Gestion de la contention

Probabilité pour l'Informatique et Évaluation de Performances

Jean-Marc Vincent

MESCAL-INRIA Project
Laboratoire d'Informatique de Grenoble
email Jean-Marc.Vincent@imag.fr

Ensimag 2ième Filière ISI

Outline

Outline

Queues

Queues are among simplest dynamic systems, but are still the source of many open problems.
Tasks do not have any constraints, sizes and arrival times are often independent.

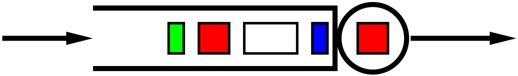

Queues

Queues are among simplest dynamic systems, but are still the source of many open problems.
Tasks do not have any constraints, sizes and arrival times are often independent.

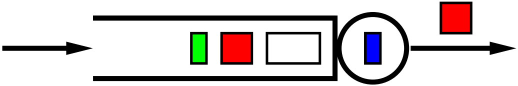

Queues

Queues are among simplest dynamic systems, but are still the source of many open problems.
Tasks do not have any constraints, sizes and arrival times are often independent.

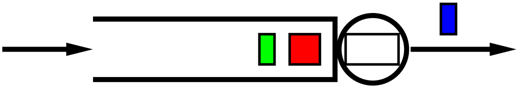

Queues

Queues are among simplest dynamic systems, but are still the source of many open problems.
Tasks do not have any constraints, sizes and arrival times are often independent.

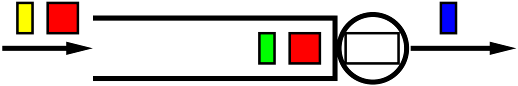

Queues

Queues are among simplest dynamic systems, but are still the source of many open problems.
Tasks do not have any constraints, sizes and arrival times are often independent.

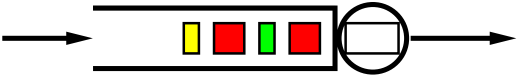

Queues

Queues are among simplest dynamic systems, but are still the source of many open problems.
Tasks do not have any constraints, sizes and arrival times are often independent.

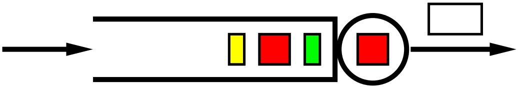

Quizz from Harchol-Balter(2013)

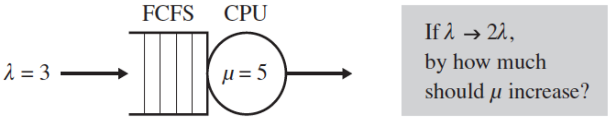

Figure 1.2. A system with a single CPU that serves jobs in FCFS order.

Answers

(a) Double the CPU speed
(b) More than Double the CPU speed
(c) Less than double the CPU speed

Performance Modeling and Design of Computer Systems: Queueing Theory in Action Mor Harchol-Balter, Cambridge University Press, 2013

Quizz (2)

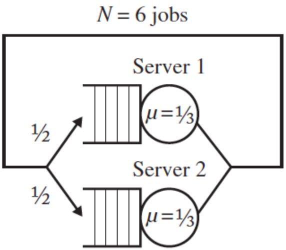

Figure 1.3. A closed batch system.

Question You replace server 1 with a server that is twice as fast (the new server services jobs at an average rate of 2 jobs every 3 seconds). Does this "improvement" affect the average response time in the system? Does it affect the throughput?

Quizz (3)

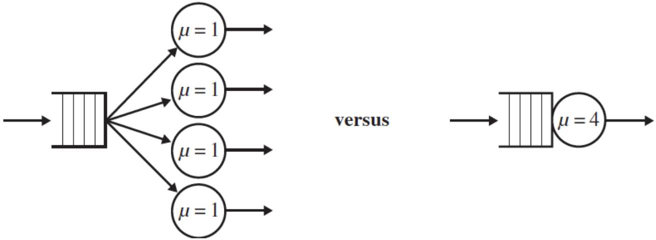

Figure 1.5. Which is better for minimizing mean response time: many slow servers or one fast server?

Quizz (4)

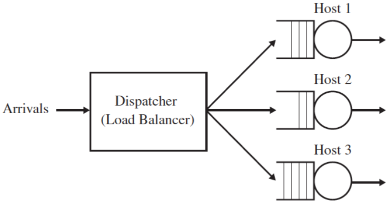

Figure 1.6. A distributed server system with a central dispatcher.

Task assignment policies

- **Random**
- **Round-Robin**
- **Shortest-Queue**
- **Size-Interval-Task-Assignment (SITA)**
- **Least-Work-Left (LWL)**
- **Central-Queue**

Question: Which of these task assignment policies yields the lowest mean response time?

Quizz (4)

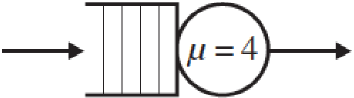

Scheduling strategy

- **Random**
- **Non-Preemptive Last-Come-First-Served (LCFS)**
- **First-Come-First-Served (FCFS)**

Question: Which of these non-preemptive service orders will result in the lowest mean response time?

Kendall's notation

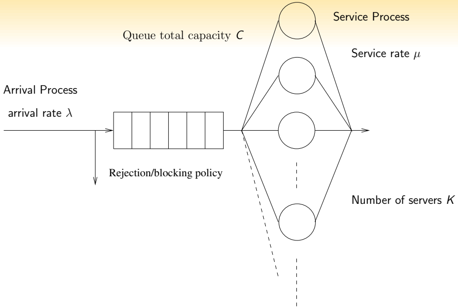

Notation : $A/S/K/C/Disc$

- A : arrival process
- B : service process
- K : number of servers
- C : total queue capacity (including currently served customers)
- $Disc$: Service discipline (FIFO, LIFO, PS, Quantum, Priorities,...)

Process types

Arrival or service process

- M : memoryless (exponential distribution);
- D : deterministic (constant);
- U : uniform distribution;
- Erl_k : Erlang distribution (sum of exponentially distributed RV, $\gamma(k, \lambda)$);
- H_k : Hyper-exponential distribution
- GI : general independent (given arbitrary distribution) independence between inter arrivals or services
- G : general

Usually service times and inter arrival processes are independent

Service discipline

- *FIFO* : First In First Out
- *LIFO* : Last In First Out
  - pre-emptive or non pre-emptive
  - resume, restart with same service, restart with new service
- Quantum : round-robin policy
  - PS : asymptotic
- Priority
  - pre-emptive or non pre-emptive
  - resume, restart with same service, restart with new service
- Adaptive discipline

David George Kendall

Born: 15 January 1918 in Ripon, Yorkshire, England
Died: 23 October 2007 in Cambridge, England
One highlight was his pioneering work of 1949 on stochastic (or random) processes for population growth. Another was his classic 1951 paper on queuing theory, which was motivated by the scheduling problems of aircraft and runways during the Berlin air lift of 1948-49. A third was a series of penetrating studies, with G.E.H. Reuter, of Markov processes (roughly, random processes without memory).
The Independent/MacTutor History
Another biography by G. Grimmet

System variables

User variables

- Input rate λ or inter-arrival δ
- Service time σ or S (service rate μ)
- Waiting time W
- Response time R (in some books W)
- Rejection probability

Resource variables

- Resource utilisation (offered load) ρ
- Queue occupation N
- System availability

Outline

One Server Queue load

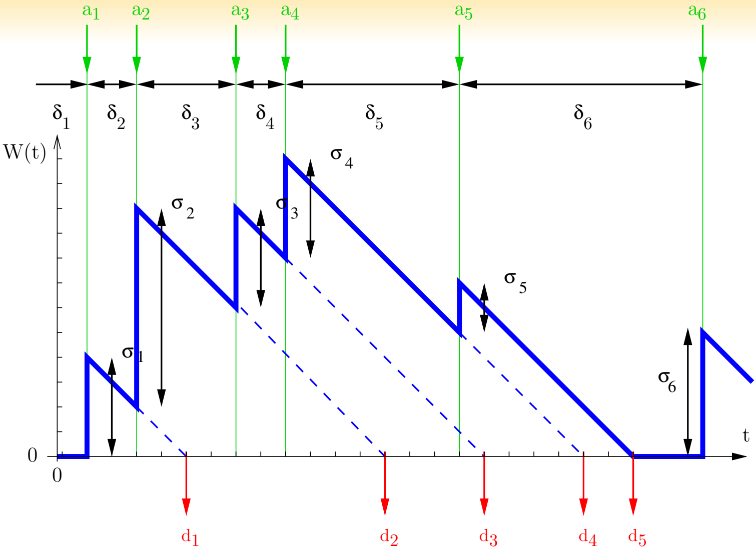

Lindley's formula

W_n is the waiting time of the n-th task. It is a dynamical system of the form $W_n = \varphi(W_{n-1}, X_n)$ with $X_n = \sigma_{n-1} - \delta_n$ and φ defined by the

Lindley's equation:

$$W_n = \max\{W_{n-1} + X_n, 0\} .$$

- FIFO scheduling
- Non-linear evolution equation

Stability of the $G/G/1$ queue

$$
\begin{aligned}
W_n &= \max\left\{W_{n-1} + X_n, 0\right\}, \\
&\quad \textit{expanding the relation} \\
&= \max\left\{\max(W_{n-2} + X_{n-1}, 0) + X_n, 0\right\}, \\
&\quad \textit{property } \max\{\max\{a,b\},c\} = \max\{a,b,c\} \\
&= \max\left\{W_{n-2} + X_{n-1} + X_n, X_n, 0\right\}, \\
&= \max\left\{W_{n-3} + X_{n-2} + X_{n-1} + X_n, X_{n-1} + X_n, X_n, 0\right\}, \\
&= \max\left\{W_0 + X_1 + \cdots + X_{n-1} + X_n, \cdots, X_{n-1} + X_n, X_n, 0\right\}, \\
&= \max\left\{X_1 + \cdots + X_{n-1} + X_n, \cdots, X_{n-1} + X_n, X_n, 0\right\}, \\
&\quad \textit{renumbering of variables iid} \\
&\sim \max\left\{X_n + \cdots + X_2 + X_1, \cdots, X_2 + X_1, X_1, 0\right\}, \\
&\stackrel{def}{=} M_n.
\end{aligned}
$$

$$
W_n =_{st} M_n = \max\left\{X_n + \cdots + X_1, M_{n-1}\right\} .
$$

Stability of the $G/G/1$ queue (2)

$$W_n =_{st} M_n = \max\{M_{n-1}, X_1 + \cdots + X_n\}.$$

M_n is a non-decreasing sequence
Either $M_n \longrightarrow M_\infty$ or $M_n \longrightarrow +\infty$

Stability

- $\mathbb{E}X = \mathbb{E}(\sigma - \delta) < 0$ The system is **Stable**

  $$M_\infty =_{st} \max(M_\infty + X, 0).$$

  Functional equation on the distribution

  $$\mathbb{P}(M_\infty < x) \stackrel{def}{=} F(x) = \int F(x-u)dF_X(u).$$

  Condition : $\mathbb{E}\sigma < \mathbb{E}\delta$ or $\lambda < \mu$
- $\mathbb{E}X = \mathbb{E}(\sigma - \delta) > 0$ The system is **Unstable**

Depends only on service and inter-arrival expectation

Outline

Little's Formula

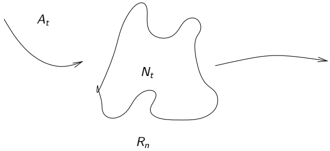

Assumptions

$$\lim_{t\to+\infty} \frac{A_t}{t} = \lambda, \quad \lim_{t\to+\infty} \frac{1}{t}\int_0^t N_s ds = \mathbb{E}N \text{ and } \lim_{n\to+\infty} \frac{1}{n}\sum_{i=1}^{n} R_i = \mathbb{E}R,$$

Little's Formula

$$\mathbb{E}N = \lambda \mathbb{E}R.$$

Little's Formula (proof)

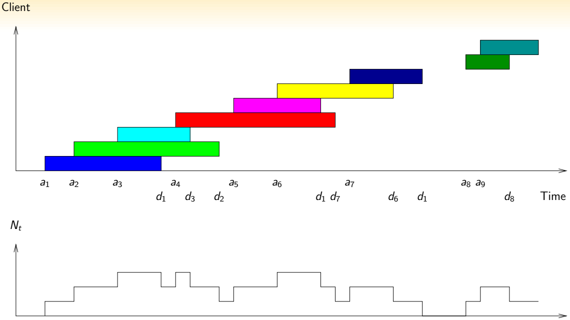

$$\frac{1}{T}\int_0^T N_s ds = \frac{A_T}{T}\frac{1}{A_T}\sum_{i=1}^{A_T} R_i.$$

$T \to \infty$ implies $\mathbb{E}N = \lambda \mathbb{E}R$.
The equality is valid if the system is empty at time T ($N_T = 0$), what is changed in the proof if $N_T \neq 0$?

Little's Formula (proof)

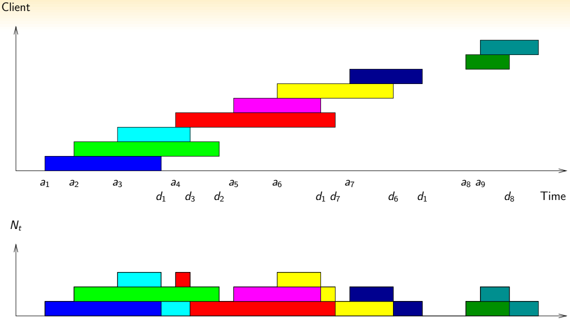

$$\frac{1}{T}\int_0^T N_s ds = \frac{A_T}{T}\frac{1}{A_T}\sum_{i=1}^{A_T} R_i.$$

$T \to \infty$ implies $\mathbb{E}N = \lambda \mathbb{E}R$.
The equality is valid if the system is empty at time T ($N_T = 0$), what is changed in the proof if $N_T \neq 0$?

Little's Formula (proof)

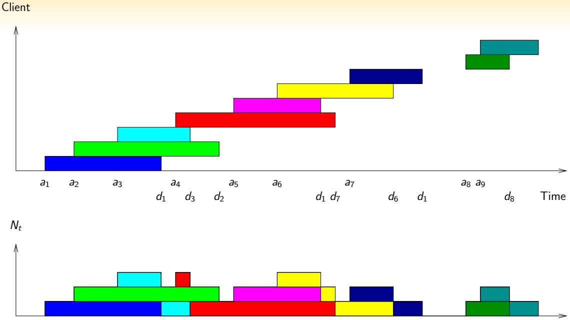

$$\frac{1}{T}\int_0^T N_s ds = \frac{A_T}{T}\frac{1}{A_T}\sum_{i=1}^{A_T} R_i.$$

$T \to \infty$ implies $\mathbb{E}N = \lambda\mathbb{E}R$.
The equality is valid if the system is empty at time T ($N_T = 0$), what is changed in the proof if $N_T \neq 0$?

Little's Formula (examples)

Server utilization

$$\rho = \lambda \mathbb{E} S$$

$M/M/1$ queue

$$\rho = \frac{\lambda}{\mu}, \mathbb{E} N = \frac{\rho}{1-\rho}$$

Response time

$$\mathbb{E} N = \lambda \mathbb{E} W$$

Valid for all service disciplines

Outline

Birth and Death models

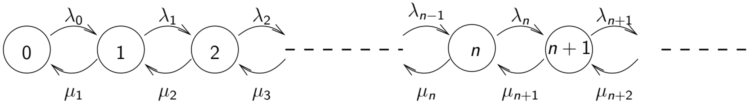

- Stationary distribution

$$\pi_n = \pi_0 . \frac{\lambda_0 \lambda_1 \cdots \lambda_{n-1}}{\mu_1 \mu_2 \cdots \mu_n}$$

- Stability condition

$$\sum_n \frac{\lambda_0 \lambda_1 \cdots \lambda_{n-1}}{\mu_1 \mu_2 \cdots \mu_n} < \infty$$

- Model of many systems
  queueing systems, reliability models, epidemic models, systems of particles, ...

The classic M/M/1 queue

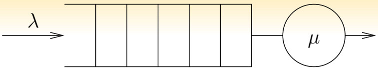

$\{X_t\}_{t \in \mathbb{R}}$, number of clients in the queue at time t is a Birth-and-Death process

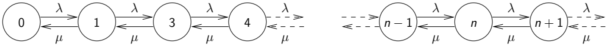

Formula

- **Load** of the queue $\rho = \frac{\lambda}{\mu}$ (utilization of the server)
- **Stability** of the queue $\rho < 1$
- **Steady-state** of the queue $\pi_n = (1-\rho)\rho^n$ (geometric distribution)
- **Mean number of clients** in the queue $\overline{N} = \frac{\rho}{1-\rho}$
- **Mean response time** of clients $\overline{W} = \frac{1}{\mu - \lambda}$
- **Overflow probability** $\mathbb{P}(X \geqslant n) = \rho^n$
- **Buzy period** $\overline{B} = \frac{1}{\mu - \lambda}$

Performances

Mean number of clients $\mu = 1$

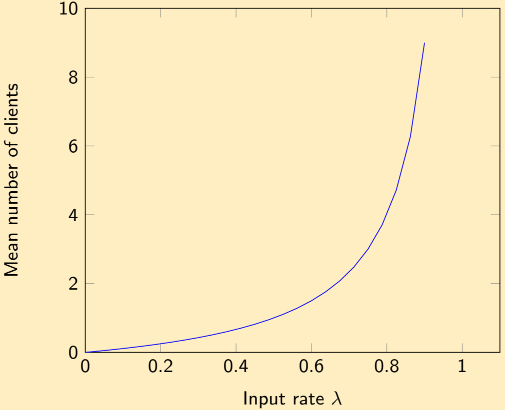

Performances

Average response time $\mu = 1$

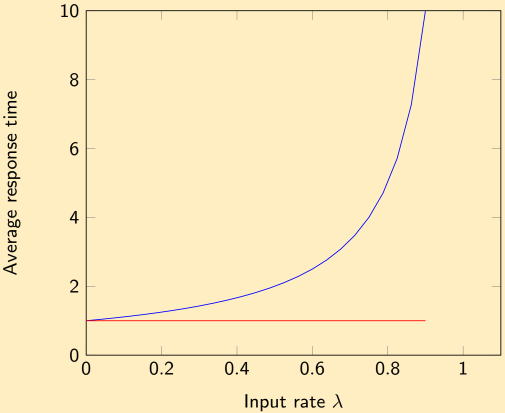

With a FIFO scheduling $W \sim \mathcal{E}(\frac{1}{\mu-\lambda})$

Outline

Synthesis

The basic building block : the queue

The Kendall's notation

$$A/S/K/C/Disc$$

Stability Condition and Average Identities

$$\rho = \frac{\lambda}{\mu} < 1.$$

Little's formula

$$\overline{N} = \lambda \overline{W}.$$

Formula : Markovian simple queue

- **Load** of the queue $\rho = \frac{\lambda}{\mu}$ (utilization of the server)
- **Stability** of the queue $\rho < 1$
- **Steady-state** of the queue $\pi_n = (1-\rho)\rho^n$ (geometric distribution)
- **Mean number of clients** in the queue $\overline{N} = \frac{\rho}{1-\rho}$
- **Mean response time** of clients $\overline{W} = \frac{1}{\mu-\lambda}$
- **Overflow probability** $\mathbb{P}(N \geqslant n) = \rho^n$
- **Buzy period** $\overline{B} = \frac{1}{\mu-\lambda}$