

Notas:	



Hay un gran número de sistema que requieren modelar diferentes colas. Una transaccion puede requerir servicios de uno o mas recursos o servicios

En el caso general se necesitan modos de analizar sistemas de colas a fin de evaluar el comportamiento de sistemas informáticos compuestos por múltiples recursos (CPU, memoria, unidades I/O, etc.). En la nomenclatura de los modelos de colas, esto corresponde a decir que dentro de una transacción se ha de esperar a múltiples colas, o incluso realizar dentro de la misma transacción diferentes requerimientos a un mismo recurso. Por ejemplo, la transacción puede requerir ejecutar deferentes secciones de código en la CPU, realizar diferentes lecturas y escrituras sobre el disco, y transmitir y recibir diferentes datos por la unidad de I/O.

La resolución de un circuito de colas no se puede basar en maquinas de estados, ya que el estado de un sistema de colas es el producto cartesiano de los estados de cada una de las colas, y ello introduce una complejidad inabordable. Hay dos soluciones:

-Utilizar una herramienta de simulación: Con ella se pueden realizar resultados muy detallados y precisos, pero presenta el problema de requerir una gran capacidad de computo, y además se obtienen resultados sin saber a que características del modelo se deben.

-Utilizar métodos analíticos: Para obtener soluciones hay que realizar suposiciones simplificadoras que conduce a cálculos aproximados, pero a cambio de ello, los requerimientos de computación son bajos, y se obtienen relaciones entre los resultados y los parámetros de los modelos.



Un sistema de colas puede clasificarse como un conjunto de colas que modelan los recursos del sistema y un conjunto de arcos que describen las secuencias de requerimientos de una a otras. Las reglas que determinan el flujo de los eventos son probabilísticas y definen la probabilidad de que en una transacción, después de acceder a un recurso, aceda a uno u otro recurso.

Las magnitudes que se utilizan son:

- $X_{\rm i}$: flujo de requerimientos (throughput medidos como requerimientos/s) que se realizan en la cola i.

- S_i: Tiempo de servicio en el recurso de la cola i

- R_i: Tiempo de estancia en la cola i

- $V_i\!\!:$ Tasa de visitas. Número de veces que dentro de una misma activación de la transacción se accede al recurso i

- qi: Numero medio de clientes en espera en la cola i.

- p_i: probabilidad de que un evento que llega un punto de bifurcación siga una

rama.

- D_i=V_iS_i: Demanda del servicio i por la transacción.



The simplest way to classify a queueing network is either open or closed. An **open queueing network** has external arrivals and departures, as shown in Figure (1). The jobs enter the system at "In" and exit at "Out." The number of jobs in the system varies with time. In analyzing an open system, we assume that the throughput is known (to be equal to the arrival rate), and the goal is to characterize the distribution of number of jobs in the system.

A **closed queueing network** has no external arrivals or departures. As shown in Figure 2, the jobs in the system keep circulating from one queue to the next. The total number of jobs in the system is constant. It is possible to view a closed system as a system where the Out is connected back to the In. The jobs exiting the system immediately reenter the system. The flow of jobs in the Out-to-In link defines the throughput of the closed system. In analyzing a closed system, we assume that the number of jobs is given, and we attempt to determine the throughput (or the job completion rate).

It is also possible to have **mixed queueing networks** that are open for some workloads and closed for others. Figure 3 shows an example of a mixed system with two *classes* of jobs. The system is closed for interactive jobs and is open for batch jobs. The term *class* refers to types of jobs. All jobs of a single class have the same service demands and transition probabilities. Within each class, the jobs are indistinguishable.



Sistemas con múltiples workload que se componen de varios tipos de clientes que fluyen con diferentes throughput, requieren de los mismos recursos servicios con tiempos diferentes, y que visitan las colas con diferentes tasas.



Poisson Processes: If the interarrival times are IID and exponentially distributed, the number of arrivals n over a given interval (t, t + x) has a Poisson distribution, and therefore, the arrival process is referred to as a **Poisson process** or a **Poisson stream.** The Poisson streams are popular in queueing

theory because the arrivals are then memoryless as the interarrival time is exponentially distributed. In addition, Poisson streams have the following properties:

(a) Merging of k Poisson streams with mean rate λ_i results in a Poisson stream with mean rate $\lambda = \sum \lambda_i$

(**b**) If a Poisson Stream is split into k substreams such that the probability of a job going to the *i*th substream is p_{i} , each substream is also Poisson with a mean rate of $p_i \lambda$.

(c) If the arrivals to a single server with exponential service time are Poisson with mean rate λ , the departures are also Poisson with the same rate λ , provided the arrival rate λ is less than the service rate μ .

(d) If the arrivals to a service facility with *m* service centers are Poisson with a mean rate λ , the departures also constitute a Poisson stream with the same rate λ , provided the arrival rate λ is less than the total service rate Here, the servers are assumed to have exponentially distributed service times.



Considerese un circuito simple compuesto por tres colas M/M/1 en seri con tiempos de servicio S1=1.0 s, S2=2.0 s y S3=3.0 s. y que recibe requerimientos de una transacción que se ejecuta con una frecuencia de $\lambda=0.1$ invocación/s. Debido a las propiedades de las distribuciones de Poisson, se puede tratar como tres colas independientes, todas ellas (en este caso) solicitadas con el mismo throughput λ .

En la tabla se calculan la utilización, el tiempo de estancia, y la longitud de la espera en cada cola.

En la última línea se comprueba que la fórmula de Little se puede aplicar globalmente al sistema como si fuese una cola.

Si las colas fuesen de múltiples servidores, o fuesen múltiples colas, habría que utilizar las correspondientes ecuaciones para calcular el tiempo de estancia R y el número medio de elementos en las colas.



Jackson (1963) showed that the above method of computing joint probability is valid for any arbitrary open network of *m*-server queues with exponentially distributed service times (see the Figure). In particular, if all queues are single-server queues, the queue length distribution is exponencial. However, it is not correct to assume that each queue becomes an independent M/M/1 queue with a Poisson arrival process. In general, the internal flow in such networks is not Poisson. Particularly, if there is any feedback in the network, so that jobs can return to previously visited service centers, the internal flows are not Poisson. It is surprising that even though the flows are not Poisson, the queues are separable and can be analyzed as if they were Jackson's results were later extended to closed networks by Gordon and Newell (1967). They showed that any arbitrary closed networks of *m*-server queues with exponentially distributed service times also have a product form solution. showed that product form solutions exist for an even broader class of networks. This class consists of networks satisfying the following criteria:

1. Job Flow Balance: For each class, the number of arrivals to a device must equal the number of departures from the device.

2. *One-Step Behavior*: A state change can result only from single jobs entering the system, moving between pairs of devices in the system, or exiting from the system. This assumption asserts that

simultaneous job moves will not be observed.

3. *Device Homogeneity*: A device's service rate for a particular class does not depend on the state of the system in any way except for the total device queue length and the designated class's queue length. This assumption implies the following:

(a) *Single-Resource Possession*: A job may not be present (waiting for service or receiving service) at two or more devices at the same time.

(b) No Blocking: A device renders service whenever jobs are present; its ability to render service is not controlled by any other device.

(c) *Independent Job Behavior*: Interaction among jobs is limited to queueing for physical devices; for example, there should not be any synchronization requirements.

(d) Local Information: A device's service rate depends only on local queue length and not on the state of the rest of the system.

(e) *Fair Service*: If service rates differ by class, the service rate for a class depends only on the queue length of that class at the device and not on the queue lengths of other classes. This means that the servers do not discriminate against jobs in a class depending on the queue lengths of other classes.

(f) *Routing Homogeneity*: The job routing should be state independent. In the last condition, the term *routing* is used to denote a job's path in the network. The routing homogeneity condition implies that the probability of a job going from one device to another device does not depend upon the number of jobs at various devices.



En la red simple de la figura, los clientes llegan a una frecuencia λ de Poisson. Dado que a la salida hay una bifurcación que hace retornar a la entrada de la cola, se verifica $\lambda_1 = \lambda + p \lambda_1$. Despejando, resulta la tasa de visita V =1/(1-p) y la demanda de servicio D=VS.

La utilización del recurso es U= $\lambda_1 S$.

Aplicando la ecuación de una cola M/M/1, $R_v = S/(1 - \lambda_1 S)$

El tiempo de respuesta global del sistema es R=VR_v=VS/(1- V λ S)=D/(1- λ D)



Open queueing network models are used to represent transaction processing systems such as airline reservation systems or banking systems. In these systems, the transaction arrival rate is not dependent on the load on the computer system. The transaction arrivals are modeled as a Poisson process with a mean arrival rate ». In this section, we present equations for an exact analysis of such systems. We assume that all devices in the computer system can be modeled as either fixed-capacity service centers (single server with exponentially distributed service time) or delay centers (infinite servers with exponentially distributed service time).



The figure shown represents a queueing model of a file server consisting of a CPU and two disks, A and B. Measurements on a distributed system with six clients systems making requests to the file server produced.









Notas:





Notas:	





Mean-Value Analysis (MVA) allows solving closed queueing networks in a manner similar to that used for open queueing networks. As the name implies, it gives the mean performance. The variance computation is not possible using this technique. Mean-value analysis applies to networks with a variety of service disciplines and service time distributions. However, we will initially limit our discussion here to fixed-capacity service centers. Delay centers are considered later. Load-dependent service centers are considered also later. Given a closed queueing network with *N* jobs, Reiser and Lavenberg (1980) showed that the response time of the ith device is given by

$$R_i(N) = S_i[1 + Q_i(N-1)]$$

Here, $Q_i(N-1)$ is the mean queue length at the ith device with N-1 jobs in the network. Notice the similarity between this equation and la equación habitual en colas simples $R_i(N) = S_i[1 + Q_i]$ for open queueing networks. As before, one way to explain this equation is to consider what happens when we add a tagged Nth job to the network with N-1 jobs. On arrival at the *i*th service center, the tagged job sees $Q_i(N-1)$ jobs ahead (including the one in service) and expects to wait $Q_i(N-1)S_i$ seconds before receiving service. Including the service time for itself, the job should expect a total response time of $S_i[1 + Q_i(N-1)]$. This is not an operational law. It assumes that the service is memoryless, an assumption that is not operationally testable.

Given the performance for N-1 users, is sufficient to allow us to compute the performance for N users. Since the performance with no users (N = 0) can be easily computed, performance for any number of users can be computed iteratively, as we show next.



Notas:			



Notas:



Mean-value analysis is a recursive algorithm. Computation of performance with N jobs in the network requires knowledge of performance with N-1 jobs. Since performance with N=0 is trivially known, we always start the analysis with N=0 and compute the performance for N=1,2,... successively. For small values of N, this procedure is not computationally too expensive. However, for large values of N, particularly if the performance for smaller values of N is not required, it would be preferable to avoid this recursion.

Several approximate analysis techniques have been developed with this goal. Here we describe one such approximation technique known as **Schweitzer's approximation**. It avoids the recursion in MVA by appropriately estimating the queue lengths with *N* jobs and computing the response times and throughputs. The values so computed can be used to recompute the queue lengths, and if the previous estimate was good, the new computed value would be close to the estimate.

The approximation, due to Schweitzer (1979), is based on the assumption that as the number of jobs in a network increases, the queue length at each device increases proportionately. For example, doubling the

number of jobs in the network will result in doubling the number of jobs at each device. Analytically:

In particular, this implies

or

The MVA equations can therefore be written as follows:



Notas:











The basic MVA algorithm is quite powerful and elegant. It is applicable across a wide set of performance models. It has been the focus of much research and several extensions have been developed. These include:

- Multi-class networks
- Networks with load dependent servers
- Networks with open and closed classes of customers

Several approximation techniques have also be adapted to MVA to model systems having non-product form features, including first-come-first-serve and priority multi-class networks. The extension of MVA to product form, load-independent, multi-class networks is found in references.

However, even with its widespread applicability, there are limitations and shortcomings surrounding MVA. For example:

•MVA does not provide the steady state probabilities of individual system states. Thus, if it were important to know the probability that there were, say, less than five customers at a device in order to meet some QoS criteria, MVA would not be helpful. It would be necessary to revert to solving the global balance equations. As its name implies, MVA only provides the mean values of various performance metrics, not the associated distributions.

•MVA does not provide transient analysis information. For instance, if it were necessary to know how long it would take the system to recover from a temporary overload in one sector of the system and return to "steady state" behavior, MVA would be inadequate.

•MVA does not model state dependent behavior. For example, consider the modeling of a simple routing protocol, with two paths between a particular source and destination, where the customers (i.e., message packets) select the path least busy (i.e., dependent on the particular system state). Although a Markovian model can be easily constructed and solved for this situation, MVA is of little use.

•MVA solves product form networks. As a result, MVA is not directly applicable to non-product form situations. These restrictions exclude certain device service distributions (e.g., Gaussian, uniform, constant), certain device queuing disciplines (e.g., priority, multi-class FCFS), and certain device loading strategies (e.g., shortest queue routing, deterministic routing).