technique has been applied - the tollgate problem. At this point I had a fancy slide showing a tollgate, but somewhere it got lost. So, picture, if you will, a single tollbooth on a turnpike - I think you're all familiar with that - and at this tollbooth we find a line of cars, and there are also cars approaching the tollbooth. The basic components of a queuing situation; customers and a customer input process, service points and a service mechanism, and some sort of queue discipline. Each of these is illustrated in the tollgate problem. The customers are motorists who arrive in some random way at the tollgate which is the service point. And they pay a toll before they may pass this point.

Their queue discipline is to line up on a first-come first-served basis. In this particular problem we assume only one gate. Service times at the tollgate will depend on whether the motorist has his money out, how much change is required, whether or not he stalls his car, and other factors. This summer I had a flat tire right at the tollgate. They wouldn't let me stay.

Both the intervals of time between new arrivals and service time must be treated as random variables. And the frequency distribution must be determined for each of these. In analyzing this sort of queuing problem one is interested in such information as the waiting times of the customers, the number of customers in line, and what fraction of time the servor is idle.

The next slide shows a particular set of conditions, and the probabilities with which there will be various numbers of customers at the tollgate at any given time. In this example an average of six cars per minute arrive at the booth. And the toll collector can handle an average of ten cars per minute. For arrival and service